

Value Function Approximation for Direct Control of Switched Power Converters

Nicholas Moehle

Stephen Boyd

Abstract—We consider the problem of controlling switched-mode power converters using model predictive control. Model predictive control requires solving optimization problems in real time, limiting its application to systems with small numbers of switches and a short horizon. We propose a technique for using off-line computation to approximate the model predictive controller. This is done by dividing the planning horizon into two segments, and using a quadratic function to approximate the optimal cost over the second segment. The approximate model predictive algorithm minimizes the true cost over the first segment, and the approximate cost over the second segment, allowing the user to adjust the computational requirements by changing the length of the first segment. We conclude with two simulated examples.

I. INTRODUCTION

Model predictive control (MPC) is emerging as a powerful tool for controlling switched-mode power converters. A model predictive control scheme uses a planning model to predict the effects of control decisions on the converter over some finite planning horizon. It then compares the outcomes of these decisions using a cost criterion, choosing an optimal control input sequence, *i.e.*, one that minimizes the cost over the horizon. The controller then uses the first control input in the trajectory, waits a specified time, then repeats the process. Model predictive control is now widely used in many control applications, particularly those that are sampled slowly enough that the required minimization problem can be solved in the sample time.

Finite-control-set MPC (FCS-MPC) refers to the case in which the control input takes on a finite number of values. This is particularly relevant for switched-mode power converters, where the control inputs typically represent allowable switch configurations. In this case, solving the optimization problem only requires evaluating the objective function a finite number of times and choosing the control inputs that achieve the lowest cost. When the total number of control inputs over the planning horizon is small, this can be done quickly on modest computational platforms, which is attractive for embedded control.

Unfortunately, the computational cost of solving the FCS-MPC problem increases exponentially with the length of the planning horizon, which means that

checking all combinations of control input trajectories becomes very difficult for discrete horizon lengths of more than, say, four or five. In practice, this means FCS-MPC is only useful for converters with simple dynamics, for which short planning horizons (and myopic controllers) suffice, or converters with powerful processors available for control. Even worse, many converter topologies have dynamics that inherently require a long planning horizon for FCS-MPC to work well, such as those with inverse response (*i.e.*, non-minimum phase behavior). For example, to increase the steady-state output voltage of an ideal boost converter, it may be necessary to charge the inductor, which can have the short-term effect of *decreasing* the output voltage. A long planning horizon may be necessary for the controller to predict the benefits of charging the inductor before increasing the output voltage. We illustrate this behavior for a boost converter in §VI. (For further discussion on non-minimum phase behavior in power converters, see [EM07, Ch. 9] and [KGM14].)

In this paper, we provide a technique for *approximately* solving the optimization problems arising in FCS-MPC with a long planning horizon, with very low computational cost. This allows the resulting approximate FCS-MPC controller to make near-optimal decisions, with long prediction horizons, at only a tiny fraction of the computational cost of traditional long-horizon FCS-MPC.

To do this, we split the planning horizon into two segments. The dynamics and cost function over the first segment are represented exactly, and the problem is solved by exhaustive search, as in traditional FCS-MPC. The optimal cost over the second (typically much longer) segment is approximated using a multivariate quadratic function, called the *approximate value function*. This function is determined during a pre-solve phase (*i.e.*, during controller synthesis), using a desktop computer or server. The controller itself simply compares a relatively small number of possible input sequences with a short horizon, and chooses the best one, taking into account the approximate value function. Our method therefore shifts computational burden from real-time execution to an off-line, precomputation phase, which allows us to execute FCS-MPC-like controllers with very long horizon lengths on very modest

computational hardware.

The resulting control algorithm, which we call *value function approximation model predictive control*, which we shorten to *approximate model predictive control* (A-MPC) in this paper, is very similar to short-horizon FCS-MPC, differing only by the addition of an extra cost term (given by the approximate value function), which represents the approximate long-term cost. This resemblance is a practical and pedagogical advantage, as our method can be explained and implemented as a simple modification of short-horizon FCS-MPC that imparts better dynamics on the closed-loop system, since it has the advantage of a long planning horizon.

The price of doing this is decreased flexibility: because deriving an approximate value function often requires extensive offline precomputation, our formulation does not handle gross changes in the objective function or dynamics, although some moderate changes, such as changing the desired state value, can be easily accommodated. However, we suspect that such large-scale changes in the operating requirements are rare, and in any case, can be handled by synthesizing several controllers, and switching between them as desired. (This is similar to gain scheduling linear controllers.)

A. Related work

a) FCS-MPC for power converters: A good introduction to FCS-MPC is given by Kuoro [KCV⁺09] and Cortés [CKK⁺08], which include many references to practical applications as well as tricks that can be used with FCS-MPC (*e.g.*, shaping the switching spectrum or controlling reactive power). These papers primarily cover short-horizon FCS-MPC, which is easily solved by brute-force enumeration of all possible switching trajectories.

For long-horizon FCS-MPC, enumerating all switching trajectories is not computationally feasible, and more sophisticated techniques are used to solve (or approximately solve) the optimization problem. For example, branch and bound uses lower bounds on the optimal cost to identify and remove suboptimal parts of the decision tree (see [BM07]). Simple lower bounds for the FCS-MPC problem are proposed by Geyer [Gey11], by simply assuming zero future cost starting from any node in the decision tree. Another general technique is move blocking (*i.e.*, artificially limiting the number of control input trajectories to obtain a suboptimal but practical solution); an application to boost converter control can be found in [KGM12]. Another recent approach is to use a sphere decoding algorithm to solve the FCS-MPC optimal control problem [GQ13]. Additionally, techniques from hybrid system theory can be used to explicitly write the solution of the FCS-MPC problem as a lookup table with regions defined

by linear or quadratic inequalities. (see [BBM] for details). Unfortunately, the same exponential blowup that occurs in long-horizon FCS-MPC can appear in explicit solutions of hybrid MPC, *i.e.*, it is possible that there are an exponential number of regions in the lookup table. There have been many applications of this approach to power converters; see [QAG14] and [GPM05], and references therein.

b) Other control techniques: Most traditional techniques for controlling power converters are based on linear control. Because power converters are switched nonlinear systems, it is first necessary to linearize the system around an operating point, and use PWM, or a similar modulation scheme, to convert the continuous control decisions of the linear controller to discrete switching decisions. For an introduction to these techniques, see [KSV91] and [EM07]. For a discussion of continuous MPC techniques (and others) applied to controller power converters, see [WCY⁺14]. Deadbeat control policies attempt to drive the system state to a desired value as quickly as possible, often using a short prediction horizon. A good summary can be found in [QAG14]; applications to power converters can be found in the references of [KCV⁺09].

Many control methods are based on explicitly defining the switching boundaries in state space, usually defined using first- or second-order equations; this idea originated with Burns and Wilson [IW76], [IW77]. For some simple converters, these boundaries can be derived as solutions to certain optimal control problems using Pontryagin's maximum principle. For example, the solution to the minimum-time state transfer problem for a buck converter has a switching boundaries defined by second-order equations. However, this approach becomes unwieldy for even slightly more complicated systems (see [DKDG⁺14] for a discussion).

c) Dynamic programming: Recovering optimal control policies through value function approximation, which is the basis of the approach in this paper, is the subject of approximate dynamic programming, and has been studied extensively in operations research, control, robotics, and other fields (see [Ber05] and [Pow07] for an introduction). The general method of exactly solving an optimal control problem over some time steps, and using an approximation of the value function for the rest of the horizon, is well known in the dynamic programming literature, with names such as roll-out policy or look-ahead policy. Adaptive versions of these methods, which change the approximate value function during operation of the system, are possible; algorithms of these types are the subject of reinforcement learning; see [SB98] and [Ber05].

Our algorithm relies on quadratic approximate value functions, obtained by simple function fitting tech-

niques; other methods for obtaining quadratic approximate value functions for switched-affine systems can be found in [Ran06] and [WOB15]. Fast evaluation of control policies involving quadratic value functions (as well as much more general problems) can be found in [WB11], [MWB11], and [WB10].

d) Mixed-integer nonlinear programming: The optimal control problem solved in an FCS-MPC control policy is a special case of mixed-integer nonlinear programming (MINLP). MINLP problems are difficult to solve in general, but algorithms exist that appear to work well in practice. Several special cases, such as mixed integer linear programming [Wol98], and mixed integer second-order cone programming [BS13], have dedicated solvers (see Gurobi [Inc15], Mosek [AA00], and CPLEX), which often outperform general MINLP solvers. In many cases, small problems (such as the FCS-MPC problem for power converters) can be solved on a desktop in a few seconds. Furthermore, these solvers can often be interfaced through modeling languages, such as CVXPY [DCB14], and CVX [GB14], [GB08], which increase accessibility and decrease development time.

B. Contributions

In this paper we present a tractable approximate method for solving long-horizon predictive control (FCS-MPC) problems for switched-mode power converters. We provide a theoretical justification for our method, as well as practical approach for computing this added extra cost term. We give simulated results for controlling a boost converter control and an inverter, and we show that our method performs markedly better short-horizon predictive control and performs comparably to long-horizon predictive control, at a tiny fraction of the computational cost.

We contrast our method with other methods to achieve long-horizon FCS-MPC. Many methods attempt to solve the optimal control problem on the fly, which often necessitates a fast control processor. Our method, by contrast, can be viewed as a simple modification of short-horizon (or even one-step) FCS-MPC. This allows our method to work with very low computational cost, so that our method can be executed in microseconds or less on modest computational platforms, while still achieving some of the benefits of long-horizon MPC methods.

C. Outline

In §II, we present the theoretical basis for our method, starting with the mathematical optimal control problem to be solved by the controller, and our general approach for approximately solving it. In §III we show

how to numerically derive the controller by approximating the value cost as a quadratic function. In §VI, we demonstrate our approach on a boost converter example. In §VII, we conclude with some ideas for extending our method, and alternative techniques for computing the extra cost term.

II. MODEL PREDICTIVE CONTROL

e) Switched dynamical systems: We consider the switched dynamical system

$$x_{t+1} = f(x_t, u_t), \quad t = 0, 1, \dots, \quad (1)$$

with state $x_t \in \mathbf{R}^n$ and (finite) control input $u_t \in \{1, \dots, K\}$. The subscript t denotes (sampled) time.

In a converter, the state typically consists of physical variables such as inductor currents and capacitor voltages, but it can also include other, non-physical variables such as a running sum of errors. The input corresponds to the allowed switch configurations, and K is the number of allowed switch positions. The dynamics function f gives the state at the next sample time, given the current state and choice of switch configuration.

f) Control: Our task is to choose a *control policy* $\phi : \mathbf{R}^n \rightarrow \{1, \dots, K\}$ that maps the system state x_t at time t to a control input u_t , so that

$$u_t = \phi(x_t).$$

We design the policy according to some objectives, such as maintaining the state near some desired value, or minimizing the converter power loss. The policy divides the state space into K (possibly disconnected) regions, so that if the state is in the i th region, the controller chooses input i . The boundaries of these surfaces are often known as switching surfaces or switching boundaries; when the state of the system crosses a switching surface, the input (*i.e.*, switch configuration) changes. Implicit here is that we can measure the state; when this is not the case, the state can be estimated using, for example, a Kalman filter.

g) Model predictive control: Under a model predictive control policy ϕ_{mpc} , the control input is chosen by solving an optimal control problem: to find $\phi_{\text{mpc}}(z)$, we first solve the finite-horizon optimal control problem

$$\begin{aligned} & \text{minimize} && \sum_{t=0}^{T-1} g(\tilde{x}_t) + h(\tilde{x}_T) \\ & \text{subject to} && \tilde{x}_{t+1} = f(\tilde{x}_t, \tilde{u}_t) \\ & && \tilde{x}_0 = z. \end{aligned} \quad (2)$$

The decision variables are the predicted system states $\tilde{x}_0, \dots, \tilde{x}_T$ and the planned control inputs $\tilde{u}_0, \dots, \tilde{u}_{T-1}$. The integer T is the horizon; it determines how far ahead in time we plan when we choose the current input. The function $g : \mathbf{R}^n \rightarrow \mathbf{R}$ is the stage cost function, and the function $h : \mathbf{R}^n \rightarrow \mathbf{R}$ is

the terminal cost function. Then $\phi_{\text{mpc}}(z)$ is chosen to be \tilde{u}_0 , where \tilde{u} is any solution of the optimal control problem (2).

To specify an MPC policy, the designer must choose the horizon T , the stage cost function g , and the terminal cost h . The main difficulty of implementing the MPC policy ϕ_{mpc} is quickly solving (2), which is (in general) a mixed-integer nonlinear program (MINLP). In principle the problem can be solved by checking the objective value for each possible control input trajectory, *i.e.*, for each of the K^T possible trajectories, and simply simulating the dynamics for each choice of input sequence. However, for large values of K and T (larger than, say $K = 3$ and $T = 5$), this approach is not practical. In this case, (2) is very difficult to solve exactly on an embedded processor at sub-millisecond speeds, required for real-time control of switched power converters. We note, however, that many effective algorithms exist to solve MINLP problems reasonably quickly (if not at sub-millisecond time scale) on desktop computers or servers, which we will return to later.

A large value of the planning horizon T allows a receding-horizon controller to consider the impact of current actions far in the future, and almost always results in a superior control law. Using a short planning horizon results in a myopic controller, which can seriously degrade performance. Unfortunately, our desire for long horizon directly contradicts our desire to easily solve (2) using low-cost, embedded processors, because a large value of T makes (2) much more difficult to solve numerically. This effectively limits receding-horizon controllers based on (2) to converters with simple dynamics (for which myopic controllers perform well) or converters with powerful processors available for control. In the following sections we discuss an alternative approach to approximately solve (2), with a long planning horizon T , in a manner suitable for real-time, embedded control. We will do this by moving most of the computational burden offline (*i.e.*, to controller synthesis, not execution), using value function approximation.

A. Value functions

We will use a standard trick in dynamic programming to express the problem (2) as a similar one, with a shorter time horizon $\tau < T$. Define $V_\tau(z)$ as the optimal value of the problem

$$\begin{aligned} & \text{minimize} && \sum_{t=\tau}^{T-1} g(\tilde{x}_t) + h(\tilde{x}_T) \\ & \text{subject to} && \tilde{x}_{t+1} = f(\tilde{x}_t, \tilde{u}_t) \\ & && \tilde{x}_\tau = z, \end{aligned} \quad (3)$$

with variables u_τ, \dots, u_{T-1} , $\tilde{x}_\tau, \dots, \tilde{x}_T$. Thus $V_\tau(z)$ is the optimal value of the control problem (2), starting from state z at time τ . It is also called the (Bellman) value function.

We can use the value function to rewrite (2) as

$$\begin{aligned} & \text{minimize} && \sum_{t=0}^{\tau-1} g(\tilde{x}_t) + V_\tau(\tilde{x}_\tau) \\ & \text{subject to} && \tilde{x}_{t+1} = f(\tilde{x}_t, \tilde{u}_t) \\ & && \tilde{x}_0 = z, \end{aligned} \quad (4)$$

with variables $\tilde{x}_0, \dots, \tilde{x}_\tau$ and $\tilde{u}_0, \dots, \tilde{u}_{\tau-1}$.

Note that the new problem (4) has the same form as the MPC problem (2), but with shorter horizon length τ and terminal cost V_τ . We can solve it by evaluating the cost of K^τ input sequences, a far smaller number than the original problem, which requires evaluating K^T input sequences. In the general case this advantage is an illusion, since the cost of evaluating $V_\tau(\tilde{x}_\tau)$ is $K^{T-\tau}$ evaluations of value input sequences. But if the value function V_τ were known explicitly, with a simple form, solving (4) would be easier than solving (2). For example, if n is very small, say, 2 or 3, we could evaluate V_τ on a fine grid of points offline and store the values in a table.

If we take $\tau = 1$, we have

$$\phi_{\text{mpc}}(z) = \underset{u}{\operatorname{argmin}} V_1(f(z, u)).$$

In this case, the control policy can be carried out using just K evaluations of V_1 and f . If V_1 were easy to evaluate (which it is not in general), this would give a very simple way to evaluate the MPC policy.

B. Approximate value functions

Unfortunately, computing, storing, and evaluating V_τ is difficult in all but a few special cases. We therefore settle for an approximate or surrogate value function \hat{V}_τ which approximates V_τ . We consider the problem

$$\begin{aligned} & \text{minimize} && \sum_{t=0}^{\tau-1} g(\tilde{x}_t) + \hat{V}_\tau(\tilde{x}_\tau) \\ & \text{subject to} && \tilde{x}_{t+1} = f(\tilde{x}_t, \tilde{u}_t) \\ & && \tilde{x}_0 = z, \end{aligned} \quad (5)$$

which is an approximation to (4) with the same decision variables. This optimization problem is the basis for the approximate MPC (A-MPC) policy, ϕ_{ampc} . We define $\phi_{\text{ampc}}(z)$ as an optimal value of \tilde{u}_0 for (5). If we have $\hat{V}_\tau = V_\tau$ (*i.e.*, \hat{V}_τ is a perfect approximation), we recover the MPC policy, regardless of the shortened horizon length τ . Alternatively, if we take $\tau = T$ and $\hat{V}_T(z) = h(z)$, we recover the MPC policy.

III. OBTAINING AN APPROXIMATE VALUE FUNCTION

In order to use an approximate MPC policy, we first need to synthesize an approximate value function \hat{V}_τ . In practice, good control can be obtained with very simple forms for \hat{V}_τ , even when \hat{V}_τ is not a particularly good estimate of V_τ . (Indeed, the success of short-horizon predictive control suggests that even a greedy policy, obtained using $\hat{V}_\tau = g$, works well in some cases.)

h) *Quadratic functions*: We limit our attention to quadratic value function approximations, *i.e.*, functions of the form

$$\hat{V}_\tau(x) = (x - x_{\text{des}})^T P (x - x_{\text{des}}) + r, \quad (6)$$

where $P \in \mathbf{R}^{n \times n}$ is a symmetric positive definite matrix, $r \in \mathbf{R}$ is a scalar, and x_{des} is a desired system state, representing an area of the state space with low cost. The scalar r has no effect on the A-MPC control policy, but is included since we will choose it and P so that $\hat{V}_\tau \approx V_\tau$.

For example, one simple quadratic approximate value function is given by

$$\hat{V}_\tau(x) = (x - x_{\text{des}})^T P_{\text{energy}} (x - x_{\text{des}}), \quad (7)$$

where the symmetric matrix P_{energy} is such that $x^T P_{\text{energy}} x$ gives the total energy stored in the converter. Control policies based on the energy function appear to have good stability properties; for details, see [SV92] and [MK92].

i) *Function fitting*: We first choose state vectors $x^{(1)}, \dots, x^{(N)}$, which are points at which we desire an accurate approximation. Then for each state vector $x^{(i)}$, we compute the value function $V_\tau(x^{(i)})$ by solving problem (2) with $z = x^{(i)}$. This is done offline, possibly using substantial computational resources. We then choose the parameters P and r of the approximate value function as the solution to the least-squares approximation problem

$$\begin{aligned} \text{minimize} \quad & \frac{1}{N} \sum_{i=1}^N \left(V_\tau(x^{(i)}) - \hat{V}_\tau(x^{(i)}) \right)^2 \\ & + \lambda (P - \alpha P_{\text{energy}})^2. \end{aligned} \quad (8)$$

The decision variables are the approximate value function parameters P and r , as well as the scaling factor α . Note that P and r enter linearly in \hat{V}_τ , which allows us to solve this problem as a least-squares problem. The parameter $\lambda > 0$ is a regularization parameter; if we choose λ to be small, the optimal P and r almost minimize mean-squared error between the approximate value function and the sampled values, thus returning a good value function approximation at the points. By choosing λ large, we force P to be close to (a scaled version of) P_{energy} ; this tends to make the resulting control more robust, especially if we have few samples.

With values of P and r that solve (8), we expect that $\hat{V}_\tau(z) \approx V_\tau(z)$ in areas of the state space that are well represented by the state samples. Because the flexibility (and therefore, accuracy) of the approximate value function is a limited resource, it is important to choose state samples that are likely to be encountered in normal operation, and not to include samples far from

relevant areas of the state space, since including them may decrease accuracy in relevant areas.

We can explicitly add the constraints that P be positive definite, and $\alpha \geq 0$; the approximation problem above is no longer a least-squares problem but it is still convex and readily solved [BV04]. But when the parameter λ and the sample points are appropriately chosen, these constraints are typically satisfied automatically. We also note that since we are approximating the value function, we do not need to evaluate $V_\tau(x^{(i)})$ exactly; good approximations are good enough.

IV. SUMMARY OF APPROXIMATE MPC METHOD

Here we summarize the steps needed to use the approximate MPC policy.

First, the parameters of the underlying MPC policy are determined. This means that a dynamics model (1) of the system is derived and the cost function and horizon length are determined. The dynamics model should be simple enough, and the horizon length modest enough, for problem (2) to be solved fairly quickly on a desktop computer, so that many (hundreds or thousands) of samples can be generated reasonably quickly.

Second, the reduced horizon length τ is chosen, based on the computational capabilities of the embedded control platform (*i.e.*, a large value of τ results in a more accurate approximation of the MPC problem, but also in a higher computational burden, as the controller will need to evaluate K^τ different trajectories each time step.) Sample states likely to be encountered in normal operation are then chosen, and the value function $V_\tau(x^{(i)})$ is evaluated for each sample $x^{(i)}$ by solving (3). Once the value function has been evaluated for all sample points, we choose the regularization tradeoff parameter λ and solve the approximation problem (8) to obtain the approximate value function parameters P and r .

Once the approximate value function is obtained, the control policy can be evaluated by solving the short-horizon policy (5), implementing the optimal first control input u_0 , waiting until the next sampling time, and repeating.

V. EXTENSIONS

A. Switching cost

In this section, we describe how to include switching costs, *i.e.*, a cost associated with switching from one control input to another, so that (2) becomes

$$\begin{aligned} \text{minimize} \quad & \sum_{t=0}^{T-1} (g(\tilde{x}_t) + \ell(u_{t-1}, u_t)) + h(\tilde{x}_T) \\ \text{subject to} \quad & \tilde{x}_{t+1} = f(\tilde{x}_t, \tilde{u}_t) \\ & \tilde{x}_0 = z. \end{aligned} \quad (9)$$




Fig. 1: Boost converter model.

where ℓ is a switching cost function, and the parameter u_{t-1} represents the control input most recently implemented. In theory, this problem can be put into the form of problem (2) by encoding the switch position as a state variable.

However, a more convenient way to handle the switching cost is to include it in the first τ time steps and simply ignore it during the last $T - \tau$ time steps. By doing this, the process of deriving the value function (see §III) remains the same, but the control policy is modified to reflect the switching costs:

$$\begin{aligned} & \text{minimize} && \sum_{t=0}^{\tau-1} (g(\tilde{x}_t) + \ell(u_{t-1}, u_t)) + \hat{V}_\tau(\tilde{x}_\tau) \\ & \text{subject to} && \tilde{x}_{t+1} = f(\tilde{x}_t, \tilde{u}_t) \\ & && \tilde{x}_0 = z. \end{aligned} \quad (10)$$

Trivial modifications to this setup allow us to include many of the standard tricks for FCS-MPC can be applied directly, such as using a frequency-based switching cost to enforce a constant switching frequency. (See [KCV⁺09], [CKK⁺08], for details.) We expect this approach to work well for converters in which the switching dynamics do not have a significant long-term effect, *e.g.*, converters with a high switching frequency.

VI. BOOST CONVERTER EXAMPLE

We consider the ideal boost converter shown in figure 1. The system state is

$$x_t = \begin{bmatrix} i_{L,t} \\ v_{C,t} \end{bmatrix}.$$

where $i_{L,t}$ and $v_{C,t}$ are the inductor current and capacitor voltage at discrete epoch t . We (approximately) discretize the converter dynamics to have the form (1).

We use the cost function to penalize deviation of the capacitor voltage from a desired value:

$$g(x_t) = |v_{C,t} - v_{\text{des}}|.$$

With this cost function, and with the discretized dynamics, problem (2) can be written as a mixed-integer linear program (MILP).

We used the parameters of [KGM13], so that $V_{\text{dc}} = 10$ V, $L = 450$ μH , $R_L = 0.3$ Ω , $C = 220$ μF , and $R_{\text{load}} = 73$ Ω . We took $v_{\text{des}} = 30$ V (instead of $v_{\text{des}} = 15$ V, as in [KGM13]), because this increased




Fig. 2: For the boost converter example, the 100 sampled points of the value function are shown on the left, and the fitted quadratic approximate value function \hat{V}_τ is shown on the right, as a function of the inductor current i and the capacitor voltage v .

target voltage means that the controller must charge up the inductor at the cost of temporarily increasing voltage error. (For $v_{\text{des}} = 15$ V, the controller can simply turn off the switch, charging the inductor and capacitor simultaneously until the output voltage reaches the desired value.)

Following §III-0i, we randomly generated 100 initial states by sampling uniformly over a box with a lower bound of (0 A, 0 V) and an upper bound of (10 A, 50 V), which are reasonable limits for normal operation of the converter. We then solved problem (2) for each of the initial states, to within one percent accuracy. Using these sample points, we solved the regularized approximation problem (8) to obtain an approximate value function, with parameter $\lambda = 100$, using a positive semidefinite constraint on P . The resulting value function is shown in figure 2 along with the 100 randomly sampled values. To evaluate the ADP policy, we must explicitly define a desired state vector x_{des} , which we take to be $(v_{\text{des}}/R_{\text{load}}, v_{\text{des}})$.

In figure 3 we show the closed-loop response of the system under FCS-MPC controllers with horizon lengths of 1, 5, and 30, with initial condition $x_0 = (0, 0)$ (*i.e.*, zero inductor current and capacitor voltage). To evaluate the FCS-MPC control policy, (2) was solved to one percent accuracy. With a horizon length of $T = 1$, the controller does not reach the desired steady-state value of 30 V and for $T = 5$, the desired steady-state value is reached, but only slowly. For $T = 30$, the desired steady-state value is reached very quickly; however, using such a long planning horizon is not practical. (In our desktop implementation, each MPC evaluation took around 10 minutes on average.)

In figure 4, we show the response using the A-MPC controller, using $\tau = 1$ and using the approximate value function derived from function fitting. We can see that the A-MPC controller outperforms the MPC controllers with short horizons ($T = 1$ and $T = 5$), and is on par with the long-horizon MPC policy ($T = 30$).




Fig. 3: Closed-loop trajectories for the boost converter example with the MPC policy, for different values of the planning horizon T . The desired output voltage value is 30 V.




Fig. 4: Closed-loop trajectories for the boost converter example, using the A-MPC policy. The desired output voltage value is 30 V.

VII. CONCLUSION

We propose a method for approximately solving long-horizon FCS-MPC problems that avoids extensive computation on the control processor. Our method relies on estimating a quadratic approximate value function, which we do by sampling.

We conclude by noting some other ways to obtain a quadratic approximate value function \hat{V} . For example, if we minimize a quadratic cost function over an approximate, linearized system, a quadratic value function can be obtained by solving a Riccati equation; this method can be used instead of the function fitting method described in §III-0i. Another possible way involves solving linear matrix inequalities corresponding to a relaxed Bellman equation; for details, see [Ran06] or [WOB15].

REFERENCES

[AA00] E. D. Andersen and K. D. Andersen. The Mosek interior point optimizer for linear programming: An im-

plementation of the homogeneous algorithm. In *High Performance Optimization*, pages 197–232. Springer, 2000.

- [BBM] F. Borrelli, A. Bemporad, and M. Morari. *Predictive Control for Linear and Hybrid Systems*. Preprint available at http://www.mpc.berkeley.edu/mpc-course-material/MPC_Book.pdf.
- [Ber05] D. P. Bertsekas. *Dynamic programming and optimal control*, volume 1. Athena Scientific Belmont, MA, 3 edition, 2005.
- [BM07] S. Boyd and J. Mattingley. Branch and bound methods. http://stanford.edu/class/ee364b/lectures/bb_notes.pdf, 2007. Course notes for EE364b, Stanford University.
- [BS13] H. Y. Benson and U. Saglam. Mixed-integer second-order cone programming: A survey. *Tutorials in Operations Research*, pages 13–36, 2013.
- [BV04] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [CKK⁺08] P. Cortés, M. P. Kazmierkowski, R. M. Kennel, D. E. Quevedo, and J. Rodríguez. Predictive control in power electronics and drives. *Transactions on Industrial Electronics*, 55(12):4312–4324, 2008.
- [DCB14] S. Diamond, E. Chu, and S. Boyd. CVXPY: A Python-embedded modeling language for convex optimization, version 0.2. <http://cvxpy.org/>, May 2014.
- [DKDG⁺14] S. V. Dhople, K. A. Kim, A. D. Domínguez-García, et al. Time-optimal control in DC-DC converters: A maximum principle perspective. In *Proceedings of the IEEE Applied Power Electronics Conference and Exposition*, 2014.
- [EM07] R. W. Erickson and D. Maksimovic. *Fundamentals of power electronics*. Springer Science & Business Media, 2007.
- [GB08] M. Grant and S. Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008.
- [GB14] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, March 2014.
- [Gey11] T. Geyer. Computationally efficient model predictive direct torque control. *IEEE Transactions on Power Electronics*, 26(10):2804–2816, 2011.
- [GPM05] T. Geyer, G. Papafotiou, and M. Morari. Model predictive control in power electronics: A hybrid systems approach. In *Conference on Decision and Control, European Control Conference*, pages 5606–5611. IEEE, 2005.
- [GQ13] T. Geyer and D. E. Quevedo. Multistep direct model predictive control for power electronics—part 1: Algorithm. In *Energy Conversion Congress and Exposition*, pages 1154–1161. IEEE, 2013.
- [Inc15] Gurobi Optimization Inc. Gurobi optimizer reference manual, 2015.
- [IW76] W. W. Burns III and T. G. Wilson. State trajectories used to observe and control dc-to-dc converters. *IEEE Transactions on Aerospace and Electronic Systems*, 12(6):706–717, 1976.

- [IW77] W. W. Burns III and T. G. Wilson. Analytic derivation and evaluation of a state-trajectory control law for DC-to-DC converters. In *IEEE Power Electronics Specialists Conference*, pages 70–85. IEEE, 1977.
- [KCV⁺09] S. Kouro, P. Cortés, R. Vargas, U. Ammann, and J. Rodríguez. Model predictive control—a simple and powerful method to control power converters. *Transactions on Industrial Electronics*, 56(6):1826–1838, 2009.
- [KGM12] P. Karamanakos, T. Geyer, and S. Manias. Direct model predictive current control of dc-dc boost converters. In *15th International Power Electronics and Motion Control Conference*, 2012.
- [KGM13] P. Karamanakos, T. Geyer, and S. Manias. Direct model predictive current control strategy of DC–DC boost converters. *Journal of Emerging and Selected Topics in Power Electronics*, 1(4):337–346, 2013.
- [KGM14] P. Karamanakos, T. Geyer, and S. Manias. Direct voltage control of dc–dc boost converters using enumeration-based model predictive control. *Power Electronics, IEEE Transactions on*, 29(2):968–978, 2014.
- [KSV91] J. G. Kassakian, M. F. Schlecht, and G. C. Verghese. *Principles of power electronics*. Addison-Wesley Reading, USA, 1991.
- [MK92] P. Midya and P. T. Krein. Optimal control approaches to switching power converters. In *Power Electronics Specialists Conference*, pages 741–748. IEEE, 1992.
- [MWB11] J. Mattingley, Y. Wang, , and S. Boyd. Receding horizon control: Automatic generation of high-speed solvers. *IEEE Control Systems Magazine*, 31(3):5265, June 2011.
- [Pow07] W. B. Powell. *Approximate Dynamic Programming: Solving the curses of dimensionality*, volume 703. John Wiley & Sons, 2007.
- [QAG14] D. E. Quevedo, R. P. Aguilera, and T. Geyer. Predictive control in power electronics and drives: Basic concepts, theory, and methods. In *Advanced and Intelligent Control in Power Electronics and Drives*, pages 181–226. Springer, 2014.
- [Ran06] A. Rantzer. Relaxed dynamic programming in switching systems. *IEE Proceedings, Control Theory and Applications*, 153(5):567–574, 2006.
- [SB98] R. S. Sutton and A. G. Barto. *Introduction to reinforcement learning*, volume 135. MIT Press Cambridge, 1998.
- [SV92] S. R. Sanders and G. C. Verghese. Lyapunov-based control for switched power converters. *IEEE Transactions on Power Electronics*, 7(1):17–24, 1992.
- [WB10] Y. Wang and S. Boyd. Fast model predictive control using online optimization. *IEEE Transactions on Control Systems Technology*, 18(2):267–278, March 2010.
- [WB11] Y. Wang and S. Boyd. Fast evaluation of quadratic control-Lyapunov policy. *IEEE Transactions on Control Systems Technology*, 19(4):939946, July 2011.
- [WCY⁺14] L. Wang, S. Chai, D. Yoo, L. Gan, and K. Ng. *PID and Predictive Control of Electrical Drives and Power Converters Using MATLAB/Simulink*. John Wiley & Sons, 2014.
- [WOB15] Y. Wang, B. O’Donoghue, and S. Boyd. Approximate dynamic programming via iterated Bellman inequalities. *International Journal of Robust and Nonlinear Control*, 25(10):1472–1496, 2015.
- [Wol98] L. A. Wolsey. *Integer programming*, volume 42. Wiley New York, 1998.