Optimal Transient Real-Time Engine-Generator Control in the Series-Hybrid Vehicle

(article starts on next page)

DSCC2019-8964

OPTIMAL TRANSIENT REAL-TIME ENGINE-GENERATOR CONTROL IN THE SERIES-HYBRID VEHICLE

Jonathan Lock
Department of Electrical Engineering
Chalmers University of Technology
Gothenburg 41258
Sweden
Email: lock@chalmers.se

Rickard Arvidsson
Department of Electrical Engineering
Chalmers University of Technology
Gothenburg 41258
Sweden
Email: arvidssr@chalmers.se

Tomas McKelvey
Department of Electrical Engineering
Chalmers University of Technology
Gothenburg 41258
Sweden
Email: tomas.mckelvey@chalmers.se

ABSTRACT

We study the dynamic engine-generator optimal control problem with a goal of minimizing fuel consumption while delivering a requested average electrical power. By using an infinite-horizon formulation and explicitly minimizing fuel consumption, we avoid issues inherent with penalty-based and finite-horizon problems. The solution to the optimal control problem, found using dynamic programming and the successive approximation method, can be expressed as instantaneous non-linear state-feedback. This allows for trivial real-time control, typically requiring 10–20 CPU instructions per control period, a few bytes of RAM, and 5–20 KiB of nonvolatile memory. Simulation results for a passenger vehicle indicate a fuel consumption improvement in the region of 5–7% during the transient phase when compared with the class of controllers found in the industry. Benchtests, where the optimal controller is executed in native hardware, show an improvement of 3.7%, primarily limited by unmodeled dynamics. Our specific choice of problem formulation, a guaranteed globally optimal solution, and trivial real-time control resolve many of the limitations with the current state of optimal engine-generator controllers.

1 Introduction

Engine-generator (gen-set) systems are used in a wide range of applications, from the drivetrain in a series-hybrid passenger vehicle to the larger diesel-electric locomotive and natural gas power plants. Regardless of application a gen-set fundamentally consists of a combustion engine mechanically connected to an electric machine, whose purpose is to convert a combustible fuel into electrical power. One particularly relevant attribute of a well-designed gen-set is it's fuel efficiency, i.e. how efficiently consumed fuel is converted to electrical power. In this paper we will consider the gen-set of a series-hybrid passenger vehicle and specifically a method of improving its fuel efficiency (though this method can easily be extended to other applications). In particular, we will focus on methods that include transient operation, i.e. where the gen-set engine speed is allowed to vary over time.

One method of generating a fuel-efficient control scheme is to formulate the objective as an *optimal control problem*, and then use a method from the field of optimal control to solve the problem. The specific formulation of the optimal control problem and the method used to solve it is crucial. In this paper we will require the solution to

1. accurately capture our intent of minimizing fuel consumption,
2. be (close to) the globally optimal solution, and
3. allow for real-time control with low computational demand.

Previous work in transient gen-set control covers a range

of different optimal control formulations, solution methods, and with varying degrees of real-time applicability.

If we first consider the formulation of the optimization problem, some have elected to use quadratic penalty-based formulations that balance the quadratic gen-set losses and quadratic deviation from requested power [4, 6]. Though a conventional formulation, this results in the gen-set generally never delivering exactly the requested power. Furthermore, this formulation does not generally minimize fuel consumption (as the minimand is quadratic losses and power deviation, not consumed fuel). One alternative formulation is to apply strict equality constraints to the delivered power at all times [3]. This is suitable for applications where there is a true requirement to deliver a given power at every point in time, however, as described in their paper this is at cost of potential infeasibility as well as reducing the range of permissible control actions, virtually guaranteeing a higher fuel consumption compared to the case where delivered power is allowed to vary during the transient phase. A third option is to use an explicit minimum-fuel problem formulation [5]. This choice unequivocally matches our intent, but has the potential of leading to a problem formulation that is difficult to numerically solve.

The specific method used to solve the optimization problem is another aspect relevant to consider. The gradient descent method [4] is a traditional method, but requires a differentiable problem formulation and is not guaranteed to return a globally optimal solution. The multiple shooting and sequential quadratic programming method [5] is more sophisticated, but is still subject to the same fundamental limitations. An alternate approach is the use of stochastic optimization, e.g. a genetic algorithm [6]. This avoids the need for a differentiable problem formulation and is less likely to be trapped in a local optimal solution, but is typically computationally demanding and has no global optimality guarantee. One attractive option is to use a method based on Pontryagin's maximum principle or dynamic programming, as these typically are ensured to give the globally optimal solution. An example of this is seen in [3], though the specific method used requires hours of computational time to solve a problem that covers 10 seconds of operation.

Finally, it is worth considering how well (if at all) solution methods can translate into real-time control. One straightforward method is to limit equilibrium operation to a small number of discrete points [6], allowing for storing pre-computed trajectories and applying them during runtime. Though effective, this scales poorly with the number of equilibrium points, as we will typically need to store the optimal control trajectories over time for every combination of equilibrium points. Alternatively, domain-specific knowledge can be used to implement real-time methods

FIGURE 1: Series-hybrid vehicle drivetrain, the gen-set consists of the combustion engine and electrical machine.

FIGURE 2: Gen-set dynamic model overview.

that approximate the behavior of the computed optimal solutions [3, 4]. Though this can lead to very efficient real-time control methods, we are no longer guaranteed optimality while also requiring extensive knowledge of the specific problem to be solved.

In this paper we will introduce a method that addresses many of the current limitations. In particular, we will formulate the optimization problem as an explicit fuel consumption minimization problem, constrained to deliver a given average electrical power over an infinite time horizon which we solve using dynamic programming, ensuring a (close-to) globally optimal solution. The infinite time horizon is useful both as it ensures us that fuel consumption is balanced between the transient and stationary phase, as well as giving a solution that can be formulated as state-feedback. This ultimately allows for implementing real-time optimal control using a bare minimum of computational power (on the order of 5–20 KiB of nonvolatile memory, a few bytes of RAM, and 10–20 CPU instructions per control period).

2 Problem formulation

In this paper we will study the gen-set of a series-hybrid vehicle, as illustrated in Fig. 1. In particular, we will model the gen-set as a combustion engine directly connected to an electrical machine, with dynamics that arise from the moment of inertia of moving parts (illustrated in Fig. 2).

In the gen-set, we model the combustion engine (CE) as consuming fuel at a rate

$$\dot{m}(t) = f_C(\omega(t), \tau_C(t)), \qquad (1)$$

i.e. a function of the instantaneous torque $\tau_C(t)$ and crankshaft angular speed $\omega(t)$. Similarly, the electric machine (EM) delivers or consumes electric power as

$$P_E(t) = f_E(\omega(t), \tau_E(t)) \tag{2}$$
$$= -\tau_E(t)\omega(t) - f_{E,\text{loss}}(\omega(t), \tau_E(t)), \tag{3}$$

for a given instantaneous torque $\tau_E(t)$ and crankshaft speed $\omega(t)$. We will view τ_E and τ_C as control inputs that should ultimately be selected in a way that minimizes the fuel consumption of the gen-set.

Note that we assume f_C and f_E are instantaneous functions of the current crankshaft velocity and their respective torques, without any additional dynamics. This instantaneous formulation allows for determining them by using standard experimentally obtained equilibrium maps (i.e. by experimentally measuring the fuel massflow and delivered electrical power respectively at a wide range of operating points). We will return to the validity of this assumption in Subsection 4.1. Furthermore, note that (3) implies that we can view the EM as an ideal mechanical to electrical converter (with power given by $-\omega(t)\tau_E(t)$), with all non-frictional losses lumped into $f_{E,\text{loss}}$. For brevity, we will no longer explicitly state the time-dependence of the previous terms.

We lump all frictional components in the CE and EM together to form $\tau_f(\omega)$, which is typically non-linear. This implies that we can view τ_C and τ_E as indicated torques rather than crankshaft torques. Similarly, we lump together all inertial terms forming a net moment of inertia J. This gives a net continuous-time dynamic gen-set model

$$\dot{\omega} = \frac{1}{J}\left(\tau_C + \tau_E + \tau_f(\omega)\right), \tag{4}$$

i.e. a non-linear first-order dynamic system. Note that the nonlinearity of τ_f implies that (4) is generally a non-linear differential equation.

This specific choice of modeling and sign notation implies that a negative EM torque will apply a retarding torque to the crankshaft and deliver electrical power. Similarly, for positive EM torques the EM will tend to accelerate the crankshaft and thus consume electrical power. In steady-state operation, where the gen-set delivers electrical power, we will thus have $\tau_C > 0$, $\tau_E < 0$, and $P_E > 0$.

Assume that we will ultimately implement real-time control using a digital controller with a fixed sample rate t_s, where the control signals τ_C and τ_E are held constant between sample times (zero-order hold). For convenience,

we will use the notation

$$\omega_k \equiv \omega(kt_s) \tag{5a}$$
$$\tau_{C,k} \equiv \tau_C(kt_s) \tag{5b}$$
$$\tau_{E,k} \equiv \tau_E(kt_s) \tag{5c}$$

to indicate the discrete-time crankshaft speed and torques respectively. Using a numerical ODE solver we can solve (4) for a given initial condition and torques over time t_s, allowing us to introduce a discrete-time dynamic equation

$$\omega_{k+1} = f\left(\omega_k, \tau_{C,k}, \tau_{E,k}\right) \tag{6}$$

where f is given by the ODE solver.

We can now introduce the optimization problem we ultimately wish to solve as

$$(\tau_C^*, \tau_E^*) = \underset{\tau_C, \tau_E}{\operatorname{argmin}} \ \lim_{N \to \infty} \sum_{k=0}^{N} f_C\left(\omega_k, \tau_{C,k}\right) \cdot t_s \tag{7a}$$

subject to

$$\omega_{k+1} = f\left(\omega_k, \tau_{C,k}, \tau_{E,k}\right) \tag{7b}$$
$$\frac{1}{N}\sum_{k=0}^{N} f_E\left(\omega_k, \tau_{E,k}\right) = P_{\text{tgt}} \tag{7c}$$
$$g\left(\omega_k, \tau_{C,k}, \tau_{E,k}\right) \leq 0, \tag{7d}$$

i.e. over an infinite horizon minimize fuel consumption, while delivering a given average power P_{tgt}, and demanding the CE and EM torques and crankshaft velocity lie in a permissible range given by g. Note that though the sum in (7a) grows arbitrarily large with increasing N, the minimand (τ_C^*, τ_E^*) is for this problem well-defined [1], as $f_C \cdot t_s$ is finite (the CE can only consume a finite amount of fuel in finite time).

Note that the choice of an infinite horizon ($N \to \infty$) obviates the need for determining a control and prediction horizon commonly found in finite-horizon problem formulations (e.g. traditional model predictive control (MPC)).

In this paper we have studied a conventional direct-injection turbocharged gasoline combustion engine coupled with a permanent-magnet synchronous electric machine. These are representative examples of gen-set components

that could be found in a typical light-duty series-hybrid. The functions f_C and f_E are approximated by empirically measuring the physical fuel flow and delivered electrical power respectively for a large set of equilibrium operating points. Similarly, τ_f is determined by measuring the torque required to motor the unloaded crankshaft (i.e. τ_C, τ_E are set to 0) while maintaining a constant speed using an additional external EM. The function g is defined to give box constraints for the permissible state and control signals as

$$g\left(\omega, \tau_C, \tau_E\right) = \begin{cases} 0 & \text{if } 0 \leq \omega \leq 750 \\ & \text{and } 0 \leq \tau_C \leq 390 \\ & \text{and } -150 \leq \tau_E \leq 150 \\ 1 & \text{else.} \end{cases} \quad (8)$$

Illustrations of the combustion engine's normalized indicated specific fuel consumption, the electric machine's efficiency (defined as $\eta_{\text{EM}} \equiv (|f_E| - f_{E,\text{loss}})/|f_E|$, and the net friction τ_f are shown in Fig. 3.

3 Method

Minimum Horizon Dynamic Programming (MHDP), a method recently developed by the authors [2], is a method based on dynamic programming and the successive approximation method for solving general average-constrained infinite-horizon nonlinear optimal control problems of form

$$\left(x^*, u^*\right) = \underset{x,u}{\text{argmin}} \lim_{N \to \infty} \sum_{k=0}^{N} c\left(x_k, u_k\right) \quad (9a)$$

subject to

$$x_{k+1} = f_d\left(x_k, u_k\right) \quad (9b)$$

$$\frac{1}{N} \sum_{k=0}^{N} f_a\left(x_k, u_k\right) = \alpha \quad (9c)$$

$$g\left(x_k, u_k\right) \leq 0 \quad (9d)$$

where $x_k \in \mathbb{R}^n$ and $u_k \in \mathbb{R}^m$ are the k'th state and control variables respectively. Ultimately, MHDP determines the optimal controls u^* in the sense that (9a) is minimized while respecting the system dynamics and constraints.

It can be shown that under some mild assumptions, for problems of form (9a), that the optimal control u^* to apply at any given time solely depends on the current system

(a) CE fuel consumption.

(b) EM electrical power.

(c) Net crankshaft friction.

FIGURE 3: Plots of normalized CE fuel consumption, EM efficiency, and τ_f. As the CE and EM are modeled as frictionless devices the crankshaft torques can be views as indicated rather than brake-specific torques.

state [2]. Essentially, there exists a function $f_{\mathrm{sf}} : \mathbb{R}^n \to \mathbb{R}^m$ such that $u_k^* = f_{\mathrm{sf}}(x_k)$. Typically f_{sf} is referred to as an *optimal control policy* or *optimal non-linear state feedback*. MHDP generates an approximation of f_{sf} and returns a look-up table with the optimal controls to apply for a given set of discrete state values. Once f_{sf} is known this allows for near-trivial on-line control, as the controller can simply look up the tabulated state value (typically using some form of interpolation if the current system state lies between tabulated values) and apply the resulting control signal.

Though a detailed description of the MHDP method is beyond the scope of this paper, an overview of the parameters required by the method can aid in understanding and interpreting the generated results. Given a non-linear (but well-behaved) cost c, system dynamics f_d, averaging function f_a, a demanded average α, and general constraints g MHDP can return the optimal control policy associated with (9). Furthermore, MHDP requires selecting a given state and control quantization; more densely sampled points will give a more accurate approximation of f_{sf}, but generating the state-feedback law will take longer as well as requiring more memory in the real-time controller (as the optimal control policy table will contain more entries). Finally, an additional consequence of the MHDP method is that the generated optimal control policies are not explicitly parameterized by the average constraint α (P_{tgt} in our application), but by another scalar parameter λ analogous to a Lagrange multiplier. We resolve this in this paper by searching for values of λ that give the desired average power levels.

Returning to our specific problem, solving (7) with MHDP will generate the optimal control policies

$$(\tau_C^* (\omega, P_{\mathrm{tgt}}), \tau_E^* (\omega, P_{\mathrm{tgt}})), \tag{10}$$

which we can view as the optimal torques τ_C^* and τ_E^* to apply tabulated by different crankshaft speeds ω and the target-power P_{tgt}.

A block diagram illustrating the final controller construction is shown in Fig. 4. As is shown, the requested power P_{tgt} is fed into the tabulated control laws τ_C^*, τ_E^* along with the current crankshaft speed ω, giving a closed-loop non-linear state-feedback controller.

For our specific problem (7) the real-time system need only perform a two-dimensional interpolation operation every sample. This typically takes on the order of 10–20 CPU instructions and (using the parameters to be presented shortly in Table 1) requires a look-up-table consuming approximately 13 KiB (as $N_{P_{\mathrm{tgt}}} \cdot N_\omega = 6560$ entries are required, using 16-bit words gives ≈ 13 KiB), both of which are trivially performed in virtually any embedded system.


FIGURE 4: Net closed-loop block diagram.

TABLE 1: MHDP parameters used for (7).

PARAMETER	VALUE	UNIT	DESCRIPTION
t_s	$50 \cdot 10^{-3}$	s	Sample time
$\Delta\tau$	4	Nm	Torque quantization
$N_{P_{\mathrm{tgt}}}$	80	-	Number of evenly spaced equilibrium power levels.
N_ω	82	-	Number of evenly spaced discretized states.

4 Results

We have solved (7) using MHDP (taking approximately 12 hours with a standard desktop PC) with the configuration listed in Table 1, where the state and control signals have been discretized by evenly spacing them in the range allowed by (8). The specific choice of parameters in Table 1 was chosen to give a torque and crankshaft speed quantization that is on par with the measurement accuracy of the system.

Illustrations of the optimal CE and EM torques for representative target powers are shown in Fig. 5a and Fig. 5b. These plots illustrate the low computational demand of the on-line controller — the optimal control is given simply by consulting the stored torque corresponding to the current crankshaft speed and desired power.

Studying Fig. 5a and Fig. 5b allows us to come to some conclusions about the optimal control policy. For instance, for the 37 kW and 62 kW target power the EM will accelerate the gen-set when the crankshaft speed is less than approximately 200 rad/s. We speculate that this is due to the slightly lower efficiency of the CE at low speeds, as is seen in Fig. 3a, and it is thus optimal to spend as little time

Optimal CE torques

Optimal EM torques

FIGURE 5: Optimal CE and EM torques for representative target powers.

as possible in this region. Furthermore, for the specific case of a 0 kW setpoint, $\tau_C = 0$, i.e. the combustion engine is set to consume no fuel. In the event that $\omega \neq 0$ then $\tau_E < 0$, which we can interpret as the optimal policy being to decelerate the crankshaft to zero speed while converting as much as possible of the crankshaft's stored kinetic energy to electrical energy.

In order to evaluate the performance of the optimal controller generated by MHDP we have chosen to compare it to a traditional gain-scheduled proportional controller. In this paper we will use a controller with control law

$$
\begin{bmatrix} \tau_{C,P}\left(\omega,\omega_{\mathrm{EQ}}^*\right) \\ \tau_{E,P}\left(\omega,\omega_{\mathrm{EQ}}^*\right) \end{bmatrix} = \mathrm{clamp}\left(\begin{bmatrix} \tau_{C,\mathrm{EQ}}^*\left(\omega_{\mathrm{EQ}}^*\right) \\ \tau_{E,\mathrm{EQ}}^*\left(\omega_{\mathrm{EQ}}^*\right) \end{bmatrix} + \begin{bmatrix} k_p\left(\omega_{\mathrm{EQ}}^*\right) \\ -k_p\left(\omega_{\mathrm{EQ}}^*\right) \end{bmatrix} \left(\omega - \omega_{\mathrm{EQ}}^*\right) \right) \quad (11a)
$$

$$
\omega_{\mathrm{EQ}}^* = f_\omega^*\left(P_{\mathrm{tgt}}\right). \quad (11b)
$$

Here, $f_\omega^*\left(P_{\mathrm{tgt}}\right)$ is a function that gives the equilibrium crankshaft speed that consumes the least fuel for a given

target power P_{tgt}. $\tau_{C,\mathrm{EQ}}^*$ and $\tau_{E,\mathrm{EQ}}^*$ are functions that return the torques needed to keep the system stationary at the equilibrium speed (which can be viewed as feed-forward terms), clamp is the clamping function that limits $\tau_{C,P}$ and $\tau_{E,P}$ to the permissible values given in (8), and k_p is a parameterized controller gain (which can be viewed as a proportional feed-back gain term). Due to the system's dynamics (4), the structure of (11) ensures that the gen-set will asymptotically approach the optimal operating point but will likely be sub-optimal during the transient approach.

The specific value of the controller gain k_p is critical for the performance of the closed-loop system. With small values a large amount of time is spent operating in regions with low efficiency, while large values will tend to apply very large torques which also typically reduces efficiency. In order to give a fair comparison between the optimal and proportional controllers, we have selected k_p to give a step response time from standstill ($\omega = 0$) to the target speed ($\omega = \omega_{\mathrm{EQ}}^*$) that takes an equally long time as the optimal solution (see Fig. 6). Note that we have now implicitly made our proportional controller to some degree optimal — it is unlikely that a real-world controller would happen to be optimally tuned and would presumably consume even more fuel than the results shown below.

Note that this specific choice of a gain-scheduled proportional controller ensures that the gen-set will deliver the same electrical power and consume fuel at the same rate when at equilibrium operation. In all following comparisons we will thus only consider the transient phase, as it is only in this region where the optimal controller has any potential for improvement.

In all the following results, we have chosen to use the piecewise-constant setpoint function

$$
P_{\mathrm{tgt}}\left(t\right) = \begin{cases} P & 0 \leq t \leq T \\ 0 & \mathrm{else}, \end{cases} \quad (12)
$$

where T is chosen to be large enough to bring the gen-set to near-equilibrium operation. We have chosen this specific setpoint function as this allows us to fairly compare gen-set controllers. A full drive-cycle simulation would require choosing a specific torque-split controller (i.e. when and how much power to draw from the battery and gen-set respectively), the choice of which will greatly influence the total system behavior. Furthermore, we can note that only steps in power will give a truly optimal control signal, as the problem formulation (9) assumes a constant average power constraint. We speculate however that slowly-varying power demands will also be near-optimal (as no significant dynamics are excited). To limit the scope of this paper we will thus only consider the simple power step (12).

FIGURE 6: Step response from standstill to equilibrium crankshaft speed for optimal and proportional controllers. For fairness, the proportional controller gain is set to make the system reach equilibrium at the same time as the optimal controller.

In Fig. 7 we can study the operating points for the CE and EM for the optimal controller (7) and the traditional controller (11). We can see that though both controllers reach the same equilibrium operation point (indicated by x_{EQ}), their paths to the equilibrium point are different. In particular, in Fig. 7a we can see that the proportional controller generates a trajectory that happens to pass through a region with poor efficiency ($\omega = 75$, $\tau_C = 280$), a region the optimal controller avoids entirely.

As both controllers bring the gen-set to the same operating point in the same time we can numerically compare them by studying their effective efficiency (i.e. delivered electrical energy per consumed fuel mass)

$$\eta = \frac{\int_0^\infty f_E \, dt}{\int_0^\infty f_C \, dt}. \tag{13}$$

Note that as $P_{\text{tgt}} = 0$ for $t > T$, (13) can be viewed as a measure of the efficiency of the gen-set during solely transient operation, and indicates the gen-set efficiency both for an increase and decrease in target power.

As the controllers do not necessarily deliver the same amount of energy during the transient phase, the time T in (12) is adjusted separately for each controller until the total delivered power is identical.

In essence, both controllers will start and end at the same crankshaft speed, reach the same equilibrium speed at the same time, and deliver the same amount of energy

over the whole cycle. The controller's freedom lies in the choice of transient torques, as well as when electrical power is delivered.

Fig. 8 shows the efficiency improvement during the transient phase of operation, with fuel savings on the order of 2–7%. We believe that the large variation in efficiency gains is primarily due to the traditional controller sometimes performing fairly well (e.g. at 52 kW) and sometimes poorly (e.g. for ≤ 45 kW) compared to the optimal controller. We can see an indication of this in Fig. 7, where in Fig. 7a the proportional controller's trajectory passes through an inefficient region, while in Fig. 7c the proportional controller's trajectory happens to be closer to the optimal trajectory.

4.1 Bench-test results

In this section we verify the previous results by implementing and evaluating the optimal controller in a physical, real-time, control system.

We have performed this test using a light-duty plug-in hybrid electric vehicle with dynamometers mounted on all wheels as shown in Fig. 9. We use the same problem formulation as in the previous section, i.e. solve for (7), (8), with an additional constraint limiting the electrical power to 15 kW due to the vehicle's relatively small inverter. Mathematically we can demand this by adding

$$|f_E| \leq 15 \cdot 10^3$$

as a constraint to (8).

Notably, due to the minimal computational load of implementing the optimal control policy as given by MHDP, the entire control loop is performed natively in the vehicle's engine control unit without any support from external hardware (as shown in Fig. 4). In these tests we have generated state-feedback control laws for 102 discrete crankshaft speeds over the gen-set's operating range (using linear interpolation to determine τ_C and τ_E for intermediate crankshaft speeds) and two power levels ($P_{\text{tgt}} = P$ and $P_{\text{tgt}} = 0$). τ_C^* and τ_E^* were stored as 32-bit floating point values, giving a total nonvolatile memory consumption of 1.6 KiB. Tests have shown that the on-line controller consumes on the order of 10–15 CPU cycles per sample, which is similar to that of the existing (suboptimal) controller. We have compared the optimal controller with a controller of form (11), which is comparable to controllers used in the industry.

We have tested the vehicle's gen-set for a setpoint power of $P = 15$ kW, with a corresponding equilibrium crankshaft speed of 205 rad/s. Due to resource limitations only a single (15 kW) power level was evaluated. In Fig. 10 we can view ensemble plots of the CE and EM torque for several
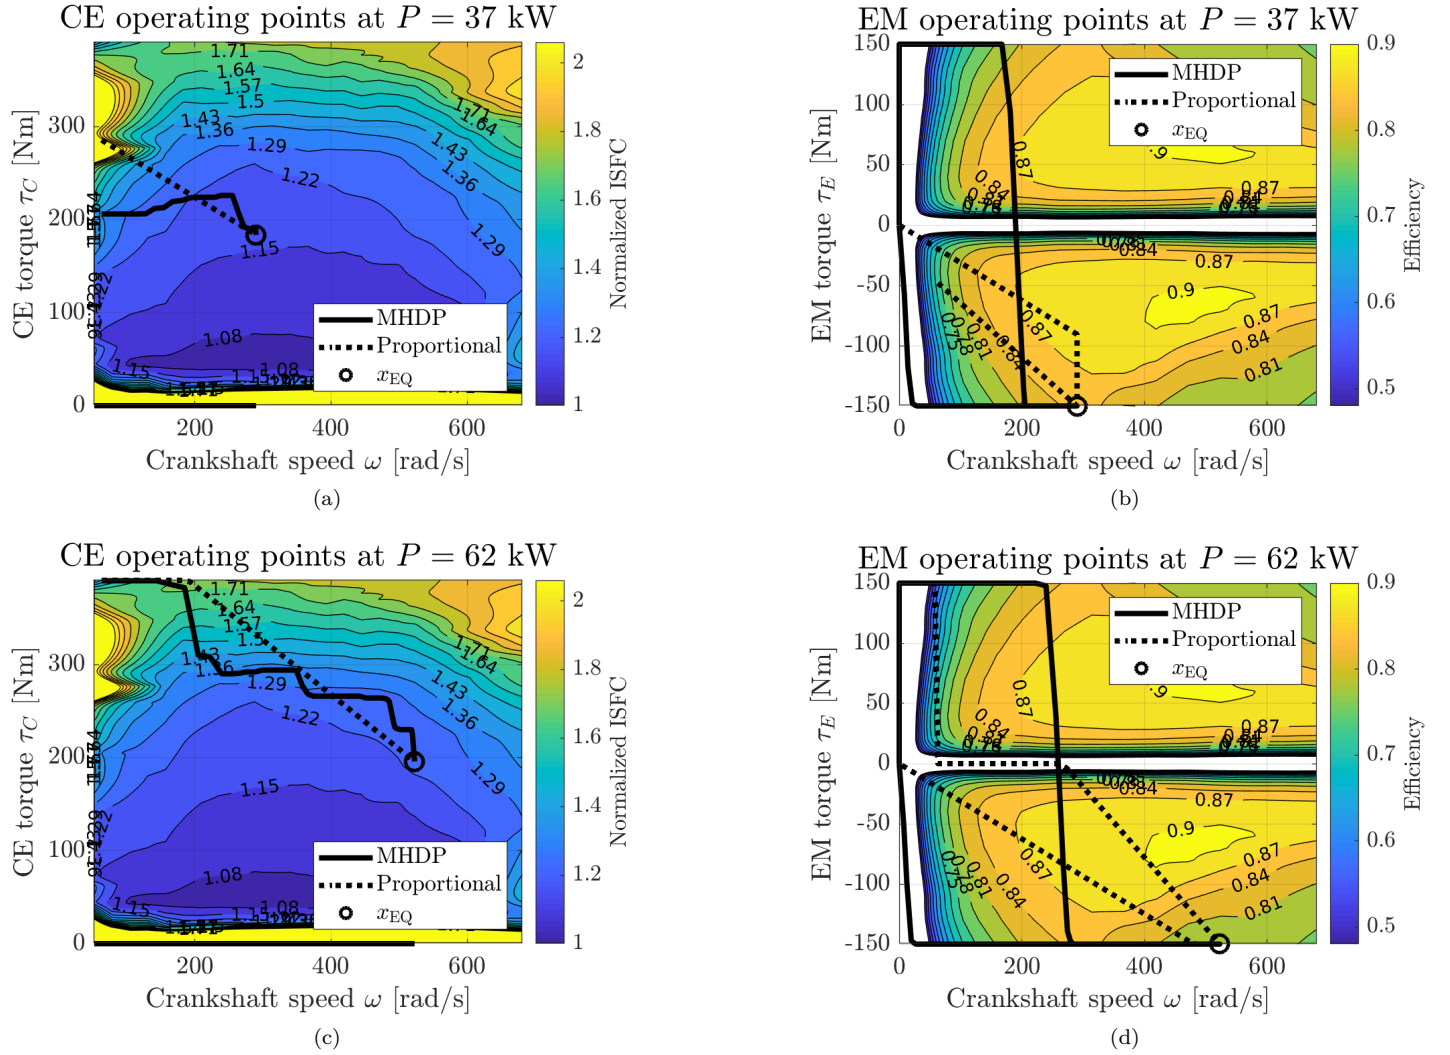
FIGURE 7: Typical CE and EM transient operating point trajectories.

repeated tests, with an associated average efficiency improvement of 3.7%. The difference in equilibrium speed in Fig. 10c is an artifact of the test-bench, but has little effect on the results as the CE efficiency is virtually identical for $\omega = 200$ and $\omega = 210$. Notably, the structure of the proportional controller gives a solution that tends to load the CE more at speeds where it is inefficient ($0.5 \leq t \leq 1$).

There are several potential avenues of further improving the fuel consumption of the optimal controller beyond the results seen in this test setup. One significant limitation is the gen-set's EM inverter power constraint. This limits the permissible operating points to a small portion of the theoretical speed/torque space. For these tests we also have allowed the engine to rest at idle between steps to P (rather than allowing it to reach a standstill) to avoid activating the existing engine start-up routines. Furthermore, there are several unmodeled sources of dynamics in this physical setup, the turbocharger among others. Finally, the functions f_C, f_E, and τ_f only accurately model equilibrium operation, while we have implicitly assumed that they hold during transient operation. We ultimately have a trade-off between model accuracy and the computational load of the controller (whose non-volatile memory usage scales exponentially with the number of state variables).

FIGURE 8: Efficiency improvement (delivered energy per consumed fuel mass) for the optimal controller compared to the proportional controller.

FIGURE 9: Test-bench setup. A conventional passenger vehicle is used with the same CE and EM as shown in Fig. 3.

FIGURE 10: Ensemble plots of CE and EM torque and crankshaft speed ω from bench-tests.

5 Conclusions

We have shown that an optimal gen-set controller that takes transient dynamics into account can effectively be implemented and run in real-time. Minimizing the consumed fuel using a dynamic programming method with an infinite horizon ensures that the optimal control policy captures the true intent of the problem studied in this paper; minimize fuel consumption while delivering a requested equilibrium power. This also avoids issues with penalty-based methods, does not require any gradients to be known, and is guaranteed to be close to the globally optimal solution. Using the MHDP method, which is based on dynamic programming and the successive approximation method, allows us to off-line easily construct a state-feedback control law with very low on-line computational demand (on the order of 10–20 CPU instructions and 5–20 KiB of consumed memory).

Simulation results show that a gen-set in a typical light-duty vehicle can expect a fuel consumption improvement on the order of 5–7% during transient operation when compared to current industry-standard controllers. Bench-tests displayed a fuel consumption improvement of 3.7%, indicating that the simulation results are fairly representative. This reduction of fuel consumption may be further increased by using more sophisticated models, in particular a combustion engine fuel consumption model that takes additional transient dynamics, such as a turbocharger, into account.

Relevant future work includes studying the efficacy of the proposed method for the more ultimately relevant case of a whole-vehicle drive cycle. The choice of torque-split controller (i.e. when and how much power is drawn from the gen-set and battery respectively) as well as the selected drive cycle will heavily influence the net effectiveness of the proposed gen-set controller, and should thus be prudently chosen.

Acknowledgements

References

[1] Dimitri P. Bertsekas. *Dynamic Programming: Deterministic and Stochastic Models*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1987.

[2] Jonathan Lock and Tomas McKelvey. Constant-setpoint infinite-horizon nonlinear optimal control using dynamic programming. 2019. Available at `https://doi.org/10.5281/zenodo.2613415`.

[3] Tomas Nilsson and Anders Fröberg. Optimal operation of a turbocharged diesel engine during transients. *SAE International Journal of Engines*, 5(2):571–578, April 2012.

[4] Carlos E. Nino-Baron, Abdul R. Tariq, Guoming Zhu, and Elias G. Strangas. Trajectory optimization for the enginegenerator operation of a series hybrid electric vehicle. *IEEE Transactions on Vehicular Technology*, 60(6):2438–2447, July 2011.

[5] Martin Sivertsson and Lars Eriksson. Optimal transient control trajectories in diesel-electric systems – part I: Modeling, problem formulation, and engine properties. *Journal of Engineering for Gas Turbines and Power*, 137(2):021601, September 2014.

[6] Xi Zhang, Zixian Wu, Xiaosong Hu, and Wei Qian. Trajectory optimization-based auxiliary power unit control strategy for an extended range electric vehicle. *IEEE Transactions on Vehicular Technology*, 66(12):10866–10874, December 2017.