An Optimal Velocity Planning Scheme for Vehicle Energy Efficiency through Probabilistic Prediction of Traffic Signal Timing

Grant Mahler Ardalan Vahidi

***Abstract*—The main contribution of this paper is the formulation of a predictive optimal velocity planning algorithm that uses probabilistic traffic Signal Phase And Timing (SPAT) information to increase a vehicle's energy efficiency. We introduce a signal phase prediction model which uses historically-averaged timing data and real-time phase data to determine the probability of green for upcoming traffic lights. In an optimal control framework, we then calculate the best velocity trajectory that maximizes the chance of going through greens. Case study results from a multi-signal simulation indicate that energy efficiency can be increased with probabilistic timing data and real-time phase data. Monte-Carlo simulations are used to confirm that the case study results are valid, on average. Finally, simulated vehicles are driven through a series of traffic signals, using recorded data from a real-world set of traffic-adaptive signals, to determine the applicability of these predictive models to various types of traffic signals.**

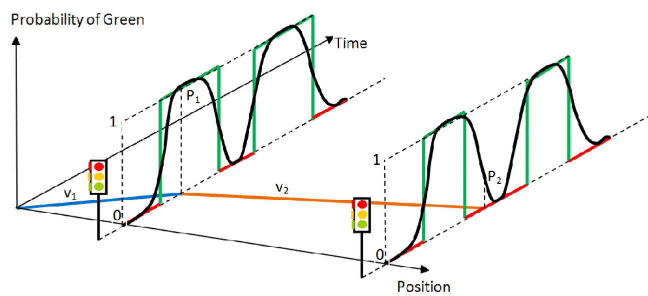

Fig. 1. Schematic of velocity planning based on probability of green for two consecutive lights

I. Introduction

A significant amount of fuel is spent by vehicles slowing down, sitting behind, and accelerating away from traffic signals [1], [2], [3]. With Corporate Average Fuel Economy standards set to rise, new technologies must be developed to meet the more stringent standards. Avoidance of red signals could improve vehicle specific fuel economy, reduce emissions, and help automotive manufacturers meet these new standards. While cities can spend time and money improving the timing of traffic signals [4], new research in velocity advisory algorithms suggests that it is possible to avoid red traffic signals through intelligent usage of traffic signal phase and timing information [5], [6], [7], [8], [9], [10], [11]. This benefit comes without the cost imposed by significant changes to infrastructure or production vehicles.

A velocity planning algorithm which guided a driver through multiple traffic signals given full and exact knowledge of the future state of traffic signals has previously been published by our group [5]. This algorithm could be implemented as a smart phone application displaying a suggested velocity to the driver, or it could provide the reference velocity to the adaptive cruise control system of a car. Simulation results using the algorithm in [5] indicated approximately 24% increase in fuel economy when passing through a series of simulated traffic signals. Koukoumidis et al. have presented an iPhone application to guide a driver through a light [7]. An underlying assumption in both papers is that accurate Signal Phase And Timing (SPAT) information is readily available. Unfortunately due to timing drift in fixed-time lights, and ever-changing traffic conditions for actuated and adaptive lights, accurate SPAT information is difficult to obtain.

In this paper we propose a method which utilizes either the base timing plan or historically averaged timing data (for example a 24 hour average), in complement with real-time signal phase data, to produce probabilistic SPAT predictions. A velocity planning algorithm then uses the prediction to reduce the chance of idling at a red light. A schematic is shown in Figure 1. This method can be implemented today as a result of its reliance only on real-time data instead of full-horizon data, and potentially overcomes issues of timing drift and unknown traffic conditions at actuated and adaptive signals.

In [7], a machine learning technique, Support Vector Machine (SVM), was used to predict the current phase of the light based on information from windscreen mounted iPhone cameras. Each iPhone utilized an image analysis algorithm to determine the current phase of the signal, and generated an ad-hoc wireless network to distribute this information to other iPhones running the same application. While SVM is one potential method of prediction, the application's reliance on ad-hoc wireless networks and image processing to provide signal timing has shortcomings. In [12], several traffic signal phase length prediction algorithms were presented, with the goal of utilizing the information in vehicle efficiency applications. However, the accuracy of the proposed methods as phase length predictor algorithms was deemed insufficient by the thesis author. In [13], the authors looked at clustering of velocity profiles from drivers who previously visited a road segment to determine traffic signal state estimations;

G. Mahler (corresponding author) is with the Department of Mechanical Engineering, Clemson University, Clemson, SC 29634 gmahler@clemson.edu.
A. Vahidi is with the Department of Mechanical Engineering, Clemson University, Clemson, SC 29634 avahidi@clemson.edu.

but the authors do not make phase length predictions. In [6], an iPhone application to re-route drivers around red lights was presented. Unfortunately the prediction method was not explicitly indicated. In current literature, there exists no comprehensive framework for velocity planning when inexact or incomplete traffic signal phase and timing information is available.

The goal of this paper is to fill the gap in current research, by developing SPAT prediction and probabilistic velocity planning algorithms that are applicable to both fixed time and actuated signals, in the presence of both exact and inexact red/green split information. An optimal control formulation for the velocity planning problem is presented in Section II. Section III describes a procedure for predicting, probabilistically, the future phase of a signal, based on its current phase and averaged timing data. A full array of simulations in Section IV expands on our initial results in [14] and statistically evaluates fuel economy gains attainable with our proposed methods versus cases with no signal information (baseline minimum efficiency) and full horizon signal information (maximum attainable efficiency). In particular in Section IV-A, we present a motiving simulation case study consisting of three consecutive traffic signals with our selection of phase and timing configurations. Next in Section IV-B we move to a simulation containing a large number of test cases, involving randomized traffic signal timings of known red/green split; this builds statistically significant evidence for use when the driver has access to exact red/green split information akin to a driver passing through a series of fixed time lights. To further demonstrate the applicability of proposed methods, in Section IV-C we show results for a series of consecutive traffic signal phase and timing configurations matching a series of real-world intersections. In these simulations we utilize actuated traffic signal timings, pre-recorded from an actual street. The goal of using the signal phase and timings from the real world is two-fold: first it allows the evaluation of the algorithm in the presence of actuated signals and inexact red/green splits, and secondly it reduces author influence on the road geometry and signal timings.

II. Optimal Velocity Planning

Our goal is to find a velocity profile which reduces the total energy consumption during a trip based on full or partial SPAT information. One can formulate this problem as an energy (fuel) minimization problem, but this requires inclusion of dynamic models of a specific vehicle and its propulsion system (combustion engine, etc.) to relate energy use to the velocity profile (for example, see Rakha et al.[15]). To avoid the ensuing computational complexity and to decouple the choice of optimal speed from a vehicle's make and model, we simplify the cost function by penalizing in it a weighted sum of the total trip time and all of the acceleration and decelerations, instead of total energy use. The underlying assumptions in this choice are that idling at a traffic light and excessive accelerations and decelerations induced by a traffic light cost energy with no benefits to the driver. Other factors such as motion constraints imposed by red intervals, road speed limits, and the fact that very low velocities will be unacceptable to consumers, can be accounted for by constraining the solution space. We evaluate the fuel economy for a specific vehicle model a-posteriori, by feeding the optimal velocity profile to a high-fidelity dynamic model of the vehicle.

We first describe, in Section II-A, the scenario when deterministic and accurate SPAT information over the entire planning horizon is available. When the phase and timing of upcoming signals are uncertain, a probabilistic term can be added to the cost function, as described in Section II-B.

A. Planning with Deterministic SPAT Information

To obtain a best achievable energy efficiency baseline, we first solve the optimal control problem assuming full and deterministic knowledge of signals' phase and timing over the planning horizon. The following cost function is used:

$$J = \sum_i \left[w_1 \frac{t_{i+1} - t_i}{\Delta t_{min}} + w_2 | \frac{a_i}{a_{max}} | + c(x_i, t_i) \frac{1}{\varepsilon} \right] \tag{1}$$

where J is the total cost and is indexed over position x with index i, $t_{i+1} - t_i$ is the time required for a vehicle to cover the distance between steps x_i and x_{i+1} given the velocity at x_i and the acceleration a_i; Δt_{min} is the minimum time to complete the step if starting and ending at the maximum velocity and is used as a scaling factor, a_i is the constant acceleration assumed during step i, and a_{max} is the maximum allowed acceleration. The constants w_1 and w_2 are weighting terms. Motion constraints imposed by a red interval are imposed as a soft constraint by inclusion of the term $c(x_i, t_i)\frac{1}{\varepsilon}$ in the cost function[1]. The value of $c(x_i, t_i)$ is zero except for spatiotemporal intervals when a light is red in which case its value is set to one, and ε is a very small constant (for example 10^{-6}), such that idling at red is discouraged.

The vehicle kinematics, realized by the following two-state equations, are imposed as equality constraints. Here x is the independent variable, velocity v and time t are the two states, and acceleration a is the input:

$$\begin{cases} \frac{dv}{dx} = \frac{a}{v} \\ \\ \frac{dt}{dx} = \frac{1}{v} \end{cases} \tag{2}$$

Discretizing the above equations with a constant sampling interval of $\Delta x = x_{i+1} - x_i$ and with a zero-order hold on acceleration, we obtain:

$$\begin{cases} v_{i+1} = \sqrt{(v_i)^2 + 2a_i \Delta x} \\ \\ t_{i+1} = t_i + \frac{2\Delta x}{v_i + \sqrt{(v_i)^2 + 2a_i \Delta x}} \end{cases} \tag{3}$$

We also enforce the hard inequality constraints: $v_{min} \leq v_i \leq v_{max}$ and $a_{min} \leq a_i \leq a_{max}$. Here v_{min} and v_{max} are the road speed limits and can also include lowest speed acceptable to a driver; $-a_{max}$ and a_{max} are the feasible bounds for deceleration and acceleration.

The above optimal control problem is solved numerically using Deterministic Dynamic Programming (DDP) and based

[1] Note that in simulations, a low level controller verifies and can override the recommendation of the velocity planner, if the planner makes a recommendation which would pass through a red light.

on the discretization on position, time, and velocity[16]. The DDP is solved by value function iterations for each stage, backwards. Using Bellman's principle of optimality, one only has to solve for one control input, here v_i. The trajectory is found recursively, instead of attempting to find the whole velocity trajectory at once.

B. Planning with Probabilistic SPAT Information

Because perfect full-horizon SPAT information is generally not available, a solution which takes advantage of currently available information and technologies is preferable. The goal is a solution that, given imperfect or incomplete phase and timing information, is still able to increase the energy efficiency by taking advantage of available data. Because of imperfect or incomplete starting data, this resulting energy efficiency is expected to be lower than the case with full-horizon information.

The cost function in (1) is modified to the following to take into account the probabilistic nature of SPAT information:

$$J = \sum_i \left[w_1 \frac{t_{i+1}-t_i}{\Delta t_{min}} + w_2 \left| \frac{a_i}{a_{max}} \right| + c(x_i,t_i) \left| \log_e \left(p(x_i,t_i) \right) \right| \right] \quad (4)$$

All parameters and variables in (4) are the same as those described for (1); the only new variable is $p(x_i,t_i)$ which represents probability of green at time t_i for a light situated at position x_i. Therefore higher costs are assigned to solutions that pass through time intervals where probability of green is lower. At the limit when probability of green at x_i,t_i is zero, $log_e(p(x_i,t_i)) = \infty$ and passing through a red would be discarded. Where $p(x_i,t_i) = 1$, this term of the cost function drops to zero and increases the likelihood that the corresponding velocity will be selected. The probability of green for each light can be generated based on real-time and/or historical information as described in Section III. Minimization of the cost function (4) with the equality and inequality constraints described in the previous subsection, remains a deterministic optimal control problem. The problem is solved using DDP but in a receding horizon manner; as new information becomes available, the DDP is re-solved taking into account the updated information over the remaining trip horizon.

III. Prediction

There can be much uncertainty in the phase and timing of a traffic signal which makes predicting its future state quite challenging. For fixed-time traffic signals which do not respond to traffic conditions and operate only on a timing table, we have confirmed the finding that the traffic signal clock drifts significantly during a 24 hour period. Therefore, it is not possible to know with certainty the start of greens and reds, even for fixed-time signals. The level of uncertainty is higher for actuated and adaptive traffic signals which do respond to traffic conditions. Although they have a base timing table, the timings of actuated and adaptive lights may change according to traffic conditions, rendering not only the start of reds and greens but also the phase lengths uncertain.

Due to the aforementioned uncertainties, it is difficult to determine the start and duration of greens deterministically. Therefore in this paper we employ a probabilistic prediction framework to handle the case with partial or uncertain information. We focus on cases where only i) the current phase (color) and ii) the average red and green lengths for a signal are known. We use this information to predict the probability of a green over the planning horizon.

Access to the current phase of the traffic signal is a major technological hurdle. However, solutions have been proposed and implemented in [6], [17] that could address this problem. Other approaches, including those that rely on Dedicated Short Range Communication (DSRC), can be found in [5], [18], [19].

Obtaining a base timing plan for a traffic signal is not trivial either. Direct access to signal timing plans is prohibitively difficult due to hundreds of local and federal entities that manage the more than 330,000 traffic lights across the United States [4]. To overcome these problems, it is possible to combine historical data, operating logic of signalized intersections, infrastructure sensor data, and even crowd source information to generate an average timing table. This can be done for different times in a day (rush hour/midday) and days of a week (weekday/weekend). The outcomes are average cycle times, and percentage of green and red in each travelling direction for each signal. Mere knowledge of such a baseline schedule, obtained offline and using only historical data, has statistical value even when the signal clock time is unknown.

Let us denote the state of a light by $\ell(t)$ which can assume two values, g and r, representing green and red respectively. We are interested in determining the probability of a light being green at time $t+t_p$ conditioned on its current color at time t. To form this conditional probability function, we assume the durations of green and red are known to be t_g and t_r on average. We also assume the traffic signal operates cyclically[2] and as a result the total cycle time[3] is fixed and equal to $t_g + t_r$. In this formulation, we assume the arrival of vehicles at the intersection to be uniformly distributed; if the arrival distribution of vehicles at an intersection is known (for example, in [20]), that distribution may be used as a weighting function in place of the uniform assumption. Using relatively straight-forward probabilistic reasoning, the chance of a green light in t_p seconds, given a green at current time t can be found to be:

$$P[\ell(t+t_p) = g | \ell(t) = g] = \begin{cases} \frac{t_g - t_m}{t_g} & t_m \le t_r, \quad t_m \le t_g \\ \frac{t_g - t_r}{t_g} & t_r \le t_m \le t_g \\ 0 & t_g \le t_m \le t_r \\ \frac{t_m - t_r}{t_g} & t_g \le t_m, \quad t_r \le t_m \end{cases} \quad (5)$$

where $t_m = mod(t_p, t_g + t_r)$ is the residue of division of t_p by $t_g + t_r$. In other words, because the signal clock is assumed to be periodic, the resulting conditional probability is also going to be a periodic function of time with the same period. Similarly, the chance of a green light in t_p seconds, given a red at time t is:

[2]This is true for many traffic signals; even many of those that react to traffic can theoretically have a fixed cycle time.

[3]We include the yellow time with red time; for safety reasons we do not make recommendations which would guide a driver through a yellow light.

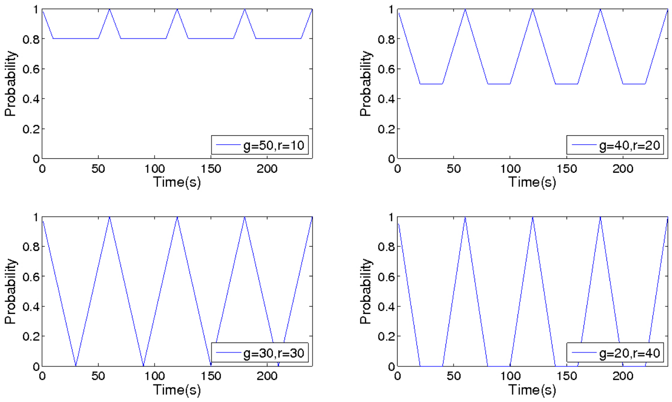

Fig. 2. Conditional future probability of green given that the light is currently green, for four different light timing patterns. In all patterns the total cycle time is 60 seconds, with the lengths of green and red indicated in the legends. The time axis is t_p as described in Equations 5 and 6.

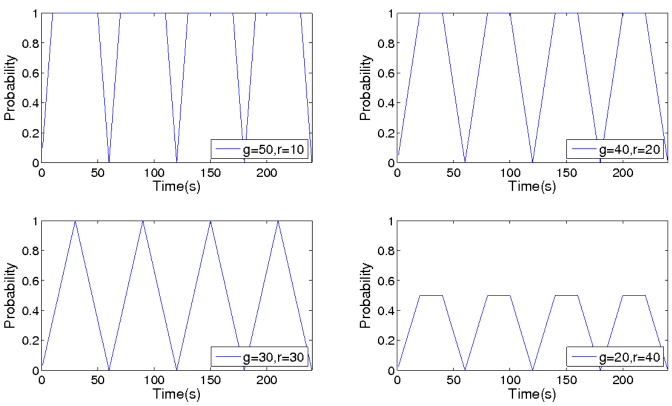

Fig. 3. Conditional future probability of green given that the light is currently red, for four different light timing patterns. In all patterns the total cycle time is 60 seconds, with the lengths of green and red indicated in the legends. The time axis is t_p as described in Equations 5 and 6.

$$P[\ell(t+t_p)=g|\ell(t)=r]=\begin{cases}\frac{t_m}{t_r} & t_m \le t_r, \quad t_m \le t_g \\ 1 & t_r \le t_m \le t_g \\ \frac{t_g}{t_r} & t_g \le t_m \le t_r \\ \frac{t_g+t_r-t_m}{t_r} & t_g \le t_m, \quad t_r \le t_m \end{cases} \tag{6}$$

Figures 2 and 3 show several probabilistic prediction examples with different splits between red and green but with the same cycle length. These are visualizations of the probabilities used in the probabilistic simulation cases described next.

IV. Simulations

A series of simulations were run in increasing order of complexity and increasing applicability to real world applications. The first step involved a simple case study, whereby the efficacy of the algorithm was examined in a generic scenario. The second step was to implement the algorithms as part of a Monte-Carlo simulation, whereby in each simulation the consecutive signal timing configurations were randomly adjusted to simulate spatio-temporal effects similar to that of varied intersection geometries. In the third step, traffic signal spacing and timing configurations were adapted from a semi-urban environment, further validating the applicability of the algorithms.

We evaluate three levels of SPAT information (none, deterministic, and probabilistic) in all of the following studies. In many of the simulations, a vehicle which is unaware of the future phase of traffic signals would have to alter the vehicle velocity for some of the traffic lights and stop at some of them. In many of the simulations, a driver with full SPAT information and sufficient space and time is able to avoid coming to a stop at any of the traffic signals. The real time information case, with probabilistic models, often falls somewhere in the middle.

In the scenario with probabilistic information, the optimal control problem in Section II-B was solved in a receding horizon manner and once per sample time using DDP. To simulate the uncertainty in the phase and timing of an actual traffic signal, a random number generator could be used to slightly and randomly shift the start of a green and change phase durations. At each sample time, a prediction of the probability of green was made for the remainder of the trip using only the current color of the lights and an assumed and fixed green/red split ratio as described in Section III; this prediction was fed to the DDP algorithm at each sample time. The recent behavior history of a signal was not accounted for in the prediction stage.

A maximum speed limit of 20 meters/second is enforced in all simulations, corresponding to an arterial road. The simulated driver is required to start at zero velocity and a terminal constraint is enforced such that the driver ends at zero velocity.

In all simulations, the penalty weights in the cost function J are set equally to empirically derived values of w_1=1/8, w_2=1/8. Weighting factors values are derived empirically such that simulated vehicles complete the distance in a reasonable time without violating red lights. This involved several, but not necessarily exhaustive, iterations. The value of ε is set at 10^{-6}. To solve the DDP, the solution space is discretized to distances of 20 meters, time increments of 1 second, and velocity steps of 1 meter/second. In this discretization grid choice, we have tried to maintain the computational time and memory requirements at a reasonable level without noticeably influencing the solution.

AUTONOMIE, a high fidelity vehicle simulation environment developed by Argonne National Laboratory, was used in calculating fuel economy. High fidelity models of a production car were assembled and a causal forward looking simulation was run by feeding each velocity profile to the software. The velocity profiles created previously were used as drive cycles; a driver model attempts to follow the drive cycle while respecting constraints on engine power, engine transient response time, acceleration or braking traction and other variables. In the motivating case study, MATLAB's ODE4 Runge/Kutta solver was used with a step size of 0.01 seconds and the simulation was stopped at the time the vehicle reached its destination.

In simulations where it was not computationally feasible to run all cases through a full high-fidelity AUTONOMIE simulation cycle, a simplified vehicle model was developed using efficiency maps taken from AUTONOMIE and a simplified gear shifting logic [4]. The simplified vehicle model is the same two-wheel-drive, automatic transmission, conventional-engine vehicle. This vehicle had a total mass of 1580 kg, an engine producing a peak of 115 kW, and a constant electrical load of 200 W. The velocity profiles generated by the dynamic program were fed to this model to calculate the fuel economy for each case. The simplified model provides a significant reduction in computational time when calculating the fuel economies for large numbers of simulation cases and we believe the fuel economy numbers will remain directionally valid, if not in the absolute sense[5].

A. Motivating Case Study

A case study was run as a motivating first step and involved a single simulation of a set of three consecutive signal timing configurations. The velocity profiles for the case studies with

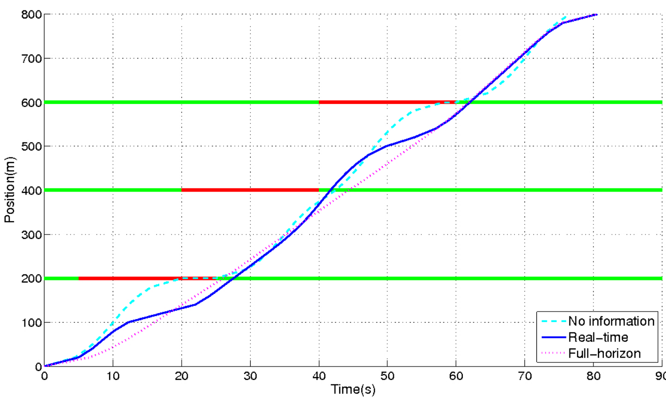

Fig. 4. Case Study: Comparison of velocity profiles of a driver with no information, real-time information, and full horizon information.

no advanced information, with read-time only information, and with full horizon information are shown in Figure 4. A DP solution for the full horizon information case has a smooth velocity profile, an uninformed driver must stop and start at lights, and a driver with access to real-time signal information is able to partially smooth her/his velocity profile . In Figure 4 the uninformed driver is required to come to a complete stop twice, for a combined total of about 7 seconds of idle time.

To help visualize how the decisions are made as the car receives more information and progresses along, time-lapses of the real-time information simulation in Figure 4 are shown in six subplots of Figure 5. In these subplots the information about the future color of a light is only probabilistic and is visualized by a red to green color spectrum. As the simulated driver approaches a traffic signal, the probability prediction becomes both more confident (i.e. probability becomes bimodal around either 1 or 0), and more relevant to the simulated driver. Bright red indicates the probability of green is near 0. Bright (neon) green indicates the probability of green is near 1. Dark reds and dark greens indicate probabilities in the middle. Because probability is only one of the terms in the cost function, at times it may appear as though the driver is moving

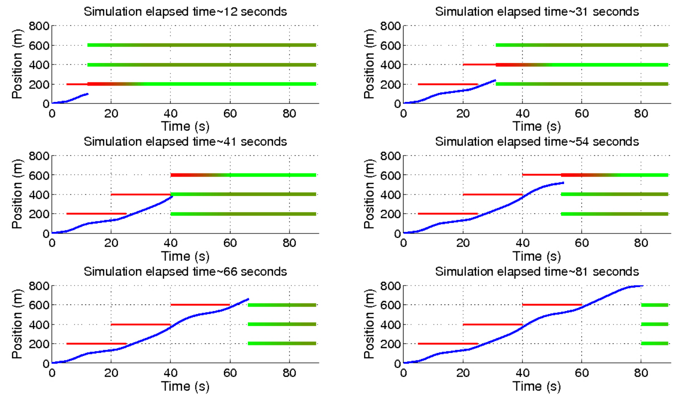

Fig. 5. Snapshots of a running simulation for Case Study 1. Probability prediction is indicated by the red to green color spectrum. See the video at the following link [21].

aggressively towards a light with a high probability of red; the color spectrum shown is only the value of probability of green, and does not reflect the total cost function. A simulation movie for the first case study can be found at [21].

TABLE I
FUEL ECONOMIES IN MILES PER GALLON (MPG). THE RESULTS SHOW THE IMPACT INFORMATION CAN HAVE ON THE ENERGY EFFICIENCY.

	Case Study
No Information	10.3
Real Time Information	16.6
Full Information	31.5

The fuel economy for each scenario was evaluated in AUTONOMIE and using the full vehicle model. The results presented in Table I are promising. With only real-time phase data and the probabilistic prediction model, a 61% increase in fuel economy over an uninformed driver is observed. This corresponds to 29% of the potential benefit of having full and exact future knowledge of SPAT information. Note that because the total simulation distance is only 800 meters with three signals, the fuel economy differences may be more exaggerated than average gains expected over driving cycles where traffic signals are less frequent.

B. Monte-Carlo Simulations

While the results of the preceding case study is promising, it is not clear if significant improvement in average fuel economy can be gained with the proposed algorithm, if relative offsets in the three signal timings are varied. In other words, it remains to verify that fuel economy gains were not solely a result of author-designed signal offsets. Therefore in this section we evaluate a statistically significant number of cases with randomly generated timing offsets; this is a variant on a Monte-Carlo experiment.

For these Monte-Carlo simulations, drivers with access to the three levels of information were run, in which the start of red phases were randomized within a window of sufficient length for the driver to complete the route. The total cycle length, and length of each red were kept constant. Also the proportion of red to green times across all simulations were constrained to be the same (this ratio is the average used for

[4]For example, the effects of engine start and stop transients on fuel economy were not modeled in the author-derived fuel economy calculations.

[5]The directionality of fuel economy results utilizing this simplified model matches those computed using AUTONOMIE as in [14]

the simulated signals, and could match a 24 hour or any other temporal average for a specific traffic signal). The start of the red phase of each signal was uniformly varied within the cycle so long as the full length of red was preserved. The start of each red of a traffic signal was chosen independently of the start of red of the next traffic signal. Three thousand simulated cases, with three traffic signals per simulation, with a simulation length of eight hundred meters were run (1000 simulations for each level of information: no information, real-time information, and full information).

The fuel economy for each of the The Monte-Carlo simulations was obtained by feeding the resulting velocity profile to our simplified vehicle model. The averaged results for each information level are summarized in Table II. The results indicate that for the road conditions described and with only real-time information and the probabilistic models, an average of 16% increase in fuel economy could be expected, representing approximately 62% of the benefit of full and exact traffic signal timing information.

TABLE II
MONTE-CARLO SIMULATION RESULTS REFLECT THE POSITIVE INFLUENCE OF INFORMATION, ON AVERAGE, ON FUEL ECONOMY.

	Mean(MPG)	Standard Deviation(MPG)
No Information	25.9	5.0
Real Time Information	29.9	3.7
Full Information	32.5	3.0

It was determined that the chosen 1000 simulations per information level is sufficient and captures reasonably well the average fuel economy for each information level. This can be seen in Figure 6 where the cumulative average fuel economy over the number of simulations is shown for each information level. For example, at $X = 100$, the Y value represents the average fuel economy of the first 100 simulations (for each level of information). By the time 1000 runs were simulated for each level of information, only very minor changes in average MPG occurred with the addition of more cases.

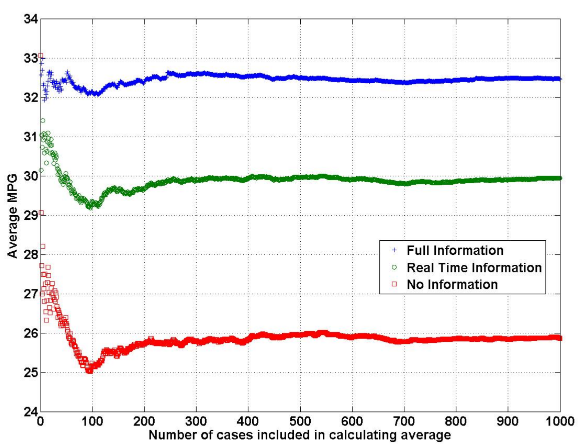

Fig. 6. Cumulative average fuel economy for the simulated 3 signal corridors.

Figure 7 shows histograms of Monte-Carlo fuel economy traces. As apparent in the figure and also shown in Table II, the standard deviation of the case where the driver has no information is highest, where the case where the driver has full information is lowest. One possible explanation of this is that the driver with no information has significantly different fuel economy when the driver has to stop at all lights versus no lights, whereas the driver with full information is able to achieve more uniform fuel economy results by smoothing the velocity and avoiding stopping.

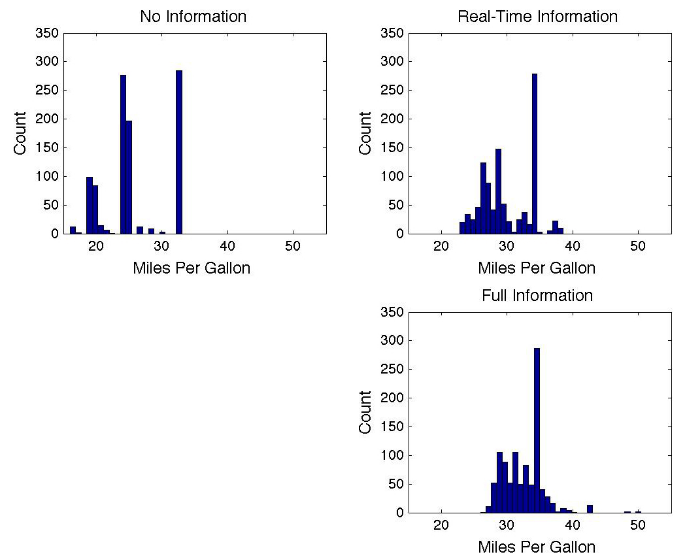

Fig. 7. Histogram of Monte-Carlo simulation results for the three levels of information.

The driver with no information is occasionally able to be faster and more efficient than the driver with some or full information. This situation may occur as a result of the no-information driver accelerating at full speed, and just passing before a light turns red. When this occurs, the driver is able to significantly reduce total trip time, which in turn reduces fuel usage. Because the real time or full information drivers have cost functions (Equations 1 and 4) which place some weight on acceleration, the optimal solution is rarely maximum acceleration. The driver with no information applied full acceleration and full deceleration as necessary, reflecting an aggressive driver.

C. Simulations Using Recorded Timings from Arterial Adaptive Lights

With Monte-Carlo simulations indicating positive relationships between future information about traffic signals and fuel economy, we turn to more realistic examples in which we use signal timing and geometry of actual intersections. Toward this goal, we have obtained the timing of signals of three intersections from the city of Fremont in California.The lengths of the green phases in the direction of travel for these simulations for a 24 hour period can be found in Figure 8. From this Figure, it is clear that the traffic signals are not fixed, they respond to traffic conditions. Use of actual signal timings and actual offsets between multiple intersections, reduces any unintended bias that may have been present in our Monte Carlo simulation design. Moreover, we show in this section that while our proposed algorithms were developed for fixed-time signals, they are robust to variance in nominal traffic signal timing and could potentially be used even in the presence of actuated traffic signals.

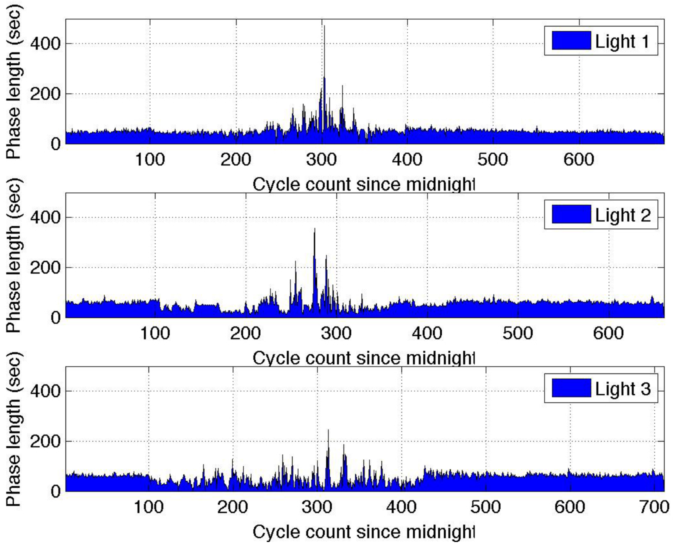

Fig. 8. Histories of relevant phases for each light along the chosen real-world route, for every cycle over a 24 hour period (midnight to midnight).

A vehicle was simulated driving through the three traffic signals every 10 minutes over the 24 hours yielding a total of 144 simulated drives per level of information. The real-world distance between the signals is preserved in the simulation, such that the simulated vehicle has to cover the same distance using the same traffic signal timing offsets as a real driver would encounter. The total simulation distance is 1320 meters. The first light occurs 520 meters into the simulation. The second light occurs 280 meters later. The third and final light occurs at 1200 meters from the start. The resolution of the dynamic programming algorithm was kept similar to simulations in Section IV-B. Velocity resolution remains at 1 meter per second, distance resolution remains at 20 meters, time resolution remains at 1 second. The traffic signal phase and timings are taken from recorded data from the city and are merely played back into this simulation. No other vehicles are considered to be on the road - the only obstacles the vehicle routing algorithm must avoid are the lights themselves. For the purposes of prediction, the real-time simulation is given a 24 hour average of red and green lengths, though if more relevant averages (for example a short-term average, a time of day average, or other statistical means) are available, they may improve the performance of this case.

TABLE III
FUEL ECONOMY RESULTS FROM RECORDED REAL-WORLD TRAFFIC SIGNAL TIMINGS WITH SIMULATED VEHICLES MOVING BETWEEN THE LIGHTS REFLECT THE POSITIVE INFLUENCE OF INFORMATION.

	Mean(MPG)	Standard Deviation(MPG)
No Information	31.7	3.1
Real Time Information	33.7	3.0
Full Information	34.5	3.6

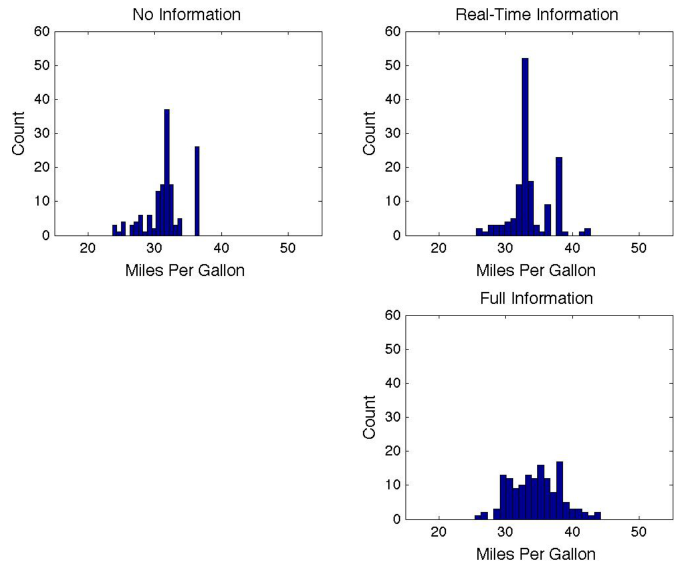

Fig. 9. Histogram of fuel economies for real-world traffic signal timings.

Fuel economy was calculated in the same manner as the Monte-Carlo simulations, due to the calculation time of AUTONOMIE for the number of drive cycles considered. The results of these fuel economy calculations can be found in Figure 9, and the means of those simulations found in Table III. The maximum error between the velocity profile generated and the velocity profile followed by the fuel economy calculations was 3.6%. Figure 9 and Table III both confirm the positive influence of information on fuel economy.

The real-world data used in Section IV-C provides another step towards implementability in comparison to the Monte Carlo simulations of Subsection IV-B. The use of recorded traffic signal data reduces the possibility of author-induced errors, e.g. inadvertently creating red or green waves. The simulation indicates that drivers with access to real-time information were able to improve fuel economy over drivers with no information by approximately 6%. This accounts for roughly 70% of the potential gains available through access to full and exact future knowledge of traffic signal timing. The particular implementation chosen here allows finding the new optimal trajectory based on the most current data by recalculating the cost to go and control matrices at each position step. This is one technique for dealing with unexpected traffic, pedestrians crossing out of cross-walks, and other disturbances.

In the real-world data, at some times throughout the day the base timing plan of some of the traffic signals have cycle lengths of 60 seconds. The decision to simulate a vehicle every 10 minutes (a multiple of the 1 minute nominal cycle length) is therefore a potentially problematic choice. However, as the traffic signals are adapting to traffic conditions, the cycle length and splits adjust, reducing the chance of aliasing of results. This ensures that the simulations are not 144 repetitions of the same cycle (the variation in simulations can be confirmed by reviewing Figure 9). Additionally, the timing plans of the various traffic signals appear to be lacking in synchronization, as even late at night and early in the morning, when neither vehicles nor pedestrians make calls, the lights drift with respect to each other.

In general, in a coordinated series of signals, under actuated or adaptive control logic, the offsets and other signal timing parameters may play an increasingly important role - for example in creating a green wave. If the proposed probabilistic models are able to assist a driver in joining a green wave, this may have a positive effect on fuel economy for that driver.

V. Conclusions

This paper statistically evaluated velocity planning algorithms which minimize idling time behind red lights based on probabilistic traffic signal timing models that we proposed. Three cases were evaluated - vehicles with no information about upcoming traffic signals, vehicles with real-time information, and vehicles with full and exact future knowledge of traffic signal timings. Drivers today fit into the first case - the least efficient. Drivers of the future may fit into the third and most efficient case, if infrastructure and technologies develop to provide this information. The middle case is feasible today, and obtains much of the potential benefit obtainable via knowledge of upcoming traffic signal timing.

Real time knowledge with probabilistic models where the driver encountered fixed-time lights yielded an optimistic 61% increase in a motivating case study, and a 16% increase in average fuel economy across 1000 multi-signal simulations of fixed time signals. The same models in combination with real time information yielded a 6% increase in fuel economy for actuated signals. These reflect technologies which could be feasibly implemented with little or no infrastructure changes and with only software updates to current production vehicles.

Future work will focus on the network wide effect of informed vehicles on uninformed vehicles, penetration rates, and fleet efficiency.

VI. Acknowledgement

The authors thank Dr. Chen Zhang for his assistance regarding shaping the DP discretization. This work was supported by the National Science Foundation grant number CMMI-0928533, and by BMW Group.

References

[1] "Traffic Congestion and Urban Mobility," http://tti.tamu.edu/infofor/media/topics/congestion_mobility.htm, Accessed:2010-01-25.

[2] S. M. Chin, O. Franzese, D. L. Greene, H. L. Hwan, and R. C. Gibson, "Temporary losses of highway capacity and impacts on performance: Phase 2," Tech. Rep., U.S. Department of Transportation, Oak Ridge National Laboratory, 2004.

[3] M. Sivak, "Effects of vehicle fuel economy, distance travelled, and vehicle load on the amount of fuel used for personal transportation in the U.S.: 1970-2010," Tech. Rep. UMTRI-2013-10, University of Michigan Transportation Research Institute, February 2013.

[4] U.S. Department of Transportation, "Intelligent transportation systems for traffic signal control: Deployment benefits and lessons learned," http://ntl.bts.gov/lib/jpodocs/brochure/14321_files/a1019-tsc_digital_n3.pdf, Accessed:2013-04-29.

[5] B. Asadi and A. Vahidi, "Predictive cruise control: Utilizing upcoming traffic signal information for improving fuel economy and reducing trip time," *IEEE Transactions on Control Systems Technology*, vol. 19, pp. 707–714, 2010.

[6] J. Apple, P. Chang, A. Clauson, H. Dixon, H. Fakhoury, M. Ginsberg, E. Keenan, A. Leighton, K. Scavezze, and B. Smith, "Green Driver: AI in a microcosm," in *AAAI Conference on Artificial Intelligence*, San Francisco, CA, August 2011.

[7] E. Koukoumidis, L.-S. Peh, and M. Martonosi, "Signalguru: Leveraging mobile phones for collaborative traffic signal schedule advisory," in *9th International Conference on Mobile Systems, Applications, and Services (MobiSys)*, Bethesda, Maryland, June 2011.

[8] P. Schuricht and B. Baker, "Efficiency-increasing driver assistance at signalized intersections using predictive traffic state estimation," in *Proceeding of 14th International IEEE Conference on Intelligent Transportation Systems*, 2011.

[9] S. Mandava, K. Boriboonsomsin, and M. Barth, "Arterial velocity planning based on traffic signal information under light traffic conditions," in *Proceeding of 12th International IEEE Conference on Intelligent Transportation Systems*, St. Louis, MO, October 2009.

[10] J. Wollaeger, S. Kumar, S. Onori, D. Filev, U. Ozguner, G. Rizzoni, and S. Di Cairano, "Cloud-computing based velocity profile generation for minimum fuel consumption: A dynamic programming based solution," in *American Control Conference*, Montreal, Canada, June 2012.

[11] J. Stevanovic A. Stevanovic and C. Kergaye, "Impact of signal phasing information accuracy on green light optimized speed advisory systems," in *Transportation Research Board Annual Meeting 2013*, Washington, D.C., January 2013.

[12] V. Khakhutskyy, "Signal phase and timing prediction for intelligent transportation systems," M.S. thesis, Der Technischen Universitat Munschen, 2011.

[13] M. Kerper, C. Wewetzer, A. Sasse, and M. Mauve, "Learning traffic light phase schedules from velocity profiles in the cloud," in *5th IFIP International Conference on New Technologies, Mobility and Security*, Istanbul, Turkey, May 2012.

[14] G. Mahler and A. Vahidi, "Reducing idling at red lights based on probabilistic prediction of traffic signal timings," in *American Control Conference*, Montreal, Canada, June 2012.

[15] H. Rakha, K. Ahn, K. Moran, B. Saerens, and E. Van Den Bulck, "Virginia Tech Comprehensive Power-Based Fuel Consumption Model: Model Development and Testing," *Transportation Research Part D: Transport and Environment*, vol. 16, no. 7, pp. 492–503, October 2011.

[16] D. E. Kirk, *Optimal Control Theory: An Introduction*, Dover Publications, 2004.

[17] Federal Highway Administration, "Investigating advanced traffic signal control," http://www.fhwa.dot.gov/advancedresearch/pubs/11044/11044.pdf, Accessed:2013-04-29.

[18] Audi, "Audi travolution: efficiently through the city," http://www.audiusanews.com/newsrelease.do?id=1812, Accessed:2013-04-29.

[19] U.S. Department of Transportation, "Connected vehicle: Vehicle-to-vehicle and vehicle-to-infrastructure technology test bed," http://www.its.dot.gov/connected_vehicle/technology_testbed2.htm, Accessed:2011-08-24.

[20] C. Day and D. Bullock, "Computational efficiency of alternative arterial offset optimization algorithms," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2259, pp. 37–47, 2011.

[21] G. Mahler, "Real-time traffic signal data driver simulation," http://www.youtube.com/watch?v=STZiyzi0k5M, Accessed:2013-04-29.

Grant Mahler received his B.S. in mechanical engineering from Northwestern University in 2008. He is currently a Ph.D. candidate in mechanical engineering at Clemson University in Clemson, South Carolina. He was previously an intern at the BMW Information Technology Center in Munich, Germany, as well as a visiting scholar at the BMW Information Technology Research Center in Spartanburg, South Carolina. He is currently a visiting researcher at the BMW Group Technology Office USA in Mountain View, California, where his research focus remains on connected vehicle technologies.

Ardalan Vahidi is currently an Associate Professor with the Department of Mechanical Engineering, Clemson University, Clemson, South Carolina. He received the Ph.D. degree in mechanical engineering from the University of Michigan, Ann Arbor, in 2005. the M.Sc. degree in transportation safety from George Washington University, Washington, DC, in 2002, and B.S. and M.Sc. degrees in civil engineering from Sharif University, Tehran, Iran, in 1996 and 1998, respectively. He has been a visiting scholar at the University of California, Berkeley and a visiting researcher at the BMW Group Technology Office USA in 2012-2013. His current research interests include control of vehicular and energy systems, and connected vehicle technologies.