Energy minimization for an electric bus using a genetic algorithm

Downloaded from: <https://research.chalmers.se>, 2023-04-26 12:33 UTC

Citation for the original published paper (version of record):

Torabi, S., Bellone, M., Wahde, M. (2020). Energy minimization for an electric bus using a genetic algorithm. *European Transport Research Review*, 12(1).
<http://dx.doi.org/10.1186/s12544-019-0393-1>

N.B. When citing this work, cite the original published paper.

ORIGINAL PAPER

Open Access

Energy minimization for an electric bus using a genetic algorithm

Sina Torabi^{*} , Mauro Bellone and Mattias Wahde

Abstract

Background and methods: This paper addresses, in simulation, energy minimization of an autonomous electric minibus operating in an urban environment. Two different case studies have been considered, each involving a total of 10 different 27 km bus routes and two different average speeds. In the proposed method, the minibus follows an optimized speed profile, generated using a genetic algorithm.

Results: In the first case study the vehicle was able to reduce its energy consumption by around 7 to 12% relative to a baseline case in which it maintains a constant speed between stops, with short acceleration and deceleration phases. In the second case study, involving mass variation (passengers entering and alighting) it was demonstrated that the number of round trips that can be completed on a single battery charge is increased by around 10% using the proposed method.

Keywords: Energy minimization, Autonomous vehicles, Electric minibuses, Speed profile optimization, Genetic algorithms

1 Introduction

In recent years, to meet the high demands of reducing greenhouse gas emissions in transportation, the public transportation sector has started to include more electric vehicles (such as city buses) in their fleets. In addition to providing lower emission and environmental pollution compared to vehicles equipped with internal combustion engines, electric vehicles offer more high-level control opportunities to increase their energy efficiency and longevity.

This study has been carried out in the framework of the Sohjoa Baltic project,¹ which involves pilot studies in several cities of the Baltic sea region, namely Kongsberg, Helsinki, Tallinn, Gdansk, and Vejle. The project aims at improving the first-last mile part of public transportation to provide a ubiquitous and efficient service by using autonomous electric minibuses. The initial study within this project has identified several gaps to be overcome, see Ainsalu et al. [1], one of which is the fact that energy efficiency in electric buses must be improved for them to provide a stable service as current vehicles require

frequent and time-consuming recharging cycles. Indeed, the current technology typically requires vehicles to stop for between 6 and 8 h to recharge. Thus, extending the range of such vehicles (on a single battery charge) is a highly relevant issue.

This paper is focused on energy optimization for an autonomous urban electric bus with regenerative braking, operating under a drive cycle involving frequent stops with short distance between. The method used here is based on the speed profile optimization (SPO) framework introduced and successfully applied in earlier work; see Torabi and Wahde [2]. The main contribution of this paper is an extension of that method to the case of urban electric minibuses with regenerative braking.

This paper is organized as follows: Sect. 2 provides a brief overview of related work. The vehicle model, and the concepts of road and speed profiles are described in Sect. 3, along with implementation details. In Sect. 4, the procedure used for evaluating the speed profiles and the optimization of the speed profiles are presented along with a description of the baseline case against which the optimized speed profiles are compared. The road profile used in the simulations is introduced in Sect. 5. The results are presented in Sect. 6 and are discussed in Sect. 7. Finally, the conclusions are presented in Sect. 8.

¹see www.sohjoabaltic.eu

* Correspondence: sina.torabi@chalmers.se
Chalmers University of Technology, Gothenburg, Sweden

2 Related work

Current work on reducing energy consumption of vehicles and improving their energy efficiency can be divided into three general categories, namely improving drivetrain design, developing new drivetrain systems (such as in hybrid and fully electric engines), and optimizing the utilization of existing vehicle technologies such as, but not limited to, energy-efficient driving (or eco-driving) that can be achieved by following an optimized speed profile.

For heavy-duty vehicles (HDVs) with combustion engines, there has been much work on reducing fuel consumption (see e.g. Henzler et al. [3], Hellström et al. [4], and Torabi and Wahde [2]) something that involves generating energy efficient speed profiles, i.e. the speed of the vehicle as a function of its position along the road. These strategies result in fuel savings from 3% to 12% on average (relative to the case where the vehicle drives using standard cruise control). On the other hand, for hybrid electric vehicles (HEVs) energy optimization strategies typically involve efficient management of the energy flow between the batteries and the fuel sources (see e.g. Wirasingha et al. [5], Moura et al. [6], and He et al. [7]), velocity trajectory optimization (see e.g. Gue et al. [8] and Mensing et al. [9]) and battery charge sustaining strategies (see e.g. Mets et al. [10] and Bashash et al. [11]). For fully electric vehicles, however, only few methods have been proposed for generating optimized speed (or torque) profiles (see e.g. Dib et al. [12], Rohkämper et al. [13], and Jiang et al. [14]) similar to those optimized for vehicles equipped with internal combustion engines.

The problem of energy-efficient driving is commonly addressed using optimal control. With this formulation, the problem of finding an optimized speed profile for a given road is often solved numerically using dynamic programming (DP) both for vehicles with combustion engines, see e.g. Henzler et al. [3], and for hybrid and fully electric vehicles, see e.g. Mensing et al. [9], Guo et al. [8], and Dib et al. [12]. Speed trajectory optimization for HEVs has been carried out in various ways. For example, in Mensing et al. [9], the optimized velocity is found by solving a single-objective optimization problem in which the cost function is a weighted sum of the fuel consumption and the battery state-of-charge (SOC). Their approach was compared to real-life drive cycles for a specific route with various weight settings (in the objective function) during optimization, resulting in fuel savings of 10% with similar battery discharge over the considered cases. In Gue et al. [8], an energy management strategy for an HEV using a bi-level model predictive control (MPC) is proposed, where the torque split ratio, gear shift schedule, and the velocity trajectory is optimized for a given road. The proposed bi-level MPC improves the computational efficiency of the

common DP-based methods while resulting in similarly energy-efficient speed trajectories. This is achieved by first generating an optimized speed trajectory in an outer loop and then determining the optimal torque and gear scheduling in an inner loop so that the vehicle can track the speed trajectory as closely as possible.

Similar DP-based speed trajectory optimization approaches have been introduced for fully electric vehicles as well; see e.g. Dib et al. [12], Jiang et al. [14], Maamria et al. [15], and Rios-Torres et al. [16]. Energy-efficient driving strategies developed for fully electric vehicles and HEVs are typically considered as driving assistance systems that provide reference speed recommendations for a human driver to follow. The energy savings obtained using these strategies (i.e. when the driver follows the recommended speed trajectory) are often in the range of 9% to 14% relative to the case in which a human driver controls the vehicle's motion during the drive cycle without any recommendations. Even though the reported energy savings are high in DP-based methods, they are not directly comparable to the autonomous urban bus driving case considered here, since they involve a human driver for defining the baseline case, and it is unclear how that person drove in the absence of recommendations (for their baseline case). Moreover, DP-based methods suffer from the curse of dimensionality and are thus computationally heavy. Additionally, in order to use DP algorithms, one should discretize the problem so as to ensure that the optimization problem is solvable in a reasonable time. Therefore, DP-based methods are often used with large discretization steps and short optimization horizon which decreases the performance of the obtained solutions.

An alternative approach is to solve the optimization problem using stochastic optimization methods such as genetic algorithms. In Torabi and Wahde [2] for example, the fuel savings obtained by their speed profile optimization (SPO) method using genetic algorithms compares favorably to those obtained by methods based on dynamic programming for HDVs equipped with internal combustion engines. Moreover, the SPO framework can be used without discretization of the problem while also providing longer optimization horizon.

3 Model and representation

This section starts with a description of the electric vehicle model used to calculate the energy consumption of an electric vehicle and its battery state-of-charge. Next, the road and speed profiles are introduced.

3.1 Vehicle model

The longitudinal vehicle model used here is implemented based on a computationally efficient energy consumption

model proposed in Genikomsakis and Mitrentsis [17]. This model, in addition to transforming the traction power requirement at the wheels into battery power requirements using a generic physics-based vehicle model, provides an estimation of the battery's state-of-charge, while also considering the resistive forces such as aerodynamic drag and rolling resistance. Moreover, the model considers regenerative braking system into the battery state of charge model. The energy consumption of the vehicle is calculated by integrating the following differential equation:

$$dE/dt = P_{tot} \quad (1)$$

where E is the energy consumption and P_{tot} is the total power provided by the battery which is computed as

$$P_{tot} = P_b \eta_{RTE}, \quad (2)$$

$$P_b = P_m + P_{aux}. \quad (3)$$

Here, P_b denotes the battery's output power. P_m represents the required power by the electric motor whereas P_{aux} is the auxiliary power, i.e. the power needed to operate the vehicle's auxiliary systems (a value that is constant throughout the simulation). Moreover, η_{RTE} (slightly smaller than 1) represents the battery's round trip efficiency, the precise value of which depends on whether the battery is in charging or discharging mode; see Eq. (6) below. The battery's state-of-charge (S) is updated as follows:

$$dS/dt = -P_b/C, \quad (4)$$

where C denotes the battery's capacity (in Ws). Finally, the electric motor's required power is calculated by considering the longitudinal dynamics of the vehicle and its engine efficiency:

$$P_t = v \times (m(1 + C_i)a + F_{aero} + F_g + F_{roll}) \quad (5)$$

$$P_m = \begin{cases} P_t \eta_r, & \text{if } P_t < 0 \\ P_t / \eta_t, & \text{if } P_t > 0 \end{cases} \quad (6)$$

where P_t is the traction power required from the electric motor, v denotes the vehicle's speed, m is the mass, C_i is the correction factor for the rotational inertia force, and a is the vehicle's acceleration. The terms F_{aero} , F_g and F_{roll} represent the aerodynamic force, gravity force, and the rolling resistance force acting on the vehicle, respectively. Finally, η_r and η_t are the electric motor's efficiency in regenerative (charging) mode and in traction (discharging) mode, respectively. For a full description of the model and its parameter values, see Genikomsakis and Mitrentsis [17].

3.2 Road profile

In order to evaluate speed profiles in the simulation, a road profile is needed that represents the elevation as a function of longitudinal distance along the road. The road profiles have been defined using composite Bézier curves, i.e. a sequence of Bézier splines, as in Torabi and Wahde [2]. Since only the longitudinal motion is considered, a road profile can be modelled using two-dimensional composite Bézier curves where the two dimensions are the longitudinal position and elevation along the road. Thus, the road profiles are represented in the following form:

$$(s, z) \equiv (s_i(u), z_i(u)) = P_0(1-u)^3 + 3P_1u(1-u)^2 + 3P_2u^2(1-u) + P_3u^3. \quad i = 0, \dots, n-1 \quad (7)$$

where the vectors P_i are two-dimensional control points, u is a parameter ranging from 0 to 1, and n is the total number of splines. With this representation, at any given point s along the road, the corresponding spline index and the u -value can be determined. Therefore, the elevation can be written as $z = z(s)$.

3.3 Speed profile

A speed profile is represented using two-dimensional composite Bézier curves, similar to the road profile representation defined in Eq. (7) above, where the two dimensions are the longitudinal position and the longitudinal speed of the vehicle. An example is given in Fig. 1, in which the stops can be clearly seen. The speed profile provides a desired speed at any given position along the road, for the vehicle to follow using a PID controller. Note that, when the vehicle is at a stop, and is about to leave, the desired speed is artificially inflated to a small positive value, in order to start the motion. Once the vehicle's speed is positive, it then follows the speed profile.

4 Methods

This section first describes the procedure for evaluating a given speed profile. Next, the speed profile optimization method is described, and then the baseline method against which the speed profile optimization results are compared.

4.1 Evaluation of a speed profile

Assuming that a speed profile is available, the energy consumption of a vehicle can be measured by simulating the motion of the vehicle as it follows the speed profile. A simulation environment was written (in C#. NET) implementing the vehicle model described in Sect. 3.1 and the road and speed profiles defined in Sects. 3.2 and 3.3. To evaluate a speed profile performance in terms of energy consumption over a given road profile, the following procedure is used: At each time step, the current

Fig. 1 Top panel: A 2 km road profile with slope variations within $\pm 1\%$. The boxes represent the bus stops along the route with the final stop at the travelled distance of 2 km. The figure shows the complete profile for an entire drive cycle, from start to end, and back again. Bottom panel: A corresponding optimized speed profile, for a case with an average speed between stops of 10 km/h; see Sect. 6.1

longitudinal location of the vehicle along the road is determined, and it is then used to obtain the desired speed (at the same location) from the speed profile. Thus, the speed profile acts as a lookup table. The desired speed is then fed to the PID controller as its set speed. Note that the set speed varies over a road profile. Therefore, due to the controller's non-zero response time, the vehicle will not follow the set speed exactly, and neither is it required to do so in the method presented here. The procedure is repeated, time step by time step, until the vehicle reaches the end of the road profile.

In this work, it is assumed that the motion of the vehicle takes place in a dedicated lane so that it is able to follow its speed profile. It is further assumed that there are no traffic lights along the bus routes. However, if traffic lights were present, they could be dealt with as additional stops; see Sect. 7 for detailed discussion.

4.2 Speed profile optimization

The speed profile optimization is carried out with respect to energy consumption only, using a genetic algorithm (GA). A speed profile must fulfil certain constraints, namely (i) the instantaneous maximum speed must not exceed an upper limit (the speed limit), (ii) the average speed should not be below a certain user-defined threshold, and (iii) the final battery state of charge (for the electric vehicle) should be above a threshold to prevent battery damage.

As in any standard GA, the algorithm maintains a population of M individuals (here 20), each defining a speed profile. The speed profiles are encoded in chromosomes that specify the speed values (i.e. the second dimension of the control points in Eq. (7)), with the requirement of positional continuity (C0) and derivative

continuity (C1), which together ensure that the speed profile is continuous and smooth after decoding. In order to meet these two requirements, the following conditions must hold:

$$P_{i+1,0} = P_{i,3} \quad \text{for } i = 0, \dots, n-1, \quad (8)$$

$$P_{i+1,1} - P_{i+1,0} = P_{i,3} - P_{i,2} \quad \text{for } i = 0, \dots, n-1, \quad (9)$$

where $P_{i,j}$ is the j^{th} control point of the i^{th} spline. Moreover, the speed at the start and the end of a profile should be set to 0. The first spline thus requires three parameters, the following $n-2$ splines two parameters each, and the final spline one parameter. Therefore, each chromosome contains a total of $3 + 2(n-2) + 1 = 2n$ parameters.

For a given individual, the decoding procedure results in a speed profile that is evaluated as described in Sect. 4.1. The evaluation results in a single number, namely the energy consumption E_{end} obtained when the vehicle reaches the end of the profile. If the maximum speed or the final battery state-of-charge constraints are violated, the speed profile is assigned zero fitness. If not, the fitness measure f is taken as

$$f = \begin{cases} \frac{1}{E_{\text{end}}} e^{-c_v \left(\frac{\bar{v}}{\bar{v}_d} - 1 \right)^2}, & \text{if } |\bar{v} - \bar{v}_d| < 2 \frac{km}{h} \\ 0, & \text{if } |\bar{v} - \bar{v}_d| > 2 \frac{km}{h} \end{cases} \quad (10)$$

where $\bar{v}$ is the actual average speed, $\bar{v}_d$ is the desired average speed (see Sect. 6), and c_v is a positive constant.

Selection is carried out using standard tournament selection with a given tournament size (typically 2–5) and a tournament probability of p_t (around 0.7–0.8). Single-point crossover with a probability of p_{cross}

(around 0.7–0.9) is used. Once two new individuals are formed through selection and crossover, mutation is applied with probability of p_{mut} (around 0.05–0.1) with either full-range mutation (with probability of 0.5) or creep mutation (with probability of 0.5) with a creep rate of 10% of the full range. In order to make sure that the best speed profile found so far is not lost when forming the next generation (in the GA), this profile is passed unchanged to the next generation, a concept known as *elitism*.

4.3 Baseline method

The baseline case roughly mimics a typical driving sequence performed by a bus driver, albeit executed in a more precise manner (with a PID controller), see e.g. Pourabdollah et al. [18]. In this case, a speed profile is defined with constant speed (referred to as the cruising speed) between stops. The first and last splines of a baseline speed profile define the acceleration and deceleration phase, respectively. The cruising speed was set so as to achieve a given average speed; see Sect. 6 below for numerical values. An example of a baseline case is given in Fig. 2 below.

5 Road data

Ten synthetic road profiles of 2 km length were generated corresponding to short urban bus routes. The generated bus routes were defined in such a way that their slope variations were within $\pm 1\%$ for five routes (Route group I) and within $\pm 3\%$ for the other five (Route group II). There were four stops in total, at 0, 500, 1500, and 2000 m from the starting point; see also the top panel of Fig. 1. The generated road profiles were similar to a bus route in Kongsberg, Norway (considered within the framework of the Sohjoa Baltic project) which is roughly 2 km long, with a total of 4 stops, and with small slope variation. The road profiles were defined with one spline for every 500 m, which is sufficient given the mild slope variation.

6 Simulation results

In order to evaluate the performance of the speed profile optimization method, two case studies have been

considered. For each of the road profiles, the desired average speed was set either to 10 km/h or 15 km/h, which are suitable values for cases where the distance between the bus stops are rather short; Oskarbski et al. [19]. The maximum speed was set to 40 km/h. Moreover, in all cases below, the vehicle's battery initial state of charge was set to 95% at the start of the trip and the mass of the vehicle (when empty) was set to 2000 kg. In order to form a baseline case against which the optimization algorithm can be compared, the energy consumption of the bus was computed using the method described in Sect. 4.3 with same average speed, initial battery state of charge, and mass. The performance is typically measured either as the total energy saving (i.e. E in Eq. (1); see also Dib et al. [12]), or the battery usage savings (i.e. S in Eq. (4); see also Mensing et al. [9]). Those two values differ slightly, due to the battery round-trip efficiency coefficient; see Eq. (2). Here, both measures will be given. In all cases (speed profile optimization and the baseline case), the speed profiles consisted of $n = 40$ splines so that each spline covered 50 m.

6.1 Case study I

In this case study, the optimization was carried out for a single round trip, i.e. total of 4 km driving, for both of the bus route groups. For the purpose of comparison, the mass of the bus was fixed to the initial value (2000 kg) both for the baseline case and for the speed profile optimization.

The results obtained from the simulations (henceforth: Case study I) are presented in Table 1 for each of the bus route groups (Groups I and II). The second column from the left shows the required average speed between consecutive bus stops, and the two right-most columns show the energy savings and the battery usage savings, respectively. Both the energy savings and the battery usage savings are calculated relative to the energy consumption and battery usage of the vehicle when following the speed profiles defined using the baseline method. The results obtained from the simulations show that by following the optimized speed profiles, the total energy consumption of the vehicle is reduced by an average of

Fig. 2 An example of a baseline speed profile for an entire driving cycle (from start to end, and back again). In this case, the average speed between stops was set to 10 km/h

Table 1 Simulation results obtained for Case study I. The second column from the left shows the average speed of the vehicle between stops (set to either 10 or 15 km/h). The two right-most columns show the energy savings and battery usage savings relative to the energy consumption and battery usage of the baseline case, respectively

Bus route	Average speed (km/h)	Energy savings (%)	Battery usage savings (%)
Group I	10	12.62 ± 2.36	12.44 ± 2.69
	15	10.19 ± 1.69	10.50 ± 1.57
Group II	10	9.50 ± 1.74	9.49 ± 1.95
	15	7.40 ± 1.73	7.39 ± 1.95

around 11.5% and 8.5% for bus routes with slopes within $\pm 1\%$ (Group I) and $\pm 3\%$ (Group II) respectively. Moreover, the battery usage of the vehicle is decreased by similar values relative to the battery usage of the baseline method.

The bottom panel of Fig. 1 shows an example of an optimized speed profile for one of the routes in Group I, in a case with 10 km/h average speed between stops. Note that even though the speed profile gives an impression of high speed variability, exaggerated by the compressed horizontal axis, it should be noted that the (magnitude of the) maximum acceleration is never above 0.6 m/s^2 . Moreover, with 10 km/h average speed, the travel over the first 500 m (i.e. between the start and the first stop) takes 3 min, implying that the 3–4 mild acceleration and deceleration cycles in that interval should not cause discomfort to the passengers. Similar observations can be made for the other intervals.

6.2 Case study II

An additional case study (Case study II) was considered as well in which the vehicle continues its operation until the battery state-of-charge drops below a certain threshold (here 20%). In this experiment, which is an extension of Case study I, the vehicle's motion was simulated using both the speed profile optimization method and the baseline method. In order to emulate a typical bus driving scenario, the mass of the vehicle was varied at each bus stop in the following way

$$m = m_0 + n \times 75 \text{ kg} \quad (11)$$

where m_0 is the initial mass (2000 kg) and n is the number of passengers chosen randomly from 0 to 6 (the maximum number for the minibus considered here) multiplied by the average weight of a passenger (here considered to be 75 kg). In order to make a fair comparison, the mass variation at stops was identical for both methods (SPO and baseline). The time spent at the bus stops (which affects the total energy consumption) was not included in the analysis, since it is the same for both methods.

The optimization was carried out for both bus routes (Groups I and II), stop-by-stop, generating one speed profile for each segment (i.e. the stretch of road between two stops) for every possible mass value, and for each average speed value. As before, the vehicle's average speed was set either to 10 km/h or 15 km/h between stops. At every stop during a round trip, an optimized speed profile was selected by considering the change in the vehicle's mass and the required average speed. The results obtained from the simulations are presented in Table 2. The results show that by using the SPO method, the electric vehicle is able to increase the number of completed round trips by around 12% relative to the baseline method. Moreover, the energy and battery usage savings when using SPO are similar to the values obtained in Case study I.

Table 2 Simulation results obtained for Case study II. The second column from the left shows the average speed of the vehicle between stops (set to either 10 or 15 km/h) whereas the third and the fourth columns give the average energy consumption of the vehicle per round trip using the SPO method and the baseline method, respectively. The fifth and sixth columns from the left show the number of completed round trips before the vehicle's battery state-of-charge drops below 20%. The two right-most columns show the average battery usage of the vehicle

Bus route	Average speed (km/h)	Average energy consumption per round trip (kWh)		Number of round trips before discharge		Average battery usage per round trip (%)	
		SPO	baseline	SPO	baseline	SPO	baseline
Group I	10	1.013 ± 0.157	1.141 ± 0.186	37 ± 1	32 ± 1	2.03 ± 0.10	2.26 ± 0.03
	15	1.015 ± 0.124	1.115 ± 0.167	42 ± 1	37 ± 1	2.01 ± 0.13	2.23 ± 0.19
Group II	10	1.074 ± 0.181	1.216 ± 0.193	36 ± 1	32 ± 1	2.08 ± 0.10	2.26 ± 0.04
	15	1.120 ± 0.163	1.193 ± 0.183	43 ± 1	38 ± 1	2.18 ± 0.09	2.34 ± 0.03

7 Discussion

The results presented in Case study I show that the SPO method improves the efficiency of an electric bus by reducing the energy consumption of the vehicle by around 11.5% and 8.5% for bus routes with slope variation within $\pm 1\%$ (Group I) and $\pm 3\%$ (Group II), respectively, in a single round trip. Moreover, as was shown in Case study II, with the SPO method, the electric bus is able to carry out around four more round trips, compared to the baseline method introduced above, before the vehicle's battery needs recharging (when the state-of-charge drops below a certain threshold).

The reported energy savings in Sect. 6.1 (Case study I) are within the range commonly found in the literature. However, our results are not directly comparable to those values since our baseline case is different, and more stringent: Here, the results are compared to a (quite energy-efficient) baseline case, involving autonomous operation with PID control. In other works, e.g. Dib et al. [12] the comparison is made with a (largely unspecified) human driving cycle, which in all likelihood would not be as energy-efficient as our baseline case, with its PID control, see e.g. Mersky and Samaras [20] and Pourabdollah et al. [18]. Furthermore, as was shown in Torabi and Wahde [2], for vehicles with combustion engines, the SPO method outperforms DP-based methods when both method types are applied to the *same* situation with *identical* baseline cases.

Unlike DP-based methods (and MPC methods in particular), the SPO method does not require online iterative calculations. Once a speed profile has been generated for a given road profile, it can be used without any further modifications (assuming that the vehicle is able to follow the profile; see the next paragraph). Moreover, the SPO based method can be used in cases where longer optimization horizons must be considered (10 km for example) with larger speed range for the vehicle (highway driving, for instance), situations that would require heavy computations, and long running times, if DP-based methods were to be used.

Here it has been assumed that the vehicle is able to follow the speed profile without any interference from the surrounding traffic. This is a valid assumption for an increasing number of urban bus routes for which dedicated lanes are being built in city centers. Another assumption made here is that the motion of the vehicle is not interrupted by any traffic lights. However, the results can be generalized to routes involving traffic lights: For a given bus route, with known traffic light positions, a set of speed profiles can be generated to handle the binary case of traffic lights being either green or red. With n traffic lights between two stops, one could thus optimize 2^n speed profiles (noting again that the optimization can

be done a priori, before driving), covering all possible traffic light states, and then let the vehicle switch dynamically between profiles during driving, depending on the encountered situations. Even though the number of speed profiles grows rapidly with n , it should be no problem to handle realistic values of n , especially for the short bus routes considered here. A more difficult problem, which has not been explicitly modeled here, is the case in which the bus has to slow down or even stop due to unpredictable events (e.g. pedestrians crossing the road). In such cases, which can perhaps largely be avoided by careful city planning, one could simply revert to following the speed profile once the obstacle is gone, thus again accelerating to reach the cruising speed. Any such event would negatively impact the energy savings (regardless of the method used) but would be necessary since safety is a more important consideration.

Moreover, in Case study II, it was assumed that the vehicle's mass is known, including the weight of the passengers, which can increase the total weight by up to 40–50% for a bus. This can for example be achieved by combining statistical data on the typical passenger weight with onboard sensors measuring the number of passengers entering or exiting (or even a camera-based system). A bank of speed profiles can then be generated for a set of different total masses, and the vehicle can then switch between profiles (much as in the traffic-light case considered above), depending on its current estimated mass. Alternatively, the speed profile (for a given mass) could be generated remotely, on demand, and then uploaded to the bus before it leaves the stop. Generating a speed profile for a 500 m stretch between two stops typically takes a few minutes on a standard desktop computer, a time interval that could be reduced to seconds using a faster, cloud-based computer resource noting also that the genetic algorithm is almost completely parallelizable.

As for the length of the acceleration (and deceleration) phases (i.e. the first and last spline; see Sect. 4.3) in a speed profile, several different values were tested before settling on a value of 50 m, which provided the best (i.e. most challenging) results for the baseline case against which our results are compared. With smaller values, the acceleration would be too high, thus resulting in excessive energy usage, and with larger values the cruising phase would be very short, again leading to increased energy usage (as well as a risk that the top speed would have to exceed the allowed maximum speed).

The average speed was chosen to be either 10 or 15 km/h in order for the autonomous bus to comply with the regulations imposed in this project. However, even without those regulations, at least the higher of those speeds is a realistic average speed for a bus operating in an urban environment [19]. Nevertheless, the average

speed can be set to higher values as well, in order to follow a specific timetable. With a different average speed setting, one would then have to rerun the simulations to calculate the energy savings.

A possible extension of this work would be to adapt the fitness measure (in the optimization) so that it increases the emphasis on regenerative braking, thus resulting in larger energy savings, but possibly at the expense of slightly larger acceleration and deceleration. Another topic for future work would be to migrate the optimization procedure to a cloud-based server, as well as running tests in real minibuses using the infrastructure and vehicles available in the Sohjoa Baltic project.

8 Conclusion

For an autonomous electric bus with regenerative braking operating in an urban environment, the proposed method provides energy savings of around 7 to 12% relative to a baseline case with a constant reference speed between stops. The savings translate to an increase of around 10% in the number of round trips that can be executed on a single battery charge, corresponding roughly to an extra hour of operation for the case considered here. The results also demonstrate the applicability of the SPO method, originally developed for vehicles with combustion engines, to electric vehicles.

Acknowledgements

The funding support of the EU Baltic Sea Interreg program through the Sohjoa Baltic project (grant number R073) is gratefully acknowledged.

Declarations

The data used during the current study are available from the corresponding author upon request.

Authors' contributions

All authors read and approved the final manuscript.

Funding

Open access funding provided by Chalmers University of Technology.

Competing interests

The authors declare that they have no competing interests.

Received: 3 July 2019 Accepted: 17 December 2019

Published online: 06 January 2020

References

1. Ainsalu, J., Arffman, V., Bellone, M., Ellner, M., Haapamäki, T., Haavisto, N., Josefson, E., Ismailogullari, A., Lee, B., Pilli-Sihvola, E., Madland, O., Madzulis, R., Muur, J., Mäkinen, S., Nousiainen, V., Rutanen, E., Sahala, S., Schönfeldt, B., Marek Smolnicki, P., Soe, R. M., Sääski, J., Szymańska, M., Vaskinn, I., & Åman, M. (2018). State of the art of automated buses. *Sustainability*, 10(9), 3118–3152.
2. Torabi, S., & Wahde, M. (2017). Fuel consumption optimization of heavy-duty vehicles using genetic algorithms. In *2017 IEEE Congress on Evolutionary Computation (CEC)* (pp. 29–36).
3. Henzler, M., Buchholz, M., & Dietmayer, K. (2015). Optimal parameter selection of a model predictive control algorithm for energy efficient driving of heavy-duty vehicles. In *IEEE Intelligent Vehicles Symposium (IV)* (pp. 742–748).

4. Hellström, E., Åslund, J., & Nielsen, L. (2010). Design of an efficient algorithm for fuel-optimal look-ahead control. *Control Engineering Practice*, 18(11), 1318–1327.
5. Wirasingha, S. G., & Emadi, A. (2011). Classification and review of control strategies for plug-in hybrid electric vehicles. *IEEE Transactions on Vehicular Technology*, 60(1), 111–122.
6. Moura, S. J., Fathy, H. K., Callaway, D. S., & Stein, J. L. (2011). A stochastic optimal control approach for power management in plug-in hybrid electric vehicles. *IEEE Transactions on Control Systems Technology*, 19(3), 545–555.
7. He, Y., Chowdhury, M., Pisu, P., & Ma, Y. (2012). An energy optimization strategy for power-split drivetrain plug-in hybrid electric vehicles. *Transportation Research Part C: Emerging Technologies*, 22, 29–41.
8. Guo, L., Gao, B., Gao, Y., & Chen, H. (2017). Optimal energy management for HEVs in eco-driving applications using bi-level MPC. *IEEE Transactions on Intelligent Transportation Systems*, 18(8), 2153–2162.
9. Mensing, F., Trigui, R., & Bideaux, E. (2012). Vehicle trajectory optimization for hybrid vehicles taking into account battery state-of-charge. In *2012 IEEE vehicle power and propulsion conference* (pp. 950–955).
10. Mets, K., Verschueren, T., Haerick, W., Devellder, C., & De Turck, F. (2010). Optimizing smart energy control strategies for plug-in hybrid electric vehicle charging. In *2010 IEEE/IFIP Network Operations and Management Symposium Workshops* (pp. 293–299).
11. Bashash, S., Moura, S. J., Forman, J. C., & Fathy, H. K. (2011). Plug-in hybrid electric vehicle charge pattern optimization for energy cost and battery longevity. *Journal of Power Sources*, 196(1), 541–549.
12. Dib, W., Chasse, A., Mouline, P., Sciarretta, A., & Corde, G. (2014). Optimal energy management for an electric vehicle in eco-driving applications. *Control Engineering Practice*, 29, 299–307.
13. Rohkämper, S., Hellwig, M., & Ritschel, W. (2017). Energy optimization for electric vehicles using dynamic programming. In *2017 International Conference on Research and Education in Mechatronics (REM)* (pp. 1–5).
14. Jiang, X., Chen, L., Xu, X., Cai, Y., Li, Y., & Wang, W. (2018). Analysis and optimization of energy efficiency for an electric vehicle with four independent drive in-wheel motors. *Advances in Mechanical Engineering*, 10(3) (pp. 1–9).
15. Maamria D., Gillet K., Colin G., Chamailard Y. and Nouillant C., (2016). On the use of Dynamic Programming in eco-driving cycle computation for electric vehicles. In *2016 IEEE Conference on Control Applications (CCA)* (pp. 1288–1293).
16. Rios-Torres, J., Sauras-Perez, P., Alfaro, R., Taiber, J., & Pisu, P. (2015). Eco-driving system for energy efficient driving of an electric bus. *SAE International Journal of Passenger Cars - Electronic and Electrical Systems*, 8, 79–89.
17. Genikomsakis, K. N., & Mitrentsis, G. (2017). A computationally efficient simulation model for estimating energy consumption of electric vehicles in the context of route planning applications. *Transportation Research Part D: Transport and Environment*, 50, 98–118.
18. Pourabdollah, M., Björkvik, E., Furer, F., Lindenberg, B., & Burgdorf, K. (2017). Fuel economy assessment of semi-autonomous vehicles using measured data. In *2017 IEEE Transportation Electrification Conference and Expo (ITEC)* (pp. 761–766).
19. Oskarbski, J., Birr, K., Miszewski, M., & Żarski, K. (2015). Estimating the average speed of public transport vehicles based on traffic control system data. In *2015 IEEE International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS)* (pp. 287–293).
20. Mersky, A. C., & Samaras, C. (2016). Fuel economy testing of autonomous vehicles. *Transportation Research Part C: Emerging Technologies*, 65, 31–48.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.