A Two-Speed Synchronous Traffic Protocol for Intelligent Intersections: From Single-Vehicle to Platoon Crossing

DANIEL MARKERT and ALEJANDRO MASRUR, Department of Computer Science, TU Chemnitz, Germany

With progress in cooperative and autonomous driving, there is an increasing interest in intelligent intersections to replace conventional traffic lights and, thereby, improve traffic efficiency. To avoid accidents in such safety-critical systems, a traffic protocol needs to be implemented. In this article, we are concerned with synchronous traffic protocols, i.e., those that synchronize the arrival time of vehicles at the intersection. In particular, such protocols are normally conceived for homogeneous vehicles of approximately the same size/length. However, these do not extend well to heterogeneous vehicles, i.e., they lead to unviable requirements on the road infrastructure. To overcome this limitation, based on the observation that large/overlength vehicles like buses and trams are less frequent than passenger vehicles, we propose an approach that treats them as exceptions (rather than the rule) leading to a much more efficient design. In contrast to approaches from the literature, we implement a two-speed policy—with a high speed for drive-through and a low speed for turn maneuvers—and analyze both single-vehicle as well as fairness-based platoon crossing. To conclude, we perform detailed comparisons illustrating the benefits by the proposed approach.

CCS Concepts: • **Networks** → *Cyber-physical networks*; • **Computer systems organization** → **Embedded and cyber-physical systems;**

Additional Key Words and Phrases: Intelligent intersections, traffic automation, traffic protocols, vehicle throughput

ACM Reference format:
Daniel Markert and Alejandro Masrur. 2023. A Two-Speed Synchronous Traffic Protocol for Intelligent Intersections: From Single-Vehicle to Platoon Crossing. *ACM Trans. Cyber-Phys. Syst.* 7, 2, Article 10 (April 2023), 21 pages.
https://doi.org/10.1145/3571289

1 INTRODUCTION

Nowadays, traffic lights around the globe are gradually including features such as broadcasting their time-to-green [20], prioritizing emergency vehicles [12] or public transportation [19], and so on, to improve traffic efficiency. With advances in cooperative and autonomous driving, these are expected to further evolve into intelligent intersections that dynamically coordinate vehicles in all possible directions, improving efficiency (i.e., throughput) while guaranteeing safety via a

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - MA 6480/3-1.
Authors' address: D. Markert and A. Masrur, Department of Computer Science, TU Chemnitz, Strasse der Nationen 62, Chemnitz, Saxony, Germany, 09111; emails: {daniel.markert, a.masrur}@cs.tu-chemnitz.de.

2378-962X/2023/04-ART10 $15.00
https://doi.org/10.1145/3571289

traffic protocol. In such settings, vehicles are the physical components (part of the real world) and the traffic protocol is a cyber component, which collects information from vehicles (over wireless communication such as IEEE 802.11p or cellular) and processes it, constituting a cyber-physical system.

In this article, we mainly focus on the traffic protocol itself, which is only a—but probably the most important—part of this cyber-physical system. Traffic protocols can be basically asynchronous or synchronous. Asynchronous traffic protocols schedule vehicles one by one as they arrive at the intersection, i.e., no synchronization of arrival times, and are rather suitable for low-traffic scenarios. On the contrary, synchronous traffic protocols enforce arrival patterns and aim to interleave vehicles crossing in all possible directions, being more effective under high-traffic situations.

Both synchronous and asynchronous traffic protocols can either be implemented in a centralized, i.e., with a **road side unit (RSU)** collecting and processing data, or a decentralized manner, i.e., with vehicles directly exchanging data among themselves. However, synchronous traffic protocols are by nature easier to implement in a centralized manner.

In this article, we are concerned with centralized synchronous traffic protocols, and, hence, we assume that vehicles exchange data with an RSU. To ensure safety, vehicles on the same lane need to be separated by sufficient distance, which allows interleaving vehicles on perpendicular lanes through the intersection. Clearly, inter-vehicle separations depend on vehicles' physical dimensions, in particular, on their lengths.

Available approaches such as **Ballroom Intersection Protocol (BRIP)** account for this dependency by assuming that the largest/longest possible vehicle in the intersection represents the default case [4], which works well for homogeneous traffic, e.g., with only passenger cars. However, under heterogeneous traffic including, e.g., buses, trams, and the like, this leads to stringent space requirements, greatly restricting applicability. That is, disproportionally wide lanes are required, which either implies changing or excluding most existing intersections.

Contributions. This article presents the following advancements with respect to our previous work [16, 18]:

- We propose a traffic protocol called **LTR (Left Turn, Through, Right Turn)**, where *overlength vehicles* are treated as an exception rather than the default case. This significantly relaxes the resulting space requirements on intersections.
- In addition, in contrast to the approaches from the literature, we introduce a two-speed policy, which allows vehicles to drive through the intersection at a higher speed. Similarly, this way, vehicles perform turn maneuvers at a more *passenger-friendly*, i.e., lower, speed.
- Finally, we propose a platoon crossing strategy, i.e., where a group of vehicles cross at once in one direction, to increase the throughput of LTR while retaining its benefits.

The proposed LTR segments the intersection center (similar to BRIP) into square *sectors* of constant size, in particular, fixed by the lane width. Contrary to BRIP, however, a vehicle does not necessarily have to fully fit into one of these sectors, but it can extend over multiple adjacent sectors, incurring what we call an *overlength penalty.* This temporarily increases the necessary inter-vehicle separation on the affected lanes, but it also allows for a considerably higher space efficiency compared to BRIP. Given that overlength vehicles are less frequent, our results show that the proposed single-vehicle crossing LTR reduces the required intersection size by 75% while outperforming BRIP in terms of throughput for overlength vehicles of 9.6 m onward. Further, the proposed platoon crossing outperforms BRIP in terms of throughput already for overlength vehicles of 7.5 m onward (e.g., 150% more throughput is possible when vehicles are as long as 15 m) with the same space efficiency as shown by detailed comparisons.

Structure. The remainder of this article is organized as follows: In Section 2, we discuss existing related publications, while Section 3 introduces the basic assumptions and definitions required for Section 4, where the proposed two-speed LTR protocol is introduced, including our platoon crossing strategy. In Section 5, we then present our analysis and comparison results with regard to space efficiency and throughput. Finally, Section 6 concludes the paper.

2 RELATED WORK

To the best of our knowledge, the first traffic protocol/approach for intelligent intersections was proposed with **AIM (Asynchronous Intersection Management)** in [8], [9], and [10], a **FCFS (First Come, First Served)** reservation strategy to schedule vehicles through an intersection (in an asynchronous fashion). To that goal AIM divides the intersection into a grid of $n \times n$ tiles, where n is termed *granularity* [10]. When approaching the intersection, a vehicle requests its intended trajectory, which translates into a sequence of tiles that need to be traversed (at specific points in time). This is then only granted by AIM, if the requested trajectory does not contain any tile, which has been already reserved for another vehicle (at the same point in time).

If a vehicle's trajectory is rejected, this has to stop and wait for its turn. AIM is quite effective under low traffic, but it rather leads to a stop-and-go behavior in high-traffic situations. For this reason, in this article, we focus on synchronous algorithms, which are more suitable for high traffic.

Extensions to the AIM were presented in [13] for multi-intersection settings and in [3] with autonomous driving in focus. However, these extensions remain asynchronous in nature and, hence, inherit the advantages and disadvantages discussed above.

The first synchronous traffic protocol to be proposed is **BRIP (Ballroom Intersection Protocol)** [4], which is designed to increase the overall throughput at the intersection by enforcing different four-way crossing patterns. To this end, several intersection layouts (i.e., different combinations of lanes and directions) are considered. Similar to AIM, the intersection is divided into square sectors with sides equal to $S_{BRIP} \geq L + W$, with L and W being vehicle length and width. Since BRIP assumes the same size for all vehicles, L and W inadvertently represent the dimensions of the longest possible vehicle allowed at the intersection. For example, this leads to BRIP requiring a two-lane intersection to be at least 52 m wide to accommodate a truck with a length of 10 m and a width of 3 m, i.e., it requires 13 m per lane. As a consequence, BRIP is not suitable fort most existing intersections, but rather requires infrastructural modifications.

Recently, BRIP was extended into two new protocols called CSIP [1] and DSIP [2]. **CSIP (Configurable Synchronous Intersection Protocol)** increases inter-vehicle distances to allow for better guarantees of collision avoidance when considering positioning errors by GPS, while **DSIP (Distributed Synchronous Intersection Protocol)** is a decentralized traffic protocol for mixed-traffic environments, i.e., when human participants are also present. While both these protocols clearly outperform conventional traffic lights, they also introduce additional constraints, effectively increasing inter-vehicle distances. This reduces their throughput when compared to BRIP. On the other hand, similar to BRIP, they are still based on stringent space requirements as stated before.

The above approaches are basically intended for single-vehicle crossing. In [14], vehicles going in the same direction may spontaneously group into platoons, which are then scheduled through the intersection in an asynchronous fashion. This allows for a reduction in travel and wait time as well as the overall communication overhead to the RSU compared to single-vehicle crossing. In this article, we follow a similar approach in proposing our platoon crossing strategy; however, in contrast to [14], our approach is synchronous and takes vehicles with varying lengths into account.

Further, in [15], a vehicle grouping technique is proposed. In particular, vehicles are grouped into pre-determined zones, taking, among others, their length, width, and arrival direction into account.

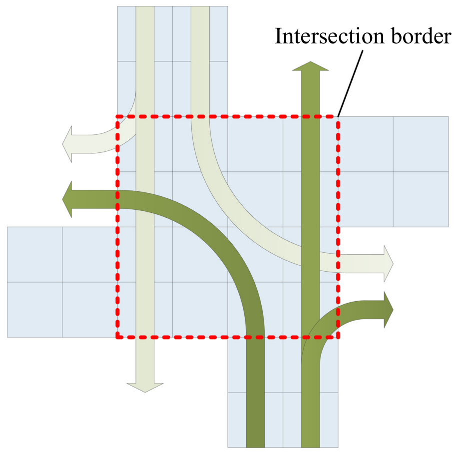

Fig. 1. Basic intersection layout of the proposed LTR protocol: with one left-turn and one shared through/right-turn lane. The underlying squares highlight the concept of sectors used throughout this article. Furthermore, all figures and behaviors are based on right-hand traffic.

These groups are then scheduled to cross together. In contrast, our work attempts to interleave vehicles from different arrival directions at the same time. Additionally, we considered the waiting time incurred on conflicting lanes to ensure fairness as well as the impact on throughput.

In [5], policies based on stop signs were introduced for autonomous intersection management and later refined into a **platoon-based autonomous intersection management (PAIM)** [6]. This is based on a time-triggered 4-phase plan (i.e., one phase for each direction). PAIM was shown to result in a reduction in travel delay and fuel consumption and to increase throughput with respect to conventional traffic lights. Contrary to our work, neither space requirements nor fairness-related metrics are discussed. In addition, since the four phases are fixed in time and order, it is difficult to adapt PAIM to scenarios of heterogeneous vehicles traveling in random directions.

3 ASSUMPTIONS AND DEFINITIONS

While the concepts in this article apply to both left-hand and right-hand traffic, for simplicity and consistency, we assume right-hand traffic in all figures and explanations.

As previously mentioned, the center of the intersection is divided into square, same-size sectors with a side length S equal to the width of a lane as depicted in Figure 1. We therefore assume that all lanes have identical widths in the intersection center.

Further, to allow describing distances smaller than S, we define σ as a *sector fraction* of S, in particular, $\sigma = \frac{S}{k}$ for any positive integer k. In this article, for ease of exposition, we set $\sigma = 1$ m limiting S to multiples of 1 m, however, other values are also possible.

We denominate the distance from front bumper to front bumper of two consecutive vehicles on the same lane by *vehicle period* P. Therefore, the inter-vehicle separation is given by $D = P - \lceil \frac{L}{\sigma} \rceil$, with L being the leading vehicle's length. Assuming that vehicles always fit within one S, both vehicle period P and inter-vehicle separation D can be expressed as multiples of S to keep them independent of absolute values.

A synchronous traffic protocol can be easily implemented using a centralized approach, i.e., where an RSU collects from and distributes data to vehicles.[1] We set the intersection's *region of*

[1]A decentralized implementation is also possible, but considerably more challenging. For example, one of the approaching vehicles can become coordinator collecting and distributing data. In this case, a new coordinator needs to be found for each set of vehicles crossing the intersection.

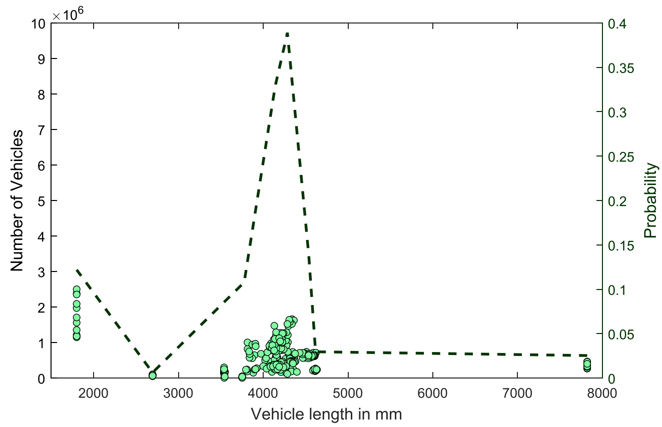

Fig. 2. Number of vehicles sold in Europe and their corresponding length distribution (up to around 8 m) from [17]. Note that buses and trams may be considerably longer (in some cases around 20 m).

influence as a radius (e.g., 200 m) from the center, in which the RSU enforces desired speeds on all vehicles. As previously mentioned, we propose a two-speed policy where vehicles are assigned one of two speeds, namely a low V_{LO} and a high speed V_{HI}, with V_{LO} being the standard speed. On the other hand, V_{HI} is intended only for drive-though maneuvers. Note that, whereas a vehicle can also drive through the intersection at V_{LO}, it cannot turn either right or left at a speed of V_{HI}. In this article, for ease of exposition, we set $V_{HI} = 1.5V_{LO}$; however, other values are also possible. We discuss this in detail in Section 4.1.1.

To allow for comparisons independent of a specific speed, we refer by cycle C to the time required to cover a distance equal to the sector size S at the standard speed V_{LO}. Clearly, it takes a vehicle $1/k$-th of a cycle to cover a distance equal to σ at V_{LO}.

Depending on the length L of a vehicle, it may not fit entirely into one sector. While regular passenger vehicles have a length of around 5 m, there also exist exceptionally long vehicles (e.g., buses, trucks, etc.) [11]. In this article, we refer to them as *overlength vehicles*, requiring additional space that we denominate *overlength penalty* O.

$$O = \max\left(0, \left\lceil \frac{L-S}{\sigma} \right\rceil\right) = \max\left(0, \left\lceil \frac{L}{\sigma} \right\rceil - k\right). \quad (1)$$

While increasing the sector size S reduces the overlength penalty (if $S \geq L$, the overlength penalty is $O = 0$), this can lead to space-hungry intersections. In this case, overlength vehicles are treated as the default case. Instead, in this article, we opt to incorporate overlength penalty into the traffic protocol to account for the low probability of occurrence of overlength vehicles (see Figure 2). That is, the overlength penalty O needs to be added to the vehicle period P to ensure sufficient inter-vehicle separation D.

Without loss of generality, in this article, we assume uninterrupted vehicle flow following the patterns of the corresponding protocols, as well as a constant *baseline speed* of $V_{LO} = V_{BRIP} =$ 30 km/h.

This allows us to obtain an upper bound on throughput, which we wanted to show for all approaches the same. If, on the other hand, traffic flow is sparse, there will be empty slots. However, this affects all protocols the same and has no impact on the presented comparison.

Further, in contrast to the proposed protocol, the approaches from the literature assume the same speed for all maneuvers, including turns. Hence, we selected a speed ($V_{LO} = V_{BRIP}$) that is as large as possible, maximizing throughput while still ensuring passenger-friendly turns.

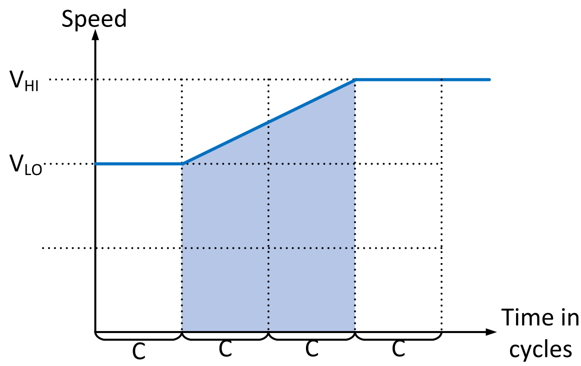

Fig. 3. Since $V_{HI} = 1.5V_{LO}$, vehicles can transition from one speed to the other in a time equal to $2C$, traversing a distance equal to $2.5S$ in the process, with $S = V_{LO} \cdot C$.

4 PROPOSED PROTOCOL

LTR is a synchronous traffic protocol and, hence, it enforces arrival patterns at an intersection with the aim of increasing vehicle throughput. In contrast to similar approaches from the literature, LTR does not require modifying the existing road infrastructure (e.g., when overlength vehicles are allowed at the intersection as discussed previously). We analyze both single-vehicle as well as platoon crossing.

4.1 Single-Vehicle Crossing

To introduce LTR's basic behavior, we consider the intersection layout of Figure 1 consisting of a shared through/right-turn lane and a dedicated left-turn lane.[2] Let us first assume that every vehicle fits entirely into one of the depicted sectors.

In this setting, LTR enforces arrival patterns with either two or four vehicles at the same time depending on where vehicles are heading to (i.e., their drive directions). We denominate this as *two-way synchronization* and *four-way synchronization,* respectively. In particular, in a two-way synchronization, only vehicles on opposing lanes cross simultaneously. That is the case, for example, when too many vehicles go in the same direction, etc.

4.1.1 Transition between Speeds. Since the system allows for two different speeds (V_{LO} for turn and V_{HI} for drive-through maneuvers), transitions between them need to be defined in advance. For simplicity, we set $V_{HI} = 1.5V_{LO}$ and assume that it takes any vehicle two cycles (i.e., $2C$) to transition from one speed to the other at a constant acceleration/deceleration — see Figure 3.[3]

As a result, the distance traveled in that process is equal to:

$$d = V_{LO}(2C) + \frac{1}{2} \cdot \frac{V_{LO}}{4C}(2C)^2 = 2.5S,$$

where $\frac{V_{LO}}{4C}$ is the absolute value of the acceleration/deceleration (i.e., the speed changes by an amount $\frac{V_{LO}}{2}$ in a time equal to $2C$), and again S is given by $V_{LO} \cdot C$.

4.1.2 Drive Through Only/Drive Through and Right Turns. Due to the synchronous nature of the protocol, all vehicles arrive at the border of the intersection simultaneously. As previously

[2]Note that LTR can be extended to more complex settings as well.
[3]Note that, in principle, any other relation between V_{HI} and V_{LO} is possible. However, the transition between speeds will then differ from Figure 3.

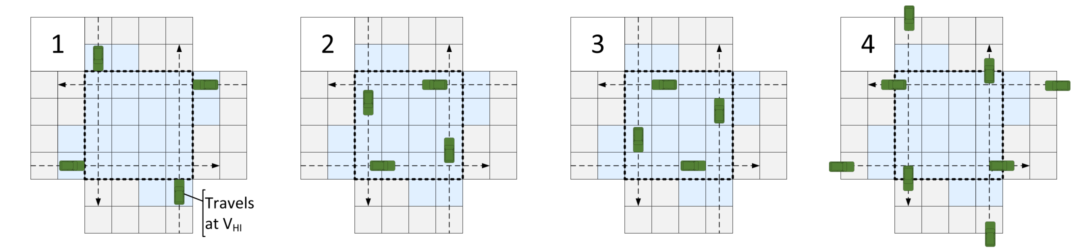

Fig. 4. Driving through in a *four-way synchronization*: This maneuver is divided into 4 one-cycle steps. Since vehicles travel at V_{HI}, they traverse $1.5S$ during $1C$. yielding a vehicle period $P = 6S$. The total throughput is 4 vehicles every 4 cycles.

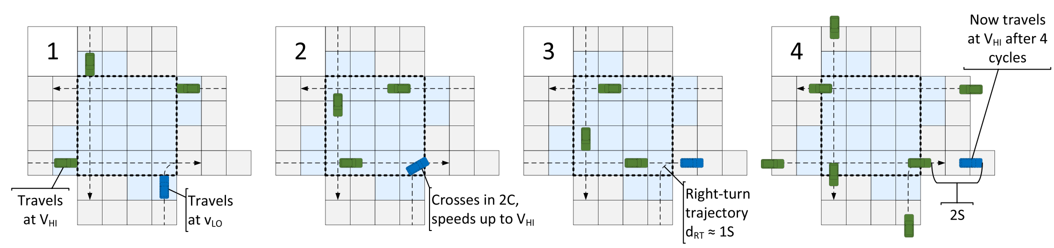

Fig. 5. Driving through with one vehicle turning right.

discussed, vehicles cover a distance equal to S in one cycle when traveling at V_{LO}. However, since vehicles driving through travel at $V_{HI} = 1.5V_{LO}$, they traverse a distance of $1.5S$ per cycle. Therefore, the crossing process can be divided into 4 different steps (one per cycle), as depicted in Figure 4, yielding a vehicle period P of $6S$ (i.e., 4 cycles at V_{HI}).[4] A shorter vehicle period can potentially lead to accidents with vehicles crossing perpendicularly, while a longer vehicle period decreases throughput.

Note that it is possible that any of the vehicles in Figure 4 turns right instead of driving through. Such a vehicle arrives at the intersection at the same time with the others (see step 1 in Figure 5), but it turns right at V_{LO}[5] traversing an arc-like trajectory $d_{RT} = \frac{\pi}{4}S \approx 0.79S$ (i.e., 1/4-th of a circumference with radius $r = S/2$, assuming the vehicle travels from the middle of its current lane to the middle of its target lane. At V_{LO}, this maneuver takes around $0.79C$ to complete, however, we approximate this to one full cycle for simplicity. The vehicle then transitions to V_{HI}, which takes a time $2C$ and a distance of $2.5S$ (see again Figure 3), leading to vehicle period $P = 2S$ to its following vehicle on the same lane, see step 4 in Figure 5.

Note that any single vehicle in the drive-through setting can be substituted with a vehicle turning right without increasing the vehicle period. Moreover, if all vehicles turn right, one can reduce the vehicle period even further, as described later when discussing right and left turns in Figure 8.

Finally, when adding another through/turn-right lane to this setting, vehicles on the outermost such lane will have to arrive at the intersection border with a delay of one cycle (with respect to vehicles on the innermost lane). Otherwise, these will collide with vehicles turning right from the innermost through/turn-right lane. In general, there will have to be an additional one-cycle delay per through/turn-right lane added to the setting.

[4]While a vehicle period $P = 5S$ would also be possible, we selected $6S$ to preserve safety, since we later combine drive-through with other maneuvers.

[5]Note that a vehicle turning right at the intersection may approach the intersection at V_{HI}, however, it has to slow down to V_{LO} for the turn maneuver.

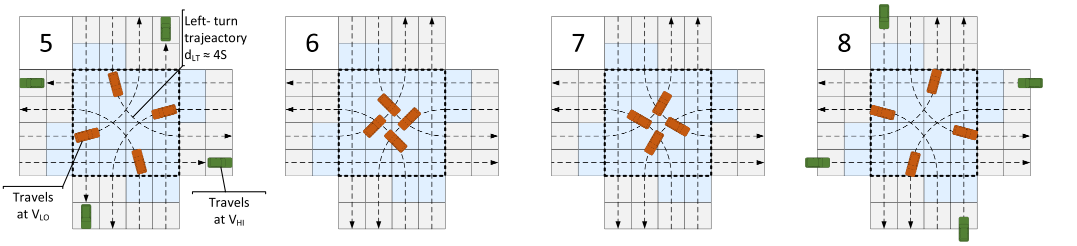

Fig. 6. Combined drive-through and left-turn maneuvers in a *four-way synchronization*. Note that a throughput of eight vehicles every eight cycles is possible. When a *two-way synchronization* is required/used instead (see explanation below), the throughput drops to six vehicles every eight cycles. This figure illustrates the last four four steps (i.e., steps 5 to 8). Steps 1 to 4 are identical to those in Figure 4 (with the only exception that approaching vehicles in step 4 turn left instead of driving through).

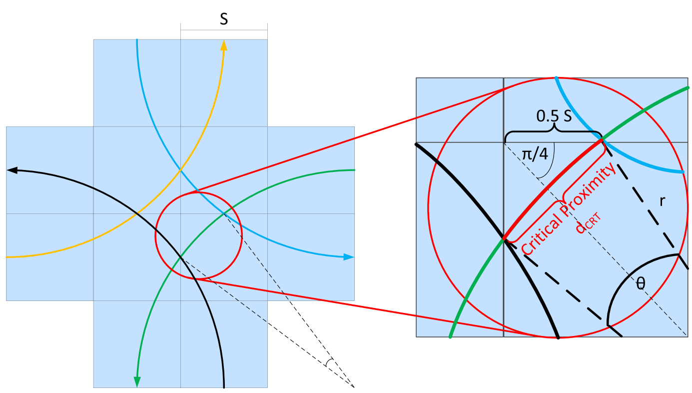

Fig. 7. Trajectories of vehicles turning left. If one vehicle exceeds a given maximum length equal to the critical proximity d_{CRT}, a two-way synchronization where only opposing lanes cross simultaneously must be enforced instead.

4.1.3 Drive Through and Left Turns. Having a set of vehicles driving through requires four cycles, as described above. Once these vehicles have cleared a sufficiently large portion of the intersection center, such that collisions can be avoided, left-turn maneuvers can start as depicted in Figure 6. Since left turns describe a curved path as shown in Figure 7, vehicles will require a non-integer number of cycles to exit the intersection. Specifically, their curved trajectory corresponds to 1/4-th of a circumference with radius $r = 2.5S$, assuming again that vehicles travel from the middle of their current lane to the middle of their target lane. Hence, the covered distance in sectors can be calculated as follows:

$$d_{LT} = \frac{\pi}{4}(2 \cdot 2.5S) = \frac{5\pi}{4}S = 3.9270S. \tag{2}$$

This leads to left-turning vehicles exiting the intersection center 0.073 cycles to the end of step 8 in Figure 6. While this desynchronizes their trajectory, it does not compromise safety due to left-turning vehicles leaving the intersection center earlier (before the 8th step ends). In addition, they still have a sufficient distance to preceding vehicles in the same direction (>1.9S). For simplicity, we assume approximate $d_{LT} = 4S$.

Ideally, each set of left-turning vehicles consists of four vehicles in a four-way synchronization. However, since vehicles turning left on perpendicular lanes have overlapping trajectories, these

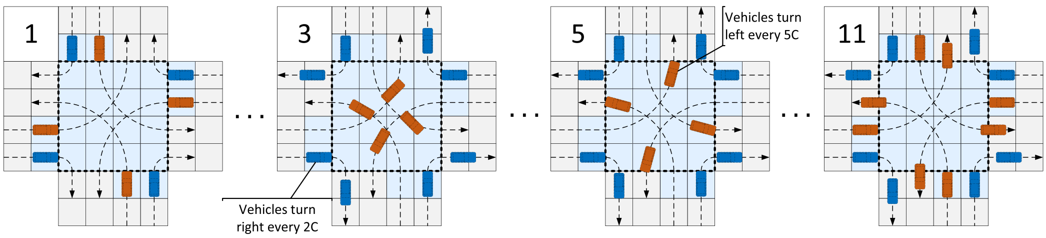

Fig. 8. Right and left turns in a *four-way synchronization*. Since trajectories do not interfere with each other, this combination results in the highest possible throughput. Vehicles making right turns can cross in sets of four every two cycles. In addition, vehicles turning left can cross in sets of two to four every five cycles. This results in 24 to 28 vehicles every 10 cycles, depending on whether two-way or a four-way synchronization needs to be used for left turns.

approach up to a distance we denote *critical proximity* d_{CRT}:

$$d_{CRT} = r \cdot \theta = 2.5S \cdot \theta = 0.7095S, \tag{3}$$

where θ can be calculated through the Law of Sines, see Figure 7. That is:

$$\frac{r}{sin(\frac{\pi}{4})} = \frac{d}{sin(\frac{\theta}{2})}.$$

Therefore, if any of the vehicles turning left has $L \geq 0.7S$ (i.e., is longer than 3.5 m assuming $S = 5$ m), a four-way synchronization will lead to collisions. In that case, only opposing lanes can perform left-turn maneuvers simultaneously since their trajectories never coincide, leading to a two-way synchronization.

In summary, interleaving vehicles driving through and turning left allows for a total throughput of six to eight vehicles per eight cycles. That is, four vehicles drive through in the first four cycles (steps 1 to 4 as per Figure 4), whereas either two or four vehicles turn left over a time of five cycles (steps 4 to 8 as per Figure 6).

4.1.4 Right and Left Turns. As shown in Figure 8, since there are no conflicting trajectories, right and left turns do not interfere with each other. Similar to the previous case, vehicles turning left can cross every five cycles in either sets of four, provided that all vehicles are shorter than $0.7S$, or sets of two, otherwise. In addition, right turns allow for four vehicles to cross every two cycles and, hence, we have a total throughput of 24 to 28 vehicles per 10 cycles, consisting of 20 vehicles turning right and 4 to 8 vehicles turning left.

4.1.5 Considering Overlength Vehicles. All previous considerations assume that vehicles always fit entirely into one sector. Extending this to vehicles longer than S is achieved by introducing an overlength penalty which adds an appropriate number of sector fractions σ to the vehicle period P following Equation (1).

As an example, a vehicle with $L = 7.5$ m in an intersection with $S = 5$ m and $\sigma = 1$ m (i.e., $\sigma = \frac{S}{5} = 0.2S$) leads to $O = 3\sigma$. In a drive-through-only setting, this increases the set's vehicle period from $P = 4S$ to $P = 4.6S$ (i.e., 4 full cycles and 3 additional σ, with $\sigma = 0.2S$), reducing throughput from 4 vehicles per 4 cycles to 4 vehicles per 4.6 cycles.

Other crossing patterns are affected a similar way, i.e., the vehicle throughput is reduced by some amount that depends on the overlength penalty.

In particular, note that overlength vehicles unavoidably lead to a two-way synchronization in left-turn settings (i.e., they always exceed the critical distance as per Equation (3)) and may further

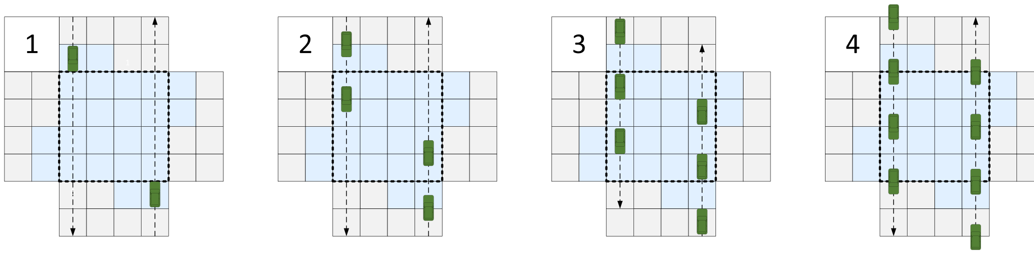

Fig. 9. Drive-through platoon crossing. After an initialization phase of four cycles, six vehicles cross the intersection every four cycles with a vehicle period $P = 2S$.

enforce a two-way synchronization in drive-through-only settings when $O \geq 10\sigma$, i.e., $2S$ with the values of S and σ assumed before.

4.2 Platoon Crossing

To improve throughput, we propose modifying the previously described single-vehicle crossing by grouping vehicle into platoons. This allows reducing the vehicle period in the direction of crossing, but it increases the waiting time of vehicles with conflicting trajectories. Next, we discuss the different platoon crossing regimes and transitions between them.

4.2.1 Drive-Through Platoon Crossing. If vehicles fit in one sector of length S, as shown in Figure 4, we can reach a throughput of 4 vehicles per 4 cycles under single-vehicle crossing. However, having platoons (i.e., chains of vehicles) that are almost always longer than $2S$ enforces a two-way synchronization, where only vehicles on opposing (and not on perpendicular) lanes are allowed to cross at a time.

On the other hand, platoon crossing allows multiple vehicles to cross (in the same direction) in the least possible time as depicted in Figure 9. The intelligent intersection then switches to allow vehicles to cross in a different direction.[6]

There are three different cases leading to a state transition from a drive-through regime: approaching vehicles that turn left either in the same or perpendicular direction or vehicles driving through in the perpendicular direction. The different possible transitions are displayed in Figure 10.

4.2.2 Left-turn Platoon Crossing. As discussed above, a throughput of 4 vehicles per 5 cycles can be achieved for single vehicles turning left, if vehicles are shorter than $0.7S$. However, again, having platoons that are typically longer than $2S$ enforce a two-way synchronization, i.e., only vehicles on opposing lanes can cross at a time. On the other hand, this allows us to increase the throughput to 2 vehicles per 3 cycles as depicted in Figure 11. Necessary transitions to a left-turn platoon crossing regime are shown in Figure 11.

Similar to before, there are three different situations forcing a state transition from a left-turn regime: approaching vehicles that turn left in the perpendicular direction or vehicles driving through in either the same or the perpendicular direction, see Figure 12.

4.2.3 Considering Right Turns. Since vehicles turning right do not have overlapping trajectories with vehicles turning left, they can coexist under both single-vehicle and platoon crossing. On the other hand, right-turn and drive-through trajectories overlap, however, it is possible to

[6]Clearly, this is similar to conventional traffic lights that change after some time. However, the proposed platoon crossing allows for much shorter inter-vehicle separations achieving a considerably higher efficiency.

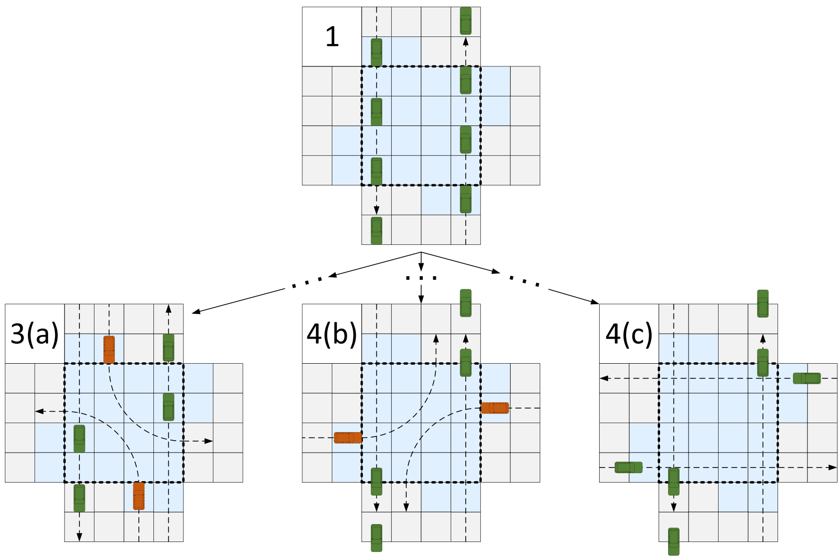

Fig. 10. Transition from drive-through platoon crossing to (a) a left-turn regime from the same direction (b) a left-turn regime from the perpendicular direction (c) a drive-through regime from the perpendicular direction. Due to the increased speed $V_{HI} = 1.5V_{LO}$, transitioning between perpendicular drive-through regimes requires a non-integer amount of cycles ($(4 + 1/3)C$) for the next set of vehicles to reach the intersection center.

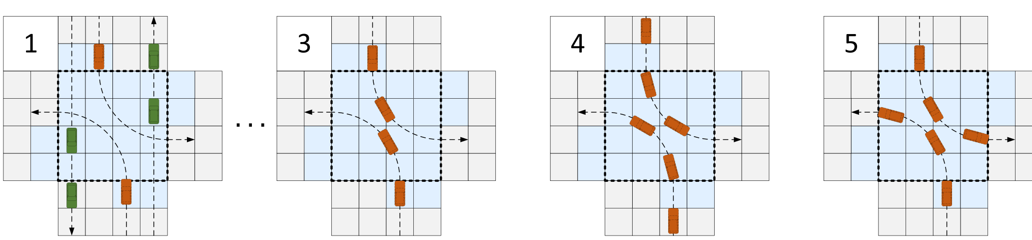

Fig. 11. Left-turn platoon crossing. After an initialization phase of 5 cycles, 2 vehicles cross the intersection every 2 cycles with a vehicle period $P = 2S$.

interleave them under single-vehicle crossing using a vehicle period of $P \geq 6S$ (see Figure 4) to avoid collisions. Since drive-through platooning schemes use a vehicle period of $P = 2S$, interleaving drive-through and right-turn maneuvers is not possible unless an outer dedicated lane is available for right turns, in which case these are completely independent of the rest of the intersection.

4.2.4 Maximum Blocking Time. To ensure fairness by avoiding long waiting times, the intersection switches from one platoon regime to another based on what we call maximum blocking threshold, denoted by B_{max} and expressed in cycles. That is, having $B_{max} = 25$ indicates that the current platoon regime (e.g., on the southbound/northbound lanes) will run for 25 cycles before a transition is performed (at the next appropriate point in time) to allow the vehicles on the currently inactive lanes (in this example, the eastbound/westbound) to cross.

Finally, all considerations towards overlength discussed in Section 4.1 also apply to the platoon crossing regimes introduced before.

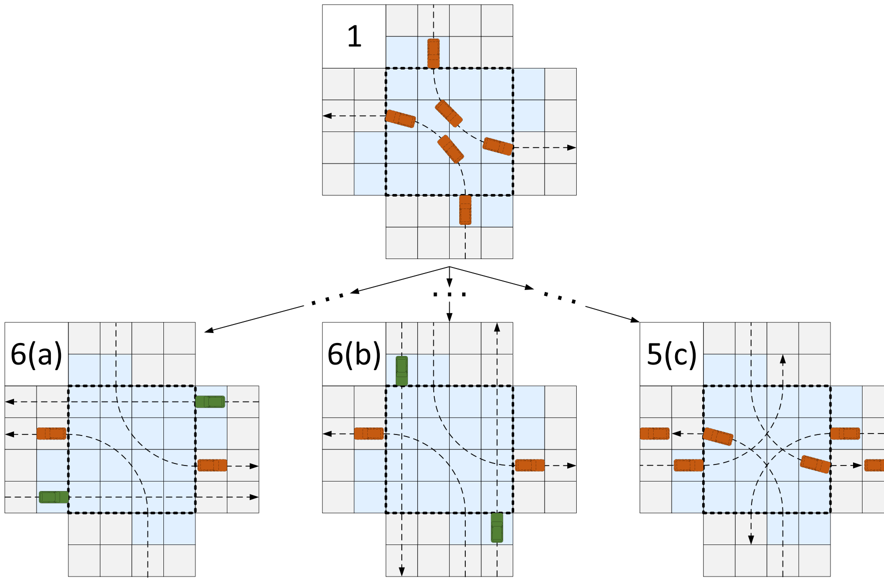

Fig. 12. Transition from a left-turn platoon crossing to (a/b) a drive-through regime from either direction (requiring 6 cycles) or (c) a left-turn regime from the perpendicular direction (requiring 5 cycles). Similar to before, step 1 in this figure corresponds to step 5 in Figure 11, step 5(c) corresponds to step 1(c) in Figure 11 and step 6(a/b) in this figure corresponds to step 1 in Figure 9, respectively.

4.3 The Algorithm

All vehicles in the intersection are handled by the RSU, which periodically runs the algorithm presented below. Basically, the algorithm (i) monitors planned vehicles for deviations from their intended trajectory so as to guarantee safety and (ii) assigns unplanned (i.e., newly arrived) vehicles to either an existing set or (iii) creates a new set that accommodates an unplanned vehicle if there is no match.

To this end, the RSU loops over all vehicles in the intersection's region of influence, including vehicles that are not yet assigned to a vehicle set (i.e., that have arrived after the last execution of the algorithm). It then request updates from all vehicles in the system (line 3), e.g., position, speed, etc.

If a planned vehicle (already part of a set) deviates from its intended trajectory, the RSU triggers a fall-back mechanism which brings the intersection to a fail-safe state to avoid collisions, e.g., all red lights (line 7).

For unplanned vehicles, the RSU checks all existing vehicle sets for empty slots that are compatible with the vehicle, i.e., the set has a slot the vehicle can realistically reach and that matches the vehicle's direction (line 18). If compatible, the vehicle is assigned to the corresponding set (line 20). If no sets are compatible, a new set is created for the vehicle (line 30).

Complexity. Since the algorithm checks all existing vehicle sets (lines 12 to 27) for all unplanned vehicles, it has a linear complexity of $O(n)$, where n is the number of vehicles at the intersection. This is because we can have one vehicle per set and the number of unplanned (i.e., newly arrived) vehicles is upper-bounded by the number of lanes, e.g., 8 for a two-lane intersection (2×4, i.e., two lanes in every direction).

5 COMPARISON AND EVALUATION

As discussed in Section 2, while BRIP has been extended into CSIP [1] and DSIP [2], the additional constraints these protocols introduce reduce their throughput when compared to BRIP. On the

ALGORITHM 1: LTR control loop for all vehicles

```
foreach (v in Vehicles) /* Cycle through all vehicles */
do
    Update(v); /* Get current vehicle information for v */
    if (v.Planned == true) then
        if (⌈v.CurrentPosition - ExpectedPosition(v)⌉ > threshold) /* Deviation from expected position
          */
        then
            Fallback(); /* If deviation exceeds critical threshold, engage fall-back mechanism */
        else
            continue; /* Check next vehicle */
        end
    else
        foreach (s in Sets) /* Vehicle sets are sorted, from crossing time from earliest to latest
          */
        do
            if (s.Full == true) /* Set is full set, i.e., no empty slots */
            then
                continue; /* Check next set */
            else
                if (IsCompatible(s, v) /* Direction of v matches the direction of an empty slot in
                  s */
                then
                    Assign(s, v); /* Assign v to s) */
                    v.Planned=true; /* Flag vehicle as planned  */
                    break; /* Step out of inner foreach-loop, go to next vehicle */
                else
                    continue; /* No match, go to next set */
                end
            end
        end
        if (v.Planned==false) /* v does not match any existing set */
        then
            s_new = CreateNewSet(v); /* Create new set that fits the trajectory of vehicle v */
            Assign(s_new, v); /* Assign v to new set s_new */
            v.Planned = true; /* Flag vehicle as planned */
            Sets.Append(s_new); /* Add s_new to list of sets */
        else
            continue; /* Vehicle is assigned - go to next vehicle. */
        end
    end
end
```

other hand, they retain the same space requirements as BRIP. In this section, for ease of exposition, we compare and evaluate the proposed LTR with BRIP alone based on best-case and worst-case performance. Note that, all comparisons are also valid for CSIP and DSIP.

To this end, as mentioned in Section 3, we assume an uninterrupted vehicle flow following the patterns by the corresponding protocol and consider a constant *baseline speed* of

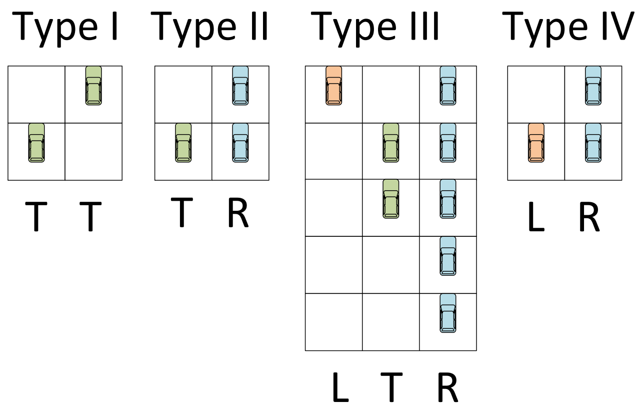

Fig. 13. Overview of arrival patterns for the different BRIP types [4]. In this picture, T stands for through lane, whereas R and L stand for right- and left-turn lane respectively.

$V_{LO} = V_{BRIP} = 30\ km/h$. This way, we can express the throughput by vehicles per minute (instead of vehicles per cycle). This also leads to $V_{HI} = 1.5V_{LO} = 45\ km/h$ for vehicles driving through under the proposed LTR.[7]

5.1 The BRIP Protocol

As already mentioned, BRIP has different types corresponding to different intersection layouts as depicted in Figure 13.

One of the core assumptions of BRIP is that the sector length has to be at least $L + W$, where L and W are the length and width[8] of the largest vehicle in the intersection. Since the definition of a cycle is bound to the sector length, we distinguish between the previously discussed cycles in LTR and *BRIP cycles*, which are considerably larger. Furthermore, BRIP synchronizes vehicles in all directions with the same constant speed. For settings with both vehicles driving through and vehicles turning right/left, this requires the chosen speed to be low enough to allow for safety and comfort, even for the vehicles driving through. This affects all BRIP types except for Type I.

While BRIP was conceived for same-size vehicles only, the possibility considering overlength vehicles was contemplated by changing arrival patterns and allowing for more empty arrival slots [4]. However, this implies nontrivial modifications, certainly diminishing throughput, and remains unsolved.

BRIP Type I is a two-lane setup for drive-through maneuvers only. BRIP Type I can accommodate one vehicle per direction and BRIP cycle. Type II consists of a mix of drive through and right turns, similar to Type IV, which combines left turns and right turns. These latter two types yield a throughput of three vehicles per direction every two BRIP cycles. Type III allows for 8×4 vehicles in the different directions every 5 BRIP cycles, however, considering a three-lane setup instead of only two lanes as in the previous cases.

5.1.1 Concerning Right Turns. BRIP Types II, III, and IV have a dedicated right-turn lane. All three types allow for one vehicle per BRIP cycle and direction to turn right. However, the remaining lanes in these three cases are completely independent of right turns, as depicted in Figure 14 for

[7] Clearly, the higher the baseline speed considered, the higher the throughput. However, note that this comparison and the corresponding conclusions drawn remain valid independent of the selected speed.

[8] To this end, we have set the width of a vehicle to be 30% of its length in our analysis, but no more than 3 m (i.e., all vehicles longer than $L = 10\ m$ are considered to be as wide as $W = 3\ m$), which is true for most commercial vehicles [7, 11].

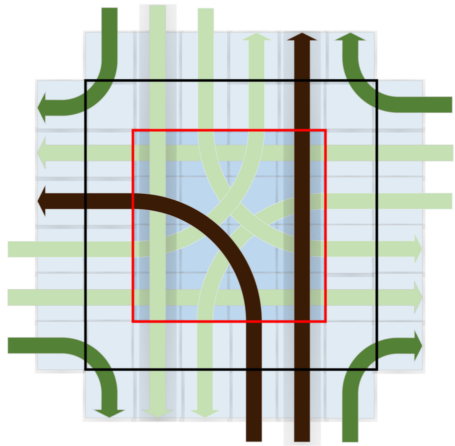

Fig. 14. Intersection setup for BRIP Type III. Note that left-turn and drive-through maneuvers are confined to a two-lane space as marked by the red square. Right turns do not impact the remaining lanes.

the case of Type III. As a result, we separate the analysis of right turns from the rest of the lanes. We first perform a detailed comparison of our proposed LTR with BRIP Type I and Type III where right turns are not considered and then analyze the effect of a dedicated right-turn lane on the different protocols separately. Note that Type III is a generalization of Type II and Type IV. So Type II and IV are not explicitly included in this comparison.

5.2 Intersection Size

We first investigate how the proposed LTR[9] and the existing BRIP affect the intersection size.

We consider that vehicles can be as short as 3 m, e.g., motorbikes, small vehicles, and the like, and as long as 30 m, which accounts for public transportation vehicles such as buses or trams, and the like.

As the LTR's sector length S depends on the width of a lane, we consider different variants in this comparison. These are LTR_3 with $S = 3$ m, LTR_5 with $S = 5$ m, and LTR_7 with $S = 7$ m. In contrast, BRIP's sector length depends on the size of the largest possible vehicle allowed at the intersection, i.e., $S_{BRIP} \geq L + W$.

As discussed above, both LTR and BRIP Type I consist of two-lane roads going in all directions (i.e., north, south, east, and west). These yield an intersection side/length of $4S$ and $4S_{BRIP}$, hence, leading to a surface (i.e., space requirement) of $16S^2$ and $16S_{BRIP}^2$ respectively.

Figures 15(a) and 15(b) depict the space and area requirements on the intersection infrastructure by BRIP and the LTR variants discussed above. As expected, we can observe that BRIP's space requirements rapidly increase with the length of the longest vehicle allowed at the intersection. For example, for vehicles with $L_{max} = 15\ m$, BRIP requires the intersection to provide a length of $80\ m$, which translates in $6400\ m^2$. Clearly, the longer vehicles are, the less existing intersections comply with BRIP's space requirements.

In contrast to this, LTR does not have any inherent space requirements, but is rather parameterized to current lane width in an existing intersection. For example, LTR_5 requires the intersection width to be 20 m (i.e., 400 m^2) independent of the vehicles' lengths.

[9]Note that LTR in both its variants, i.e., single-vehicle and platoon crossing, leads to the same space requirements on the intersection. As a result, we do not distinguish between them in this comparison.

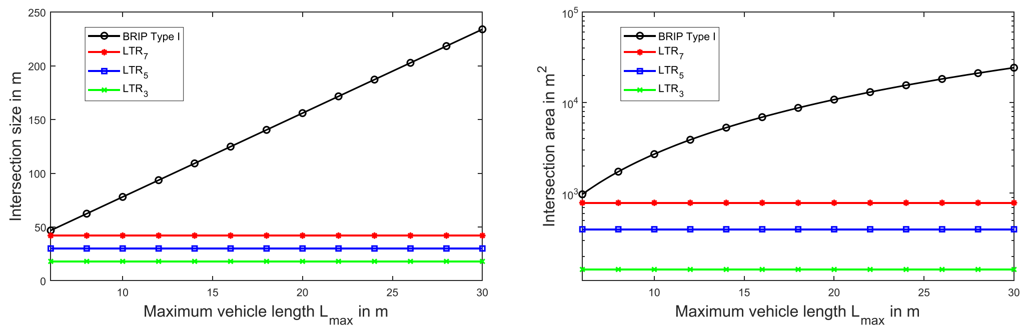

(a) Intersection size.

(b) Intersection area.

Fig. 15. Intersection size and area versus vehicle length: BRIP Type I against LTR variants. As the LTR's sector length S depends on the width of a lane, we consider different variants. These are LTR_3 with $S = 3$ m, LTR_5 with $S = 5$ m, and LTR_7 with $S = 7$ m.

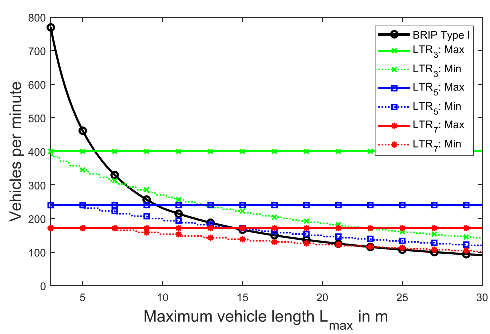

Fig. 16. Throughput versus maximum vehicle length: BRIP Type I against LTR variants at $V_{Hi} = 45$ km/h for all vehicles. As mentioned above, we considered the variants LTR_3, LTR_5, and LTR_7 for sector sizes 3 m, 5 m, and 7, respectively.

Note that the space requirements as per Figure 15(a) correspond to BRIP Type I. Due to needing 3 lanes (instead of only 2), BRIP Type III requires 50% wider intersections (leading to an area increase of 125%).

5.3 Throughput for Single-Vehicle Crossing

In this section, we compare the throughput of BRIP Types I and III against LTR for single-vehicle crossing as introduced in Section 4.1.

5.3.1 BRIP Type I. Note that, since BRIP Type I only consists of vehicles driving through, both BRIP and LTR utilize the higher speed $V_{BRIP} = V_{HI} = 1.5V_{LO}$ for the entire maneuver. As shown in Figure 16, for settings consisting of only short vehicles, BRIP Type I substantially outperforms LTR in all its variants in terms of throughput. With longer vehicles being allowed at the intersection, this disparity grows smaller until LTR starts outperforming BRIP.

For a given maximum vehicle length L_{max}, LTR leads to a maximum and a minimum throughput, depending on the frequency with which these vehicles actually appear at the intersection. If these long vehicles are possible, but never cross the intersection, LTR results in the best/maximum possible throughput (i.e., the best case). If only these long vehicles cross the intersection, LTR results in the worst/minimum possible throughput (i.e., the worst case).

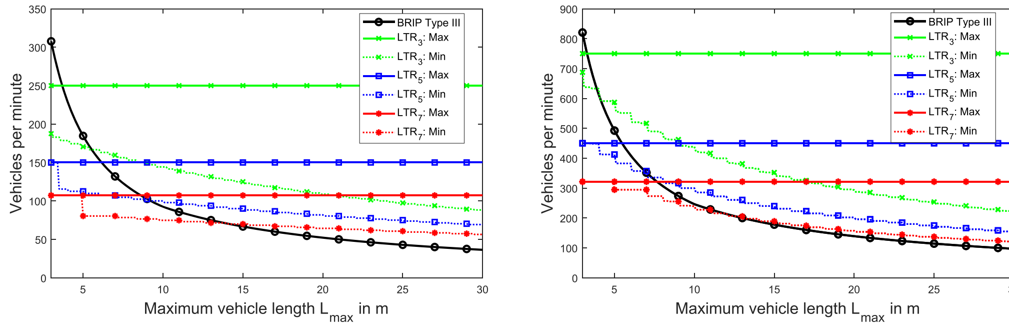

(a) Disregarding right turns. (b) Considering right turns.

Fig. 17. Throughput versus maximum vehicle length when disregarding or considering right turns: BRIP Type III against single-vehicle crossing LTR variants for $S = 5$ m. Since vehicles drive through and turn left, all vehicles in BRIP have the same constant speed $V_{LO} = 30$ km/h. Under LTR, vehicles drive through at $V_{HI} = 45$ km/h. Again, LTR_3, LTR_5, and LTR_7 are considered for sector sizes 3 m, 5 m, and 7 m, respectively.

On average, one can expect a throughput that is in between these maximum and minimum curves of the corresponding LTR variant. For example, LTR_5 provides a throughput between 170 and 240 vehicles per minute for vehicles as long as $L_{max} = 15\ m$. BRIP Type I allows for around 167 vehicles per minute in this case (independent of how often long vehicles cross). That is, when considering vehicles as long as $L_{max} = 15\ m$, LTR_5 roughly breaks even with BRIP in the worst case, but it allows for up to 40% more throughput in the best case. In particular, LTR_5 starts outperforming BRIP Type I for $L_{max}^{BC} > 9.6$ m in the best case and $L_{max}^{WC} > 14.5$ m in the worst case. From now on, we will denote these bounds by $L_{max} > [L_{max}^{BC}, L_{max}^{WC}]$, i.e., in this case: $L_{max} > [9.6\text{ m}, 14.5\text{ m}]$.

5.3.2 BRIP Type III. Since Type III combines turn and drive-through maneuvers, BRIP is bounded by the turn speed V_{LO}, since, in contrast to LTR, it does not allow for a two-speed regime.

Independent of whether right turns are considered or not (in Figures 17(a) and 17(b) respectively), we observe a similar behavior as in the previous comparison against BRIP Type I. That is, for settings with only short vehicles, BRIP Type III outperforms LTR in all its variants. However, LTR starts outperforming BRIP Type III as longer vehicles start being possible/allowed at the intersection.

For instance, when disregarding right turns in Figure 17(a), LTR_5 leads to throughputs between roughly 90 and 150 vehicles per minute when considering that vehicles of up to $L_{max} = 15$ m are allowed at the intersection. Similar to before, the minimum throughput results when only 15 m vehicles cross the intersection, whereas the maximum throughput results when 15 m vehicles are possible, but seldom present at the intersection. In contrast to this, BRIP Type III reaches a performance of around 67 vehicles per minute if vehicles as long as 15 m can cross, independent of whether these vehicles are present or not at the intersection. Specifically, LTR_5 starts outperforming BRIP Type III in this case for $L_{max} > [6.1\text{ m}, 9.2\text{ m}]$, i.e., for vehicle lengths greater than 6.1 m in the best case and for vehicle lengths greater than 9.2 m in the worst case.

Taking right turns into consideration substantially increases throughput numbers due to the constant uninterrupted flow of right turns on a dedicated lane. Here, as shown in Figure 17(b), LTR_5 leads to throughputs between 240 and 450 vehicles per minute for the same maximum vehicle length $L_{max} = 15$ m, whereas BRIP reaches a throughput of roughly 178. When considering right turns, LTR_5 starts outperforming BRIP Type III for $L_{max} > [5.4\text{ m}, 7.4\text{ m}]$.

Note that for $L_{max} = 5\ m$, BRIP outperforms LTR, since BRIP always allows for a four-way synchronization (i.e., vehicles crossing in all four directions at the same time). More specifically,

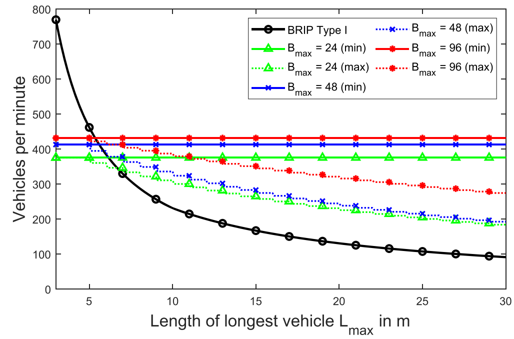

Fig. 18. Throughput versus maximum vehicle length when considering right turns: BRIP Type I against LTR under platoon crossing for different B_{max} and $S = 5$ m. Solid lines represent best-case throughputs, while dotted lines represent worst-case throughputs.

fixing the vehicle length to 5 m changes BRIP's sector size to something more than 5 m, i.e., length plus width of the vehicle. BRIP then assumes that the intersections are always able accommodate its necessary sector size.

In contrast to this, LTR's sectors do not depend on vehicle length and are rather determined by the (existing) lane width. As previously mentioned, we have considered lane widths of 3 m, 5 m, and 7 m, which results in the LTR variants LTR_3, LTR_5, and LTR_7.

As a result, LTR does not always allow for a four-way synchronization, due to the lack of space in smaller intersections (with narrower lanes). For example, LTR does not allow for a four-way synchronization for left turns with vehicles longer than $\approx 0.7S$ (as described in Section 4.1.3) and, hence, has a worse performance than BRIP.

5.4 Throughput for Platoon Crossing

Now, we consider LTR's platoon crossing as proposed in Section 4.2. For ease of exposition, in the following analysis, we show LTR_5's behavior for different maximum blocking thresholds B_{max}. The conclusions drawn from it are, however, also valid for the other to LTR variants introduced before.

5.4.1 BRIP Type I. A comparison of BRIP Type I with the LTR_5 under platoon crossing can be seen in Figure 18.

Just as before, we consider both the best case (i.e., the longest possible vehicles are rarely present) and the worst case (all vehicles are as long as the longest possible vehicle). As noted previously, LTR outperforms BRIP in terms of throughput under single-vehicle crossing for $L_{max} > [9.6 \text{ m}, 14.5 \text{ m}]$. Under platoon crossing with $B_{max} = 24$, these values improve to $L_{max} > [6.2 \text{ m}, 6.7 \text{ m}]$. Similarly, LTR outperforms BRIP under platoon crossing for $B_{max} = 48$ and $L_{max} > [5.6 \text{ m}, 6.1 \text{ m}]$, whereas $L_{max} > [5.5 \text{ m}, 5.6 \text{ m}]$ results for $B_{max} = 96$. That is, the greater B_{max}, the lesser difference there will be between the best and the worst case.

Now, for $L_{max} = 15$ m, BRIP allows for 167 vehicles per minute in a Type I setting. Our platoon crossing LTR with $B_{max} = 48$ can reach a throughput of 282 to 412 vehicles per minute, i.e., 65% to 145% higher throughput than BRIP. Since B_{max} is the maximum number of cycles vehicles on conflicting lanes have to wait to cross, $B_{max} = 48$ results in a waiting time $t_{wait} = 48\frac{S}{V_{LO}} = 28.8$ s for $S = 5$ m and $V_{LO} = 30$ km/h.

5.4.2 BRIP Type III. In Figure 19(a), disregarding right turns, we can observe bounds of $L_{max} > [7.4 \text{ m}, 16.2 \text{ m}]$ for $B_{max} = 24$, $L_{max} > [5.9 \text{ m}, 8.4 \text{ m}]$ for $B_{max} = 48$ and $L_{max} > [5.4 \text{ m}, 6.4 \text{ m}]$ for

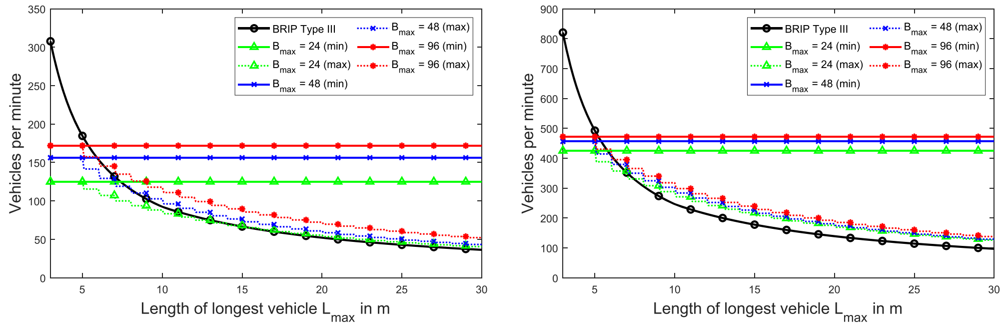

(a) Disregarding right turns.

(b) Considering right turns.

Fig. 19. Throughput versus vehicle length when disregarding or considering right turns: BRIP Type III against platoon crossing LTR variants for $S = 5$ m. Solid lines represent best-case throughputs, while dotted lines represent worst-case throughputs.

$B_{max} = 96$. When comparing to Figure 17(a), we can observe that LTR under platoon crossing outperforms the single-vehicle variant for a sufficiently high B_{max}. When considering vehicles that are as long as 15 m, LTR allows for between 76 to 156 vehicles per minute for $B_{max} = 48$ (which corresponds to a waiting time of 28.8 s on conflicting lanes). For the same setting, BRIP allows only 67 vehicles per minute.

The same holds true when considering right turns, see Figure 19(b). In particular, we can observe that we now have $L_{max} > [5.8 \text{ m}, 7.5 \text{ m}]$ for $B_{max} = 24$, $L_{max} > [5.4 \text{ m}, 6.5 \text{ m}]$ for $B_{max} = 48$, and $L_{max} > [5.2 \text{ m}, 6.2 \text{ m}]$ for $B_{max} = 96$. When we again consider vehicles of at most 15 m, LTR can achieve between 226 and 456 vehicles per minute for $B_{max} = 48$ (i.e., a waiting time of 28.8 s on conflicting lanes), as compared to 178 vehicles per minute by BRIP, i.e., allowing for an increase in throughput of roughly 25% to 150%.

As mentioned before, for $L_{max} = 5$ m, BRIP outperforms LTR for most cases due to allowing for a four-way synchronization at all times.

While vehicle periods (i.e., bumper-to-bumper distances) between the two approaches are similar (e.g., $2S_{LTR}$ and $2S_{BRIP}$ for driving through on one lane), this consistent four-way synchronization gives BRIP an initial advantage in terms of throughput. However, with increasing vehicle lengths, S_{BRIP} also increases while S_{LTR} remains constant, which shifts the comparison in favor of LTR.

Summary. To summarize, BRIP provides very high throughput numbers and reasonable space requirements on the intersection for short (e.g., less than 5 m) and homogeneous vehicles. However, with increasing vehicle lengths, BRIP's space requirements increase drastically, which makes it rather unsuitable for most existing intersection.

In contrast to this, the proposed LTR requires substantially less space both for single-vehicle and platoon crossing alike. Moreover, when heterogeneous vehicles are allowed at the intersection, LTR's throughput starts outperforming BRIP. The greater the difference in length between the longest and the standard vehicle at the intersection, the higher LTR's throughput will be with respect to BRIP.

6 CONCLUDING REMARKS

In this article, we proposed a synchronous traffic protocol for intelligent intersection called LTR. In contrast to approaches from the literature, LTR is space efficient and, hence, it allows for heterogeneous traffic (in particular, vehicles with varied lengths) without modifying the existing road infrastructure. LTR divides the intersection into sectors of the same size, which does not depend on vehicle dimensions (i.e., in particular, length), but rather on the lane width. This way, long vehicles diverging from the standard case are treated as an exception, i.e., as overlength vehicles, and inter-vehicle separations are adapted accordingly.

In addition, we introduced a two-speed policy to differentiate between slower turn maneuvers and faster drive-through maneuvers without sacrificing throughput or passenger comfort. To the best of our knowledge, this is the first synchronous traffic protocol that allows for different speeds depending on the maneuver.

Furthermore, LTR can be parameterized to allow for platoon crossing in one direction and, thereby, increasing throughput. On the other hand, to control the waiting time of vehicles crossing in other directions, we introduced the concept of *maximum blocking threshold* that limits the number of cycles (i.e., the time) which a given platoon regime can be executed before enforcing a transition to let vehicles with conflicting trajectories cross.

Finally, we have included a detailed comparison of the proposed variants of LTR with BRIP, the most prominent traffic protocol from the literature, and could show that LTR improves the vehicle throughput under heterogeneous traffic without requiring changes to the existing road infrastructure.

REFERENCES

[1] Shunsuke Aoki and Ragunathan Rajkumar. 2019. CSIP: A synchronous protocol for automated vehicles at road intersections. *ACM Transactions on Cyber-Physical Systems* 3, 3 (2019).

[2] Shunsuke Aoki and Ragunathan Rajkumar. 2019. V2V-based synchronous intersection protocols for mixed traffic of human-driven and self-driving vehicles. In *Proceedings of the 2019 IEEE 25th International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA)*. IEEE.

[3] Tsz-Chiu Au, Shun Zhang, and Peter Stone. 2014. Semi-autonomous intersection management. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems*. International Foundation for Autonomous Agents and Multiagent Systems.

[4] Reza Azimi, Gaurav Bhatia, Ragunathan Rajkumar, and Priyantha Mudalige. 2015. Ballroom intersection protocol: Synchronous autonomous driving at intersections. In *Proceedings of the 2015 IEEE 21st International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA)*. IEEE.

[5] Masoud Bashiri and Cody H. Fleming. 2017. A platoon-based intersection management system for autonomous vehicles. In *Proceedings of the 2017 IEEE Intelligent Vehicles Symposium (IV)*. IEEE.

[6] Masoud Bashiri, Hassan Jafarzadeh, and Cody H. Fleming. 2018. PAIM: Platoon-based autonomous intersection management. In *Proceedings of the 2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE.

[7] Directive (EU) 2015/719 of the European Parliament and of the Council of 29 April 2015 - Maximum Authorised Dimensions in National and International Traffic. 2021. https://eur-lex.europa.eu/eli/dir/2015/719/oj. (July 2021).

[8] Kurt Dresner and Peter Stone. 2004. Multiagent traffic management: A reservation-based intersection control mechanism. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 2*. IEEE Computer Society.

[9] Kurt Dresner and Peter Stone. 2008. A multiagent approach to autonomous intersection management. *Journal of Artificial Intelligence Research* 31 (2008).

[10] Kurt M. Dresner and Peter Stone. 2007. Sharing the road: Autonomous vehicles meet human drivers. In *Proceedings of the International Joint Conferences on Artificial Intelligence (IJCAI)*, Vol. 7.

[11] European Vehicle Market Statistics - Pocketbook 2016/2017. 2017. http://www.theicct.org/sites/default/files/publications/ICCT_Pocketbook_2016.pdf. (June 2017).

[12] Amnesh Goel, Sukanya Ray, and Nidhi Chandra. 2012. Intelligent traffic light system to prioritized emergency purpose vehicles based on wireless sensor network. *International Journal of Computer Applications* 40, 12 (2012), 36–39.

[13] Matthew Hausknecht, Tsz-Chiu Au, and Peter Stone. 2011. Autonomous intersection management: Multi-intersection optimization. In *Proceedings of the 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE.

[14] Qiu Jin, Guoyuan Wu, Kanok Boriboonsomsin, and Matthew Barth. 2013. Platoon-based multi-agent intersection management for connected vehicle. In *Proceeings of the 16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*. IEEE.

[15] Harsha Vardan Maddiboyina and VA Sankar Ponnapalli. 2019. Mamdani fuzzy-based vehicular grouping at the intersection of roads for smart transportation system. *International Journal of Intelligent Engineering Informatics* 7, 2-3 (2019).

[16] Daniel Markert and Alejandro Masrur. 2019. Space-efficient traffic protocols for intelligent crossroads. In *Proceedings of the 2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE.

[17] Daniel Markert, Philip Parsch, and Alejandro Masrur. 2017. Using probabilistic estimates to guarantee reliability in crossroad VANETs. In *Proceedings of the 6th ACM Symposium on Development and Analysis of Intelligent Vehicular Networks and Applications*. ACM.

[18] Daniel Markert, Philip Parsch, and Alejandro Masrur. 2020. Impact of probabilistic vehicle estimates on communication reliability at intelligent crossroads. *Microprocessors and Microsystems* 78 (2020).

[19] Fraser McLeod and Nick Hounsell. 2003. Bus priority at traffic signals - evaluating strategy options. *Journal of Public Transportation* 6, 3 (2003), 1.

[20] Michael Zweck and Michael Schuch. 2013. Traffic light assistant: Applying cooperative ITS in european cities and vehicles. In *Proceedings of the 2013 International Conference on Connected Vehicles and Expo (ICCVE)*. IEEE, 509–513.

Received 16 July 2021; revised 29 April 2022; accepted 5 November 2022