



Hierarchical Control of Electric Bus Lines

Downloaded from: <https://research.chalmers.se>, 2025-07-03 06:01 UTC

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

Lacombe, R., Gros, S., Murgovski, N. et al (2020). Hierarchical Control of Electric Bus Lines. IFAC-PapersOnLine, 53(2): 14179-14184. <http://dx.doi.org/10.1016/j.ifacol.2020.12.1040>

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

Hierarchical Control of Electric Bus Lines[★]

Rémi Lacombe^{*} Sébastien Gros^{**} Nikolce Murgovski^{*}
Balázs Kulcsár^{*}

^{*} Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden

(e-mail: {lacombe,nikolce.murgovski,kulcsar}@chalmers.se)

^{**} Department of Engineering Cybernetics, Norwegian University of Science and Technology, Trondheim, Norway

(e-mail: sebastien.gros@ntnu.no)

Abstract: In this paper, we propose a hierarchical control strategy for a line of electric buses with the double objective of minimizing energy consumption and providing regular service to the passengers. The state-space model for the buses is formulated in space rather than in time, which alleviates the need for integer decision variables to capture their behavior at bus stops. This enables us to first assemble a fully-centralized multi-objective line problem in the continuous nonlinear optimization framework. It is then reassembled into a hierarchical structure with two levels of control in order to improve on scalability and reliability. This new supervisory structure consists of a centralized line level controller which handles the time headway regularity of the buses, and of decentralized bus level controllers which simultaneously manage the energy consumption of each individual bus. Our method demonstrates good battery energy savings and regularity performances when compared to a classical holding strategy.

Copyright © 2020 The Authors. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0>)

Keywords: Hierarchical control, Bus bunching, Electric vehicles, Headway regularity, Optimal control, Public transport

1. INTRODUCTION

Bus networks in general are known to be very vulnerable to disturbances from their surroundings, which can in turn substantially increase passengers waiting times. Newell and Potts (1964) first proved that buses have a natural tendency to bunch if delayed, primarily due to the accumulation of passengers at stops. More recently, online control strategies have been proposed, benefiting from the rapid growth of GPS technology and vehicle-to-vehicle communications. Some are primarily based on real-time information, such as those which elaborate on Daganzo (2009) and Bartholdi and Eisenstein (2012). Other approaches rely on model-based predictions instead (Varga et al., 2018). However, most of these works focus solely on reducing bus bunching, and as such run the risk of sacrificing potential savings in energy consumption.

Due to a reduced dependence on fossil fuels, electric buses are a promising solution to reduce the environmental impact of transport systems (Lajunen et al., 2016). But battery and charging constraints are still hindering their market penetration today as they entail additional cost for transit operators. One way to mitigate this is to reduce the energy consumed by the buses when operated. Energy-saving driving strategies have been extensively investigated for heavy-duty vehicles in all kind of driving environ-




Fig. 1. Representation of the bus line considered in this paper. The electric buses are running on a circular route and the controllers can estimate and adapt their predicted time headways (noted H_1, \dots, H_n).

ments (Murgovski et al., 2016; Held et al., 2018). However, works on this topic have focused almost exclusively on trucks, and to the best of our knowledge only Varga et al. (2019) have addressed the energy-optimal bus line control problem. In their work, they formulate the line problem in the time domain and in a fully-centralized way. This may cause some reliability and scalability issues, which we try to address here.

In this paper, we focus on a circular route which is serviced by a single line of electric buses. These buses are made to keep a constant time headway provided by

[★] The authors acknowledge the contribution of Transport Area of Advance at Chalmers University of Technology. The paper has been partially supported by Energimyndigheten through the project ‘Operational Network Energy Management for Electrified Buses’ (46365-1).

the transit agency between them. Such regularity-based operating policies are common in an urban context where small headways between consecutive buses mean that passenger arrivals can be considered as random (Fan and Machemel, 2009). Buses are allowed to run at different speeds between stops, both as a way to absorb deviations from the goal headway and to adapt to local settings (location of bus stops, road topography) to spare battery energy when possible. They are however not allowed to dwell at stops longer than needed to pick passengers up in order to avoid disturbing the surrounding traffic. In addition, we make the assumption that the buses are operated automatically, i.e. without the supervision of a human driver. This enables them to react faster and more accurately to the commands and thus improves their energy-saving potential.

The main contributions of this paper are twofold. The first contribution is to assemble the multi-objective bus line optimization problem as a nonlinear program (NLP) thanks to a spatial formulation of bus dynamics. Choosing a temporal formulation for the dynamics would have required adding integer variables for the decisions of buses at stops, thus resulting in mixed-integer programs, which are notoriously harder to solve. The second contribution is to formulate the bus line problem in a hierarchical control framework. The proposed hierarchical structure contains several parallelizable finite horizon sub-problems which are solved at the bus level, and a simple constrained quadratic program (QP) at the line level. This allows for lower computation times compared with a fully-centralized approach. It also improves the reliability of the control as the commands can be computed directly in each bus independently from the rest of the line, thus making the system more resilient to e.g. communication errors.

2. BUS DYNAMICS

In this section, the longitudinal dynamics are first described at the bus level, and are then combined at the line level. For the rest of this paper, we consider a setting where n buses are running on a circular route of total length L with q bus stops.

2.1 Longitudinal Dynamics

For an electric bus $i \in \mathbb{I}_{[1,n]}$, the equations of motion along a fixed route are

$$\dot{s}_i(t) = v_i(t), \quad (1a)$$

$$\dot{v}_i(t) = \frac{1}{m_i(s_i)} (F_{m,i}(t) - F_{b,i}(t) - F_{d,i}(v_i) - F_{r,i}(s_i)), \quad (1b)$$

where s_i is the bus position and v_i its velocity at time t , m_i is its mass, which is space-dependent, and $F_{m,i}$ is the force created at the wheels by the electric motor. The force $F_{b,i}$ is generated by the friction brakes, while the aerodynamic drag $F_{d,i}$ and a force $F_{r,i}$ gathering the rolling resistance and the gravitational effects can be expressed as

$$F_{d,i}(v_i) = \frac{1}{2} \rho A_{\text{bus}} c_a v_i^2, \quad (2a)$$

$$F_{r,i}(s_i) = g m_i(s_i) (\sin \alpha(s_i) + c_r \cos \alpha(s_i)), \quad (2b)$$

where ρ is the air density, A_{bus} is the frontal area of the vehicle, c_a is the aerodynamic air drag coefficient and c_r

is the rolling resistance coefficient. The function α is the road gradient, which depends on the position of bus i .

Since m_i is also space-dependent, due to uneven passengers loads between stops, it is more practical to express (1) in the space coordinate system s in order to remove the nonlinearity in the dynamics (Murgovski et al., 2016). We now consider the travel time t_i of bus i as a state instead of its position. The kinetic energy for a unit mass $E_i(s) = \frac{1}{2} v_i^2(s)$ is chosen as a state instead of v_i to simplify the expressions. The modified dynamics read

$$\frac{dt_i}{ds} = \frac{1}{\sqrt{2E_i(s)}}, \quad (3a)$$

$$\frac{dE_i}{ds} = \frac{1}{m_i(s)} (F_{m,i}(s) - F_{b,i}(s) - \rho A_{\text{bus}} c_a E_i(s)) - g(\sin \alpha(s) + c_r \cos \alpha(s)), \quad (3b)$$

where the second equation is a reformulation of (1b). The state vector for bus i is then $x_i(s) = [t_i(s), E_i(s)]^\top$. Note here that equation (3a) imposes that the speed of bus i should be strictly positive, even though it needs to stop frequently in reality. This is addressed by enforcing very low positive speeds around bus stops, and by introducing additional delays, as explained in the rest of this section.

2.2 Speed Corridor

Some constraints must be introduced for the speed of the buses in order to e.g. capture that they should slow down around bus stops and comply with the speed limits. To do so, we create a space-dependent speed corridor, similar to Held et al. (2018), inside which the speed of each bus is constrained to be. Its upper and lower bounds are noted v_{max} and v_{min} and the constraints for a bus $i \in \mathbb{I}_{[1,n]}$ materialize as

$$\frac{1}{2} v_{\text{min}}^2(s) \leq E_i(s) \leq \frac{1}{2} v_{\text{max}}^2(s). \quad (4)$$

Both bounds are positive, even at stops, to keep the kinetic energy from becoming null. Fig. 2. displays the appearance of the speed corridor around one stop, and this can be extended to the whole route without loss of generality.




Fig. 2. Illustration of the speed corridor around a stop. The shaded area represents the feasible speeds for the buses. At the stop, the buses are forced to slow down to a small velocity v_{entry} . When they are within a distance d_{brk} of a stop, it is their acceleration which is bounded to avoid disturbing the passengers on-board. Otherwise, the buses are limited by the speed limit of the road v_{lim} , and by a minimum acceptable speed v_{low} that prevents from going too slowly.

2.3 Electric Motor and Battery

We now elaborate a model to compute the energy consumed by the buses. Since electric motors can be used both for traction and for generation, the motor torque can be positive or negative depending on the operating scenario. For an electric bus $i \in \mathbb{I}_{[1,n]}$, the motor torque $T_{m,i}$ is related to the force at the wheels through $T_{m,i}(F_{m,i}) = r_w/(M_f\eta_f)F_{m,i}$ for the former case and $T_{m,i}(F_{m,i}) = (r_w\eta_f/M_f)F_{m,i}$ for the latter. Similarly, the motor speed $\omega_{m,i}$ is proportional to the bus speed with $\omega_{m,i}(E_i) = (M_f/r_w)\sqrt{2E_i}$. Here, r_w is the wheel radius, M_f is the final gear ratio and η_f is the efficiency coefficient of the final gear, with $\eta_f < 1$.

Due to the power limitation of the electric motor, the torque is constrained by

$$|T_{m,i}(F_{m,i})| \leq T_{\max}(E_i), \quad (5)$$

where T_{\max} is the absolute value of the maximum torque that can be delivered for a given speed of the bus.

However, the motor torque is not continuously differentiable with respect to the command $F_{m,i}$, which can be problematic for solving optimization problems based on this model. A numerical work-around is to *lift* the model by adding separate force variables for each regime of the motor, such that

$$F_{m,i}(s) = F_{t,i}(s) - F_{g,i}(s), \quad (6a)$$

$$0 \leq F_{t,i}(s) \leq \frac{M_f\eta_f}{r_w}T_{\max}(E_i), \quad (6b)$$

$$0 \leq F_{g,i}(s) \leq \frac{M_f}{r_w\eta_f}T_{\max}(E_i), \quad (6c)$$

where $F_{t,i}$ and $F_{g,i}$ are the forces at the wheels for bus i when the motor is respectively working in traction or in generation. The command vector for bus i is thus $u_i(s) = [F_{t,i}(s), F_{g,i}(s), F_{b,i}(s)]^\top$. Note that $F_{t,i}$ and $F_{g,i}$ are mutually exclusive due to the subsequent problem design, as will be motivated in the next section. This yields a smoother expression for the motor torque

$$T_{m,i}(F_{t,i}, F_{g,i}) = \frac{r_w}{M_f\eta_f}F_{t,i} + \frac{r_w\eta_f}{M_f}F_{g,i}. \quad (7)$$

The internal battery power of bus i can be modeled as a nonlinear function $P_{b,i}(\omega_{m,i}, T_{m,i})$, as in e.g. Murgovski et al. (2014). The only assumption that we make on $P_{b,i}$ is that it is monotonously increasing with respect to $T_{m,i}$, i.e. that a higher motor torque draws more power from the battery, or equivalently supplies the battery with less power if $T_{m,i} < 0$.

2.4 Delays at Stops

The dwell times of a bus $i \in \mathbb{I}_{[1,n]}$ at stops are not captured by the longitudinal dynamics model (3), but they can be added to the travel time t_i as extra delay terms D_j whenever the bus reaches a stop $j \in \mathbb{I}_{[1,q]}$. Each delay D_j consists of a constant term for stopping at and leaving stop j , and a variable term for boarding the passengers currently waiting at that stop.

We note t_s the time it takes for a bus to come to a complete stop from the small velocity v_{entry} and to open its doors. It is assumed that it also takes t_s for it to close its doors and

reach v_{entry} again when leaving the stop. In addition, we consider that the passengers arrive at stop j at a constant rate λ_j , and that it takes each passenger a time b to board the bus. The general expression for D_j is then

$$D_j(t_i, s) = \mathbf{1}_{s_j}(s) (2t_s + b\lambda_j T_j(t_i)) \quad (8)$$

where s_j is the position of stop j and T_j is the elapsed time since the preceding bus has left stop j . The indicator function $\mathbf{1}_{s_j}(s)$ is 1 when $s = s_j$ and 0 otherwise.

This expression for the delays ignores the influence of the alighting passengers, since alighting is usually faster than boarding (Petit et al., 2018). In addition, we use a normalized passengers model in order to keep the problem continuous. This means that the buses have to pick up passengers at each stop, which is a reasonable assumption for buses running in a dense urban environment. Other types of stops could be included, e.g. when buses have to halt due to traffic, but we choose to ignore those in this paper in order to keep a deterministic model.

3. HIERARCHICAL BUS LINE CONTROL

In this section, we motivate our choice of treating the line level optimization problem in a hierarchical framework. The fully-centralized line optimization problem is assembled and discussed in the first subsection. The following sub-sections then focus on detailing each step of the proposed hierarchical control formulation.

3.1 Fully-centralized Line Optimization Problem

Let us assume that we are looking at a snapshot of the bus line at a given time, such as the one presented in Fig. 1. Each bus i is currently at a position p_i . The bus preceding it on the line has index $i+1$ and is at position p_{i+1} . In order to keep a schedule with a constant time headway H between the buses, each bus i should aim to travel from p_i to p_{i+1} as closely to H as possible. In what follows, we note H_i the predicted time for bus i to travel that distance, which is defined as $H_i = t_i(p_{i+1})$. Note that H_i can be seen as the current headway between i and $i+1$. Therefore, the buses do not try to keep a headway H directly at the stops, but at whatever positions they have on the snapshot. The idea is that enforcing regular time headways at each point on the route will result in constant headways being enforced at the stops too. Based on the previous models, the centralized line control problem can now be written as

$$\begin{aligned} \min_U \quad & \sum_{i=1}^n (H_i(x_i) - H)^2 \\ & + \alpha \sum_{i=1}^n (H_i(x_i) - H_{i-1}(x_{i-1}))^2 \\ & + \beta \sum_{i=1}^n \int_{p_i}^{p_{i+1}} \frac{P_{b,i}(x_i(s), u_i(s))}{\sqrt{2E_i(s)}} ds, \end{aligned} \quad (9a)$$

$$\text{s.t.} \quad \forall i \in \mathbb{I}_{[1,n]} : \quad \begin{aligned} E_i(p_i) &= \hat{E}_{0,i}, \quad t_i(p_i) = 0, \\ \forall s \in [p_i, p_{i+1}] : \end{aligned} \quad (9b)$$

$$\frac{dx_i}{ds} = f_i(x_i(s), u_i(s)) + \sum_{j=1}^q D_j(t_i, s) [1, 0]^\top, \quad (9c)$$

$$h_i(x_i(s), u_i(s)) \leq 0, \quad (9d)$$

where $U = [u_1, \dots, u_n]^\top$ is the vector containing all the command functions, which are assumed to be piecewise constant. The initial kinetic energy of bus i is noted $\hat{E}_{0,i}$. The evolution function f_i corresponds to the dynamics (3), while h_i gathers the inequality constraints from (4), (6b) and (6c). Finally, α and β are trade-off parameters to weigh the relative importance of the different objectives in the cost function.

Remark 1: The second term in the cost function includes a *look-back* feature which constrains headways of neighboring buses to be similar. This ensures that the rate of service remains somewhat homogeneous locally, even as the headways are made to converge to H by the first term. The preponderance of each is adjusted with the parameter α . The competing objective of energy minimization is added as the third term and β regulates the trade-off.

However, the nonlinear optimal control problem (9) suffers from some practical issues. Indeed, centralized computations cause it to be unadapted to an on-line setting, and vulnerable to communication losses between the buses and the central coordinator. They also forbid to run any computations in parallel to ease the computational load. Therefore, we choose to structure (9) as a bi-level optimization problem instead by splitting it into a line level problem and bus level sub-problems to increase reliability. We respectively refer to each as the high-level or the low-level hereafter. At the high-level, the quadratic terms of the cost function (9a) are gathered into a static QP, which can be solved very quickly. The energy consumption is dealt with at the low-level, and can now be computed in parallel on each separate road segment $[p_i, p_{i+1}]$ instead of on the whole route at once. The *dialogue* between the high and the low-level is summarized in Fig. 3.

3.2 Low-level Optimal Control

Let us assume that a bus $i \in \mathbb{I}_{[1,n]}$ has received from the high-level controller a travel time goal $t_{f,i}$ to reach the preceding bus at p_{i+1} . This command concerns the predicted time headway H_i , and bus i should adapt its predicted trajectory such that $H_i(x_i) = t_{f,i}$ if possible. Note that $t_{f,i}$ can be quite different from the goal headway H , e.g. if the line is currently very disturbed.

The following NLP is assembled to find the energy-optimal control of bus i that meets the travel time goal $t_{f,i}$

$$\min_{u_i} J_i(x_i, u_i) = \int_{p_i}^{p_{i+1}} \frac{P_{b,i}(x_i(s), u_i(s))}{\sqrt{2E_i(s)}} ds + C \kappa_i \quad (10a)$$

$$\text{s.t.} \quad E_i(p_i) = \hat{E}_{0,i}, \quad t_i(p_i) = 0, \quad (10a)$$

$$\kappa_i \geq H_i(x_i) - t_{f,i}, \quad \kappa_i \geq t_{f,i} - H_i(x_i), \quad (10b)$$

$$\forall s \in [p_i, p_{i+1}] :$$

$$\frac{dx_i}{ds} = f_i(x_i(s), u_i(s)) + \sum_{j=1}^q D_j(t_i, s)[1, 0]^\top, \quad (10c)$$

$$h_i(x_i(s), u_i(s)) \leq 0, \quad (10d)$$

where most functions have already been introduced for problem (9). We note J_i the cost function of this problem, κ_i is a slack variable, and C is a large constant coefficient. Note that this problem only needs to be solved locally, i.e.




Fig. 3. This diagram represents the information exchanged by the two control levels. The high-level controller sends a travel time command $t_{f,i}$ to each and every bus i of the line. Each low-level controller responds with the corresponding energy sensitivity function J_i^* . Both control levels are operated with the same frequency until convergence of the loop, as they each need updated information from the other level.

separately for each and every bus i . The low-level control can therefore always be carried out by using the latest available information, even for cases where commands from the high-level are temporarily unavailable.

Remark 2: The cost function J_i is the sum of the total electric energy consumed on the space-horizon and the slack variable. The motivation for introducing this last term is to guarantee that the problem is feasible, regardless of the travel time command $t_{f,i}$. Indeed, using this slack variable is akin to having a penalty term $|t_{f,i} - H_i(x_i)|$ in the objective function while keeping it smooth at the same time. And choosing a large coefficient C strongly penalizes solutions where the final time $H_i(x_i)$ is different from $t_{f,i}$, which imposes $H_i(x_i) = t_{f,i}$ whenever it is feasible.

Remark 3: Since (10) is an energy-minimization problem, the command function u_i is chosen to be as energy-efficient as possible. For a given feasible force requirement at the wheel, the lowest torque (and hence the lowest battery power) is achieved in (7) by having either $F_{t,i}$ or $F_{g,i}$ equal to 0 depending on the sign of the required force. This holds because we have assumed $P_{b,i}$ to be monotonously increasing with respect to the torque earlier. Therefore, $F_{t,i}$ and $F_{g,i}$ are mutually exclusive for this problem.

Remark 4: Numerical simulations are carried out to solve this optimal control problem. For bus i , we use multiple shootings to split the route segment $[p_i, p_{i+1}]$ into uniform shooting intervals. The standard Runge-Kutta method is then used to integrate the dynamics of the discretized system on each shooting interval. The shooting points corresponding to the stops are known beforehand thanks to the sampling in space, so no integer variable needs to be added. In addition, even if the indicator function from (8) makes the NLP (10) discontinuous, its discretized counterpart becomes continuous in the optimization variables since the stops locations are known. Finally, having a decentralized control at the low-level means that these problems can be solved in parallel for each and every bus. In practice, they are found to be several time faster to solve than the fully-centralized problem (9).

3.3 Sensitivity Analysis

Once the energy-optimal low-level control is known, we are interested in knowing the sensitivity of the energy

consumption to the travel time command $t_{f,i}$ for a bus $i \in \mathbb{I}_{[1,n]}$. This information is crucial to address the trade-off between energy consumption and regular time headways, and will be used at the high-level of control.

Problem (10) can be seen as a parametric optimization problem, with the travel time command $t_{f,i}$ acting as a scalar parameter (Still, 2018). Let $J_i^*(t_{f,i}^0)$ be the parametric optimal cost for a given parameter value $t_{f,i}^0$. We know from *Remark 2* that this cost is equal to the energy consumption if and only if $t_{f,i}^0$ is a feasible time command. So the sensitivity analysis is only carried out if this condition holds, which is the case most of the time in practice.

Since there is no direct way to compute the implicit function J_i^* , we will approximate it with the function \hat{J}_i^* based on the first two terms of its Taylor series instead

$$\begin{aligned} \hat{J}_i^*(t_{f,i}) &= J_i^*(t_{f,i}^0) + \frac{\partial J_i^*}{\partial t_{f,i}}(t_{f,i}^0)(t_{f,i} - t_{f,i}^0) \\ &\quad + \frac{1}{2} \frac{\partial^2 J_i^*}{\partial t_{f,i}^2}(t_{f,i}^0)(t_{f,i} - t_{f,i}^0)^2. \end{aligned} \quad (11)$$

Let \mathcal{L}_i be the Lagrange function of (10) and λ_i and μ_i the vectors of dual variables corresponding to the equality and inequality constraints respectively. Let $z_i = [w_i, \lambda_i, \mu_i]^\top$ be the solution vector, which gathers the primal and dual variables, with $w_i = [x_i, u_i]^\top$. By noting z_i^* the optimal solution, which is a function of the parameter $t_{f,i}$, and applying the chain rule we have (Still, 2018)

$$\frac{\partial J_i^*}{\partial t_{f,i}}(t_{f,i}^0) = \left(\frac{\partial \mathcal{L}_i}{\partial t_{f,i}} \right)_{z_i^*(t_{f,i}^0), t_{f,i}^0} \quad (12a)$$

$$\frac{\partial^2 J_i^*}{\partial t_{f,i}^2}(t_{f,i}^0) = \left(\frac{\partial^2 \mathcal{L}_i}{\partial t_{f,i}^2} + \nabla_{t_{f,i}}^2 \mathcal{L}_i \frac{\partial z_i^*}{\partial t_{f,i}} \right)_{z_i^*(t_{f,i}^0), t_{f,i}^0} \quad (12b)$$

Note that the first two derivatives of J_i^* are not too expensive to obtain here since the parameter $t_{f,i}$ only appears in the constraints (10b). Let n_{z_i} and n_{μ_i} respectively be the length of the solution vector z_i , and the number of inequality constraints in problem (10). For simplicity, we assume that the constraints (10b) are the last two inequality constraints. Then it holds that

$$\frac{\partial J_i^*}{\partial t_{f,i}}(t_{f,i}^0) = \mu_i(n_{\mu_i}) - \mu_i(n_{\mu_i} - 1), \quad (13a)$$

$$\frac{\partial^2 J_i^*}{\partial t_{f,i}^2}(t_{f,i}^0) = y_i(n_{z_i} - 1) - y_i(n_{z_i}), \quad (13b)$$

where $y_i = R_{z,i}^{-1}(z_i^*(t_{f,i}^0), t_{f,i}^0) e$, $e = [0, \dots, 0, -1, 1]^\top$, and where $R_{z,i}$ is the Jacobian of the KKT conditions vector with respect to the solution vector z_i (Hult et al., 2016).

Therefore, only one matrix inversion is needed to get a quadratic approximation \hat{J}_i^* of the energy consumption sensitivity for bus i . The n sensitivity functions obtained are then passed as inputs to the high-level.

3.4 High-level Controller

The role of the high-level controller is to find the best travel time commands $t_{f,1}, \dots, t_{f,n}$. This is formulated as a supervisory control problem with a static constrained QP:

$$\begin{aligned} \min_{t_{f,1}, \dots, t_{f,n}} \quad & \sum_{i=1}^n (t_{f,i} - H)^2 + \alpha(t_{f,i} - t_{f,i-1})^2 + \beta \hat{J}_i^*(t_{f,i}), \\ \text{s.t.} \quad & t_{f,i}^{\min} \leq t_{f,i} \leq t_{f,i}^{\max}, \quad \forall i \in \mathbb{I}_{[1,n]}. \end{aligned} \quad (14)$$

The approximate minimum and maximum feasible time commands $t_{f,i}^{\min}$ and $t_{f,i}^{\max}$ for any bus i are estimated from the solutions obtained by forcing the velocity of the buses to be equal to v_{\min} and to v_{\max} respectively. Thus they can provide a good approximation of the feasible set of the QP, which can be solved efficiently using dedicated toolboxes.

The travel time commands found from solving this QP are then fed in to the low-level again, and the loop continues until convergence. Note that both control levels are run with the same frequency during this process. Once the optimal solution of the overall problem is found, the bus line is updated in a MPC-like fashion with a sampling time Δt by implementing the low-level solutions from (10). This sampling time is unrelated to the low-level programs as it is common for all buses, and denotes the frequency at which the overall line problem is solved.

4. NUMERICAL RESULTS

In this section, we present some simulation results to assess how well our hierarchical control method is able to maintain a regular level of service in an energy-efficient way. To do so, we chose to compare it to a bus holding benchmark method, as is done in Varga et al. (2019). In this scenario, buses go at the maximum allowed speed in-between stops, and then wait at some predesignated control points until the headway regularity condition is met. Many transit systems operators enforce this control method in practice as it is relatively inexpensive and easy to deploy. Following our previous notations, this means that the benchmark buses are forced to run at v_{\max} .

The route considered is that of bus line 17 in Gothenburg, Sweden, with some minor adjustments. Both control methods operate from the same initial disturbed state of the bus line in 5 different scenarios, ordered from least to most disturbed (i.e. scenario 1 is the least disturbed). Since no external disturbances have been added to the model, the controlled bus line will converge back to a state of regular service in all the scenarios. Two metrics are used to monitor each method's performances: the energy consumed per distance unit ΔE and the mean value $\Delta \hat{H}$ of the absolute headway deviations last observed at each of the stops. Both of them are computed at the line level. The numerical values for the relevant parameters are summarized in Table 1. The space-dependent functions α and m_i , for $i \in \mathbb{I}_{[1,n]}$, as well as the rates λ_j , for $j \in \mathbb{I}_{[1,q]}$, are obtained directly from bus line 17 data.

Table 1. Simulation parameters. SI units: H , L , v_{\lim} , v_{low} , v_{entry} , ρ , A_{bus} , r_w , d_{brk} , t_s , b , g . Unitless: n , q , c_a , c_r , M_f and η_f .

Parameter	n	q	H	L	v_{\lim}	v_{low}	v_{entry}
Value	5	28	360	16492	13.89	11.11	1.39
Parameter	ρ	A_{bus}	c_a	c_r	r_w	M_f	η_f
Value	1.18	5.14	1	0.0047	0.49	2.8	0.98
Parameter	d_{brk}	t_s	b	g			
Value	95.5	3	1.5	9.81			




Fig. 4. Average observed headway deviations at stops with respect to time for both methods in scenario 3. The dashed line is the 15 seconds threshold. This is obtained with $\alpha = 1$ and $\beta = 0.05$.

For each set of initial conditions, the bus line is simulated for n_{iter} iterations with a sampling time $\Delta t = 40s$. Due to the different initial disturbance strengths of the scenarios, the time required to converge back to a state of regular service varies. So n_{iter} is set separately for each scenario as the number of iterations required for the deviations $\Delta \hat{H}$ of both methods to become smaller than a threshold value, which we set to 15 seconds. The trade-off parameters α and β are chosen accordingly to make the deviation of the hierarchical control method reach that threshold value in the same number of iterations as the bus holding method when possible, as is illustrated for scenario 3 in Fig. 4. The results for each scenario are presented in Table 2.

Table 2. Results.

Scenario	1	2	3	4	5
ΔE control [kWh/km]	0.649	0.635	0.646	0.639	0.627
ΔE holding [kWh/km]	0.668	0.658	0.666	0.672	0.677
Energy savings [%]	2.9	3.6	3.1	5.2	8
n_{iter}	33	39	43	61	75

The hierarchical control method consistently consumes less energy to reach the headway deviation threshold, regardless of the initial conditions of the line. More energy can in fact be saved for scenarios with larger initial disturbances. This may be the consequence of needing longer simulation times to reach a regular level of service, as some buses will reach the desired headway sooner and have more time to drive in an energy-aware fashion.

For the most disturbed cases (scenarios 4 and 5), the hierarchical control method is not able to reach the 15 seconds threshold as fast as the bus holding method, regardless of the trade-off parameters values. This is mainly due to an imbalance in what each method can do. Indeed, the bus holding method can rapidly absorb strong disturbances by having the buses wait at the control points, whereas the hierarchical control constrains the buses to be running at v_{min} or faster all the time. So it must be noted that this is not an inherent limitation of the hierarchical control strategy, as a bus-holding feature could be added to it to handle strongly disturbed lines.

5. CONCLUSION

In this paper we set to find a scalable and reliable control strategy that could solve a multi-objective bus line control

problem. We chose to avoid formulating that problem as a mixed-integer program, but preferred a classical NLP formulation. It was first assembled as a fully-centralized problem based on the bus and line models developed before. But due to practical constraints (e.g. sensitivity to communication deficiencies), we split it into a bi-level hierarchical control structure. This new control structure only needs to solve independent sub-problems at the bus level, and a single quadratic problem at the line level. When implemented on a simple test case, it demonstrates a promising potential for sparing battery energy and keeping constant headways in weakly disturbed situations.

However, it was also found that our approach should be augmented with an additional intervention method to absorb strong disturbances more quickly. Another direction for future works is to evaluate quantitatively the complexity of different problem formulations to find the most appropriate one for a future on-line implementation.

REFERENCES

- Bartholdi, J. and Eisenstein, D. (2012). A self-coordinating bus route to resist bus bunching. *Transportation Research B, Methodology*, 46, 481–491.
- Daganzo, C.F. (2009). A headway-based approach to eliminate bus bunching. *Transportation Research B, Methodology*, 43, 913–921.
- Fan, W. and Machemehl, R. (2009). Do transit users just wait for buses or wait with strategies? *Transportation Research Record*, 2111, 169–176.
- Held, M., Flärdh, O., and Mårtensson, J. (2018). Optimal speed control of a heavy-duty vehicle in urban driving. *IEEE Transactions on Intelligent Transportation Systems*, 20, 1562–1573.
- Hult, R., Zanon, M., Gros, S., and Falcone, P. (2016). Primal decomposition of the optimal coordination of vehicles at traffic intersections. In *Conference on Decision and Control (CDC)*, 2567–2573.
- Lajunen, A. and Lipman, T. (2016). Lifecycle cost assessment and carbon dioxide emissions of diesel, natural gas, hybrid electric, fuel cell hybrid and electric transit buses. *Energy*, 106, 329–342.
- Murgovski, N., Egardt, B., and Nilsson, M. (2016). Co-operative energy management of automated vehicles. *Control Engineering Practice*, 57, 84–98.
- Murgovski, N., Johannesson, L., and Egardt, B. (2014). Optimal battery dimensioning and control of a CVT PHEV powertrain. *IEEE Transactions on Vehicular Technology*, 63, 2151–2161.
- Newell, G.F. and Potts, R.B. (1964). Maintaining a bus schedule. In *Conference on Australian Road Research Board (ARRB)*. 2nd edition.
- Petit, A., Ouyang, Y., and Lei, C. (2018). Dynamic bus substitution strategy for bunching intervention. *Transportation Research B, Methodology*, 115, 1–16.
- Still, G. (2018). *Lectures on parametric optimization: an introduction*. Optimization Online.
- Varga, B., Tettamanti, T., and Kulcsár, B. (2018). Optimally combined headway and timetable reliable public transport system. *Transportation Research C, Emerging Technologies*, 92, 1–26.
- Varga, B., Tettamanti, T., and Kulcsár, B. (2019). Energy-aware predictive control for electrified bus networks. *Applied Energy*, 252.