An Intersection-Centric Auction-Based Traffic Signal Control Framework

Jeffery Raphael, Elizabeth I. Sklar, and Simon Maskell

Abstract Vehicular traffic on urban road networks is of great interest to those who monitor air quality. Combustion emissions from transport vehicles are a major contributor of air pollution. More specifically, the release of fine particulate matter which has been linked to premature deaths. Travel and idle time are two factors that influence the amount of pollution generated by traffic. Reducing idle and travel times would have a positive impact on air quality. Thus, it is increasingly crucial to manage intersections effectively, particularly in congested cities and across a range of different types of traffic conditions. A variety of market-based multi-agent traffic management mechanisms have been proposed to improve traffic flow. In many of these systems drivers "pay" to gain access to favourable road ways (e.g., minimise travel time). A major obstacle in adopting many of these mechanisms is that the necessary communication infrastructure does not yet exist. They rely on vehicle-to-infrastructure and/or vehicle-to-vehicle communications. In this work, we propose a market-based mechanism which relies on existing technology (and in some places this technology is already in use). Experimental results show that our market-based approach is better at reducing idle and travel times as compared to fixed-time signal controllers.

1 Introduction

London's *Great Smog* occurred over 60 years ago. It is estimated that over 4000 people died prematurely due to health complications brought about by the dense mix of smoke and fog [4]. The smog was a result of coal burning (most notably

---

J. Raphael (✉)
Department of Computer Science, University of Liverpool, Liverpool, UK
e-mail: jeffery.raphael@liverpool.ac.uk

AQ1
AQ2

E.I. Sklar
Department of Informatics, King's College London, London, UK
e-mail: elizabeth.sklar@kcl.ac.uk

S. Maskell
Department of Electrical Engineering & Electronics, University of Liverpool, Liverpool, UK
e-mail: s.maskell@liverpool.ac.uk

A. Alonso-Betanzos et al. (eds.), *Agent-Based Modeling of Sustainable Behaviors*,
Understanding Complex Systems, DOI 10.1007/978-3-319-46331-5_6

for heating), industrial practices of the time, and atmospheric conditions. Even though London had experienced bouts of smog in the past, the *Great Smog* launched air pollution into the spotlight as a major health hazard. Decades later, despite a number of legislative acts to curb the presence of noxious gases and particulate matter, London continues to exceed the allowable limits defined for air pollution standards [12]. There is strong evidence linking air pollution to increases in mortality and morbidity (health issues such as low birth weight, strokes and heart diseases [12]) rates [4, 16]. Motor vehicles have been identified as a major source of air pollution (specifically nitrogen dioxide and particulate matter smaller than 2.5 μm) [12, 34]. The amount of combustion emissions are directly related to long idle and travel times which can be exacerbated in congested cities. In England and Wales, over 60 % of commuters drive to work [11]. In many cities, traffic congestion does not appear to deter travelling by car. For example, over a quarter of Londoners still choose to drive to work [11], despite having access to alternative modes of transportation. Traffic congestion does not just pose a health hazard, it also costs money. Across the UK, close to £3.76 billion (€4.94 billion) [27] are lost from fuel and increased cost of delivering goods. The UK is not alone in its struggle with traffic congestion and its effect on air quality. Other cities, such as Beijing, Los Angeles, and Delhi, have similar problems. Any effort to improve travel times (and traffic throughput) would have significant benefits to air quality and a financial impact as well.

Today, traffic managers (people) are responsible for the prevention and/or reduction of traffic congestion. Setting speed limits, installing road signs, speed humps, islands, implementing movement and parking restrictions are some of the tools at the disposal of traffic managers. This includes the authority to enact transportation policies aimed at improving traffic flow. Traffic managers must be proactive and identify potential network hazards and/or issues that may disrupt traffic flow. While most measures exist for the improvement of vehicular traffic, some are put in place to maximise the synergy between pedestrian and vehicle movements. In addition, some cities now include cyclists in their traffic management schemes, with designated cycle lanes and traffic signals. Overall, traffic managers utilise their power to control different components of the transportation infrastructure to ensure the safe and efficient use of road networks.

Traffic signals (or “lights”) are one of the most common tools employed around the world to control traffic at intersections. Different municipalities have different criteria for the use of traffic signals. Traffic demand, number of conflicting vehicle manoeuvres and general delays are the most common reasons for installing traffic signals. Traffic signals manage conflicting movements by dictating which vehicle movement(s) are allowed at a given moment. A driver approaching an intersection may arrive only to discover her vehicle is not allowed through and must then wait. Although traffic signals are installed to improve the safety and flow of traffic, they nonetheless become a major source of delay in road networks.

Traffic signals are programmed to permit small groups of non-conflicting vehicle movements through an intersection for a short period of time, followed by another group of movements. Traditional traffic signals (and the most basic deployment method) repeat the same sequence of vehicle movements without changing their duration or order of execution—regardless of what is happening in real time at the intersections they control. Traffic signal timings specify the duration and order of vehicle manoeuvres at an intersection [22]. Appropriate traffic signal timings, or *phase plans*, are essential to the proper function of traffic signals. Poorly designed phase plans may lead to additional delays, traffic jams and even accidents. Although many traffic signals rely on simple fixed protocols, they are nonetheless a vital component of traffic management [1]. Historically, finding the best signal timings involved using mathematical models of traffic behaviour to determine ideal settings [22]. It later became possible to develop even better signal timings through the analysis of historical traffic data. Adaptive *Urban Traffic Controllers (UTCs)* employ information about current road conditions and determine, some in real time, the best signal settings. Adaptive UTCs attempt to harmonise the interplay between all aspects of traffic (private vehicles, public transportation, cyclist and pedestrians) in areas ranging in size from a few city blocks to entire cities. Adaptive centralised systems have been developed that apply optimisation algorithms, such as RHODES [17], OPAC [9] and SCOOT [18, 31].

In general, the traffic control problem can be stated as that of finding a policy for setting traffic signal states such that traffic flow improves while the safety of an intersection is maintained and conflicts amongst movements are resolved (including pedestrian and cyclist movements). Such a policy could take into consideration traffic conditions at the intersection and could also incorporate information from neighbouring intersections. The policy should determine which movement(s) are allowed at any given moment in time. An optimal policy could minimise travel time. Or, it could attempt to optimise other aspects of traffic, such as number of stops per vehicle or *queueing* time (i.e., how long a vehicle waits at a particular intersection before it is able to pass, which might mean waiting through multiple signal phases if traffic is heavy).

Finding the optimal traffic signal timings is a non-trivial operation for a number of reasons:

- Traffic control is geographically distributed, takes place in a dynamic environment and the interactions amongst its components are highly complex [6].
- Traffic signal timings function under rigid temporal constraints which may be represented as discrete variables. Therefore, traffic control behaves in many regards like a combinatorial optimisation problem (e.g., TSP) [19].
- Scale is always an issue with traffic control. Any reasonably sized road network will have dozens of intersections, compounding the problem of finding an optimal traffic signal timing [3, 19, 33].
- Adaptive traffic control systems that work in real time must find a solution within a very small time window in order to function properly.

Traffic consists of many independent components that are interconnected in a highly complex manner. There are vehicles, pedestrians, cyclists and traffic control devices, to name a few of the elemental components. Using mathematical models, it is difficult to capture the interacting behaviours of these individual components;

however, modelling them as a large collection of *autonomous agents* allows us to apply a wide range of methodologies designed to investigate the interplay between independent entities. For this reason, the *Multi-Agent Systems (MAS)* paradigm offers an ideal method for modelling the critical elements of traffic behaviour. The advantage of using a MAS approach over traditional mathematical models is two-fold. One, MAS does a better job of modelling the stochastic nature of traffic. Two, better models means better platforms to investigate novel solutions to traffic control. The MAS paradigm offers a flexible and inexpensive method for designing traffic control solutions [31]. There is a plethora of traffic control solutions that fall under the umbrella of MAS. Traffic control systems that are developed within a MAS framework are also easier to maintain and scale [31]. As well, the traffic domain offers many interesting challenges from a multi-agent systems perspective.

Our work focuses on solutions that utilise *market-based mechanisms*. Traffic control can be viewed as the management of a set of traffic signals in a road network in order to minimise, for example, the delay experienced by vehicles traversing the network. If traffic signals are considered agents, traffic control can be viewed as a *coordination* problem [2] where traffic signals work together to prevent congestion and keep traffic flowing. This perspective on traffic control is important as it drives our MAS design choices and sets us apart from other MAS solutions to traffic control thus far, as explained next. We propose using *auctions* to achieve coordination amongst traffic signal agents by providing a framework for the resolution of conflicts and enhancement of cooperation amongst traffic signal agents.

Our approach for controlling traffic signals has been greatly influenced by coordination efforts in *Multi-Robot Routing (MRR)* [7, 10, 13, 23]. Auctions are a form of market mechanism for resource allocation, and they can produce near optimal results in some MRR scenarios [15]. In MRR, auctions have been used to facilitate coordination amongst robots [26]; thus the same can be done with traffic signals. A common theme in the existing literature on auction-based traffic controllers is the need for a *vehicle agent*, which refers to a vehicle-borne software system responsible for tasks ranging from simple *vehicle-to-infrastructure (V2I)* communication to more demanding vehicle navigation and control. We believe that auctions can empower agents, acting either locally at a single intersection or in small groups of connected intersections, to find local solutions to traffic congestion that then emerge as global improvements in traffic performance.

There are a number of significant issues with regard to the widespread deployment of current market-based approaches to traffic management. The first issue is the development and distribution of vehicle agents. Car manufacturers will have to agree on international communication protocols, physical specifications and many other aspects of deploying vehicle agents to the millions of vehicles that are currently in use. Second, there is the current state of the transportation infrastructure worldwide. The communication systems necessary for *V2I* communication currently does not exist; and a range of issues, such as security and privacy, remain unaddressed in the traffic management domain. Lastly, there is the concept of drivers bidding for intersection usage, which introduces the issue of *fairness*. Fairness is a general term for such questions as: *Which drivers will have to pay?* and *How much will drivers have to pay?* Our overarching goal is to design a system that reaps the benefits of market mechanisms, but without its less appealing features such as driving fees and V2I communication requirements. Our approach not only does away with vehicle agents, but also does not have drivers bidding for intersection usage; hence, our approach allows us to utilise auctions without having to consider fairness at the level of each vehicle.

In this paper, we describe our theoretical and experimental work on multi-agent auction-based traffic control mechanisms. As above, our mechanisms utilise auctions without the need for vehicle agents. We demonstrate how such a system could be designed and implemented, and we ran a series of experiments to measure the effectiveness of our mechanisms. Three empirical evaluations found that our mechanisms perform better than fixed-time signal controllers. The remainder of this paper is organised as follows. Section 2 discusses other auction-based approaches to traffic control, focussing on the MAS literature. Section 3 presents our approach. Sections 4 and 5 describe our experiments and results. Finally, we close with some discussion (Sect. 6) and conclusions (Sect. 7).

2 Related Work

Dresner and Stone [8] designed a reservation-based traffic management system to reduce traffic congestion. In a reservation-based system, vehicles request time slots. The time slots are time spans when the vehicle is allowed to occupy the intersection. The reservation-based system functions on a *first-come, first-served* basis and relies on vehicle agents (autonomous cars) that have complete control of the vehicle. The authors measured the delay experienced by vehicles passing through the intersection. Dresner and Stone [8] compared their reservation-based approach to two other traffic control schemes: *overpass* and *traffic light.* "Overpass" simulates a road network with no signals; roads cross each other via bridges. "Traffic light" simulates how current fixed-time traffic signals function. Dresner and Stone [8] found that their reservation-based system did not just outperform normal traffic lights but under certain conditions eliminated delay due to intersection crossing.

Vasirani and Ossowski [30] expanded on Dresner and Stone's work and examined the performance changes to a reservation-based system where time slots were allocated using a market mechanism. The authors also proposed a market-based traffic assignment scheme using the same reservation-based system. In Dresner and Stone [8], reservations were allotted using a *first-come first-served* policy or FCFS. Vasirani and Ossowski [30] replaced FCFS with *a combinatorial auction (CA).* As drivers approach the intersection, reservations are "won" through the auction, instead of simply handed to the next arriving vehicle. In this way, a driver may express its true valuation for a contested reservation. For the market-based traffic assignment system, Vasirani and Ossowski [30] devised a protocol where route selection was accomplished through a *combinatorial traffic assignment (CTA)*. The cost of passing through an intersection continually changed depending on demand. In turn, these costs caused vehicles to select alternative or cheaper routes.

In a network with a single intersection, Vasirani and Ossowski [30] looked at the delay experienced by drivers based on the amount they were willing to "pay" to use the intersection. They were interested in finding out if drivers willing to pay more would experience less delay. They also looked at the delay experienced as traffic volume increased across the intersection. Vasirani and Ossowski [30] found that initially having a willingness to pay does decrease delay, but eventually this levels off. However, CA was found to increase overall delay. As the intensity of traffic increases, CA experienced far more delays and rejected reservations than FCFS.

The performance of CTA was studied using a simulation of a simplified road network of Madrid, Spain. The authors examined the density (number of vehicles per kilometre of roadway) and travel time to measure its performance. Vasirani and Ossowski [30] found that CTA, which used FCFS, produced a more balanced network, i.e., vehicles were better distributed throughout the road network. As both CA and CTA are extensions of [8], they too rely on vehicle agents and thus are infeasible with current transportation infrastructure.

Although Dresner and Stone's work does not directly employ a market mechanism, it does represent the state-of-the-art in terms of futuristic visions of traffic control. The reservation-based traffic management system [8] (and Vasirani and Ossowski [30] market based derivative) requires the greatest advancements in current transportation infrastructure: V2I communications and autonomous cars. Our approach on the other hand, does not require neither V2I communications nor autonomous cars; although the former could be used to improve the performance of our mechanism.

Carlino et al. [5] described a traffic control system where auctions are run at intersections to determine use. Vehicles are embedded with a software agent (the *wallet agent*) which bids on behalf of the driver. A *system agent* also bids in a manner that facilitates traffic flow beneficial to the entire transportation system—while the *wallet agent* is solely concerned with getting its occupants to their destination in the least expensive (and quickest) way. The *wallet agent* is assigned a budget to pay for trips. Carlino et al. [5] used a second-price sealed bid auction mechanism. They tested four different modes: *FIFO* (this is how a typical intersection works), *Equal* (every driver submits a bid of one), *Auction* (drivers bid an amount equal to their account balance divided by the number of intersections remaining on their trip), and *Fixed* (drivers always bid the same amount based on the value they've assigned for the trip). The authors evaluated their traffic control mechanisms in four simulated urban cities. FIFO performed the worst in three of the four cities. *Auction* (with and without the *system agent*) had the best performance. There are two important distinctions between our work and [5]. First, Carlino et al. [5] assumes vehicles have specialised software that allow drivers to effortlessly participate in the auction; we do not need require any such software. Second, although we utilise auctions in our approach, in our work the auction provides a framework for coordination and is not monetised.

Schepperle and Böhm [24] describes an intersection controller called *Initial Time-Slot Auction* (ITSA) which is *valuation-aware* (meaning the controller considers the driver's value of reducing wait time). In ITSA, as vehicles approach an intersection, they *register* with the intersection. The *intersection agent* then executes a second-price sealed-bid auction for the most current time slot that's available for usage. Here a time-slot is a window in time where a vehicle may safely use the intersection. Only the vehicles at the very front of the traffic queue participate in the auction. Schepperle and Böhm [24] utilised the FIPA Contract Net Protocol to implement the auction. Schepperle and Böhm [24] also described two variants of ITSA. In the first variant, a mechanism is included to prevent *starvation* where auctions are suspended if vehicle waiting time has reached some fixed limit. Starvation is defined as the situation where traffic is prevented from flowing in a particular direction. The other variant, ITSA+Subsidies, considers subsidies where vehicles that have not participated in an auction yet can influence the auction of the vehicles in front of them. In the subsidised variant, vehicles boost the bid of a candidate vehicle (a vehicle in front of theirs). If the candidate vehicle wins the auction, then the vehicle that subsidised its auction would be able to participate (attain a time-slot) in an auction sooner.

Schepperle and Böhm [24] used *waiting time* to measure performance. The authors defined waiting time as the difference between actual travelling time and the minimum travel time. Schepperle and Böhm [24] also examined *average weighted waiting time* where the weighted waiting time is the product of the waiting time and the driver's valuation of a reduced waiting time. They compared their traffic controller to the reservation-based system in Dresner and Stone [8]. Both ITSA and ITSA+Subsidies were able to reduce average travel time while minimising average weighted waiting time compared to FIFO, although ITSA+Subsidies was better at reducing average weighted waiting time. Drivers that had the lowest valuations, that is those drivers that did not mind waiting, fared better under ITSA+Subsidies than ITSA.

In follow-on work, Schepperle and Böhm [25] created a *valuation-aware* traffic-control mechanism which allows concurrent use of the intersection through an auction mechanism. In a valuation-aware traffic controller, the intersection takes into account the driver's value of time; but many of these systems do not allow concurrent use of the intersection. Schepperle and Böhm [25] propose two auction-based mechanisms: *Free Choice* and *Clocked.* In Free Choice, the auction winner gets to select the time slot it wants from an interval; while in Clocked, time slots are auctioned off. Schepperle and Böhm [25] concluded that Free Choice reduced the average weighted wait time by up to 38.1 %. Clocked reduced the average weighted wait time for only lower degrees of concurrency and high traffic volume. Similar to [5], Schepperle and Böhm [24, 25] assumes that cars have a vehicle agent. Again, our approach, detailed in the next section, does not involve vehicle agents or other embedded software.

Bazzan [2] constructed a decentralised method of traffic control that utilises Evolutionary Game Theory. The traffic controller facilitates coordination among intersections while minimising communication overhead. *Intersection agents* coordinate by selecting the same action (phase plan). Different phase plans favour different sets of vehicle movements. Intersection agents function in two states: local and global. In the local state, intersection agents use local information and a mixed strategy to make action selections. Intersection agents also have a payoff function which is used to update its mixed strategy. Intersection agents experience a learning phase which allows them to update their mixed strategy taking into account most recent payoffs over past payoffs. In [2], there is an entity, "Nature", that has a global view of traffic and is able to see (and process) information from a macroscopic level. Nature recognises global traffic changes and initiates the change from local to global state in intersection agents. While in the global state, intersection agents use a payment function, given to them by Nature, to update their mixed strategy.

Bazzan [2] set up a traffic scenario with vehicles travelling through a roadway with several intersections. Bazzan [2] evaluated her traffic control mechanism under three different traffic conditions. The author compared her method to a centralised traffic controller where a central computer determined the best phase plan for the traffic signals. In the case of the centralised controller, the best phase plan is the one that produced the least delay for the traffic flow (going east or west) with the heaviest volume. The author used traffic density (discretised) to measure performance. The agent-based decentralised method performed better than the centralised method in two of the traffic scenarios. In the first, east and west bound traffic had medium to high volumes of traffic. While in the second, both directions had medium to low levels of traffic. In contrast, our approach is fully decentralised at the intersection level. It does not require a global perspective (i.e., Nature) of traffic flow. Finally, in [2], intersection agents select phase plans from a closed set of phase plans while in our approach we focus on fine tuning a single phase plan in lieu of replacing it.

3 Our Approach: Multi Agent Auction Based Traffic Signalling

For our multi-agent traffic controller, we decompose each intersection into two types of agents: *intersection agents* and *traffic signal agents*. At an intersection, there is a single intersection agent and multiple traffic signal agents. The intersection agent is responsible for making adjustments to traffic signal timings and ensuring that those changes do not violate any basic traffic regulations (e.g., minimum *green times*). Traffic signal agents, on the other hand, operate on behalf of a small set of legal vehicle movements that may occur at the intersection. That is, each traffic signal agent is assigned a number of movements to manage. The traffic signal agents compete against each other for control over traffic signal timing adjustments. An intersection agent and its associated traffic signal agents work together at the intersection level to adapt signal timings in real time. The adjustments are made to improve the efficiency of the intersection and maintain its safety. Figure 1 illustrates the key components of our multi-agent traffic controller and how they are used on a global scale (image on the right) to manipulate traffic flow.

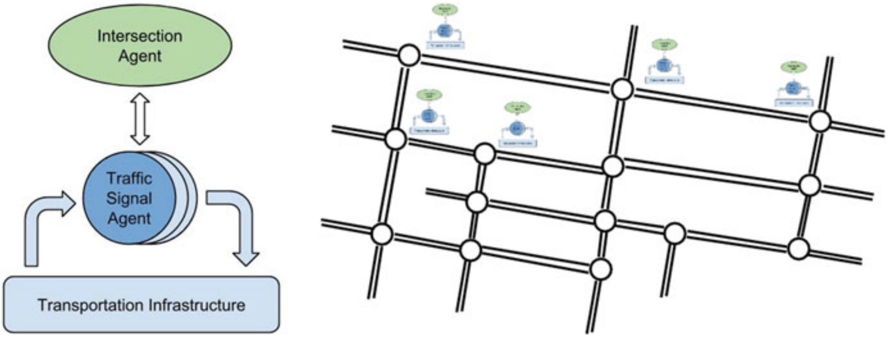

Fig. 1 Multi agent intersection controller

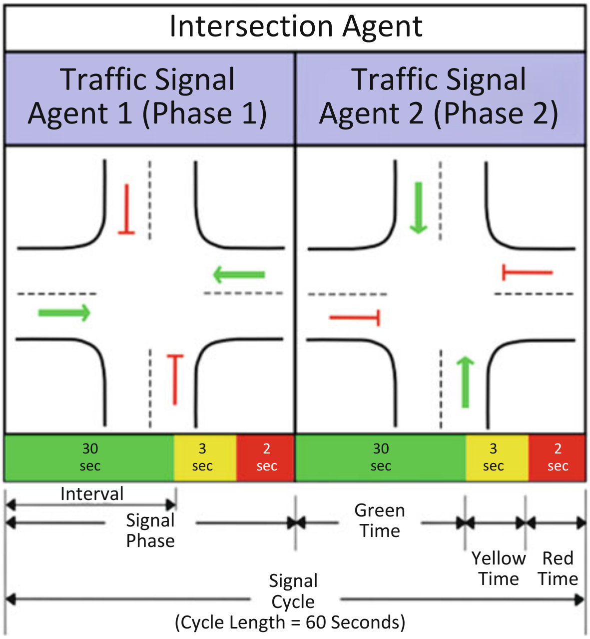

Fig. 2 Traffic signal agent and phase plan

The importance of defining our traffic signal agents in this manner is that it allows us to take advantage of in-practice methods of developing traffic signal timings. Traffic signals manage conflicting movements at an intersection by allowing and restricting movements during set time periods. A *phase plan* describes the sequence of lights a traffic signal will emit and for how long. Developing a phase plan is one of the most important first steps in the effective use of traffic signals. Without a good phase plan, a traffic signal may not be able to handle traffic demand and may even cause traffic accidents. A signal *phase* (illustrated in Fig. 2) is a portion of a traffic signal timing that is given to a set of vehicle movements [28]. A phase is a sequence of lights which includes a green *interval* followed by an amber (yellow) and then a red interval, all assigned to a single movement (or set of non-conflicting movements). The amber and red intervals are necessary so that vehicles have sufficient time to clear the intersection and come to a complete stop.

Our traffic signal agents are equivalent to traffic phases [28] in that they too represent a set of vehicle movements. Thus, for every phase in the phase plan, there is a traffic signal agent that functions on its behalf to tweak the time allotted to that phase. Together, all the phases form the signal timing for a traffic signal, while the traffic signal agents function as an intelligent counterpart to the phase. These two constructs, phase plan and traffic signal agents, address the needs of all legal vehicle movements as traffic demands change.

The design guidelines set by traffic engineers for phase plans (e.g., in the U.S., they use MUTCD [22]) provide a blueprint for determining which movements will be assigned to which traffic signal agent. Traffic engineers divide all the possible legal vehicle movements into subsets, to form phases. The most basic phase plan is the two-phase plan where each street in a standard cross junction (+) is given a phase. The two-phase signalisation plan was used in our initial work [21]. Figure 2 illustrates the relationship between our traffic signal agents and the traffic phases. As there are two phases, there are also two traffic signal agents.

There is a natural conflict that arises between traffic signal agents assigned to an intersection. Each traffic signal agent is designated to a single phase in the traffic signal timing. They compete for a slice of the limited amount of available green time in a *cycle* (see Fig. 2). Assuming the cycle length remains the same, giving more green time to one traffic signal agent means taking it away from another traffic signal agent. We needed a multi-agent interaction protocol [32] to determine an appropriate, adaptive allocation of green time to two competing entities.

As traffic flows through the intersection, auctions take place at fixed intervals which we call the *auction period*. The traffic signal agents participate in the auction and *bid* (explained below) against each other to increase the amount of green time in their respective phases. The winner is the traffic signal agent with the highest bid. The winning agent gains 5 additional seconds of green time, while the loser's green time decreases by the same amount. Although the cycle length remains the same, the amount of green time assigned to each phase changes. Note that the auction period does not (have to) match the cycle length. An auction may occur in the middle of a cycle or after a series of cycles have passed. Green time is only updated after the current traffic signal phase has completed. As a safeguard against starvation, traffic signal agents are prevented from having less than 10 s of green time. Using the taxonomy described by Parsons et al. [20], we could best categorise our auction as *single dimension*, *one-sided*, *sealed-bid*, *first-price* and *single-item*. Thus our implementation—the process that is executed for each auction—closely resembles a single-unit, seller-side English auction [20].

Traffic signal agents use *road sensors* to assess road conditions and generate an appropriate bid. Road sensors include, but are not limited to, inductive-loop vehicle detectors and cameras. The former is a loop of wire buried in the road with an

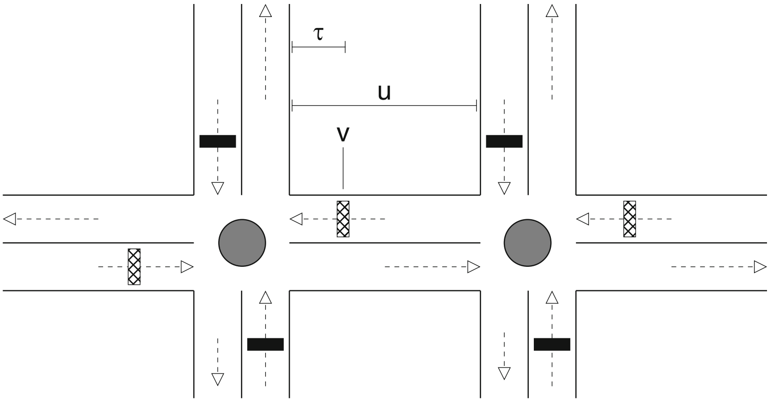

Fig. 3 Traffic Signalling Scheme. The hash-patterned *rectangles* represent the pre-existing *induction-loop* sensors for the west/east traffic signal agents; *black rectangles* for the north/south traffic signal agents. *Grey circles* indicate intersection agents (though they have no physical embodiment in the simulated system). In addition, the following parameters are indicated: **v** is the *volume* of traffic as measured by an induction-loop sensor; **u** is the occupation level between consecutive intersections; and $\boldsymbol{\tau}$ is the occupation level between the sensor and the intersection

electric current running through it and is the primary sensor used in the SCOOT system (mentioned earlier). Vehicles are detected via disruptions in the magnetic field of the wire loop caused by the metal body of the vehicle. The induction-loop sensors are located 20 m from the intersection stop line (the hash-patterned and black rectangles illustrated in Fig. 3). The vehicle detectors provide data on traffic volume, measured in vehicles per hour (*vph*).

In our initial work, we defined *saturation* as the ratio of traffic volume on a road segment to its capacity and used this as a measurement of the level of use of a phase [21]. In general, a stream of traffic that is functioning closer to its capacity is more susceptible to traffic jams and delays [29]. Given a phase p, let d_p be the measure of its *saturation*:

$$d_p = \frac{v}{c}$$

where v is the traffic volume, measured by counting the number of vehicles N (reported by vehicle detectors) that pass a point on a road segment during time interval Δt [29], computed as $v = N/\Delta t$; and c is the capacity, representing the maximum possible traffic volume on a road segment, assuming the traffic signal was always green for that movement(s) [22], computed as $c = 3600/h$ vph. Headway, h, is the average amount of time that it takes vehicles in a queue to reach the intersection. For our simulated environment reported here, headway is set to $\sim$2.54 s, resulting in a capacity of $\sim$1417 vph.

We have implemented and compared two traffic signal agents which have different bidding rules: **Saturation (SAT)** and **Saturation with Queuing (SATQ)**. These are detailed next.

Saturation (SAT) In the SAT method, the traffic signal agents compute d_p for their road segment to use as their bidding rule. In the experiments conducted here, the traffic signal agents are only concerned with the single road segment preceding the junction they manage. For example, the west/east signal agent collects volume data one block west and one block east of its location. Equation (1) defines the SAT bidding rule:

$$bid = d_p \tag{1}$$

Saturation with Queuing (SATQ) The SATQ method extends the SAT method, by augmenting its bidding rule with road occupation, u, an indication of how "full" the road is. This provides a better picture of road conditions (e.g., whether there is a queue of vehicles leading up to the road sensor) than the saturation value alone. A traffic camera could be used to obtain this data. Equation (2) defines the SATQ bidding rule:

$$bid = d_p + u \tag{2}$$

4 Experiments

We evaluated our auction-based methods using the *Simulation of Urban MObility (SUMO)* traffic simulator [14]. SUMO is an open source microscopic traffic simulator and is often used in vehicle communication research [e.g., V2I or *vehicle-to-vehicle (V2V)*]; but it has also been used to study route choice and traffic control algorithms [14]. Although it has a GUI front-end, for our experiments we treated it as a back-end server in order to complete a statistically significant number of experimental runs across a range of traffic conditions. We developed a client application to control the simulation through a TCP socket in SUMO's Traffic Control Interface (TraCI).

As a benchmark for evaluating the effectiveness of our market-based methods, we also tested a **Fixed** method of controlling traffic signals. Fixed represents the traditional approach to tuning traffic signal timings. A fixed-time traffic signal maintains the same timings or light durations throughout the day. Fixed-time signals can be classified by their cycle length. So, we evaluated three types of fixed-time signals: *short*, *medium* and *long* cycle length (tested one at a time, i.e., one per experiment). Note that the starting signal timing (base timings) for our market-based approaches was initialised to the *medium* cycle length.

4.1 Simulation Environment

For the purpose of experimentation, we used the grid-style road network shown in Fig. 4. There are 25 intersections in a 5 × 5 grid layout. *Blocks* are square shaped and measure 200^2 m. The four traffic signals in the corners of the network are deactivated. These four traffic signals control streams that run without conflicts, meaning vehicles traversing these intersections will never have to yield to one another, therefore they are set to always show green. All the other intersections are on the two-phase signal plan. The signal plan does not include dedicated turning (right or left) phases, therefore left and right turns are given lower priority than *through* movements (going straight), i.e., vehicles turning left or right must wait until it is safe to do so. Induction-loop vehicle detectors are placed on roadways (as in Fig. 3) to collect traffic flow data. In Fig. 4, the vehicle detectors are represented by the black and hashed rectangles in the inset.

Vehicles in SUMO have a single goal: reach their destination as quickly and as safely as possible. Vehicles only perform legal manoeuvres including waiting to enter the intersection box until there is ample room to pass completely through it. Vehicles try to maintain a *safe* driving speed based on several pre-set parameters such as *maximum velocity*, *deceleration* and *acceleration*. The safe speed ensures that the vehicle will always be able to safely react to changes in the speed of the vehicle in front of it. Table 1 contains the SUMO driver model parameters that were used in our simulations. Also, drivers follow set routes which are determined before the simulation begins.

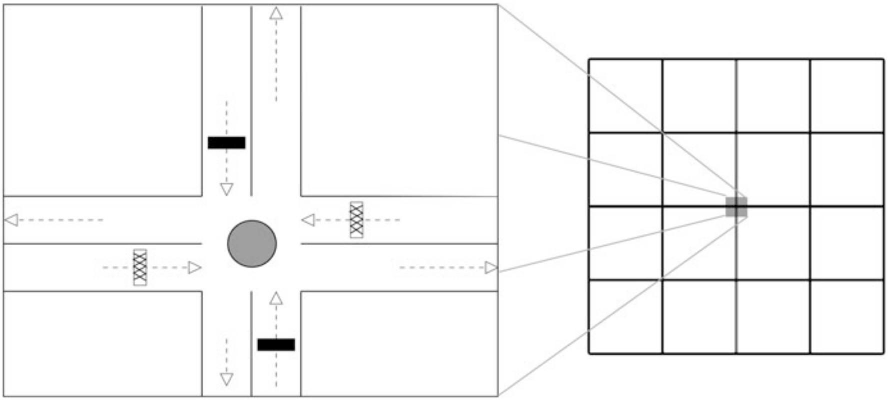

Fig. 4 Grid-based city plan with intersection layout

Table 1 SUMO driver model specifications

AQ3

Driver model	
Parameter	Value
Acceleration	0.8 m/s^2
Deceleration	4.5 m/s^2
Sigma[a]	0.5
Tau[b]	1.0
Vehicle length	5 m
Minimum gap (between cars)	2.5 m
Maximum speed	16.67 m/s

[a] Used to perturb driver behaviour
[b] Driver reaction time

4.2 Experiment Setup

The aim of the experiments we present here is to clearly map the traffic "landscape" and assess the performance of our auction-based mechanism broadly across this landscape. Four different traffic conditions were simulated:

- **Structured** is traffic that flows through the network with an identifiable path with heavy flow;
- **Unstructured** is traffic flow with no identifiable path with heavy flow;
- **Regional** is identical to Structured, except that cross traffic is kept at minimal levels; and
- **Directional** is similar to Structured, but there is a shift in the direction of the heavy flow midway through each experiment.

In the scenarios, the level of cross traffic (east versus west) was varied. The rationale behind Structured, Regional and Directional is that these represent the ideal traffic conditions where an adaptive urban controller, such as SCOOT, would be used. We were interested in how our market-based approach performs under *normal* conditions, as well as in the face of *disruptions*. We produced two types of traffic flow disruptions: *intensity* and *direction*. "Intensity" simulated a sudden increase in overall traffic volume, while "direction" simulated a change in the direction of the flow of traffic with the heaviest volume. We raised the intensity of traffic at the 1 hour mark during Structured, Unstructured and Regional traffic conditions. With the Directional traffic condition, the disruption is the change in direction of the heaviest traffic stream, which occurs at the 1 hour mark as well. Traffic scenarios ran for 3 simulated hours in SUMO (simulations ran for a maximum of 7 simulated hours). Each set of experimental conditions was repeated 30 times to attain suitable statistics.

5 Results

In this section, we present the results of our experiments. We measured performance in terms of *average travel time* and *vehicle queue length*. Results for each metric are tabulated and analysed next.

5.1 Average Travel Time

Our market-based approach significantly reduced travel time for Unstructured traffic, which was the least predictable traffic flow (see Table 2). In fact, SATQ reduced the average travel time by over 25 % compared to the best **Fixed** traffic controller. For the other traffic flows—those that presented a patterned flow—SAT and SATQ performed second best. The best average travel times for Regional, Directional and Structured were attained by the fixed-time controller with the longest cycle length (FXL). The shortest traffic signal timing (FXS) had the worst travel times by a significant margin.

The *cumulative average travel time* of SAT and SATQ with Unstructured traffic, shown in Fig. 5a, remains fairly steady, with very little change, throughout the simulation. The cumulative average never reaches above 625 s for SAT and SATQ. Initially, with Unstructured traffic, FXS and FXM provided the best travel times, but halfway through the simulation, both controllers experience a sharp and steady rise in travel time. Meantime, SAT and SATQ remain relatively unperturbed. With the Regional condition, shown in Fig. 5b, the cumulative average travel time of SAT and SATQ is nearly identical to FXL—all relatively unperturbed. With the Directional and Structured conditions (Fig. 5c, d), the results are similar, though with more separation between SAT/SATQ and FXL.

We used the percent change between SAT, SATQ and FXM to measure improvements on average travel time (shown in Table 3). We compared our auction-based mechanisms only to FXM because SAT and SATQ began with the same signal timing as FXM, thus highlighting that any differences between our approach and FXM are due to the adaptive nature of SAT and SATQ. We found that SAT and

Table 2 Average travel time over all simulation runs—mean (and standard deviation) reported

	Average travel time			
Traffic signal control	Unstructured	Regional	Directional	Structured
SAT	*623.64*(42.31)	*140.26* (6.00)	159.42 (9.36)	160.22 (8.22)
SATQ	**604.78** (32.14)	143.14 (5.95)	*158.64* (7.16)	*150.31* (8.54)
FXS	1096.26 (169.99)	322.25 (10.96)	272.38 (6.42)	250.75 (8.46)
FXM	927.47 (107.39)	183.72 (1.44)	172.60 (0.94)	165.93 (1.38)
FXL	832.71 (52.46)	**139.13** (0.38)	**131.90** (0.43)	**129.04** (0.39)

Fastest times are highlighted in bold; second-fastest times in italics

AQ4

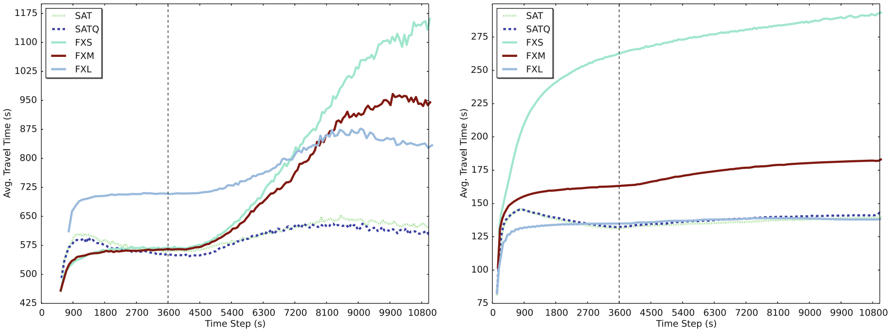

(a) Unstructured (b) Regional

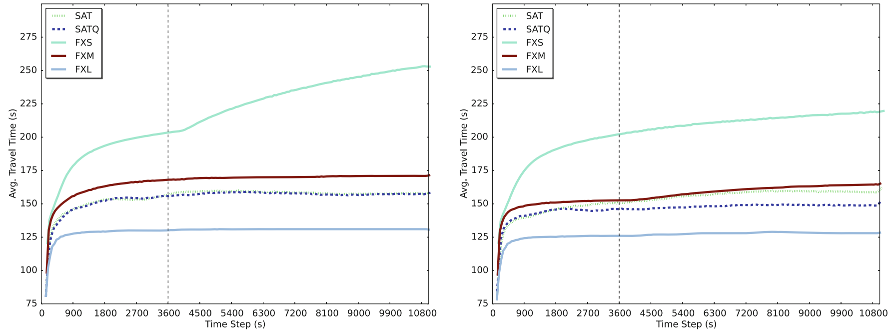

(c) Directional (d) Structured

Fig. 5 Cumulative average travel times. Disruptions occurred at the 3600 s mark. Note the y-axis range for (**a**) is larger and shifted, as compared to (**b**–**d**), further illustrating the distinctive flow exhibited by Unstructured versus the other traffic patterns

SATQ performed in a similar fashion under all four traffic conditions. SATQ's bidding rule, which utilises queue lengths, only provided a slight edge over SAT as Table 3 illustrates. SATQ, compared to SAT, reduced average travel time only slightly more than SAT under the Unstructured and Directional conditions. While SATQ did reduce travel time by nearly 10 % under the Structured condition, it increased travel time under the Regional condition. Under Unstructured conditions, both SAT and SATQ reduced average travel time by over 30 % compared to FXM. The Regional condition experienced similar reductions in travel time, over 20 %, using SAT and SATQ (see Table 3). Under Structured and Directional conditions, SATQ reduced average travel time by just under 10 %, compared to FXM. Overall, when compared to SAT, SATQ provided better travel times in three of the four traffic flows.

The average travel time of vehicles finishing their trip at each time step under the Unstructured condition was greatly reduced under the control of SATQ as compared to when they were using FXM (see Fig. 6). Although there is a slight rise in travel time around the 4500th second, SATQ quickly plateaus and eventually lowers

Table 3 Percent increase in average travel time

AQ5

Difference in average travel time				
		% change		
Traffic pattern		SAT	SATQ	FXM
Unstructured	SAT	*N/A*	3.12	−32.76
	SATQ	−3.02	*N/A*	**−34.79**
	FXM	48.72	53.36	*N/A*
Regional	SAT	*N/A*	−2.01	**−23.65**
	SATQ	2.05	*N/A*	−22.09
	FXM	30.98	28.35	*N/A*
Directional	SAT	*N/A*	0.49	−7.63
	SATQ	−0.49	*N/A*	**−8.09**
	FXM	8.26	8.80	*N/A*
Structured	SAT	*N/A*	6.59	−3.44
	SATQ	−6.18	*N/A*	**−9.41**
	FXM	3.56	10.39	*N/A*

Each *(row, col)* entry in the table is computed as *(row −col)/col*

travel time around the 8100th second. On the other hand, travel times increased dramatically under FXM after the 4500th second and remained elevated for the remainder of the simulation. Vehicles traveling under FXM during the Unstructured condition experienced a much broader range of travel times than under SATQ. We can see in Fig. 7 that with Directional, Regional and Structured conditions, average travel times fell within a very narrow band for both FXM and SATQ.

5.2 *Vehicle Queue Length*

In addition to analysing average travel time, we also measured the size of the *queue of vehicles* that formed at every time step as the simulations ran. The queue length was converted to a value, x, where $0 \leq x \leq 1$, representing the percentage of the road segment that was occupied with vehicles.[1] Figure 8 shows this occupancy measurement for the four incoming roadways at an intersection under the Unstructured condition. Around the 4000th second, FXM experiences an increase in its north-bound queue. During that same period, queues under SAT and SATQ control suffer only a slight increase in queue length. However, the opposite happens on the east-bound roadway, where there is an increase in queue length for both SAT and SATQ. That increase in queue length on the east-bound roadway corresponds to the increase seen under FXM on the north-bound roadway, illustrating that the priorities for traffic flow vary with the different mechanisms.

---

[1] This is the same as the u parameter included in the SATQ bid.

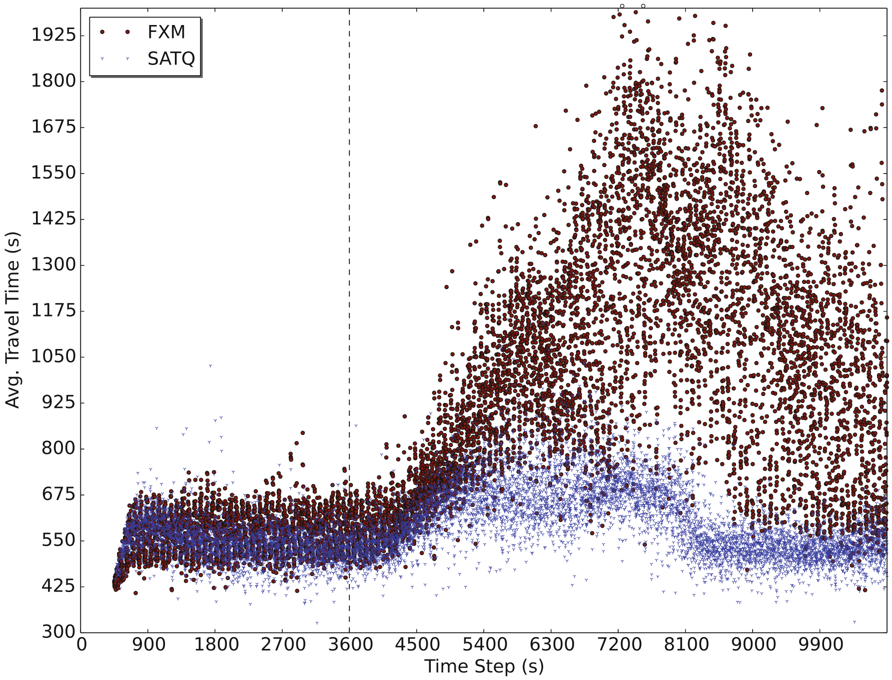

Fig. 6 A comparison of average travel times of vehicles that have completed their journey at each time step under the Unstructured condition. The first vehicles entering the simulation complete their journeys at around the 600 s time step; moving right along the *y*-axis, vehicles continue to complete journeys. The *darker dots* illustrate that the average travel times for vehicles passing through intersections managed by SATQ controllers is much more consistent than those managed by FXM. Figure 7 shows similar results under the other simulated traffic conditions

6 Discussion

Our findings demonstrate that auction mechanisms can be used to manage traffic flow effectively, without the need for vehicle agents. Our results show that under certain traffic conditions, our auction-based approach to traffic control is superior to **Fixed** time traffic signals. SAT and SATQ reduced travel times the most with Unstructured traffic. This was an important finding because the Unstructured traffic represented the sort of unpredictable traffic flow that is often found in the real world. The other traffic conditions, Regional, Directional and Structured, displayed predictable traffic behaviours, e.g., heavy traffic flow in a northerly direction. Predictable traffic behaviours, assuming they are the cause of congestion, are much easier to manage, versus unpredictable traffic flows which are more difficult to regulate. FXL outperformed SAT and SATQ on the other traffic conditions most likely because longer signal timings fair better in heavy traffic than shorter cycles [8].

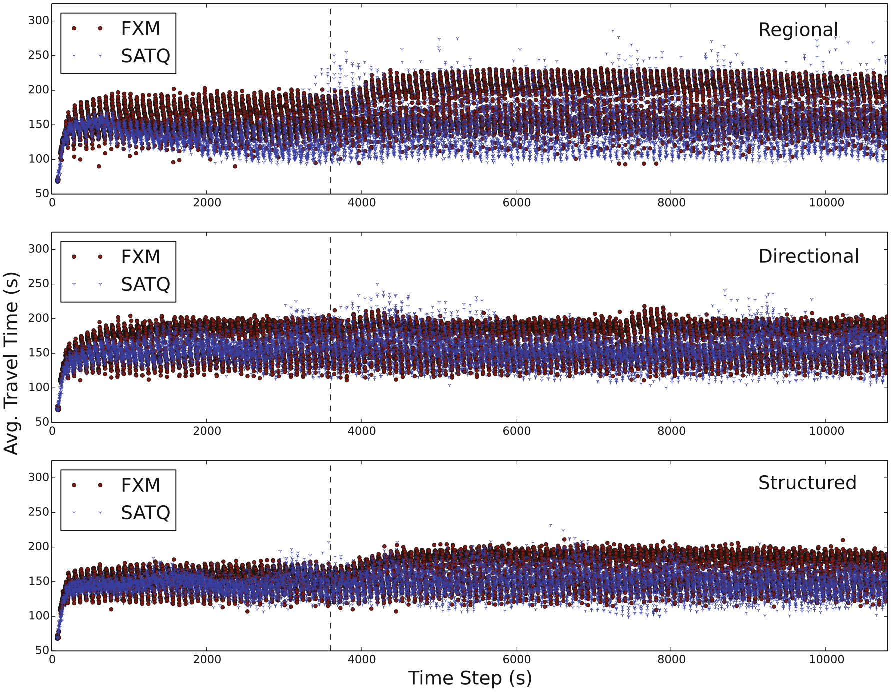

Fig. 7 A comparison of average travel times of vehicles that have completed their journey at each time step under Regional, Directional and Structured conditions. See Fig. 6 for an explanation of the plot format

Recall that SAT and SATQ have the same cycle length as FXM. Yet in Unstructured traffic, compared to FXM, SAT and SATQ reduced average travel time by over 30 % (see Table 3). This strongly suggests that the shifts in green time, caused by the auction mechanism, resulted in a reduction in average travel time. Also, there is evidence that the savings in travel time is shared by all drivers. Figures 6 and 7 show that the average travel times for all the vehicles that completed their trips fell within a narrow band. In other words, throughout the timespan of each simulation, the vast majority of vehicles experienced a reduction in travel time.

Finally, we turn to the queue measurements shown in Fig. 8. Here we can use FXM to get an idea of what queue lengths would have been like without an auction mechanism. SAT (and SATQ) had increases in queue lengths on the east/west-bound roadways with a corresponding decrease in north/south-bound queues, suggesting that in order to improve travel time, green time was shifted to the north/south-bound lanes. This means that green time was given to the roadway that needed it the most—which is the intended goal of the auction mechanism.

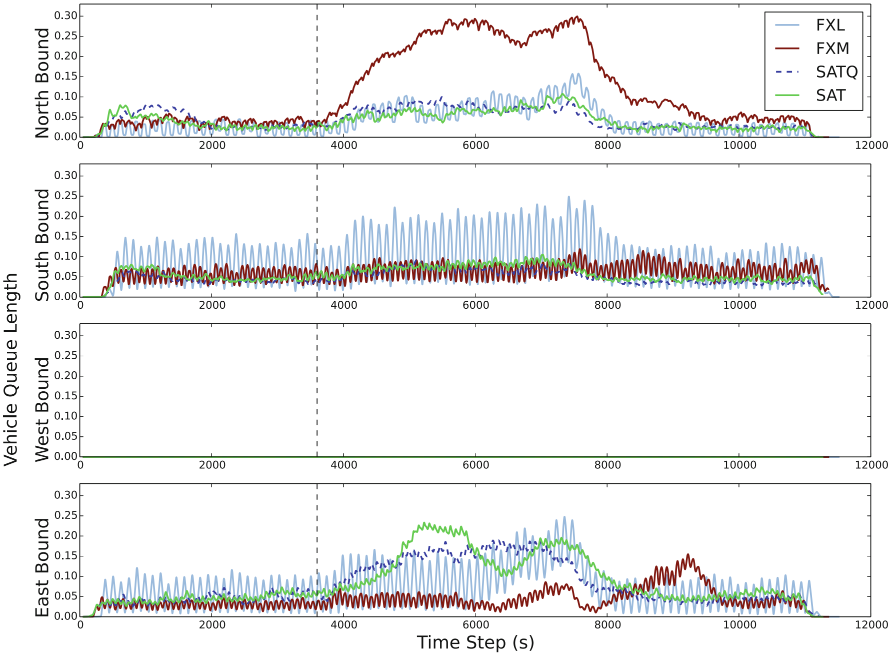

Fig. 8 Percentage of incoming roadway filled with vehicles at a single intersection within the city under the Unstructured condition. Note, for this particular intersection, the west-bound lane did not experience any traffic

Although FXL had the best average travel time under the Regional condition, a closer look revealed that SAT and SATQ behaved in a very similar manner to FXL (see Fig. 5b). This finding is interesting because the same behaviour is not seen under the Directional or Structured conditions. The difference in performance may be due to a lack of change in traffic demand. In Regional, there is very little cross traffic and the heavy flow is in a single direction. Most likely the majority of the green time was given to the heaviest flow; but with such little cross traffic, this did not cause an overall increase in travel time. Bazzan [2] utilised a similar traffic scenario in their work. However, they found the lack of demand in the opposite direction hampered their agent-based traffic controller. In contrast, our controller adapted well.

The Directional (and Structured) conditions probably experienced a similar shift in green time; but having greater cross traffic, this resulted in an overall increase in average travel time. The amount of cross traffic was not enough to influence the auction (as it remain constant), but enough to raise overall average travel times. The simulations under Directional and Structured conditions highlight the critical role the bidding rule plays in green time allocation and as a measure of traffic demand.

7 Conclusion

In conclusion, our work here and in [21] demonstrate the feasibility of a multi-agent auction-based traffic controller that does not require vehicle agents. SATQ reduced travel time by over 30 % under certain traffic conditions compared to a **Fixed** time traffic controller of initially identical cycle length. Traffic congestion costs urban areas billions of Euros in lost time, and vehicle emissions are a major source of air pollution. Auctions have been shown to be able to improve the management of intersection traffic, but thus far only when paired with vehicle agents. Our approach can be deployed in software, without the added cost of upgrading vehicles to include vehicle agents and working with existing transportation infrastructure hardware and control systems.

In this paper, we have outlined the framework for our auction-based traffic controller. Still, there is a need to further investigate the relationship between various bidding rules and traffic demand. Future work will focus on developing additional bidding rules and methods to measure a roadway's level of use, particularly under unpredictable and changeable conditions—the types of situations that currently stymie existing systems. We will also conduct experiments to compare our approach to other adaptive traffic controllers currently in use.

References

1. Bakker, B., Whiteson, S., Kester, L.J. H.M., Groen, F.C.A.: Traffic light control by multiagent reinforcement learning systems. In: Babuska, R., Groen, F.C.A. (eds.) Interactive Collaborative Information Systems. Studies in Computational Intelligence, vol. 281, pp. 475–510. Springer, Berlin/Heidelberg (2010)
2. Bazzan, A.L.C.: A distributed approach for coordination of traffic signal agents. Auton. Agent. Multi-Agent Syst. **10**(2), 131–164 (2005)
3. Bazzan, A.L.C., de Oliveira, D., da Silva, B.C.: Learning in groups of traffic signals. Eng. Appl. Artif. Intell. **23**(4), 560–568 (2010)
4. Cardiovascular Disease and Air Pollution. Technical report, Department of Health, Committee on the Medical Effects of Air Pollutants, February 2006
5. Carlino, D., Boyles, S.D., Stone, P.: Auction-based autonomous intersection management. In: Proceedings of the 16th IEEE Intelligent Transportation Systems Conference (ITSC) (2013)
6. Chen, B., Cheng, H.H.: A review of the applications of agent technology in traffic and transportation systems. IEEE Trans. Intell. Transp. Syst. **11**(2), 485–497 (2010)
7. Dias, M.B., Zlot, R., Kalra, N., Stentz, A.: Market-based multirobot coordination: a survey and analysis. Proc. IEEE **94**(7), 1257–1270 (2006)
8. Dresner, K.M., Stone, P.: Multiagent traffic management: a reservation-based intersection control mechanism. In: Proceedings of the Third International Joint Conference on AAMAS, pp. 530–537. IEEE Computer Society, Washington, DC (2004)
9. Gartner, N.H., Pooran, F.J., Andrews, C.M.: Implementation of the OPAC adaptive control strategy in a traffic signal network. In: IEEE Intelligent Transportation Systems. IEEE (2001) AQ6
10. Gerkey, B.P., Mataric, M.J.: Sold!: auction methods for multirobot coordination. IEEE Trans. Robot. Autom. **18**(5), 758–768 (2002)

11. Gomm, P., Wengraf, I.: The Car and the Commute, The journey to work in England and Wales. Technical report, The Royal Automobile Club Foundation for Motoring, December 2013
12. Knight, S., Qureshi, M., Cleverly, J., Duvall, L., Gavron, N., Jones, J., Malthouse, K.: Driving Away from Diesel, Reducing Air Pollution from Diesel Vehicles. Technical report, London Assembly, Environment Committee, July 2015
13. Koenig, S., Tovey, C.A., Lagoudakis, M.G., Markakis, E., Kempe, D., Keskinocak, P., Kleywegt, A.J., Meyerson, A., Jain, S.: The power of sequential single-item auctions for agent coordination. In: Proceedings of the 21st National Conference on Artificial Intelligence. AAAI 2006, vol. 2, pp. 1625–1629. AAAI Press, Boston, MA (2006)
14. Krajzewicz, D., Erdmann, J., Behrisch, M., Bieker, L.: Recent development and applications of SUMO - Simulation of Urban MObility. Int. J. Adv. Syst. Meas. **5**(3 & 4), 128–138 (2012)
15. Lagoudakis, M.G., Berhault, M., Koenig, S., Keskinocak, P., Kleywegt, A.J.: Simple auctions with performance guarantees for multi-robot task allocation. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 698–705. IEEE (2004)
16. Miller, B.G.: Report on Estimation of Mortality Impacts of Particulate Air Pollution in London. Technical report, Institute of Occupational Medicine, June 2010. Report P951-001
17. Mirchandani, P., Wang, F.-Y.: RHODES to intelligent transportation systems. IEEE Intell. Syst. **20**(1), 10–15 (2005)
18. Mladenovic, M., Abbas, M.: A survey of experiences with adaptive traffic control systems in North America. J. Road Traff. Eng. **59**(2), 5–11 (2013)
19. Papageorgiou, M., Diakaki, C., Dinopoulou, V., Kotsialos, A., Wang, Y.: Review of road traffic control strategies. Proc. IEEE **91**(12), 2043–2067 (2003)
20. Parsons, S., Rodriguez-Aguilar, J.A., Klein, M.: Auctions and bidding: a guide for computer scientists. ACM Comput. Surv. **43**(2), 10:1–10:59 (2011)
21. Raphael, J., Maskell, S., Sklar, E.: From goods to traffic: first steps toward an auction-based traffic signal controller. In: 13th International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS'15). Springer (2015)
22. Roess, R.P., Prasas, E., McShane, W.R.: Traffic Engineering: International Edition, 4th edn. Pearson Education, Inc., Upper Saddle River (2009)
23. Sariel, S., Balch, T.R.: Efficient bids on task allocation for multi-robot exploration. In: Sutcliffe, G., Goebel, R. (eds.) FLAIRS Conference, pp. 116–121. AAAI Press, Menlo Park, CA (2006)
24. Schepperle, H., Böhm, K.: Agent-based traffic control using auctions. In: Klusch, M., Hindriks, K.V., Papazoglou, M.P., Sterling, L. (eds.) CIA. Lecture Notes in Computer Science, vol. 4676, pp. 119–133. Springer, Berlin (2007)
25. Schepperle, H., Böhm, K.: Auction-based traffic management: towards effective concurrent utilization of road intersections. In: CEC/EEE, pp. 105–112. IEEE (2008)
26. Schneider, E., Sklar, E.I., Parsons, S., Tuna Özgelen, A.: Auction-based task allocation for multi-robot teams in dynamic environments. In: Proceedings of the 16th Towards Autonomous Robotic Systems (TAROS) (2015)
27. The economic costs of gridlock. Technical report, Centre for Economic and Business Research, December 2012
28. Transportation Research Board. Highway Capacity Manual. Transportation Research Board, National Research Council, Washington, DC (2000)
29. Treiber, M., Kesting, A.: Traffic Flow Dynamics. Springer, Berlin/Heidelberg (2013)
30. Vasirani, M., Ossowski, S.: A market-inspired approach for intersection management in urban road traffic networks. J. Artif. Intell. Res. **43**, 621–659 (2012)
31. Wang, F.-Y.: Agent-based control for networked traffic management systems. IEEE Intell. Syst. **20**(5), 92–96 (2005)
32. Wellman, M.P., Wurman, P.R.: Market-aware agents for a multiagent world. Robot. Auton. Syst. **24**(3–4), 115–125 (1998)
33. AQ7 Xu, Y., Zhang, Y., Liu, M.: Multiagent based decentralized traffic light control for large urban transportation system. Math. Probl. Eng. **2014** (2014)
34. Yim, S.H.L., Barrett, S.R.H.: Public health impacts of combustion emissions in the United Kingdom. Environ. Sci. Technol. **46**(8), 4291–4296 (2012)