Traffic Management Based on Negotiations between Vehicles

– a Feasibility Demonstration Using Agents*

Heiko Schepperle, Klemens Böhm, and Simone Forster

Institute for Program Structures and Data Organization (IPD)
Universität Karlsruhe (TH), Karlsruhe, Germany
`{schepperle,boehm,forster}@ipd.uni-karlsruhe.de`

Abstract. City governments currently make significant efforts to design traffic-control schemes. While existing solutions try to optimize the load of traffic resources, they do not consider one important aspect: Different drivers value short waiting times differently. For instance, the valuation of a truck driver who is part of a just-in-time production chain tends to be higher than the one of a family on an excursion for pleasure. Traffic-control schemes that give preference to drivers with higher valuations will increase the total utility of all drivers. In this paper we propose a new mechanism for traffic control at intersections. It is called *Time-Slot Exchange* and relies on bilateral negotiations. It tries to optimize the use of traffic resources in a valuation-aware fashion. The mechanism relies on agent-based driver-assistance systems with communication features which negotiate the right to cross an intersection at a certain time. Using simulations, we show that our new mechanism outperforms existing ones regarding average valuation-weighted waiting time.

1 Introduction

Although city governments currently make significant efforts to design traffic-control schemes, existing solutions are not sufficient [1]. They try to optimize the load of traffic resources, but this is only one aspect. Motorists typically have different valuations for shorter waiting times. A truck driver who is part of a just-in-time production chain typically has a higher valuation than a family on an excursion for pleasure. Traffic-control schemes that give preference to drivers with higher valuations will increase the total utility of all drivers. Approaches like traffic lights or main vs. side roads do not do this and do not even let drivers communicate their valuations to others.

The specifics of the traffic scenario make the design of valuation-aware mechanisms difficult. Our scenario is different from routing in communication networks, for various reasons: Vehicles have physical characteristics, e.g., acceleration or

---

* This work is part of the project DAMAST (Driver Assistance using Multi-Agent Systems in Traffic) (http://www.ipd.uni-karlsruhe.de/%7Edamast/) which is partially funded by init innovation in traffic systems AG (http://www.initag.com/).

speed limitations. Further, vehicles approaching the intersection from the same direction cannot overtake each other (*no overtaking assumption*). A left-turning vehicle prevents other vehicles from crossing the intersection at the same time. Because vehicles can arrive at the intersection at any time, the mechanism envisaged is embedded in a highly dynamic environment.

This article proposes a new mechanism for traffic control at intersections called *Time-Slot Exchange* (*TSE*). It is valuation-aware, i.e., it takes the valuations of waiting time into account. Progress in vehicle technology allows to implement mechanisms where drivers report their valuations to other drivers or to the infrastructure. *TSE* is not the first valuation-aware traffic-control mechanism – [2] has recently proposed such a mechanism called *Initial Time-Slot Auction* (*ITSA*), which is based on auctions. While *ITSA* increases total utility of drivers, it requires that all drivers use this mechanism. With *TSE* in turn the drivers can decide on their own if they want to use the mechanism or not. We deem this feature important, as it allows to counter any concerns regarding discrimination against 'poor drivers'. Fairness of such mechanisms appears to be an important concern of potential users. Further, *ITSA* lacks certain other desirable properties, as we will describe.

With *TSE*, vehicles can trade the right to cross an intersection for a certain period of time, the *time slot*, with other vehicles. When comparing *TSE* to other mechanisms, we have identified many requirements such a mechanism should fulfill, e.g., avoidance of starvation. We compare *TSE* to other mechanisms for traffic control based on these requirements. In particular, we compare it to the state-of-the-art *FIFO* mechanism described in [3, 4]. To compare the mechanisms regarding efficiency, we use simulations. For several reasons, e.g., open liability issues and inter-vehicle communication problems that remain to be solved, an evaluation of such mechanisms with real traffic is currently difficult to impossible. An important result of ours is that *TSE* outperforms *FIFO* both regarding average waiting time and average valuation-weighted waiting time. Heterogeneity of various kinds aggravates this effect significantly.

This paper has the following structure: In Section 2 we briefly discuss related work. Section 3 features definitions necessary to introduce the requirements on mechanisms for traffic control, as described in Section 4. After an overview of agent-based traffic control in Section 5 we describe our new mechanism in Section 6. An evaluation follows in Section 7, and we conclude in Section 8.

2 Related Work

In this section we briefly discuss approaches for agent-based traffic management which are important in our context. We will compare our mechanisms to the intersection-control system proposed by [3, 4], which lets vehicles reserve a time slot to cross the intersection in advance. The system only considers waiting time and is not valuation-aware.

[5] extends the reservation system by privileging emergency vehicles. The approach of [6] allows public transport vehicles to report their arrival to traffic

lights which can let them pass earlier. Both approaches prioritize only few vehicle types and ignore the valuations of all others.

As mentioned, [2] proposes another valuation-aware mechanism using auctions. [7] is a preliminary three-page version of this paper.

3 Definitions

To facilitate a description of desired properties of valuation-aware mechanisms for traffic control, we need some definitions, and we need to clarify some assumptions. The definitions are identical to the ones in [2].

In this work we limit ourselves to single intersections of roads. To take vehicles into account that are in the vicinity of the intersection, we define:

Definition 1. *The* neighborhood of an intersection *consists of the lanes of the intersection area, the incoming and the outgoing lanes.*

Definition 2. *The* travel time T_t^j *of a Vehicle* j *is the time from its first appearance in the neighborhood until it leaves the neighborhood. The* minimal travel time $minT_t^j$ *of* j *is the travel time if* j *was the only vehicle at the intersection, observed the speed limit and any constraints from the physical world, but ignored all rules concerning the right of way (i.e., crosses red lights, does not stop at stop signs, etc.). The difference of the travel time* T_t^j *and the minimal travel time* $minT_t^j$ *is the* waiting time $T_w^j = T_t^j - minT_t^j$.

We observe that the waiting time can never be negative.

Definition 3. *The* average waiting time *is* $\overline{T_w} = \frac{\sum_{j \in V} T_w^j}{|V|}$, *where* V *is the set of all vehicles.*

The average waiting time is widely used to evaluate intersection-control mechanisms. But to evaluate valuation-aware mechanisms, weighting the waiting time with the different valuations is more adequate.

Definition 4. *The* valuation $v^j(t)$ *of the driver of Vehicle* j *is the price the driver is willing to pay if he waits* t *seconds less.*

In what follows, we restrict ourselves to linear valuation functions (*linear valuation assumption*). Having said this, the valuation per second is a constant for each vehicle. We abbreviate it with v^j (j identifies the vehicle). v^j does not change over time.

Definition 5. *The* weighted waiting time *of Vehicle* j *is* $vT_w^j = v^j \cdot T_w^j$. *The* average weighted waiting time *is* $\overline{vT_w} = \frac{\sum_{j \in V} v^j \cdot T_w^j}{|V|}$, *where* V *is the set of all vehicles.*

Note that the average weighted waiting time does not depend on the price paid for an earlier time slot. In other words, if a 'rich' driver and a 'poor' driver are both short of time to the same degree, their valuation of reduced waiting time is the same and does not depend on the budget available to them.

Definition 6. *Let* b^j *denote the budget of the driver of Vehicle* j*. His* utility u^j *is* $u^j = b^j - v^j \cdot T_w^j$*. The* total utility U *is the sum of the utilities of the drivers:* $U := \sum_{j \in V} u^j$*, where* V *is the set of all vehicles.*

A driver cannot exceed his budget b^j. But this does not mean that the utility u_j of a driver cannot become negative. E.g., if its budget b^j is zero, the utility of a vehicle that must wait to cross the intersection becomes negative. Again, note that our definition of utility presumes that all motorists have the same valuation for money and only differ in the valuation of reduced waiting time.

We will use the sum of the budgets of drivers entering and leaving the neighborhood to assess the effects of the mechanisms on the budgets of the drivers.

Definition 7. *The* total entry budget B_e *is the sum of the budgets of all drivers when they enter the neighborhood:* $B_e := \sum_{j \in V} b_e^j$*.* V *denotes the set of all vehicles. The* total leaving budget B_l *is the sum of the budgets of all drivers when they leave the neighborhood:* $B_l := \sum_{j \in V} b_l^j$*.*

4 Desired Properties

To evaluate valuation-aware mechanisms we review various desirable properties of such mechanisms. [2] has described them originally, in order to evaluate *ITSA*.

A mechanism might offer incentives to drivers to report their valuations truthfully. This is a prerequisite for efficient allocations, i.e., a resource is assigned to the participant with the highest valuation.

Definition 8. *A mechanism is* incentive compatible *if reporting one's valuation truthfully is a dominant strategy, provided that all other do so as well [8].*

Efficiency, avoidance of starvation and zero-sum are desirable as well.

Definition 9. *If a traffic-control mechanism reduces the average weighted waiting time* $\overline{vT_w}$ *compared to a first-in first-out scheme we call such a mechanism* efficient*.*

We observe that an efficient mechanism increases total utility.

Definition 10. *A mechanism* avoids starvation *if it provides an upper bound of waiting time of vehicles when they arrive at the intersection.*

Clearly, this upper bound depends on the current number of vehicles in the neighborhood.

Motorists are likely to dislike mechanisms which increase the overall mobility costs, they would prefer mechanisms that are zero-sum.

Definition 11. *We call a mechanism* zero-sum *if the difference of the total entry budget* B_e *and the total leaving budget* B_l *is zero.*

We deem avoidance of starvation and zero-sum important. The reason is that these properties are prerequisites for a high degree of user acceptance.

A vehicle which is informed early about its time slot can adapt its speed in advance. The vehicle can avoid unnecessary acceleration and deceleration. This reduces fuel consumption and makes driving more comfortable. In particular, vehicles do not have to stop and start again. Thus, it is advantageous to allocate time slots as early as possible. We refer to this as *early allocation*. This property is in conflict with another desirable property: A mechanism which allocates time slots as late as possible fulfills the property *late allocation*. Late allocation is useful for vehicles with high valuations which arrive late. It lets them obtain time slots which would not be available any more otherwise. Late allocation is given if the time between the assignment of time slots to a vehicle and the point of time the vehicle starts to cross the intersection is minimized. Clearly, when designing a mechanism, one must position it somewhere between the two extremes early and late allocation.

5 Agent-Based Traffic Control

We have to build valuation-aware mechanisms upon an infrastructure which lets drivers communicate their valuations to each other. However, while driving, any distraction of the driver must be avoided. Therefore, the infrastructure must act autonomously (but in line with the preferences of the driver). Drivers should be able to configure vehicle agents in advance. To accomplish this, we use agent technology. Agents can act autonomously [9] to attain their goals, as specified by their drivers. To build the mechanism envisaged, every vehicle must be equipped with a platform which provides a standardized interface. One can use this interface to implement an agent-based driver-assistance system. There already exist applications which require such a platform in vehicles, e.g., the on-board units for trucks to collect toll on the German Autobahn (Toll Collect GmbH, `http://www.toll-collect.de`).

We call the agent hosted by a driver-assistance system *vehicle agent*. Such assistance systems can either instruct the driver or exert control over the vehicle. Our approach functions independently of the degree of interference: The driver-assistance system could just inform the driver when he may cross the intersection, or it could seize control of the vehicle, e.g., accelerate, as is the case with adaptive cruise-control systems (ACC, [10]).

An implementation of a mechanism relying solely on communication between agents residing in a vehicle seems rather ambitious for the time being. Therefore, our architecture (see Fig. 1) features stationary intersection-control units as well. Because agents can only communicate with other agents, these units have to host another agent, the *intersection agent*. It implements the intersection-control mechanism. In contrast to a vehicle agent which represents the driver, the goals

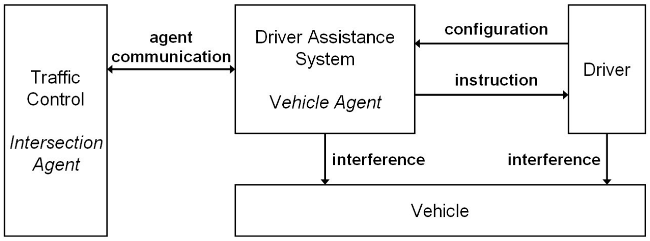

Fig. 1. Agent-based traffic control

of an intersection agent are defined by traffic planners. The intersection agent aims to increase total utility.

6 Mechanisms

The presentation of our mechanisms is based on the generic procedure for agent-based traffic control described in [2]. This procedure consists of the *initial time-slot acquisition phase* and the *subsequent time-slot acquisition phase*. In the initial time-slot acquisition phase vehicles with no time slot request one. The subsequent time-slot acquisition phase lets vehicle agents which already have a time slot obtain a better one from other vehicle agents by means of negotiation and exchange. A vehicle can trade its time slot repeatedly.

A mechanism can cover only one of the two phases, or it can be a combination of different mechanisms for the two phases and hence cover both phases.

In contrast to mechanisms for the initial time-slot acquisition phase, mechanisms for the subsequent phase have the nice characteristic that vehicle agents may, but do not have to participate. Only if the result of a negotiation increases their utility, vehicle agents will do so. This is important, as it allows to counter any concerns of potential users regarding discrimination against 'poor drivers'. Although discrimination could happen as part of the initial time-slot acquisition phase, we can at least exclude it for the subsequent phase.

In the following we describe two mechanisms for intersection control. The first one, proposed by [3, 4], is our reference point for evaluation. We then present our new valuation-aware *Time-Slot Exchange* mechanism. We study whether and to which extent the two mechanisms meet the desired properties, except for *efficiency*, which we evaluate in Section 7.

6.1 FIFO

The first mechanism uses a first-in first-out scheme and is therefore called *FIFO*. Vehicle agents entering the neighborhood interact with the intersection agent and request a time slot. The requests are answered in the order of arrival. Requesting vehicle agents do not have any time slot so far. Thus, *FIFO* belongs to the initial time-slot acquisition phase.

Different modes for the crossing of the intersection are conceivable. In contrast to [3, 4] we only distinguish three different modes. If no other vehicle may cross the intersection while a vehicle already does so we speak of the mode *intersection exclusive*. The mode *lane exclusive* prevents other vehicles from using intersection lanes which have already been assigned for the time slot in question. For instance, vehicles from opposite directions may cross the intersection at the same time if they both go straight. We achieve a more fine-grained granularity if we do not block the entire lane but only the actual spots of conflicts. Each spot of conflict is the intersection area of two lanes. We call this mode *lane shared all*. Several vehicles can cross the intersection in parallel if and only if there is only one vehicle per spot of conflict at any time.

Depending on the mode of intersection crossing, the intersection agent confirms a requested time slot if possible. If the requested time slot cannot be confirmed it looks for the earliest possible time slot after this and offers it to the requesting vehicle.

Compliance with Desired Properties. *FIFO* is not valuation-aware. Thus, incentive compatibility is not meaningful. Because a vehicle crosses the intersection after all other vehicles which have arrived before at the latest, *FIFO* avoids starvation. No vehicle pays or earns any money. This is why the budget of every vehicle does not change, i.e., *FIFO* is zero-sum. Because the only criterion is order of arrival, a vehicle arriving on short notice cannot benefit from late allocation. Thus, *FIFO* is best implemented using early allocation.

To our knowledge, no valuation-aware mechanism for intersection control exists so far, (except for [2] which is very recent). Thus, we will compare efficiency of our new valuation-aware mechanism to the one of valuation-unaware ones. *FIFO* reduces average waiting time compared to traffic lights [3]. This is why we use *FIFO* as the reference point in our evaluation. We cannot use valuation-aware mechanisms from other fields, like routing in communication networks, because they are subject to different physical characteristics. For instance, think of the restrictions incurred by a left-turning vehicle, in contrast to one that turns right.

Valuation-unaware mechanisms like *FIFO* can be combined with mechanisms for the next phase, the subsequent time-slot acquisition phase. If such mechanisms are valuation-aware, the combination is valuation-aware as well. We discuss such a combination in the following.

6.2 Time-Slot Exchange

If a vehicle agent has obtained a time slot in the initial time-slot acquisition, e.g., using *FIFO*, it may not be satisfied with it. We therefore propose a valuation-aware mechanism for the subsequent phase which can be combined with *FIFO*.

Idea. To increase the total utility of motorists, *Time-Slot Exchange* (*TSE*) lets vehicle agents negotiate the conditions for an exchange of their time slots. A vehicle agent having a late time slot may pay money to another vehicle agent and obtain its earlier time slot in return. The other vehicle agent receives the money and the time slot of the paying vehicle.

Example 1. If a vehicle agent offers 0.1 currency units per second of reduced waiting time, and the exchange takes place and reduces the waiting time by 12 seconds, the vehicle agent pays $0.1 \cdot 12 = 1.2$ currency units.

Exchange Agent. To realize *TSE*, we have extended the reference architecture of Fig. 1 with an *exchange agent* (EA). It brokers time-slot exchanges. I.e., its goal is to accomplish as many exchanges as possible. Since this goal differs from the goal of the intersection agent which aims to maximize total utility, we have modeled it as a separate agent.

Initiating Exchange. An exchange-request message sent by a vehicle agent contains the time slot currently held, the earliest time slot that is feasible, the price the vehicle agent is willing to pay for each second of reduced waiting time and the lane used to cross the intersection. The earliest time slot that is feasible is needed because a vehicle cannot make use of time slots that are too early and cannot be reached in time. The vehicle agent must specify all these parameters every time it initiates a time-slot exchange. The time slot currently held can permanently be subject to change. For instance, the earliest time slot feasible depends on the current speed and the current distance to the intersection, which may change. Even the price and the intersection lane to be used might change (even though we currently do not take this into account in our experiments).

Intersection Agent. Having received an exchange request, the exchange agent requests a list of exchange candidates (*list of candidates*) with suitable time slots from the intersection agent which knows all allocated time slots. At first sight, all time slots between the earliest time slot reported and the time slot currently held are suitable. But the intersection agent will return only vehicles if an exchange would not conflict with other reservations. This might be the case if the modes *lane exclusive* or *lane shared all* are used.

Contacting Vehicle Agents. The exchange agent iterates over the exchange candidates and informs them about the offer. Each vehicle agent has its own valuation of waiting-time reduction. Being offered a later time slot and a compensation, the potential partner now evaluates the offer. If the exchange increases its utility, i.e., the offered price per second exceeds its valuation of waiting-time reduction, the agent will agree to the exchange.

Conditions. Vehicles cannot overtake while waiting at the intersection. Thus, a vehicle agent which wants to acquire another time slot has to check two further conditions: (a) the new time slot is later than the one of the vehicle driving directly in front, and (b) the new time slot is earlier than the one of the vehicle driving directly behind if this vehicle already has a time slot. Otherwise, an exchange would harm other vehicles. They could not use their time slot any more. In our design, the exchange agent ensures these conditions for both vehicles participating in the exchange. Figure 2 shows an example of desirable, possible (but undesirable) and impossible exchange partners for a certain vehicle.

Execution. The *list of candidates* is sorted by the time slots held by the candidates. The first vehicle agent which agrees and which has a time slot which can be exchanged without harming others is the vehicle agent with the earliest suitable time slot. The exchange is executed with this vehicle.

TSE with Subsidies. [2] has shown for *ITSA* that subsidies, i.e., letting vehicles subsidize vehicles in front of them in their lane, increases efficiency significantly. We can transfer this idea to *TSE* as well: With *TSE* as described so far, the best a vehicle agent can do is to negotiate the next time slot after the time slot of the vehicle in front. This might limit efficiency. Think of a driver of a vehicle directly in front, subsequently referred to as l, with a low valuation. He does not compete for early time slots. The driver of the vehicle behind it with a high valuation, subsequently referred to as h, might not be satisfied with this. To improve the situation, the following refinement of *TSE* might help. The agent of Vehicle h can now subsidize the agent of Vehicle l. The agents of l and h would team up and request an exchange of the time slot of vehicle l. They would offer the sum of the amounts the two are willing to pay. If the exchange is accepted, the agent of Vehicle l would receive an earlier time slot, paid for by both. This now allows the agent of Vehicle h to acquire an earlier time slot for itself.

Compliance with Desired Properties. With *TSE*, neither the intersection agent nor the exchange agent pay or earn any money. Thus, *TSE* is *zero-sum*. On the other hand, pricing in *TSE* is not *incentive-compatible*. This is because a vehicle agent might offer less than its valuation to avoid unnecessarily high payments, or it might claim more than its true valuation to increase its profit. Thus, it might be possible that an exchange fails although both participating vehicle agents could benefit from it. However, a useful bilateral negotiation mechanism which is both incentive compatible and zero-sum does not exist [8]. Thus, incentive compatibility can only be achieved if we waive zero-sum and vice versa. Because a vehicle does not have to exchange time slots, *TSE avoids starvation* if the mechanism used for the initial time-slot acquisition phase has this characteristic as well. In other words, if a vehicle agent does not take part in an exchange it 'inherits' the upper bound of the waiting time from the mechanism for the initial time-slot acquisition phase. In particular, when using *FIFO* for this, the combination of *FIFO* and *TSE* avoids starvation. Finally, *late allocation* in our context means that the exchange agent looks for exchange candidates as long as possible, hoping for a vehicle arriving on short notice with a suitable time slot. Because vehicle agents initiating a request are not interested in time slots of such vehicles, and because all vehicle agents can initiate exchanges repeatedly on their own, late allocation does not have any advantages in our particular setting. Thus, *TSE* should always implement *early allocation*.

7 Evaluation

An evaluation of agent-based mechanisms in real traffic is not possible at the moment, for several reasons. We therefore use simulations to assess the mechanisms with regard to efficiency. This is in line with other research on intersection control, e.g., [3, 11, 12]. Our simulation framework uses a space-continuous and time-discrete simulation model. It simulates drivers and vehicles equipped with agent-based driver-assistance systems at intersections. We use the Java Agent

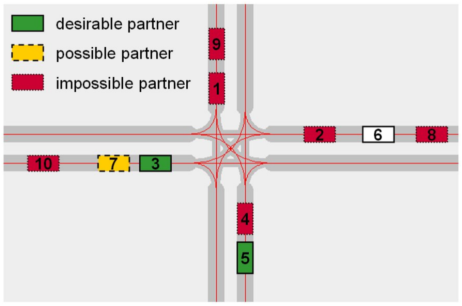

Fig. 2. Exchange partners of Vehicle 6

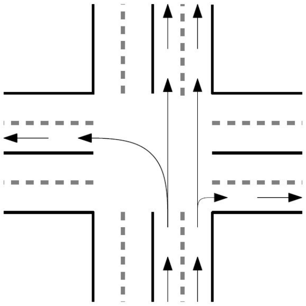

Fig. 3. Intersection layout

DEvelopment Framework (JADE, `http://jade.tilab.com/`) to implement the multi-agent system. We do not use agents for simulation, but simulate the traffic environment for agent-based driver-assistance systems. Thus agent-based simulation frameworks are not applicable.

7.1 Experimental Setup

The intersection investigated (see Figure 3) is symmetric. It consists of four directions with two incoming (right and left) and two outgoing (right and left) lanes for each direction. The left incoming lanes permit turning left or going straight (both into the left outgoing lanes) and the right incoming lanes permit turning right or going straight (both into the right outgoing lanes).

The speed limit is 50km/h in the entire neighborhood. This is the maximum speed allowed in cities in most European countries. Both incoming and outgoing lanes are 230m in length. 200m in front of the intersection, a virtual traffic sign informs driver-assistance systems that they approach an intersection which uses an agent-based intersection-control mechanism. The intersection radius is 20m. The total distance to cross the intersection in straight direction therefore is $230 + 20 + 20 + 230 = 500$m. The maximum acceleration is limited to at most 3m/s^2 and the maximum deceleration to at most 8m/s^2.

We simulate 40 minutes of each run. We leave aside vehicle behavior during the first 10 minutes to ensure that we only look at a system in steady state. Vehicles leaving the neighborhood during the following 30 minutes are taken into account to determine the average waiting time $\overline{T_w}$ and the average weighted waiting time $\overline{vT_w}$.

For each incoming lane so-called *demand sources* generate vehicles with exponentially distributed interarrival times with mean $\frac{1}{\lambda} = 36s$. Thus, 100 vehicles per hour enter the neighborhood for each incoming lane on average. To use an exponential distribution is common practice in queuing theory and in stochastic scheduling to describe arrival processes [13]. The intersection layout allows each vehicle to choose between two different directions to cross the intersection. In our simulations, both have the same probability. The valuation of each driver is also modeled using an exponential distribution. Its mean is $\frac{1}{\lambda} = \overline{v} = 0.01$. Since there are no experiences with valuation-aware mechanisms in real traffic we could not rely on any empirical data to determine this value.

800 vehicles per hour enter the neighborhood on average. Because we use the *intersection exclusive* mode for intersection crossing, only one vehicle crosses the intersection at a time. The crossing time per vehicle in our simulations is 4s. The total capacity of the intersection is 3600s/4s = 900 vehicles per hour.

The performance of a mechanism does not only depend on the average number of vehicles arriving at an intersection, but also on the distribution of interarrival times. A comparison of mechanisms is meaningful only if we compare equal arrival patterns of vehicles. We execute 25 simulation runs of *FIFO* and of *TSE* for each treatment. We use the same 25 randomly chosen numbers as seeds for the 25 runs, to compare the i^{th} run of *FIFO* to the i^{th} run of *TSE* of the same treatment. This means that a certain vehicle always has the same start time and route in the i^{th} simulation run.

7.2 Alternatives Evaluated

Our mechanism combines *FIFO* for the initial time-slot acquisition phase and *TSE* for the subsequent phase. We have not implemented the variant with subsidies of *TSE*, which lets vehicle agents subsidize vehicles in front. In the remainder of this section we use the term *TSE* for the combination of *FIFO* and *Time-Slot Exchange*.

With our implementation of *TSE*, the vehicle agent trying to initiate an exchange reveals its valuations truthfully. This leads to efficient allocations. Other strategies to determine which price to offer could be implemented as well. However, we have already mentioned that the price of an exchange does not influence our efficiency measure. Therefore, other strategies would not be meaningful.

If the budget is too low it could happen that a vehicle agent cannot afford to offer a price which corresponds to its true valuation. This would lead to biased results. To avoid this we assume in our simulations that all vehicles have unlimited budgets. In other words, all motorists have the same budget and the same valuation of money.

As mentioned before, we expect that *TSE* will have greater impact on average-weighted waiting time if there are vehicles which arrive late and have high valuations. To verify this hypothesis, we carry out four different treatments, as listed in Table 1. In each treatment we simulate both *FIFO* and *TSE* with the same parameter settings. More specifically, we examine the influence of two parameters: First, for each incoming lane the demand source generates vehicles with different valuations of waiting time, using an exponential distribution. The *mean valuation* $\overline{v}$ is the first parameter examined. Second, a virtual traffic sign marks the position where arriving vehicles can interact with the intersection agent for the first time. We denote the distance of this traffic sign to the beginning of the intersection as *begin-of-interaction distance* s. This is the second parameter examined.

Mean Valuation $\overline{v}$. Although the valuation for a vehicle is already randomly distributed, initializing a subset of the demand sources with a different mean valuation allows to examine the influence of valuations which are significantly different depending on the incoming lane of a vehicle. In the simulations with differing mean valuations the demand sources of the northern and southern incoming lanes spawn vehicles with the increased mean valuation $\overline{v} = 0.1$, while in all other cases the mean valuation is $\overline{v} = 0.01$ (see also Table 1).

Begin-of-Interaction Distance s. We evaluate the relationship between the time available for interacting with the intersection agent and the weighted waiting time by introducing vehicles arriving on short notice, i.e., vehicles which have less time to interact with the intersection agent. These vehicles will typically receive a later time slot than vehicles which are farther away. We mimic vehicles arriving on short notice by positioning the virtual traffic sign of one direction closer to the intersection. In the treatments with vehicles arriving on short notice the distance of the virtual traffic sign is $100m$ to the intersection for the North and the South, while it is 200m in the other cases (see Table 1).

Varying these parameters allows to evaluate the influence of heterogeneity. Our evaluation compares a treatment where one or two parameters are equal for every direction to one where the parameters have a different value for the northern and southern incoming lanes than for the western and eastern directions. This may be realistic, e.g., if we assume that side roads have a different infrastructure than main roads. We label simulations which have different values for a parameter depending on the direction with the name of the parameter as an index. Table 1 lists the treatments executed and shows the differing values of each parameter examined. The first value of each pair is the value of the parameter for the northern and the southern direction, and the second value is the one for the western and eastern direction.

Table 1. Executed treatments

treatment	**valuation** $\overline{v}$	**distance** s
FIFO, *TSE*	0.01/0.01	200/200
$FIFO_s$, TSE_s	0.01/0.01	100/200
$FIFO_{\overline{v}}$, $TSE_{\overline{v}}$	0.10/0.01	200/200
$FIFO_{s;\overline{v}}$, $TSE_{s;\overline{v}}$	0.10/0.01	100/200

7.3 Results

We compare the differences of average waiting time $\overline{T_w}$ and of average weighted waiting time $\overline{vT_w}$ of equally initialized runs of *FIFO* and *TSE*. For all pairs of the 25 simulation runs we can compute the mean and the 95% confidence interval (CI) of such differences. Table 2 lists the absolute and the relative differences of the average waiting time $\overline{T_w}$ and the average weighted waiting time $\overline{vT_w}$. The

Table 2. Absolute and relative differences

			T_w		vT_w	
	treatment		**mean**	**0.95% CI**	**mean**	**0.95% CI**
absolute difference	*FIFO*	*TSE*	0.567	[0.431,0.703]	0.010	[0.008,0.012]
	$FIFO_s$	TSE_s	0.658	[0.451,0.866]	0.013	[0.010,0.015]
	$FIFO_{\overline{v}}$	$TSE_{\overline{v}}$	0.626	[0.439,0.814]	0.066	[0.050,0.082]
	$FIFO_{s,\overline{v}}$	$TSE_{s,\overline{v}}$	1.150	[0.856,1.444]	0.154	[0.131,0.177]
relative difference	*FIFO*	*TSE*	0.040	[0.028,0.052]	0.073	[0.056,0.091]
	$FIFO_s$	TSE_s	0.041	[0.029,0.053]	0.083	[0.067,0.099]
	$FIFO_{\overline{v}}$	$TSE_{\overline{v}}$	0.045	[0.031,0.058]	0.089	[0.067,0.112]
	$FIFO_{s,\overline{v}}$	$TSE_{s,\overline{v}}$	0.072	[0.055,0.090]	0.157	[0.130,0.184]

absolute difference of $\overline{T_w}$ is computed as the difference of $\overline{T_w^{FIFO}}$ and $\overline{T_w^{TSE}}$. The relative difference is the absolute difference divided by $\overline{T_w^{FIFO}}$. We compute the differences of the average weighted waiting times analogously.

TSE reduces the average weighted waiting time $\overline{vT_w}$ by 0.010 units on average compared to *FIFO*, i.e., by 7.3%. The 95% confidence interval shows the reliability of these results. The relative difference is between 5.6% and 9.1% in 95% of all cases. *TSE* has a smaller but still positive influence on the average waiting times. The average waiting time is reduced by 0.567s on average, i.e., 4.0% compared to *FIFO*.

Compared to the treatments with *FIFO* and *TSE* the treatments with only one differing parameter ($FIFO_s$ and TSE_s, and $FIFO_{\overline{v}}$ and $TSE_{\overline{v}}$) only yield a slight improvement of average waiting time and of average weighted waiting time. The improvement results from a higher number of executed exchanges in the case of TSE_s and in a higher gain by a single exchange in the case of $TSE_{\overline{v}}$. This is not surprising, since TSE_s increases the chance that a vehicle receives a late time slot but could enter the intersection much earlier. With $TSE_{\overline{v}}$ the difference between the valuations of exchange partners tends to be larger. For $FIFO_{s,\overline{v}}$ and $TSE_{s,\overline{v}}$ in turn, i.e., the treatment both with differing average valuation and differing begin-of-interaction distance, the improvement is significant. The average weighted waiting is reduced by 15.7% and the average waiting time by 7.2% on average.

Our new mechanism reduces both average waiting time and average weighted waiting time, even though only 7.4% of the vehicles trade their time slot in the treatments *TSE* or $TSE_{\overline{v}}$. With the treatment TSE_s, this ratio is 8.7%. But for the treatment $TSE_{s,\overline{v}}$ it increases to 14.2%. Our explanation is that the vehicles arriving on short notice from the North or the South with high valuations have a much greater chance to find an exchange partner.

The significant reduction of average weighted waiting time with $FIFO_{s,\overline{v}}$ and $TSE_{s,\overline{v}}$ shows that the heterogeneity of the scenario influences the expected outcome of a valuation-aware mechanism. For *Time-Slot Exchange*, this is true for vehicles arriving on short notice with high valuations in particular.

8 Conclusions

Existing traffic-control schemes are not valuation-aware, even though this is expected to increase the total utility of motorists. In this paper we present the novel mechanism *Time-Slot Exchange* which has this characteristic. Using the state-of-the-art *FIFO* mechanism as a reference point for evaluation, *Time-Slot Exchange* reduces both average waiting time and average weighted waiting time significantly, in particular in the presence of vehicles arriving on short notice with high valuations.

References

1. Schrank, D.L., Lomax, T.J.: The 2005 Urban Mobility Report. Report, Texas Transportation Institute, College Station, TX (2005)
2. Schepperle, H., Böhm, K.: Agent-Based Traffic Control Using Auctions. (submitted) (2007)
3. Dresner, K., Stone, P.: Multiagent Traffic Management: A Reservation-Based Intersection Control Mechanism. In: AAMAS 2004, IEEE Computer Society (2004) 530–537
4. Dresner, K., Stone, P.: Multiagent Traffic Management: An Improved Intersection Control Mechanism. In: AAMAS 2005, ACM (2005) 471–477
5. Dresner, K., Stone, P.: Human-Usable and Emergency Vehicle-Aware Control Policies for Autonomous Intersection Management. In Bazzan, A.L.C., Chaib-draa, B., Klügl, F., Ossowski, S., eds.: Fourth International Workshop on Agents in Traffic and Transportation (ATT), Hakodate, Japan. (2006)
6. Greschner, J.T., Gerland, H.E.: Traffic Signal Priority Tool To Increase Service Quality And Efficiency. In: Proceedings APTA Bus & Paratransit Conference, Houston, TX. (2000) 138–143
7. Schepperle, H., Böhm, K., Forster, S.: Towards Valuation-Aware Agent-Based Traffic Control. In: AAMAS'07, Honolulu, Hawai'i, USA (2007)
8. Krishna, V.: Auction Theory. Academic Press (2002)
9. Wooldridge, M.J.: An Introduction to Multiagent Systems. Wiley (2002)
10. Bareket, Z., Fancher, P.S., Peng, H., Lee, K., Assaf, C.A.: Methodology for Assessing Adaptive Cruise Control Behavior. IEEE Transactions on Intelligent Transportation Systems **4**(3) (September 2003) 123–131
11. Fuerstenberg, K.C., Lages, U.: New European Approach for Intersection Safety - The EC-Project INTERSAFE. In: Intelligent Vehicles Symposium, 2005. Proceedings. IEEE. (2005)
12. Hopstock, M., Ehmanns, D., Spannheimer, H.: Development of Advanced Assistance Systems for Intersection Safety. In: AMAA 2005, 9th International Conference on Advanced Microsystems for Automotive Applications, Berlin, Germany. (2005)
13. Pinedo, M.: Scheduling – Theory, Algorithms and Systems. 2nd edn. Prentice Hall (2002)