

From Goods to Traffic: First Steps Toward an Auction-based Traffic Signal Controller

Jeffery Raphael¹, Simon Maskell², and Elizabeth Sklar¹
{jeffery.raaphael,s.maskell,e.i.sklar}@liverpool.ac.uk

¹Dept of Computer Science and ²Dept of Electrical Engineering and Electronics
University of Liverpool
Liverpool, UK

Keywords: Multi-Agent Systems, Auctions, and Traffic signal control

Abstract. Traffic congestion is a major issue that plagues many urban road networks large and small. Traffic engineers are now leaning towards Intelligent Traffic Systems as many types of physical changes to road networks are costly or infeasible. Multi-Agent Systems (MAS) have become a popular paradigm for exploring intelligent solutions to traffic management problems. There are many MAS approaches to traffic management that utilise market mechanisms. In market-based approaches, drivers “pay” to use the roadways. However, a major issue with many of these solutions is that they require technology that, as yet, does not exist or is not widely available. For example, they rely on a special software agent that resides within the vehicle. This “vehicle agent” is responsible for participating in the market mechanism and communicating with the transportation infrastructure. In this paper, an auction-based traffic controller is proposed which exploits all the benefits of market mechanisms *without the need for a vehicle agent*. Experimental results show that such a controller is better at reducing delay and increasing throughput in a simulated city, as compared to fixed-time signal controllers.

1 Introduction

Traffic congestion occurs when the volume of traffic exceeds the capacity of the infrastructure and causes traffic flow to slow. Over 60% of commuters in England and Wales drive to work [13]. In London, despite having access to an extensive public transportation network, over a quarter of Londoners still choose to drive to work [13]. During *rush hours*, traffic volume often reaches levels that severely strain current traffic management systems. Traffic volume and common work hours are just two of the many factors that can grind traffic to a halt. This type of recurring congestion pattern is responsible for 86% and 32% of traffic congestion in France and Germany respectively [1]. The cost of traffic congestion can be measured both in time and money. According to a report put out by the Centre for Economics and Business Research (CEBR), drivers in London waste around 66.1 hours a year waiting in traffic. All those hours add up and across

the UK close to €4.94 billion [2] are lost in the form of fuel and the increased cost of delivering goods. Other European countries face similar monetary losses. Traffic congestion costs France, Germany, and Spain €5.55, €7.83, and €5.5 billion respectively [2, 8]. The estimated annual cost of congestion in the EU is €111.3 billion [8]. The staggering cost of traffic congestion and its complexities make it an attractive problem to help solve.

There are many tools at the disposal of transportation departments to manage traffic flow. Traffic lights are probably the most prevalent means of controlling traffic. Other methods include stop signs and roundabouts. Although many traffic lights rely on simple fixed protocols, they are none-the-less a vital component of traffic management [4]. More advanced adaptive Urban Traffic Controllers (UTC), such as RHODES [17], OPAC [11] and SCOOT¹, have been developed in an effort to improve the performance of traffic lights [22, 18]. Adaptive UTCs use 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 cars, public transportation and pedestrians) in areas ranging in size from a few city blocks to entire cities. The majority of adaptive UTCs employ optimisation algorithms which are costly to develop, maintain and expand [22].

The fundamental nature of traffic flow makes it an ideal problem for *Multi-Agent Systems (MAS)*. Traffic control is geographically distributed, takes place in a dynamic environment and the interactions amongst its components are highly complex [7]. It is easy to see all the vehicles, pedestrians, cyclists and traffic control mechanisms as a collection of autonomous agents interacting in a large space. The MAS paradigm offers a flexible and inexpensive method for designing traffic control solutions [22]. There is a plethora of traffic control solutions that fall under the umbrella of MAS. Our work focuses on those solutions that utilise market-based mechanisms.

Our approach for controlling traffic signals has been greatly influenced by co-ordination efforts in Multi-Robot Routing (MRR) [9, 12, 14, 19]. *Auctions*, which are a form of market-based mechanism for resource allocation, can produce near optimal results in some MRR scenarios [16]. Traffic control can be viewed as a co-ordination problem [5] where traffic signals work together to maintain adequate traffic flow and minimise delays. 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 as simple as vehicle-to-infrastructure communications to more demanding vehicle navigation and control.

There are two main problems with any system that relies on vehicle agents: the development and deployment of vehicle agents and the current transportation infrastructure. Car manufacturers will have to agree on international communication protocols, physical specifications and the many other aspects of deploying vehicle agents to the millions of vehicles that are currently in use. Second, the communication infrastructure within the traffic system itself currently does not

¹ <http://www.scoot-utc.com>

exist. Our overarching goal is to design a system that reaps the benefits of a market mechanism that is able to take advantage of existing infrastructure but does not require *vehicle agents*.

We demonstrate a simple approach to such a system here. 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 (Section 6) and conclusions (Section 7).

2 Related Work

Dresner et al. [10] 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. The reservation-based system 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 et al. [10] compared their reservation-based approach to two other traffic control schemes: *overpass* and traffic light. Overpass simulates a road with absolutely no signals. Traffic light simulates how current signals functions. Dresner et al. [10] found that their reservation-based system outperformed the normal traffic light.

Vasirani et al. [21] expanded on Dresner’s work and examined the performance changes to a reservation-based system where time slots were allocated via a combinatorial auction. And they also expanded [10] to include multiple intersections to study the effects of such a market-based reservation system would have on drivers’ route choice. They viewed the space within an intersection as a resource and managed that resource using a market-based system. Vasirani et al. [21] looked at the delay experienced by drivers based on the amount they were willing to “pay” to use the intersection under various traffic densities. 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 et al. [21] found that initially having a willingness to pay does decrease delay, but eventually this levels off. As [21] is an extension of [10] it too relies on vehicle agents.

Carlino et al. [6] described a traffic control system where auctions are run at intersections to determine use. This solution assumes vehicles have an embedded agent bidding on their behalf, which is referred to as the *wallet agent*. 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). Carlino et al. [6] used a second-price sealed bid auction mechanism. They tested four different modes: *FIFO* (this is how your typical intersection works), *Equal* (every driver submits a bid of one, *Static Wallet*), *Auction* (drivers use the

Fair Wallet, and *Fixed* (drivers always bid the same amount based on the value they’ve assigned for the trip). FIFO performed the worst.

Schepperle et al. [20] 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 et al. [20] proposes 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 et al. [20] 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. Like other works of this nature, [20] assumes that cars have a vehicle agent and that intersections have an agent as well. Our approach, detailed in the next section, does not involve vehicle agents or other embedded software.

3 Our Approach: Auction-based Traffic Signalling

In this section, we describe our auction-based mechanism for traffic signalling which does not employ vehicle agents. Instead, we use an *intersection agent* (as an auction manager) and *traffic signal agents* that represent the traffic signals at each intersection—one per pair of opposing-direction traffic flows (i.e., opposing traffic light phases). Thus, at every crossroads, there is an intersection agent working in concert with two traffic signal agents to adapt the signal timing to meet traffic demands. This scheme is illustrated in Figure 1. Each intersection functions on a two-phase traffic light programme: one light phase for north/south bound traffic and the other phase for west/east bound traffic.

Our traffic signal control mechanism employs a first-price, single-item auction. As traffic flows through the intersection, auctions take place at fixed intervals which we call the *auction frequency*. The two traffic signal agents bid 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. The cycle length remains the same, but the amount of green time changes.

Note that the *auction frequency* 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 light phase has completed. As a safeguard against *starvation*, traffic signal agents are prevented from having less than 10 seconds of green time. Starvation is defined as the situation where traffic is prevented from flowing in a particular direction. *Gridlock* is defined as the situation where starvation occurs in both directions.

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 are loops of wire buried in the road

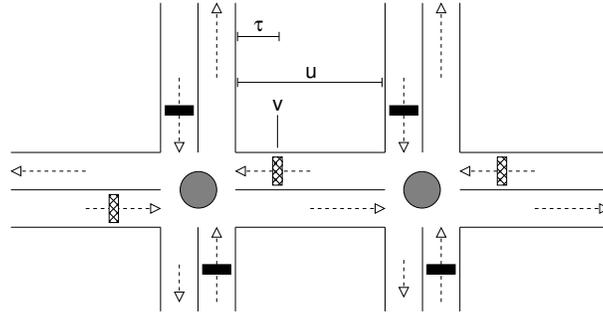


Fig. 1: Traffic Signalling Scheme. The hash-patterned rectangles represent the pre-existing *induction-loop* sensors for the west/east traffic signal agents; black rectangles for the north/south traffic signal agents. Grey circles indicate intersection agents (though they have no physical embodiment 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 τ is the occupation level between the sensor and the intersection. (See text for further explanation.)

with a current running through it and are the primary sensor used in SCOOT. Vehicles are detected via disruptions in the magnetic field of the wire loop caused by the metal body of the vehicle. Each induction-loop sensor (the hash-patterned and black rectangles illustrated in Figure 1) computes v , the number of vehicles that have crossed the induction-loop in a fixed time period. The induction-loop sensors are located 20 meters from the intersection.

We have defined two methods of traffic control: **Saturation (SAT)** and **Saturation with Queuing (SATQ)**.

SAT. In the SAT method, the traffic signal agents use the *saturation* of their road segment as a bidding rule. The *saturation* of a road segment is the ratio of the volume of traffic (here, represented as v and measured by the induction-loop sensors) to its estimated capacity c (defined by the physical road network). In the experiments conducted here, the traffic signal agents are only concerned with the single block 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 bidding rule for the traffic signal agents.

$$bid = v/c \quad (1)$$

SATQ. The SATQ method functions similarly to the SAT method, except that its bidding rule is augmented with road occupation, u , which is 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 induction-loop sensor alone. A traffic camera could be used to obtain this data. The modified bidding rule employed by the SATQ method is defined in Equation 2.

$$bid = (v/c) + u \tag{2}$$

4 Experiments

We evaluated our auction-based methods using the *Simulation of Urban MObility (SUMO)* traffic simulator [15]. SUMO is an open source microscopic traffic simulator and is often used in vehicular communications (either vehicle-to-vehicle or vehicle-to-infrastructure) research but it is also used to study route choice and traffic control algorithms [15]. Although it has a GUI front-end, for our experiments we treated it as a back-end server. We developed a client application to control the simulation using SUMO’s Traffic Control Interface (TraCI) through a TCP socket.

As a benchmark for evaluating the effectiveness of our auction-based methods, we also tested a **Fixed** method of controlling traffic signals. The Fixed method represented traditional, non-adaptive, traffic lights that display the same light sequence in every cycle. The cycle lengths ranged from 80 to 90 seconds (varying across different intersections), with each traffic signal spending at least 68% of their phase showing green.

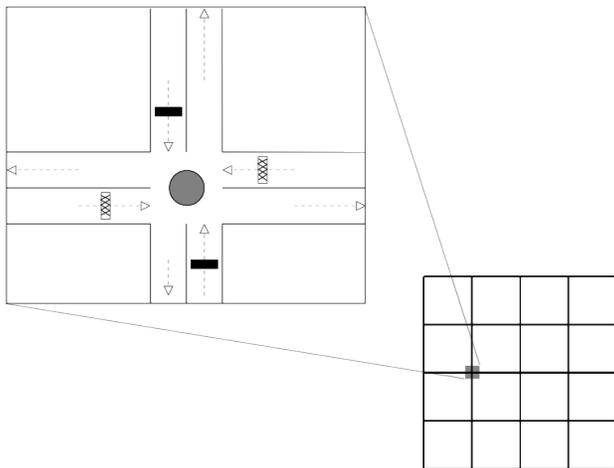


Fig. 2: Grid city

For the purpose of experimentation, to determine the effectiveness of the proposed traffic controller, a simulated “Grid” city was used, following a Manhattan-style road network (shown in Figure 2). Although simple, similar networks have

been used in other traffic experiments [3, 4]. A single Grid City block measures 200 meters. Grid City contains 25 traffic lights, but only 21 are four-way junctions (the four corners do not have opposing traffic flows). In the simulation, there are four induction-loops at each intersection, one for each traffic flow entering the intersection. The exact positions of the induction-loops are shown in the insets of Figure 2.

During each simulation run, 2,000 vehicles entered Grid City and travelled across the network. The four corners of Grid City were designated as entrance/exit points. Vehicles entered at one of four entry points and exited at another. For each of the four entry points there was a 90% probability of generating a new vehicle at any given time step. Table 1 presents the vehicle specification settings used. Each simulation run lasted a maximum of 15,000 seconds (4 hours and 10 minutes); simulations could terminate early if all vehicles reached their destination before the maximum time passed. For each traffic control method tested, 50 simulation experiments were executed.

In addition to comparing the Fixed, SAT and SATQ methods for traffic signal control, we experimented with varying the *auction frequency*. We ran fifteen sets of experiments, varying the auction frequency from 1 to 15 minutes. Note that the auction frequency remained constant throughout an experiment, and that all auctions occurred synchronously (i.e., all bidding and matching took place at the same time). Future work will explore variable auction frequencies within a simulation, as well as asynchronous auctions.

The results from the experiments are presented in Section 5.

<i>Parameter</i>	<i>Value</i>
acceleration	$0.8m/s^2$
deceleration	$4.5m/s^2$
size of vehicle (length)	$5m$
maximum velocity	$16.67m/s$
minimum gap between vehicles	$2.5m$

Table 1: Vehicle specifications

Performance was measured in three ways. The first was in terms of total trip duration: on average, how long it took for all the vehicles to complete their trips, measured in seconds. The second was *throughput* (p) which was measured in terms of *vehicles/hour*. Throughput is the estimated number of vehicles that could pass through the road network in an hour. This was calculated using Equation (3), where: n is the number of cars participating in the simulation and t is the total amount of time (in seconds) it took for all cars to complete their journeys.

$$p = \frac{n}{t} \times 3600 \tag{3}$$

The third metric was *completion rate*: the percentage of the 2000 vehicles entering the system were able to complete their journeys before the maximum simulation time had elapsed.

5 Results

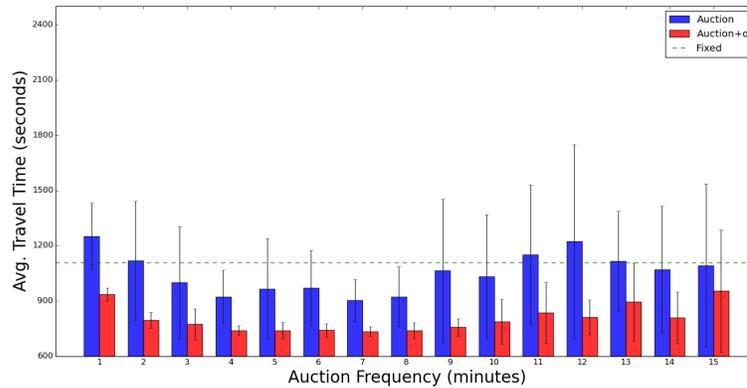


Fig. 3: Average time to complete travel plan

The results of our experiments are presented in Figures 3 and 4. Although SAT sometimes reduced travel times as much as SATQ, Figure 3 shows that SATQ consistently outperformed the other two traffic control methods. The large amount of deviation seen in SAT reflects the quality of its bidding rule. The bidding rule indirectly provides a representation of the state of traffic at the time an auction is executed. The traffic signal agents select actions that are most appropriate for the traffic conditions it perceives through its bidding rule. Traversing the solution space using the bidding rules developed thus far does not guarantee a sequence of actions that will gradually improve overall travel times from one time step to the next. If there are enough traffic signal agents that fail to accurately capture current traffic trends and act appropriately, the combined effect is an increase in overall network delay. Infrequent auctions can also lead to this disconnect between traffic signal actions and road conditions as evident in the increased deviation experienced by SATQ after the 10 minute mark in Figure 3. The lowest average travel time was achieved when the auction frequency was set to 7 minutes.

Figure 4 further supports the conclusion that SATQ was the best control method. In terms of throughput, SATQ had a statistically significant advantage over the other two methods. SATQ was able to handle nearly 50% more traffic than Fixed. Some simulations did end in gridlock using SAT or SATQ (the Fixed

controller always ran to completion). Figure 5 shows the percentage of vehicles that completed their entire trip versus the auction frequency. SATQ reached 100% completion until after the 11 minute mark. Again, we see performance issues emerge with infrequent auction frequencies.

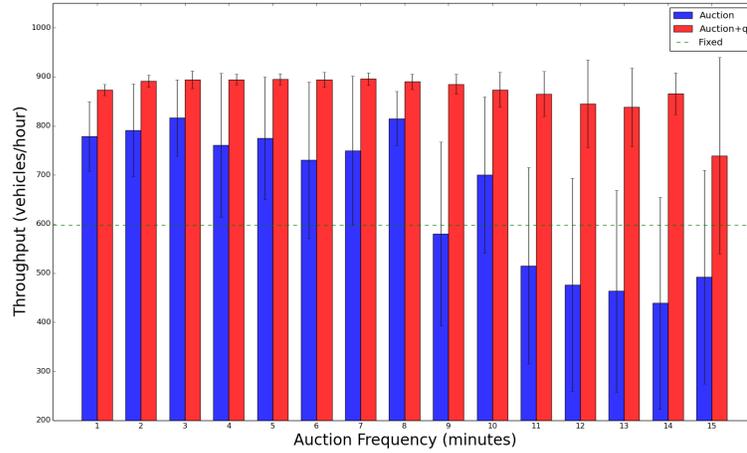


Fig. 4: Estimated throughput

6 Discussion

SATQ produced superior results as compared to a single fixed cycle. In order to get a better picture of how well SATQ worked, we compare it to the *overpass* benchmark employed by Dresner et al. [10].

<i>Control Method</i>	<i>Average Travel Time (std)</i>
<i>Overpass</i>	<i>301.30</i>
Fixed	1108.52 (23.84)
SAT	1054.23 (293.24)
SATQ	803.27 (96.90)

Table 2: Average travel time for all vehicles to complete trip in seconds. Auctions were executed every 5 minutes.

As shown in Table 2, the minimum average travel time for all vehicles to complete their trip is 301.30 seconds versus SATQ’s 818.17. The vehicles using the

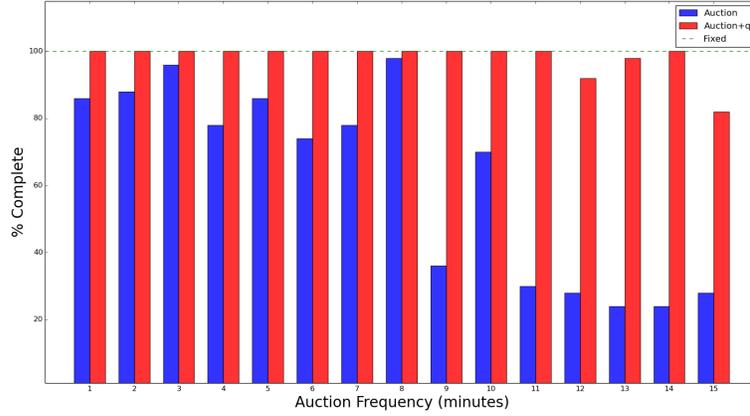


Fig. 5: Completion Rate: percentage of runs where all the vehicles reached their destination by auction frequency

Fixed method required 3.7 times more time (compared to the the lower bound) to complete their trip, while SATQ required 2.7 times more time. This in itself is impressive considering that no effort was made to optimise any global parameters in the case of SAT and SATQ. In this initial implementation, traffic signal agents behaved rather selfishly: they were concerned with improving travel time solely at their junction. These results support our belief that an appropriately designed market-based MAS can improve traffic flow. Future work will investigate traffic signal agents that consider a *neighbourhood* of intersections, not just their immediate junction.

The most likely reason as to why SATQ outperformed SAT has to do with queue formation and how inductive-loop detectors work. If and when a queue formed at an intersection and that queue surpassed the position of the induction-loop it would register traffic flow as zero. The issue with returning a zero count is that it has two meanings, either there is no traffic on the road at all or traffic is so *backed up* that a vehicle is sitting directly over the sensor. Unfortunately, one meaning suggests the signal requires less green time while the other suggests it requires more green time. The u term (Equation 2) clarifies this ambiguity. If the detector returns zero because a queue has formed, then $u \geq \tau$, where τ is the percentage of the road that would be occupied by vehicles if the queue has reached the position of the inductive-loop detector. So, the u term supports the agent's bid for more green time when a queue of sufficient size has formed. If it returns zero because there is no traffic, then $u = 0$. Interestingly, the performance gap between SAT and SATQ is a great example of how multiple sources of road data can be used in tandem to improve a traffic control mechanism.

7 Conclusion

The work presented in this paper demonstrates how an auction-based traffic controller could be implemented without the use of vehicle agents. Our approach takes advantage of road sensor devices that are currently available. We implemented two versions of our mechanism: SAT and SATQ. We tested their effects on traffic flow in a fictitious road network, Grid City. The results show that our mechanism is capable of outperforming a fixed-time signal system. Although acting locally, our *intersection agent* and signal agents are able to minimise the delay and increase the throughput of the road network. If one considers that the majority of adaptive UTCs use complex and time-consuming optimisation techniques, then our method is even more interesting. Our preliminary implementation, although simple, produced results that are quite an improvement on average travel time and throughput. Our mechanism exhibits traits that make it ideal for a real-time adaptive traffic light controller: it has minimal communications overhead, it is highly reactive to changing traffic conditions and its design is uncomplicated.

Future work will focus on the auction mechanism parameters and traffic testing scenarios. We plan to examine how including traffic signal agents from multiple connected intersections in a single auction might effect performance. We will also test our mechanisms against SCOOT in a simulation of an actual city with more intersections and realistic traffic flows.

8 Acknowledgments

Thanks to Frans Coenen for comments on earlier drafts, Simon Parsons and Liverpool City Council for research advice.

References

1. Managing urban traffic congestion. Technical report, OECD Publishing, 2007.
2. The economic costs of gridlock. Technical report, Centre for Economic and Business Research, December 2012.
3. Tiago M. L. Azevedo, Paulo J. M. de Arajo, Rosaldo J. F. Rossetti, and Ana Paula C. Rocha. JADE, TraSMAPAPI and SUMO: A tool-chain for simulating traffic light control. In *Proceedings of the Thirteenth International Joint Conference on Autonomous Agents and Multiagent Systems*, May 2014.
4. Bram Bakker, Shimon Whiteson, Leon J. H. M. Kester, and Frans C. A. Groen. Traffic light control by multiagent reinforcement learning systems. In Robert Babuska and Frans C. A. Groen, editors, *Interactive Collaborative Information Systems*, volume 281 of *Studies in Computational Intelligence*, pages 475–510. Springer, 2010.
5. Ana L. C. Bazzan. A distributed approach for coordination of traffic signal agents. *Autonomous Agents and Multi-Agent Systems*, 10(2):131–164, 2005.
6. Dustin Carlino, Stephen D. Boyles, and Peter Stone. Auction-based autonomous intersection management. In *Proceedings of the 16th IEEE Intelligent Transportation Systems Conference (ITSC)*, 2013.

7. Bo Chen and Harry H. Cheng. A review of the applications of agent technology in traffic and transportation systems. *IEEE Transactions on Intelligent Transportation Systems*, 11(2):485–497, June 2010.
8. Panayotis Christidis and Juan Nicolas Ibanez Rivas. Measuring road congestion. Technical report, European Commission, Joint Research Centre, 2012.
9. M. Bernardine Dias, Robert Zlot, Nidhi Kalra, and Anthony Stentz. Market-based multirobot coordination: A survey and analysis. *Proceedings of the IEEE*, 94(7):1257–1270, 2006.
10. Kurt M. Dresner and Peter Stone. Multiagent traffic management: A reservation-based intersection control mechanism. In *Proceedings of the Third International Joint Conference on AAMAS*, pages 530–537. IEEE Computer Society, 2004.
11. Nathan H. Gartner, Farhad J. Pooran, and Christina M. Andrews. Implementation of the OPAC adaptive control strategy in a traffic signal network. In *IEEE Intelligent Transportation Systems*. IEEE, 2001.
12. Brian P. Gerkey and Maja J. Mataric. Sold!: auction methods for multirobot coordination. *IEEE Transactions on Robotics and Automation*, 18(5):758–768, 2002.
13. Philip Gomm and Ivo Wengraf. The car and the commute, the journey to work in england and wales. Technical report, The Royal Automobile Club Foundation for Motoring, December 2013.
14. Sven Koenig, Craig A. Tovey, Michail G. Lagoudakis, Evangelos Markakis, David Kempe, Pinar Keskinocak, Anton J. Kleywegt, Adam Meyerson, and Sonal Jain. The power of sequential single-item auctions for agent coordination. In *Proceedings of the 21st National Conference on Artificial Intelligence*, volume 2 of *AAAI 2006*, pages 1625–1629. AAAI Press, 2006.
15. Daniel Krajzewicz, Jakob Erdmann, Michael Behrisch, and Laura Bieker. Recent development and applications of SUMO - simulation of urban MObility. *International Journal On Advances in Systems and Measurements*, 5(3&4):128–138, December 2012.
16. Michail G. Lagoudakis, Marc Berhault, Sven Koenig, Pinar Keskinocak, and Anton J. Kleywegt. Simple auctions with performance guarantees for multi-robot task allocation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 698–705. IEEE, 2004.
17. Pitu Mirchandani and Fei-Yue Wang. RHODES to intelligent transportation systems. *IEEE Intelligent Systems*, 20(1):10–15, 2005.
18. Milos Mladenovic and Montasir Abbas. A survey of experiences with adaptive traffic control systems in north america. *Journal of Road and Traffic Engineering*, 59(2):5–11, 2013.
19. Sanem Sariel and Tucker R. Balch. Efficient bids on task allocation for multi-robot exploration. In Geoff Sutcliffe and Randy Goebel, editors, *FLAIRS Conference*, pages 116–121. AAAI Press, 2006.
20. Heiko Schepperle and Klemens Bhm. Auction-based traffic management: Towards effective concurrent utilization of road intersections. In *CEC/EEE*, pages 105–112. IEEE, 2008.
21. Matteo Vasirani and Sascha Ossowski. A market-inspired approach for intersection management in urban road traffic networks. *Journal Artificial Intelligence Research*, 43:621–659, 2012.
22. Fei-Yue Wang. Agent-based control for networked traffic management systems. *IEEE Intelligent Systems*, 20(5):92–96, 2005.