Using Stochastic Dynamic Programming for look-ahead control of a Wheel Loader Diesel Electric Transmission

Tomas Nilsson * **Anders Fröberg** ** **Jan Åslund** ***

* *Department of Electrical Engineering, Linköping University, Linköping, Sweden, (email: tnilsson@isy.liu.se)*
** *Volvo CE, Eskilstuna, Sweden, (email: anders.froberg@volvo.com)*
*** *Department of Electrical Engineering, Linköping University, Linköping, Sweden, (email: jaasl@isy.liu.se)*

Abstract: Three Stochastic Dynamic Programming (SDP) implementations are developed for control of a diesel-electric wheel loader transmission. The implementations each use a stochastic description of the load, with the probabilities either independent of the states, dependent on previous power or on distance driven. Both the cycles used for the controller development and for the evaluation are derived from a measured sequence of cycles.
The evaluation shows that SDP can be used for control of the engine speed and that the resulting trajectories from the three implementations are very similar. The most surprising part is that the method which has constant load probability is able to adjust to the actual load. The combination of the calculation efforts and the outcomes leads to the conclusion that the constant load probability implementation is superior to the other versions.

Keywords: Vehicles, Engine control, Hybrid vehicles, Continuously variable transmission, Optimal control, Dynamic programming, Uncertainty, Stochastic variables

1. INTRODUCTION

1.1 Background

Wheel loader operation is highly transient and repetitive, and contains periods of high tractive effort at low speeds, while the engine delivers power to both the transmission and to the hydraulics. Details about common operation can be found e.g. in Filla (2008), Wang et al. (2012) and in Nezhadali and Eriksson (2013). The most common layout of heavy wheel loader powertrains is presented in Figure 1. The engine is connected directly to a variable-displacement hydraulic pump and to the drive shaft through a torque converter and an automatic gearbox.

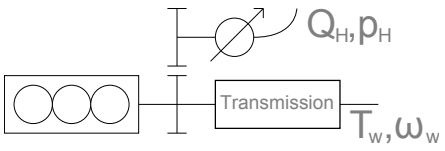

Fig. 1. The reference vehicle drivetrain setup.

In this setup the torque converter is a crucial component since it provides some disconnection between the engine and transmission speeds. This disconnection makes the system mechanically robust by preventing the engine from stalling if the bucket gets stuck, but it also causes high losses. This lack of efficiency is the reason for a desire to find other transmission concepts for wheel loaders.

1.2 On the choice of transmission

Any alternative transmission should increase the efficiency at the typical operation. The combination of low speeds, high forces and transient operation motivates the use of some type of continuously variable transmission (CVT). An evaluation, through the minimum fuel consumptions, of the standard and a hydrostatic transmission is made in Nilsson et al. (2012a). Several other papers, such as Lennevi (1995), Rydberg (1998) and Zhang et al. (2002) study the use of hydrostatic CVTs in wheel loaders, though the focus is on component control. Stein et al. (2013) analyses some different transmission solutions, including the hydrostatic device in Nilsson et al. (2012a) and a series electric hybrid, and indicates higher possible efficiency benefits in the electric solution. An electric hybrid is also considered, among other solutions, in Filla (2008). In comparing hydrostatic and electric transmissions, the low level control of a hydrostatic transmission is more complex due to slow dynamics of the hydraulic machines. Since the focus of this paper is on optimal selection of engine speed, a choice is made to use an electric transmission. The same system as in Stein et al. (2013) is used, though the system is simplified by not including a secondary energy storage.

1.3 Reasons for using stochastic dynamic programming

Some heuristic higher level CVT controller concepts can be found in Liu and Paden (1997). These however do not fully utilize the potential of the transmissions. For on-road vehicles, there have been several proposals, e.g. Hellström (2010) and Khayyam et al. (2012), for utilizing increased

computational capacity and availability of information, such as road maps and GPS data, for predictive control. This type of information is in general not available for off-road applications. Wheel loader operation is often highly repetitive, and this repetitiveness can be used for a rough prediction of future operation. One method for using an uncertain prediction for optimal control is stochastic dynamic programming (SDP). This approach is used e.g. in Johannesson et al. (2007), Kolmanovsky and Filev (2010) and McDonough et al. (2012), though these treat on-road vehicles. This paper studies the implementation of SDP for wheel loader CVT control.

2. PROBLEM

2.1 Problem statement

The problem studied in this paper is the minimization of the expected average amount of fuel needed for completing each cycle in a long series of driving cycles. For a fixed time or distance cycle, this can be formulated as a minimization of the expected average fuel flow. With U representing a set of yet unspecified control signals and M_f the total fuel use, this can be expressed as

$$\min_{U(t)} \lim_{T\to\infty} E\frac{1}{T}\int_0^T \frac{dM_f}{dt}\mathrm{d}t \tag{1}$$

while obeying the system dynamics and constraints, and following the specified driving cycle. The driving cycles of this paper consists of vehicle speed v_w and demanded power P_R, and the prediction of P_R contains uncertainties. Requiring that the engine should always be able to supply the demanded power would lead to the controller always selecting high engine and turbo speeds. Since this would be detrimental to the efficiency, a power limiter P_{max} for the actual electric load P_L is introduced according to

$$P_L = \min(P_R, P_{max}) \tag{2}$$

Using this limiter, through lowering of P_{max}, should be avoided unless necessary though. It is assumed that this limiter will be used sparsely and that the rest of the driving cycle will therefore not be affected by any usage.

2.2 System models

The complete system studied is presented in Figure 2. Any relevant dynamics and efficiencies in the electric power consuming parts are included in the desired electric power of the driving cycle. The following model, which describes the electricity producing part of the diesel electric powertrain, is identical to that used in Nilsson et al. (2012b), apart from parameter values and that a smoke limiter is added. The system consists of a turbo charged diesel engine and an electric machine. Apart from the inertia, the dynamics of the electric machine is assumed to be fast compared to the other dynamics of the system.

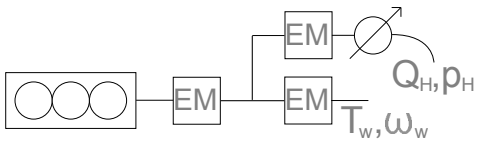

Fig. 2. The diesel-electric vehicle drivetrain setup.

The engine and generator is modeled as an inertia I_e with the speed ω_e, affected by the torque T_e and the load P_L.

$$\frac{d\omega_e}{dt}\cdot I_e = T_e - \frac{P_L}{\omega_e} \tag{3}$$

The relation between injected fuel and engine torque is described by a quadratic Willan's model, as presented in Rizzoni et al. (1999), expanded with a turbo model, according to (4). The fuel flow and fuel mass per injection are related according to (5). This model is illustrated by the engine map in Figure 3.

$$T_e = \frac{q_{lhv} n_{cyl}}{4\pi} e(\omega_e, m_f) - T_L(\omega_e) - T_{pt}(p_t, \omega_e, m_f) \tag{4}$$

$$\frac{dM_f}{dt} = \omega_e \frac{n_{cyl}}{4\pi} m_f \tag{5}$$

Here q_{lhv} is the lower heating value of the fuel, n_{cyl} is the number of cylinders, m_f is fuel mass per injection, ω_e is engine speed and e and T_L are static efficiency functions. T_{pt} is a torque loss due to low intake air pressure, caused by low turbo charger speed. The actual intake air pressure is p_t, p_{set} is a static setpoint map and the pressure offset is described by $p_{off} = p_t - p_{set}(\omega_e, m_f)$. The torque loss T_{pt} is described by

$$T_{pt} = \begin{cases} k_1(\omega_e)\cdot p_{off}^2 - k_2(\omega_e)\cdot p_{off} & \text{if } p_{off} < 0 \\ 0 & \text{if } p_{off} \geq 0 \end{cases} \tag{6}$$

while the turbo dynamics is modeled as a first order delay, with the time constant τ, for the intake air pressure according to

$$\frac{dp_t}{dt}\cdot\tau(\omega_e) = -p_{off}(\omega_e, m_f) \tag{7}$$

The engine also has a smoke limiter which sets an upper limit $T_{SL}(p_t, \omega_e)$ for the engine torque T_e.

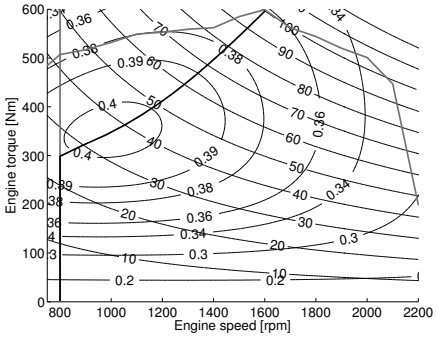

Fig. 3. Engine map with efficiency curves, output power lines with kW markings, minimum speed and maximum torque bounds and the static optimal line.

2.3 Driving cycles

The driving cycles consists of vehicle speed and electric power requirement as functions of distance driven. The requested power P_R is the sum of the transmission power P_T and the hydraulic power P_H. The cycles used were derived from data collected during a measurement sequence which consists of 34 short loading cycles. The cycles were extracted utilizing the cycle detector presented in Nilsson et al. (2014). For simplicity, the distances and vehicle speeds in each cycle were adjusted so that each driving direction change occur at the same distance in all the cycles and the maximum accelerations are unchanged.

The vehicle available for the measurement differs from the modeled system in both size and system layout. The measurement vehicle has a layout similar to that presented in Figure 1. Because of this the power required was not readily available but must be calculated from other signals and scaled to fit the modeled vehicle. The hydraulic power P_H was calculated from the hydraulic pressure and the derivative of the arm and bucket angle, which correspond to hydraulic cylinder volume changes. This power was scaled according to the maximum bucket load and cylinder volumes of the vehicles. For the transmission power P_T calculation, the cycles were partitioned into two types of operation; bucket filling and other operation. The power during the bucket filling was calculated from the torque converter characteristics according to a standard model (see e.g. Nilsson et al. (2012a)). The power at the rest of the cycle was calculated from a vehicle model according to

$$P_T = \frac{m dv_w^2}{2dt} + v_w m g c_r \tag{8}$$

in which m is vehicle mass, v_w is vehicle speed and c_r is rolling resistance. The vehicle mass consists of empty mass plus bucket load. The vehicle is assumed to be carrying a standard load weight when loaded and no weight when not loaded. The rolling resistance parameter c_r is given a typical value for this vehicle on hard packed soil. The adjusted speed and power trajectories of the 34 cycles are displayed in Figure 4. The bucket is filled during the high power episode after $5m$ and emptied at the zero speed event at $24m$. In the evaluation, 33 of these cycles are used for creating load probability functions and the last cycle is used in the subsequent simulations.

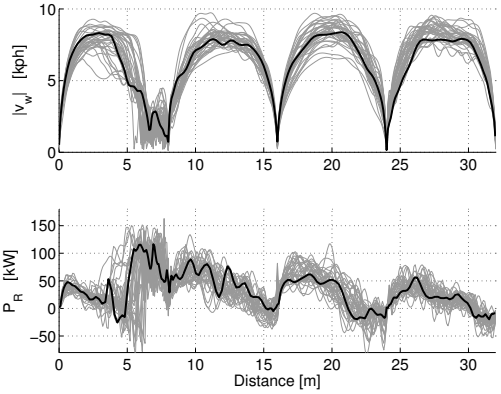

Fig. 4. The 34 driving cycles. The gray lines indicate cycles used for creating the load probability functions and the black line indicate the cycle used in the evaluation.

3. METHOD

The optimization problem is formulated as an average cost per stage problem. In this paper, stochastic dynamic programming (SDP) and the relative value iteration algorithm, as described in detail in Bertsekas (2005) and in Puterman (2005), is used for solving this problem. The same approach is also used in Kolmanovsky and Filev (2010) and in McDonough et al. (2012). The controller can be expressed as a state to control signal mapping, and the calculation of this mapping can be summarized as follows.

A cost function g, related to the target function (1), is defined. The formulation of the actual cost function is discussed later in this section. Denote the discretized states $x \in X$, and controls $u \in U$. The probability of the load being $w_l \in W$ is denoted $p(w_l)$. Select arbitrary values for J^1, an arbitrary state x^*, and an optimality condition ε. Now, update J according to

1: **while** $\rho > \varepsilon$ **do**
2: For each $x \in X$, calculate

$$J^{i+1}(x) = \min_{u \in U} \sum_{w_l \in W_k} p_k(w_l)\big(g(x_k, u, w_l) + ... + \tilde{J}^i(x_{k+1}(x_k, u, w_l))\big)$$

in which $\tilde{J}^i$ is interpolated from $J^i(x \in X)$
3: set $J^{i+1} \leftarrow J^{i+1} - J^{i+1}(x^*)$
4: calculate $\rho = \sup(J^i - J^{i+1})$
5: **end while**

The result of the algorithm is a map $J(x \in X)$ which can be used directly in a controller that at each instant selects the control action which minimizes the function

$$\hat{u} = \underset{u \in U}{\operatorname{argmin}} \sum_{w_l \in W_k} p_k(w_l)\big(g(\hat{x}_k, u, w_l) + ... + \tilde{J}(x_{k+1}(\hat{x}_k, u, w_l))\big) \tag{9}$$

in which $\hat{u}$ and $\hat{x}$ denote the actual control and state and $\tilde{J}$ is interpolated among $J(x \in X)$. Since $p(w_l)$ is a function of the states X, the controls $\hat{u}$ are functions of the states. In a real implementation the functions $\hat{u}(x \in X)$ would be precalculated to save online computational complexity. In this paper, since the controllers are only evaluated through simulations, this step has not been performed.

In this paper the load W consist of non-varying vectors of values for P_R, while the corresponding probability distributions p may depend on the states. A major part of the design of an SDP controller is the choice of states to include in X and how these should be included in the probability functions. The states used here are the two essential engine states ω_e and p_t, and a maximum of one additional state that may be used in $p(P_k|X_k)$. The probabilities can be based on e.g. previous power, as in Kolmanovsky and Filev (2010), or a position related state. Johannesson et al. (2007) implements both of these for on-road vehicles. In this paper, three probability dependencies are evaluated; independent of state $p(P_k|-)$, depending on previous power $p(P_k|P_{k-1})$ and depending on distance driven $p(P_k|s_k)$. In the later two the probabilities are assumed normally distributed due to the low number of measurements at each state P_{k-1} or s_k. The probability distributions are displayed in Figures 5 and 6.

In the general SDP problem, all state transitions may depend on both the control inputs and the stochastic load. In this problem, this is simplified in each of the three formulations. In all three setups, the change of engine speed is a control signal, as discussed in the next paragraph, and there is therefore no stochastic component to this state transition, but the resulting fuel use cause a stochastic component to the turbo dynamics. In the second setup, $p(P_k|P_{k-1})$, since the previous load is the state, the P_k state transition is directly given by the load. In the third setup, $p(P_k|s_k)$, the change of s depends on the vehicle speed, which might be described as stochastic.

Here it is decided to use a deterministic vehicle speed, not only because this reduces the number of possible load combinations, but also because the distance driven can then be discretized so that s_{k+1} is always on the grid. The interpolations of $\tilde{J}$ can therefore be made two dimensional in the J map calculations, in all three setups.

A natural choice of control signals for this system is injected fuel mass m_f, or engine torque T_e, along with the power limiter P_{max}. The large spread and quick changes in P_R means that regardless of probability function, the states $x_{k+1}(x_k, u, W)$ will get such a large spread that some component will always be outside the valid operating region, causing infinite expected cost. The only valid solution is for the controller to always keep P_{max} near zero. Since this is not acceptable, another control signal setup is proposed. It is assumed that there exist an engine speed controller which is fast compared to the update frequency of the SDP controller. The control signals in U are selected to be $\frac{d\omega_e}{dt}$ and P_{max}, and the fuel used is calculated from an inverted engine model. In the simulation, situations may now occur when a selected $\hat{u}$ is not feasible. In that case, a feasible solution is sought primarily by reducing $\frac{d\omega_e}{dt}$, and reducing P_{max} only if necessary. The signal P_{max} does not carry any natural penalty, and an artificial penalty is therefore introduced. In this paper, the cost function

$$g = M_f(\omega_e, p_t, P_L)\Delta t + \beta P_{max} \tag{10}$$

with β being a large negative constant, is used. Other penalty formulations, for example based on the probability $p(P_{max} < P_R)$, may of course also be of interest but are not investigated in this paper.

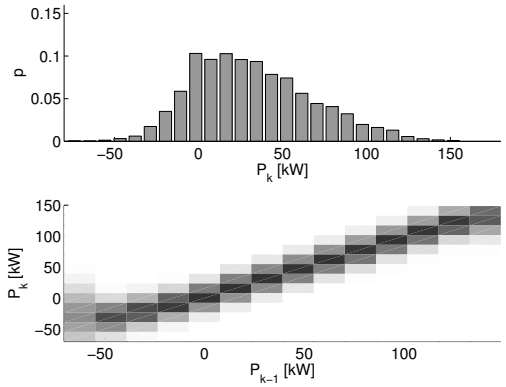

Fig. 5. The functions $p(P_k|-)$ and $p(P_k|P_{k-1})$.

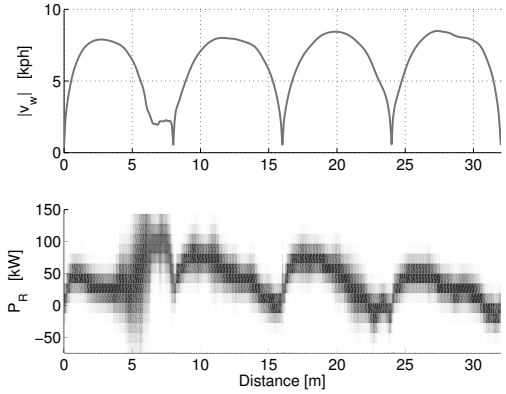

Fig. 6. The functions $v_s(s_k)$ and $p(P_k|s_k)$.

4. EVALUATION

4.1 Driving cycles

The driving cycles used are derived from a measurement sequence not specifically performed for use in this paper. Some properties of the cycles, which depend on the measurement and derivation, should therefore be pointed out.

Speed sensors often have lower accuracy at low speeds, such as those at which this vehicle operate at some points. During bucket filling, the load torque is calculated from the torque converter characteristics using the input and output axle speeds, and this may therefore be affected by the measurement uncertainty. Further, the speed signal is differentiated for use in (8), when away from bucket filling. The speed signal is filtered to reduce the impact on P_T.

In the working hydraulics, the available arm and bucket angle sensors have quite low resolution. These signals are differentiated to provide the changes of hydraulic cylinder volumes, and thus the hydraulic flow. Heavy filtering have been necessary for obtaining a realistic hydraulic flow signal. In the hydraulic pressure, the measurement is of good quality, but the instantaneous pressure is highly affected by vehicle pitch motions, which introduce a high amplitude oscillation in the hydraulic pressure. The frequency of these oscillations are in the same order of magnitude as the intended SDP control frequency. This also means that the exact choice of SDP control frequency has a big impact on the probability function $p(P_k|P_{k-1})$. The pressure signal is low pass filtered to reduce the impact of these oscillations.

4.2 Results

The iterations for calculating J, as presented in Section 3, should go on until the optimality condition ε is fulfilled. Because of the difficulty of selecting this value, the iterations was instead performed until the simulated trajectories, from each of the J^is, stopped changing. In each of the three methods, this occurred within 25 iterations. This process is illustrated with Figure 7, which displays the state trajectories for the 50 first iterations with the $p(P_k|-)$ method. The results presented in the following paragraphs are from the 25th iteration in each method, unless otherwise specified. In Figures 8 to 10, dotted is $p(P_k|-)$, dashed is $p(P_k|P_{k-1})$ and solid is $p(P_k|s_k)$.

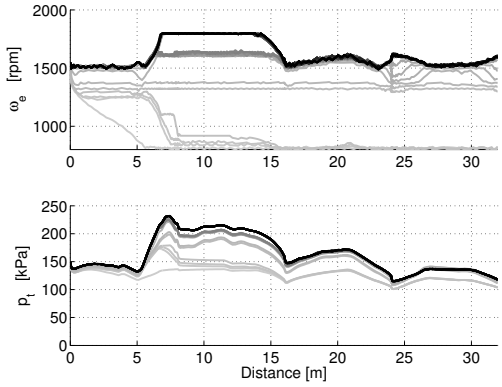

Fig. 7. State trajectories from the 50 first iterations with the $p(P_k|-)$ method. Darker is later iterations.

Figure 8 shows the states ω_e and p_t for the three methods. The maximum ω_e of the grid was $1800rpm$. The state trajectories are similar, especially in p_t, to the extent that it is difficult to distinguish the separate lines at all. The ω_e trajectory of the $p(P_k|s_k)$ method is the slight exception since this can be separated from the others, at least in parts of the cycle. It might be surprising to see that the $p(P_k|-)$ engine speed trajectory follows that of the other methods, increasing at the higher load around $10s$ into the cycle, despite only having access to a constant load probability. This change of engine speed is performed explicitly through the control signal $\frac{d\omega_e}{dt}$. The information from which the controller decides to change the engine speed is the turbo pressure p_t, which follows the injected fuel m_f, which in turn follows the load. If the turbo pressure is fixed, the controller always have the same information and will therefore always produce the same control output, causing a constant engine speed. The $p(P_k|P_{k-1})$ implementation on the other hand seem to be unable to benefit from the changing load probability distribution. The reason might be that the distribution contains is no information as to whether the trend is for increasing or decreasing power demand.

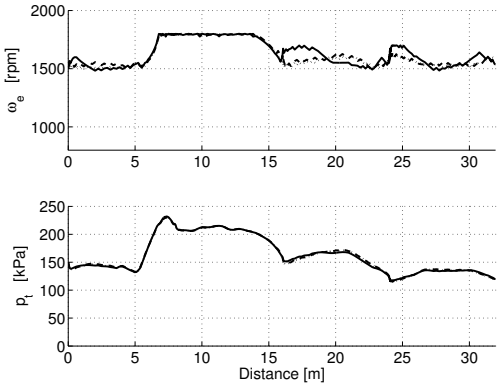

Fig. 8. Engine speed ω_e and intake air pressure p_t trajectories from the three implementations. Dotted is $p(P_k|-)$, dashed is $p(P_k|P_{k-1})$ and solid is $p(P_k|s_k)$

Figure 9 shows the control signal P_{max} for the three methods, along with the requested power P_R (gray). Just as with the states, the three signals are very similar, with that of the $p(P_k|s_k)$ method being slightly less oscillatory. In comparing to Figure 8, it can be observed that the maximum power mainly follows the turbo pressure, which indicates that the turbo pressure is the main factor affecting the choice of P_{max}.

Figure 10 shows the fuel injected per cycle m_f along with the control signal $\frac{d\omega_e}{dt}$. The injected fuel quite closely follows the actual output power $P_L = \min(P_R, P_{max})$, with the addition of inertia torque from $\frac{d\omega_e}{dt}$. This is most clearly visible just after $5s$ into the cycle, where the load P_L is at the same magnitude as at around $10s$ into the cycle, while the engine speed derivative and the amount of fuel injected is much higher. The signal $\frac{d\omega_e}{dt}$ is highly oscillatory, though slightly less in the $p(P_k|s_k)$ method. The reason for these oscillations has not yet been found.

Some of the main results from the simulations are collected in Table 1. These are total fuel used, total difference between desired and actual power, average system efficiency

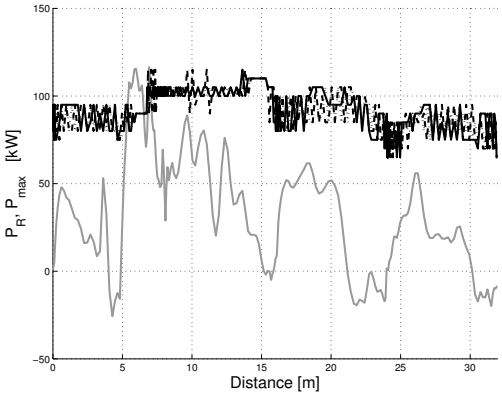

Fig. 9. Desired and maximum allowed power trajectories from the three implementations. Dotted is $p(P_k|-)$, dashed is $p(P_k|P_{k-1})$ and solid is $p(P_k|s_k)$

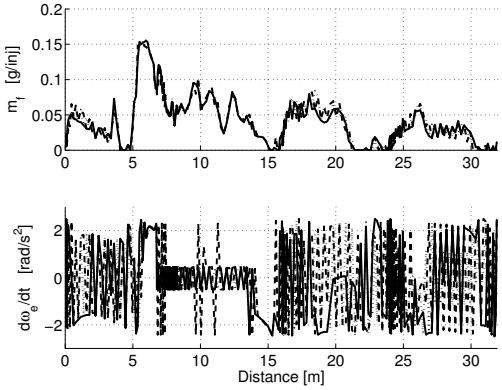

Fig. 10. Injected fuel and engine speed derivative trajectories from the three implementations. Dotted is $p(P_k|-)$, dashed is $p(P_k|P_{k-1})$ and solid is $p(P_k|s_k)$

and the magnitude of the calculation time per iteration, for each of the three methods. The solution performance values again show the close similarity between the methods. The $p(P_k|-)$ method has the best performance and the $p(P_k|P_{k-1})$ has the worst, though the differences are marginal and might very well change in either way if further iterations were to be performed, since the solutions may not have completely settled. The biggest difference between the methods are related to their relative complexity. The methods have been implemented with roughly the same resolution in P_R and exactly the same resolution in control signals and the states ω_e and p_t. The third state size however range from 1 (no third state) in the $p(P_k|-)$ method, through 15, which is the same as the P_R resolution, in the $p(P_k|P_{k-1})$ method to 250, which is related to the $25s$ duration of the cycle, in the $p(P_k|s_k)$ method. The number of states are directly proportional to the number of calculation steps, and hence the differences in the times needed for each iteration in the three methods.

Table 1. Results summary.

	$p(P_k\|-)$	$p(P_k\|P_{k-1})$	$p(P_k\|s_k)$
M_f [g]	56.4	56.8	56.9
$\int(P_R - P_L)$ $[kJ]$	21.6	22.9	21.6
mean(P_L/P_{M_f}) [%]	30.2	29.9	29.8
calculation time $[s]$	20	200	2000

5. SUMMARY AND CONCLUSIONS

This paper describes the development and evaluation of three implementations of stochastic dynamic programming (SDP) for control of the engine operating point of a diesel-electric wheel loader, operating in a short loading cycle. The driving cycles are specified by a desired electric power trajectory. The desired power trajectories should be followed unless deviating is absolutely necessary, and since deviations are assumed to be scarce, it is also assumed they do not affect the rest of the cycle. The implementations each use a stochastic description of the load, in which the probabilities are either independent of the states, depend on previous power or on distance driven. A measurement sequence with 34 loading cycles is used for the evaluation. Out of these cycles, 33 cycles are used for creating the probability distributions and the last cycle is used in the evaluations through simulations. The special characteristics of the operation, including effects from vehicle pitch dynamics, combined with the quality of some sensors, proved to have a significant effect on the calculation of the power required. It can be expected that the greatest impact would be on the implementation which use the previous power for the probability distribution.

The evaluation shows that SDP can be used for engine speed control in a diesel-electric wheel loader. It also shows that the differences between the implementations, both in the resulting trajectories and in the performance values, are extremely small. One surprising result is that the method with constant load probability is able to adjust the engine speed to the actual load. The reason is that the intake air pressure, which is related to the turbo speed, follows the injected fuel. The injected fuel in turn follows the load, and the magnitude of this pressure therefore contains information about the output power history. The differences in the sizes of the states that are added to manage the probability distributions causes proportional differences in computational time required per iteration. The combination of the computational efforts and the outcomes lead to the conclusion that the implementation with a constant load probability is superior to the implementations with probabilities depending on previous power or distance driven, for this application.

REFERENCES

Bertsekas, D. (2005). *Dynamic Programming and Optimal Control*, volume 1. Athena Scientific, 3 edition.

Filla, R. (2008). Alternative systems solutions for wheel loaders and other construction equipment. In *1st International CTI Forum Alternative and Hybrid Drive Trains*. CTI.

Hellström, E. (2010). *Look-ahead Control of Heavy Vehicles*. dissertation, Linköping University.

Johannesson, L., Åsbogård, M., and Egardt, B. (2007). Assessing the potential of predictive control for hybrid vehicle powertrains using stochastic dynamic programming. *IEEE Transactions on Intelligent Transportation Systems*, 8(1), 71–83.

Khayyam, H., Nahavandi, S., and Davis, S. (2012). Adaptive cruise control look-ahead system for energy management of vehicles. *Expert Systems with Applications*, 39(3), 3874–3885.

Kolmanovsky, I. and Filev, D. (2010). Terrain and traffic optimized vehicle speed control. In *6th IFAC Symposium on Advances in Automotive Control*. IFAC.

Lennevi, J. (1995). *Hydrostatic Transmission Control, Design Methodology for Vehicular Drivetrain Applications*. dissertation, Linköping University.

Liu, S. and Paden, B. (1997). A survey of today's cvt controls. In *Proceedings of the 36th Conference on Decision and Control*, 4738–4743. IEEE.

McDonough, K., Kolmanovsky, I., Filev, D., Yanakiev, D., Szwabowski, S., and Michelini, J. (2012). Stochastic dynamic programming control policies for fuel efficient in-traffic driving. In *American Control Conference (ACC), 2012*, 3986–3991. IEEE.

Nezhadali, V. and Eriksson, L. (2013). Modeling and optimal control of a wheel loader in the lift-transport section of the short loading cycle. In *7th IFAC Symposium on Advances in Automotive Control*. IFAC.

Nilsson, T., Fröberg, A., and Åslund, J. (2012a). Fuel potential and prediction sensitivity of a power-split cvt in a wheel loader. In *IFAC Workshop on Engine and Powertrain Control, Simulation and Modeling*, 49–56. IFAC.

Nilsson, T., Fröberg, A., and Åslund, J. (2012b). Optimal operation of a turbocharged diesel engine during transients. *SAE International Journal of Engines*, 5(2), 571–578.

Nilsson, T., Sundström, C., Nyberg, P., Frisk, E., and Krysander, M. (2014). Robust driving pattern detection and identification with a wheel loader application. *International Journal of Vehicle Systems Modelling and Testing*, 9(1), 56–76.

Puterman, M. (2005). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons Inc, Hoboken, NJ, USA, 2 edition.

Rizzoni, G., Guzzella, L., and Baumann, B. (1999). Unified modeling of hybrid electric vehicle drivetrains. *IEEE/ASME Transactions on Mechatronics*, 4, 246–257.

Rydberg, K.E. (1998). Hydrostatic drives in heavy mobile machinery - new concepts and development trends. In *International Off-Highway & Powerplant Congress & Exposition*, 981989. SAE.

Stein, G., Fröberg, A., Martinsson, J., Brattberg, B., Filla, R., and Unnebäck, J. (2013). Fuel efficiency in construction equipment - optimize the machine as one system. In *7th AVL International Commercial Powertrain Conference*. AVL & SAE.

Wang, F., Zhang, J., Sun, R., and Yu, F. (2012). Analysis on the performance of wheel loaders in typical work cycle. *Applied Mechanics and Materials*, 148, 526–529.

Zhang, R., Alleyne, A., and Prasetiawan, E. (2002). Modeling and h_2/h_∞ mimo control of an earthmoving vehicle powertrain. *Journal of dynamic systems, measurement, and control*, 124(4), 625–636.