Model Predictive Direct Torque Control – Part I: Concept, Algorithm and Analysis

Tobias Geyer, *Member, IEEE*, Georgios Papafotiou, *Member, IEEE*, and Manfred Morari, *Fellow, IEEE*

***Abstract*—This paper focuses on Direct Torque Control (DTC) for three-phase AC electric drives. A novel model predictive control scheme is proposed that keeps the motor torque, the stator flux and (if present) the inverter's neutral point potential within given hysteresis bounds while minimizing the switching frequency of the inverter. Based on an internal model of the drive, the controller predicts *several* future switch transitions, extrapolates the output trajectories and chooses the sequence of inverter switch positions (voltage vectors) that minimizes the switching frequency. The advantages of the proposed controller are twofold. Firstly, as underlined by the experimental results in the second part of this paper, it yields a superior performance with respect to the industrial state of the art. Specifically, the switching frequency is reduced by up to 50 % while the torque and flux are kept more accurately within their bounds. Moreover, the fast dynamic torque response is inherited from standard DTC. Secondly, the scheme is applicable to a large class of (three-phase) AC electric machines driven by inverters.**

***Index Terms*—AC motor drives, model predictive control, direct torque control, power electronics**

I. Introduction

In adjustable speed AC drives DC-AC inverters are used to drive three-phase AC machines as variable frequency voltage or current sources. One of the various methods used for controlling the machine's torque and speed is Direct Torque Control (DTC) [1], [2], [3], [4]. Exploiting the motor's fast stator flux dynamics, DTC directly manipulates the stator flux vector such that the desired torque (and magnitude of the stator flux) is achieved by choosing an inverter switch combination that provides the appropriate phase voltages to the motor windings. In state of the art drives, this choice is made e.g. every $T_s = 25\,\mu$s using a pre-designed switching table that – depending on the particularities of the application – addresses a number of different control objectives. These primarily concern the motor. More specifically, the electromagnetic torque and the stator flux need to be kept within pre-specified hysteresis bounds. In high power applications, where often neutral point clamped (three-level) inverters with Integrated Gate Commutated Thyristors (IGCT) are used, the control objectives are extended to the inverter and include the balancing of the inverter's neutral point.

In this paper, we present a DTC scheme that yields a significant performance improvement with respect to the state of the art without requiring excessive computational power to enable the implementation on (already existing) DTC hardware. Such a control scheme can be derived by adopting the principles of constrained optimal control with a receding horizon policy, i.e. Model Predictive Control (MPC) [5]. Specifically, the hysteresis bounds are inherited from standard DTC, whereas the DTC switching table is replaced by the following online optimization. Over a short *switching horizon* all admissible switching sequences are considered. Based on the measured (or estimated) machine quantities, for each switching sequence the evolution of the torque, stator flux and neutral point potential is predicted, using a nonlinear discrete-time model of the drive. To emulate a long *output horizon*, the predicted trajectories of the torque, flux and neutral point potential are extrapolated, and the number of time-steps is determined for which these quantities are kept within their hysteresis bounds. For each switching sequence an approximation of the average switching frequency is computed that is given by the total number of switch transitions in the sequence divided by the time duration of the extrapolated trajectory. Minimizing the switching frequency over all switching sequences yields the optimal sequence of switch transitions. Of this sequence, only the first step is applied to the drive. At the next sampling instant this procedure is repeated with new measurements thus establishing a receding horizon policy, as this method is commonly referred to.

This control approach carries several important advantages. Introducing additional control objectives (like the balancing of the switching power losses) is straightforward. As all computations are performed on-line, all quantities may be time-varying including model parameters, set points and hysteresis bounds. Even more importantly, the controller can be directly applied to a large class of three-phase AC drive. More specifically, induction machines (both squirrel-cage and ring-rotor type), synchronous and permanent magnet machines can be addressed, as well as inverter topologies such as two-, three- or five-level inverter. Yet, to simplify the exposition of the new control scheme, we focus in this paper on a specific application, where a neutral point clamped (three-level) voltage source inverter drives a squirrel-cage induction motor.

The major benefit of this controller, however, is its superior performance in terms of the switching frequency. For ABB's well-established three-level voltage source inverter ACS 6000 [6] with a squirrel-cage induction motor, the pro-

Manuscript received on April 15, 2008. This work was done at the Automatic Control Laboratory, ETH Zurich, Switzerland. This project was supported by ABB Switzerland Ltd. and by the two European Commission research projects IST-2001-33520 *Control and Computation (CC)* and FP6-IST-511368 *Hybrid Control (HYCON)*.

T. Geyer is currently with the Department of Electrical and Computer Engineering, The University of Auckland, Auckland, New Zealand (e-mail: t.geyer@ieee.org).

G. Papafotiou is now with ABB Corporate Research, Daettwil, Switzerland (e-mail: georgios.papafotiou@ch.abb.com).

M. Morari is with the Automatic Control Laboratory, ETH Zurich, Physikstrasse 3, 8092 Zurich, Switzerland (phone: +41 44 632 7626; fax: +41 44 632 1211; e-mail: morari@control.ee.ethz.ch).

posed model predictive DTC (MPDTC) scheme reduces the switching frequency over the whole range of operating points by up to 50 %, with an average reduction of 25 %, while better respecting the torque and flux hysteresis bounds. Notably, this result is independent from the power rating of the machine, which was varied from a few kVA to several MVA. All performance evaluations were carried out by applying the control scheme to a very accurate and detailed Matlab/Simulink model of the drive, which was provided by ABB to very closely resemble a physical drive and ensure a simulation set-up that is as realistic as possible. Experimental results shown in the second part of this paper confirm these simulation results to be very accurate.

However, MPC schemes are inherently computationally demanding as an underlying optimization problem needs to be solved. Yet, as will be illustrated, the proposed control approach requires a computational effort that is greatly reduced with respect to standard MPC approaches. Given the strong performance and design advantages, such an approach can be considered to be cost effective especially in the case of large drives operating in the MW region. In particular, a reduction of the switching frequency directly translates into reduced losses and therefore into energy and cost savings (in terms of operation and installation), which are significant in high power applications.

The proposed control scheme can be considered as a combination of the two DTC concepts [7], [8] we have proposed earlier. Specifically, the notion of optimality and the approximation of the average switching frequency by the number of switch transitions over a short prediction horizon were introduced in [7]; the concept of the evaluation forward in time is inherited from [8]. The key features of this new control scheme are the extrapolation, the fact that all computations are performed on-line (in [7], [8] we have pre-computed the control law off-line and stored it in a look-up table), the admissible switching sequence, and the use of a nonlinear (rather than a piecewise affine [7], [8]) prediction model for the controller synthesis.

As shown in [9], [10], the research community has recently started to consider model predictive control schemes as a way of introducing performance improvements in electrical drives. Even though DTC itself is widely interpreted as a predictive control strategy [11], [12], [13], [14], [15], it predicts only one step (one switch transition) ahead, and it lacks an internal model, a cost function and the notion of optimality, which are fundamental elements of an MPC scheme. Some of these elements are present in the more recent approaches [16], [17], [18], [19], [20], [21], [22], [23]. Yet, these schemes differ in several significant aspects from the control scheme proposed in this paper:

1) Except for [20] and possibly also [21], the prediction horizon is restricted to one and the usage of a larger (control) horizon combined with a receding horizon policy is not mentioned.
2) The DTC problem is formulated as a reference tracking problem, namely the formulated control problem tries to minimize the deviation of the torque and stator flux from their references. There are no hysteresis bounds on these variables.
3) Except for [17], [18], [21], [22], [23] a modulator is included in the control loop greatly simplifying the control problem. Thus, unlike the DTC philosophy, the switch positions are not directly chosen by the controller.
4) Only two-level voltage source inverters are considered with [13] and [23] being an exception.
5) The cost function does not emulate the switching frequency of the inverter. In combination with reference tracking and the limited number of voltage vectors (for a two-level inverter), a high switching frequency is to be expected. Particularly for the high power applications in the MVA range considered in this paper, such an approach does not seem to be applicable. [23] is an exception to this by also aiming at minimizing the switching frequency.
6) Linear (or locally linearized) models are used as prediction models. In particular for three-level inverters that include a neutral point potential with its nonlinear dynamical behavior, linear models tend to be inaccurate. Moreover, in [16], [20], [21], the cross couplings between the d and the q-axis are neglected in the machine model. Similarly, [22], [23] consider only a resistive inductive load with a back EMF rather than an electric motor.

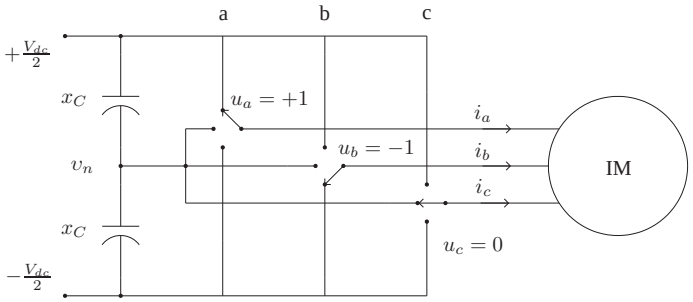

Fig. 1: Equivalent representation of a three-level voltage source inverter driving an induction motor (IM)

Here, due to the space limitation, we have provided only an overview of the literature on predictive control of motor drives that is most related to our approach. A coverage of the academic contributions to the DTC problem can be found in [24, Section 5.6] and [4].

The paper is structured as follows. Section II summarizes the physical model of the DTC drive, while Section III reformulates this model such that it can be used as prediction model in the MPC scheme. After stating the control problem in Section IV, the model predictive DTC scheme is detailed in Section V. This controller is available in two forms with control (switching) horizons $N > 1$ and $N = 1$, which differ mostly in the switching sequences' degree of freedom and the handling of the switching constraints. The control algorithms and the related computations are discussed in Section VI. Simulation results using ABB's simulation setup are presented in Section VII, and conclusions are drawn in Section VIII.

II. Physical Model

Throughout the document, we will use the normalized time scale t with one time unit corresponding to $1/\omega_b$ seconds, where ω_b is the base angular velocity. Additionally, we will use $\xi(t)$, $t \in \mathbb{R}$, to denote continuous-time variables, and $\xi(k)$, $k \in \mathbb{N}$, to denote discrete-time variables with the sampling interval $T_s = 25\,\mu$s.

All variables $\xi_{abc} = [\xi_a\ \xi_b\ \xi_c]^T$ in the three-phase system (abc) are transformed to $\xi_{\alpha\beta 0} = [\xi_\alpha\ \xi_\beta\ \xi_0]^T$ in the orthogonal $\alpha\beta 0$ stator reference frame through

$$\xi_{\alpha\beta 0} = P\,\xi_{abc}\,. \tag{1}$$

Using a stator reference frame and aligning the α-axis with the a-axis, the following transformation matrix is obtained

$$P = \frac{2}{3}\begin{bmatrix} 1 & -\frac{1}{2} & -\frac{1}{2} \\ 0 & \frac{\sqrt{3}}{2} & -\frac{\sqrt{3}}{2} \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \end{bmatrix}. \tag{2}$$

A. Physical Model of the Three-Level Inverter

The equivalent representation of a three-level voltage source inverter driving an induction motor is shown in Fig. 1. At each phase leg, the inverter can produce the three different voltages $-\frac{V_{dc}}{2}, 0, \frac{V_{dc}}{2}$, where V_{dc} refers to the total dc-link voltage. Let the integer variables u_a, u_b, $u_c \in \{-1, 0, 1\}$ denote the switch position in each phase leg, i.e. the phase state, where the values $-1, 0, 1$ correspond to the phase voltages $-\frac{V_{dc}}{2}, 0, \frac{V_{dc}}{2}$, respectively. The 27 vectors of the form $u_{abc} = [u_a\ u_b\ u_c]^T$ are transformed into the stator reference frame using (1). The resulting vectors of the form $u_{\alpha\beta 0} = [u_\alpha\ u_\beta\ u_0]^T$ are shown in Fig. 2, where they are mapped into the $\alpha\beta$ plane. The vectors in the $\alpha\beta$ plane are commonly referred to as voltage vectors, whereas we will refer to u_a, u_b, u_c as the switch positions. The actual voltages applied to the machine terminals are calculated from

$$v_{\alpha\beta 0} = \frac{V_{dc}}{2}\,P\,u_{abc}\,. \tag{3}$$

As can be seen in Fig. 1, the neutral point potential v_n depends on the state of charge of the two dc-link capacitors and is only affected when current is drawn directly from it, i.e. when one of the switch positions is zero. Introducing v_n as a state, the neutral point potential is described by

$$\frac{dv_n}{dt} = -\frac{1}{2x_C}\Big((1-|u_a|)i_{sa} + (1-|u_b|)i_{sb} + (1-|u_c|)i_{sc}\Big), \tag{4}$$

with the stator phase currents i_{sa}, i_{sb}, i_{sc} and one of the two symmetric capacitors x_C of the dc-link. Taking into account that $i_{sa} + i_{sb} + i_{sc} = 0$ it is straightforward to derive

$$\frac{dv_n}{dt} = \frac{1}{2x_C}|u_{abc}|^T\,P^{-1}\,i_{s,\alpha\beta 0}\,, \tag{5}$$

where $i_{s,\alpha\beta 0}$ is the stator current expressed in the stator reference frame, and $|u_{abc}| = [|u_a|\ |u_b|\ |u_c|]^T$ is the componentwise absolute value of the inverter switch positions. For more details about the nature of the neutral point potential and methods employed to tackle the related balancing problem, the reader is referred to [25] and [26].

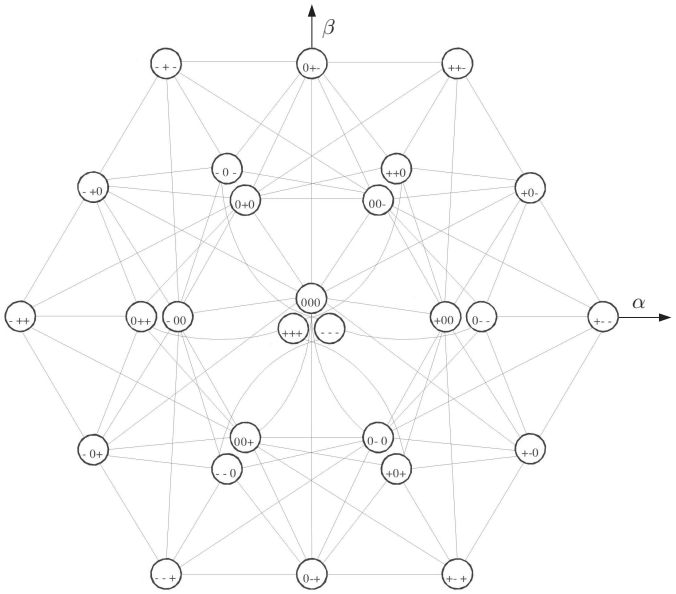

Fig. 2: Voltage vectors produced by a three-level inverter on the $\alpha\beta$ plane, the corresponding values of the switch positions u_{abc} (where '+' refers to '1' and '-' to '-1'), and the admissible switch transitions (courtesy of ABB ATDD, Switzerland)

In the inverter considered here – due to the fact that only one di/dt snubber is available in the upper and the lower half, respectively – not all switch transitions are possible, as depicted in Fig. 2. As can be seen, from $[1\ 1\ 1]^T$, for example, switching is only possible to $[0\ 1\ 1]^T$, $[1\ 0\ 1]^T$ or $[1\ 1\ 0]^T$ (and not to any of the other 23 switch positions).

B. Physical Model of the Induction Motor

The dynamics of the squirrel-cage induction motor are modelled in the stator $\alpha\beta 0$ reference frame. The α and β-components of the stator and the rotor flux linkages per second $\psi_{s\alpha}$, $\psi_{s\beta}$, $\psi_{r\alpha}$ and $\psi_{r\beta}$, respectively, and the rotor's rotational speed ω_r are used as state variables. The input voltages v_α and v_β are the stator voltages in the stator reference frame. The model parameters are the base angular velocity ω_b, the stator and rotor resistances r_s and r_r, the stator, rotor and mutual reactances x_{ls}, x_{lr} and x_m, respectively, the inertia J, and the mechanical load torque T_ℓ. Note that throughout this paper, if not otherwise stated, we are using normalized quantities, and the rotor quantities are referred to the stator circuit. The state equations are [27]

$$\frac{d\psi_{s\alpha}}{dt} = -r_s\frac{x_{rr}}{D}\psi_{s\alpha} + r_s\frac{x_m}{D}\psi_{r\alpha} + v_\alpha \tag{6a}$$

$$\frac{d\psi_{s\beta}}{dt} = -r_s\frac{x_{rr}}{D}\psi_{s\beta} + r_s\frac{x_m}{D}\psi_{r\beta} + v_\beta \tag{6b}$$

$$\frac{d\psi_{r\alpha}}{dt} = r_r\frac{x_m}{D}\psi_{s\alpha} - r_r\frac{x_{ss}}{D}\psi_{r\alpha} - \omega_r\psi_{r\beta} \tag{6c}$$

$$\frac{d\psi_{r\beta}}{dt} = r_r\frac{x_m}{D}\psi_{s\beta} + \omega_r\psi_{r\alpha} - r_r\frac{x_{ss}}{D}\psi_{r\beta} \tag{6d}$$

$$\frac{d\omega_r}{dt} = \frac{1}{J}\big(T_e - T_\ell\big)\,, \tag{6e}$$

with

$$x_{ss} = x_{ls} + x_m \tag{7a}$$
$$x_{rr} = x_{lr} + x_m \tag{7b}$$
$$D = x_{ss}x_{rr} - x_m^2 \tag{7c}$$

and the electromagnetic torque

$$T_e = \frac{x_m}{D}(\psi_{s\beta}\psi_{r\alpha} - \psi_{s\alpha}\psi_{r\beta})\,. \tag{8}$$

The length of the stator flux vector is given by

$$\Psi_s = \sqrt{\psi_{s\alpha}^2 + \psi_{s\beta}^2}\,. \tag{9}$$

Equations (6)–(9) represent the standard dynamical model of an induction motor, where the saturation of the machine's magnetic material, the changes of the rotor resistance due to the skin effect, and the temperature changes of the stator resistance are neglected.

III. Internal Model of the Controller

In this section, we derive a discrete-time model of the drive that is suitable to serve as an internal prediction model for the model predictive controller proposed in the next section. The purpose of this model is to predict the trajectory of the electromagnetic torque, the stator flux and the inverter neutral point potential over several sampling intervals in an open-loop fashion.

As the time-constant of the rotor speed dynamic exceeds the length of the prediction interval by several orders of magnitude, the rotor speed dynamics are neglected and the speed is assumed to remain constant within the prediction horizon. This allows us to treat the speed as a model parameter rather than as a state thus removing (6e) from the motor model.

The model of the inverter has one state, namely the neutral point potential, whose dynamic is described by (5) as a function of $|u_{abc}|$ and $i_{s,\alpha\beta 0}$. The α and β-components of $i_{s,\alpha\beta 0}$ are linear combinations of the stator and rotor flux components (see e.g. [27] for details), and the 0-component is always zero[1].

$$i_{s,\alpha\beta 0} = \frac{1}{D}\begin{bmatrix} x_{rr}\psi_{s\alpha} - x_m\psi_{r\alpha} & x_{rr}\psi_{s\beta} - x_m\psi_{r\beta} & 0 \end{bmatrix}^T . \tag{10}$$

We define the overall state vector of the DTC drive as

$$x = \begin{bmatrix} \psi_{s\alpha} & \psi_{s\beta} & \psi_{r\alpha} & \psi_{r\beta} & \upsilon_n \end{bmatrix}^T , \tag{11}$$

the switch positions u_a, u_b and u_c as the input vector

$$u = u_{abc} = \begin{bmatrix} u_a & u_b & u_c \end{bmatrix}^T \in \{-1, 0, 1\}^3 , \tag{12}$$

and the electromagnetic torque, the length of the stator flux and the neutral point potential as the output vector

$$y = \begin{bmatrix} T_e & \Psi_s & \upsilon_n \end{bmatrix}^T . \tag{13}$$

Combining the motor model (6a)–(6d), (8) and (9) with the model of the inverter (5) and (10), and using forward Euler approximation approach the following discrete-time model of the DTC drive is derived.

$$x(k+1) = (I + \begin{bmatrix} A & 0 \\ 0 & 0 \end{bmatrix} T_s)x(k) + \tag{14a}$$
$$+ \begin{bmatrix} B_1 \\ 0 \end{bmatrix} T_s u(k) + \begin{bmatrix} 0 \\ B_2(x(k)) \end{bmatrix} T_s |u(k)|$$
$$y(k) = g(x(k)) \tag{14b}$$

In this model, I denotes the identity matrix and $T_s = 25\,\mu$s is the sampling interval. The definitions of the matrices A, B_1 and B_2, and the vector $g(k)$ can be found in the appendix. Note that the zeros in (14) are vectors and matrices of appropriate dimensions. In (14) the first two terms capture the motor equations, while the third expression captures the dynamic of the neutral point potential.

IV. Control Problem

The DTC control objectives are to keep the three output variables, namely the torque, the length (or magnitude) of the stator flux and the neutral point potential, within given (hysteresis) bounds. At the same time, the average switching frequency of the inverter

$$f = \lim_{N\to\infty} \frac{1}{NT_s} \sum_{\ell=0}^{N} ||u(\ell) - u(\ell-1)|| \tag{15}$$

needs to be minimized, where $||\cdot||$ denotes the 1-norm.

The control problem is complicated by the fact that the control objectives comprise phenomena of very different time scales. Specifically, the control objectives relevant to the motor depend on the very fast dynamics of the stator flux, which is affected by the applied stator voltage within a few μs. On the other hand, the average inverter switching frequency needs to be evaluated over a time frame of several 100 ms. This is particularly the case in high power applications, where the switching frequencies are in the range of $200 - 400\,$Hz. This implies that every semiconductor switch is turned on roughly every $3-5\,$ms. Since the switch positions can be altered every $T_s = 25\,\mu$s, a very long horizon (a large N) of several hundred steps is necessary to capture the average switching frequency, thus leading to an intractable MPC scheme.

Apart from this, additional restrictions on the inverter switch transitions may be present resulting from the construction of the inverter (as mentioned in Section II-A). We will refer to them as *switching constraints*.

V. Model Predictive Direct Torque Control

Adopting the principles of MPC, we present a novel control methodology that considers all (admissible) switching sequences over a rather short switching horizon N, which is referred to as the control horizon in the MPC community. A switching sequence is defined as a sequence of semiconductor switch positions u_{abc} over the time-interval of length N from time-step 0 to time-step $N-1$. In a next step, based on the nonlinear discrete-time prediction model (14), the MPC scheme computes for each switching sequence the drive's response, i.e. the evolution of the output variables over the

[1]This follows from (1), taking into account that $i_{sa} + i_{sb} + i_{sc} = 0$.

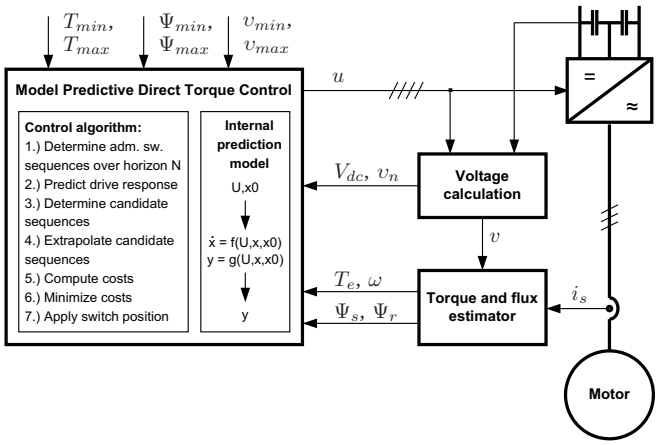

Fig. 3: Block diagram of the model predictive DTC scheme, where U denotes a switching sequences, and the initial state x_0 is composed of $x(k)$ and $u(k-1)$

switching horizon N. To emulate a long output horizon, the "promising" output trajectories are extrapolated, and the number of time-steps is determined when the first output variable hits a hysteresis bound. The cost associated with each switching sequence is determined by dividing the total number of switch transitions in the sequence by the length of the extrapolated trajectory. Minimizing this penalty yields the optimal switching sequence and the next optimal switch position to be applied to the inverter.

The model predictive DTC scheme is available in two forms with $N > 1$ and $N = 1$ differing mostly in the degree of freedom for the switching sequences and the handling of the switching constraints, and thus in the computational burden and the performance.

A. Horizon $N > 1$

For MPDTC with a switching horizon larger than one, we consider only switching sequences of length N that meet the switching constraints imposed by the physics of the inverter (see Section II-A). As shown in Fig. 3, given the current state $x(k)$, the last switch position $u(k-1)$, the bounds on the output variables, and using the nonlinear discrete-time prediction model (14) of the DTC drive, the controller computes at time-instant k the next switch position $u(k)$ according to the following procedure.

1) Given the last switch position $u(k-1)$ and taking into account the constraints on the switch transitions induced by the inverter topology, determine all switching sequences $U^i(k) = [u^i(k), \dots, u^i(k+N-1)]$ over the switching horizon N, where $i \in \mathcal{I}$ and $\mathcal{I}$ is an index set.
2) For these switching sequences, compute the drive response, i.e. compute all torque, stator flux and neutral point potential trajectories starting from $x(k)$ over the switching horizon N given by $Y^i(k) = [y^i(k), \dots, y^i(k+N)]$.
3) Determine those switching sequences that yield output trajectories that are either *feasible* at the end of the switching horizon, or *point in the proper direction* at all time-steps within the switching horizon. We refer to these switching sequences as *candidate* sequences $U^i(k)$ with $i \in \mathcal{I}_c \subseteq \mathcal{I}$. Feasibility means that the output variable lies within its corresponding bounds at time-step $k+N$; to point in the proper direction refers to the case in which an output variable is not necessarily feasible, but the degree of the bounds' violation decreases at every time-step within the switching horizon. For the case $N = 2$, Fig. 4 shows several example output trajectories that visualize these properties. The above condition needs to hold *componentwise*, i.e. for all three output variables[2].
4) For the candidate sequences, extrapolate the output trajectories from time-instant $k+N$ on linearly[3] using the samples at the time-instants $k+N-1$ and $k+N$. Derive the number of time-steps after which the first of the three output variables leaves the feasible region defined by the corresponding upper and lower bound[4]. This yields the number of time-steps n_i, $i \in \mathcal{I}_c$ this switching sequence can be applied before switching is predicted to be required again. Thus, n_i refers to the total length of the (extrapolated) output sequence.

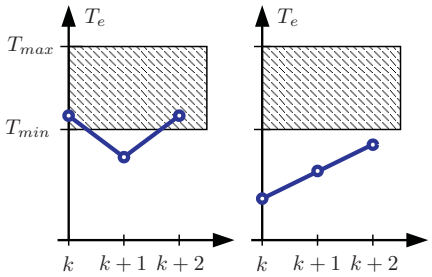

(a) Trajectories that are either feasible (left) or pointing in the proper direction (right)

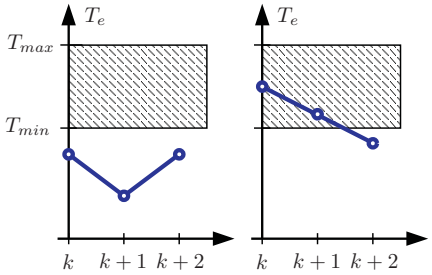

(b) Trajectories that are neither feasible (left) nor pointing in the proper direction (right)

Fig. 4: Example torque trajectories for $N = 2$ (the same concept applies to the stator flux and the neutral point potential). The feasible region between the hysteresis bounds is hatched

[2]As an example, consider the case where the torque is feasible, the stator flux points in the proper direction and the neutral point potential is feasible.

[3]Particularly during high speed operation, it is advantageous to extrapolate the stator flux quadratically using also the flux sample at $k+N+1$. The latter is computed by applying the switch position $u(k+N) = u(k+N-1)$.

[4]Note that we determine when the first output variable leaves the feasible region rather than when it hits a bound. This is done to account for situations in which an output variable lies outside its bounds but steers towards one of them.

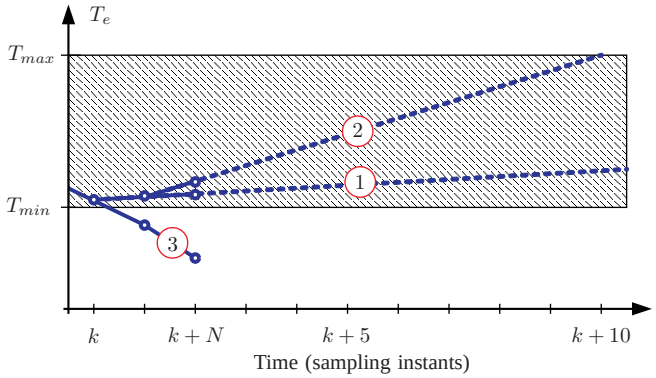

(a) Torque trajectories

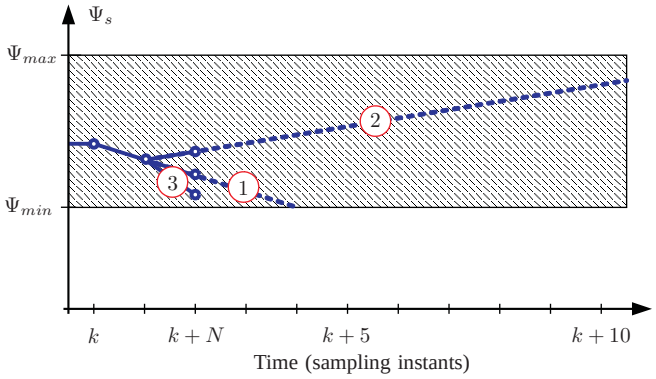

(b) Stator flux trajectories

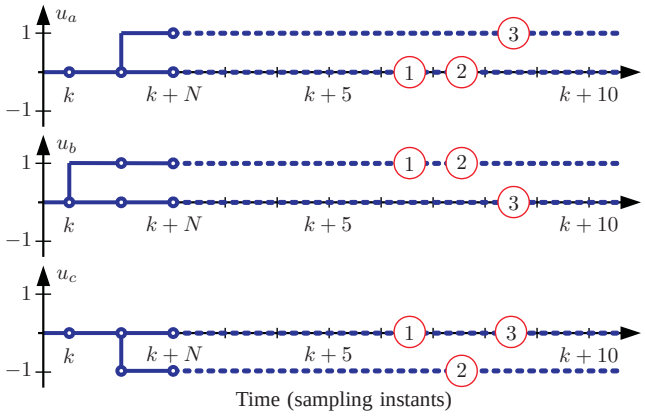

(c) Switch positions

Fig. 5: Torque and stator flux trajectories and the switch positions of Example 1 starting at time-instant k for the three switching sequences $U^1(k), U^2(k)$ and $U^3(k)$. The trajectories within the switching horizon $N = 2$ are solid, their extrapolations are dashed lines. The numbers refer to the indices of the switching sequences. The regions between the upper and lower (hysteresis) bounds are hatched

5) Compute for each candidate sequence $i \in \mathcal{I}_c$ the cost $c_i = s_i/n_i$, where

$$s_i = \sum_{\ell=k}^{k+N-1} ||u_i(\ell) - u_i(\ell-1)||_1 \qquad (16)$$

is the total number of switch transitions in the switching sequence $U^i(k)$, and n_i is the corresponding sequence length. The cost c_i is an approximation of the average switching frequency, and n_i can be interpreted as a time-varying output horizon.

6) Choose the switching sequence $U^* = U^i(k)$ with the minimal cost, where i is given by

$$i = \arg\min_{i \in \mathcal{I}_c} c_i \,. \qquad (17)$$

7) Apply only the first switch position $u^* = u^i(k)$ of this sequence.

At the next time-instant, repeat the procedure.

Example 1: To visualize the control concept, consider the example shown in Fig. 5. Assume there are the three switching sequences $U^i(k)$, $i \in \mathcal{I} = \{1, 2, 3\}$ over the switching horizon $N = 2$. According to the definition, $U^1(k)$ and $U^2(k)$ are candidate sequences, whereas $U^3(k)$ is not. Extrapolating the torque and the stator flux trajectories and determining when they leave the feasible region leads to the results summarized in Table I. Minimizing the cost yields the sequence $U^2(k)$ as the optimum. Note that this solution requires two switch transitions (one at time-instant k, the second one at time-instant $k+1$), but this investment pays out due to the longer length of the output trajectory. Without extrapolation, the controller would selected $U^1(k)$ as the optimum, since the corresponding cost expressions would be $\frac{1}{2}$ and 1 for $U^1(k)$ and $U^2(k)$, respectively. In the long run, however, this choice would be inferior compared with $U^2(k)$ thus motivating the concept of extrapolation. In this example, for the sake of simplicity, we neglect the neutral point potential, which is treated in exactly the same way as the torque and the stator flux.

Generally, the worst case computation time of the algorithm has to be accommodated in the given sampling interval. Yet, in a real world drive application, where other tasks like supervision and thermal protection are executed, the average computation time should also be kept at a minimum. When computing the next switch position, this can be achieved by first evaluating whether switching can be avoided altogether, i.e. whether the output variables are at time-step $k+N$ within their respective bounds when reapplying the last switch position for N time-steps. Only if this simple test fails, the above outlined computations need to be performed. Furthermore, bound techniques can be added to prune suboptimal branches thus avoiding the computation of the whole switching tree over N steps.

B. Horizon $N = 1$

The computational burden imposed by the model predictive DTC scheme with a switching horizon $N > 1$ might exceed the capabilities of some of the existing DTC control hardware. A further reduction of the computation time can be achieved

Sequence number i	Total length n_i of the (extrapolated) sequence	Number of switch transitions s_i	Cost c_i
1	4	1	1/4
2	10	2	1/5
3	–	–	–

TABLE I: Characteristics of the three switching sequences in Example 1

by restricting the switching horizon to $N = 1$. However, given the constraints on the switch transitions, such a short switching horizon imposes restrictions on the set of reachable voltage vectors and may lead to an infeasible control problem.

To reduce the computation time while dealing with the issue of infeasibility, we propose in this section a modified scheme that uses a switching horizon of $N = 1$ and initially ignores the switching constraints. As a result, the proposed algorithm yields an optimal switch position $u^*(k)$ that may not be directly reachable from the last switch position $u(k-1)$, since it was calculated by ignoring the switching constraints. In a last step, the constraints are reintroduced, and an admissible sequence of switch positions is calculated that leads to the optimal switch position within several time-steps. We will refer to this sequence as an *admissible switching sequence*.

More specifically, given the two switch positions μ and ν, with μ, $\nu \in \{-1, 0, 1\}^3$, the admissible switching sequence connects μ with ν via intermediate switch positions while taking into account the restrictions on the allowed switch transitions. From the fact that switching in one component (one stack of the inverter) by one switch transition at a time-step is always possible, it follows directly that an admissible switching sequence always exists, that additional switch transitions are not required and that the cost is not increased.

As the admissible switching sequence connecting two switch positions (with the lowest cost) is in general not unique, we use the following rules to narrow down the choices. (i) Follow the shortest path (in terms of time-steps); (ii) choose the switch position that yields the least number of switch transitions at time-instant k; (iii) choose the switch position that provides the most alternatives at time-instant $k + 1$. The admissible switching sequence can be stored in a look-up table, which holds only the first switch position in the sequence according to the receding horizon policy. This leads to a look-up table of dimension 27×27. Exploiting the $\frac{2\pi}{3}$ symmetry of the voltage vectors in the $\alpha\beta$ plane, the dimension of this look-up table can be reduced to 11×27 as detailed in [24].

Setting the switching horizon to $N = 1$, the control algorithm is the same as described in Section V-A with the following differences, where the numbering corresponds to the one in Section V-A.

1) Given the last switch position $u(k-1)$ and *ignoring* the constraints on the switch transitions induced by the inverter topology, 27 switching sequences of length one result.
7) From the look-up table read out of the admissible switching sequence from $u(k-1)$ to $u(k)$ the first switch position $u^*(k)$, and apply it to the inverter.

VI. Discussion and Computational Effort

The notion of the candidate switching sequences associated with output trajectories that are element-wise feasible at the end of the switching horizon or point towards the bounds, leads to the following properties. Firstly, step changes in the bounds (e.g. large steps in the torque reference) can be straightforwardly handled. In such cases, a switching sequence may not exist that moves the output variables within N time-steps back inside the bounds. Considering also sequences that move the output variables only towards – rather than inside – the bounds allows the control scheme to easily address this issue.

Secondly, excessive switching is avoided. As can be seen from Figs. 4(a) and 5, the bounds on the output variables are in general not strictly imposed by the MPDTC scheme, and output trajectories of candidate sequences may violate the bounds similar to standard DTC. As a result, one or more output variable might slightly violate a bound before a new switch position is selected. Hence, the MPDTC approach refrains from unnecessary switching when the bounds are slightly shifted, or when measurement noise and model uncertainties affect the predictions in an adverse way.

Since the proposed MPDTC scheme is intended to capture and minimize the average switching frequency, a long prediction interval is beneficial. To avoid an explosion of the related computational complexity, a short *switching (i.e. control) horizon* N (usually two or three steps), but a long *output horizon* (up to 100 steps) is used. This is achieved by extrapolating the output trajectories from the end of the switching horizon onwards until the time-instant where the first output variable hits one of its bounds. Linear extrapolation is straightforward to implement and computationally inexpensive. Note that this approach is closely related to blocking control moves in the control literature and the *multiple-rate prediction model* approach, which we introduced in [7]. As the simulation results will show, the concept of a short switching horizon combined with extrapolation allows us to greatly increase the length of the prediction interval thus enhancing the performance of the controller while keeping the computation times short.

Most importantly, the proposed control scheme can be easily adapted to different drives with different motors and inverter topologies (like two- or five-level inverters), as only the motor and inverter models need to be updated, and because there are no tuning parameters. Furthermore, possible constraints on the allowed switch transitions of the inverter can be easily incorporated by storing them in a look-up table.

A. Horizon $N > 1$

Nevertheless, when constraints on the allowed switch transitions are present, short switching horizons restrict the set of switch positions that can be reached within the switching horizon. More specifically, up to three consecutive steps are necessary to switch from one switch position to another (e.g. from $[-1 \;\; -1 \;\; 1]^T$ to $[1 \;\; 1 \;\; -1]^T$ and vice versa). Thus, a switching horizon of $N = 2$ is too short to ensure that any arbitrary switch position can be reached within the switching horizon. For $N = 2$, this occurs only very rarely, because those unreachable switch positions involve multiple switch transitions and are thus very expensive in terms of the cost expression. These issues are analyzed, visualized and solved in the second part of this paper [28] and in [24, Section 7.5].

Next, we briefly analyze the computational burden of the control algorithm. For this, we assume that all basic operations such as additions, multiplications, divisions and comparisons require one computation cycle as well as evaluating a look-up table. Possible operations for the loading or storing of variables and the execution of loops are neglected.

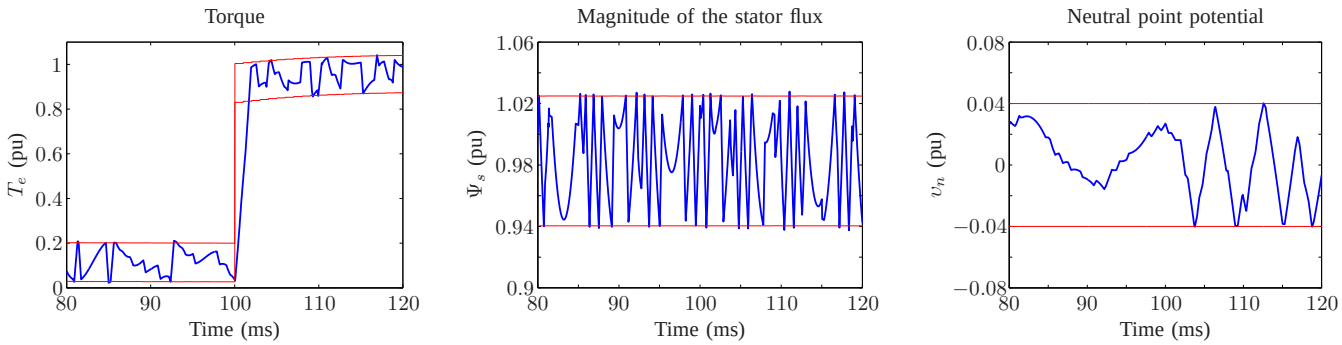

Fig. 6: Model Predictive DTC with $N = 2$: Torque step from 0.1 p.u. to 0.9 p.u. torque at $t = 100$ ms and 50 % speed. The torque, flux and neutral point potential trajectories are plotted together with their bounds over the time axis.

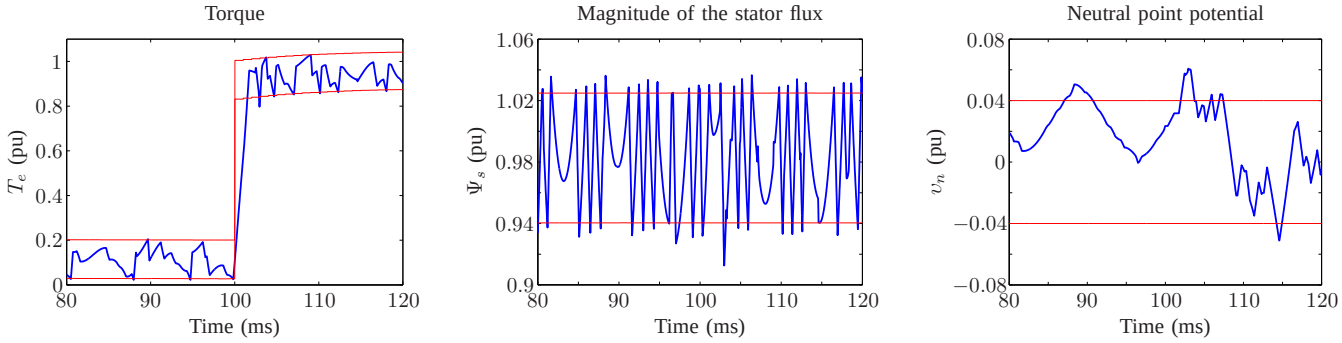

Fig. 7: Standard DTC: Torque step from 0.1 p.u. to 0.9 p.u. torque at $t = 100$ ms and 50 % speed. The torque, flux and neutral point potential trajectories are plotted together with their bounds over the time axis.

To simplify the exposition, the switching horizon is fixed to $N = 2$ hereafter. We assume that in the first step the inverter can switch from the last switch position $u(k-1)$ to s_1 voltage vectors. Subsequently, by branching on each of these s_1 switch positions we obtain a total of s_2 feasible switch positions at the second step. Then, as shown in detail in [24, Section 7.4.1], the total number of operations required to execute the control algorithm is $16 + 39s_1 + 79s_2$. For ABB's three-level inverter with the particular constraints on the admissible switch transitions, s_1 is upper bounded by 13 and s_2 is upper bounded by 121. Therefore, the upper bound on the total number of operations per control cycle is given by 10'082. Note that the computationally most expensive parts of the algorithm, which requires 70 % of the total computation power, can be easily parallelized.

B. Horizon $N = 1$

When limiting the switching horizon to $N = 1$ the same properties described above are induced, except for the following difference. Relaxing the constraints on the switch transitions and using the notion of the admissible switching sequence avoids the restrictions on the set of switch positions that can be reached within the switching horizon. More precisely, any switch position can be chosen by the optimization step, yet only the first switch position of the corresponding admissible switching sequence is applied to the inverter. This is done in accordance with the receding horizon policy.

Apart from that, bounds on the output variables are handled differently. In cases where the chosen switch position meets the switching constraints, i.e. the desired switch position can be directly applied and the admissible switching sequence is of length one, the bounds are strictly respected. However, if the admissible switching sequence comprises more than one element, the bounds are not guaranteed to be strictly respected (as it is the case for the scheme with $N > 1$).

Considering again the three-level inverter as above, the computational effort for $N = 1$ is reduced by a factor of five compared to the $N = 2$ case. Specifically, the upper bound is 1859 operations.

For a detailed analysis of the computational burden for the MPDTC algorithms with $N > 1$ and $N = 1$, the reader is referred to Sections 7.4.1 and 7.4.2 in [24], respectively. These sections also provide the pseudo code of the algorithms.

VII. Performance Evaluation

This section compares the performance of the proposed model predictive DTC schemes with ABB's ACS 6000 drive [6]. This implies that we will use ABB's DTC as a reference, to which we refer as standard DTC. The performance comparison is done through simulations that are based on a very accurate and detailed Matlab/Simulink model of the

Case study	Rated motor power	Rated motor voltage
I	1.6 MW	3.3 kV
II	6.6 MW	3.1 kV
III	15 kW	400 V

TABLE II: Motor ratings of the three case studies

drive, which was provided by ABB to ensure an as realistic as possible simulation set-up. This model includes a state estimator for the motor fluxes, and an outer (speed) control loop that adjusts the torque reference and accordingly the (time-varying) bounds on the torque. For MPDTC, the look-up table with ABB's DTC strategy is replaced by a function that solves at each sampling-instant the optimal control problem according to Fig. 3 and Section V.

Three case studies are considered comprising two medium-voltage drives and a low-voltage drive. Table II provides a rough overview of these case studies, while the detailed ratings and parameters can be found in [24, Section 7.6]. Two aspects are compared. Firstly, the torque, flux and neutral point potential are compared when applying a torque step. Secondly, the average switching frequency is compared over the whole range of operating points, and the root mean squared violation of the torque and stator flux bounds is evaluated, too. For both control schemes, the same torque and flux bounds are used. For the neutral point potential the bounds are chosen so as to reflect the behavior of ABB's control scheme, thus ensuring the comparison to be meaningful.

The evaluations are performed for the whole operating range by gridding the speed ω_r and the load torque T_ℓ at $0.1, 0.2, \ldots 1.0$ p.u.. The case of very high speed (0.9 and 1.0 p.u.) was left out for Case Studies I and II as the (total) dc-link voltage of 4294 V is too low to allow for an operation at high speed. At each operating point, the behavior of the drive was simulated over 2 s.

A. Torque Step for Horizon $N > 1$

For Case Study I and at 50 % speed a torque step is applied from 0.1 p.u. to 0.9 p.u.. Fig. 6 depicts the resulting closed-loop behavior of the torque, the stator flux and the neutral point potential with MPDTC, whereas Fig. 7 shows the corresponding trajectories resulting from ABB's DTC strategy. As can be seen, the MPDTC scheme preserves the rapid dynamic response achieved by the standard DTC approach of less than 2 ms at this operating point, while the bounds imposed on the torque, stator flux and neutral point potential are slightly better respected.

B. Average Switching Frequency for Horizon $N > 1$

First, we consider the MPDTC scheme proposed in Section V-A with a switching horizon $N > 1$. We set $N = 2$, and additionally limit the number of switching sequences by imposing an upper bound of three on the total number of switch transitions within a sequence. This bound removes a priori switching sequences with very high cost and generally does not affect the performance.

For Case Study I, Fig. 8(a) compares the average switching frequencies of standard DTC with our proposed scheme. This is done over the above defined grid of operating points. Fig. 8(b) shows the percentage-wise (relative) reduction of the average switching frequency. Averaging the data in Fig. 8(b) over all grid points yields an average switching frequency reduction of 25 % (or 60 Hz), while the maximum improvement amounts to 42 % (or 120 Hz). Over the whole range of operating points, the bounds on the torque, stator flux and neutral point potential are at least as well respected as by the standard DTC scheme, as shown in Fig. 9 (the definition of the root mean squared violation is given in Appendix A). Nevertheless, as described in Section VI-A, also the MPDTC scheme allows for slight violations of the bounds.

Case study	Average reduction of f	Maximal reduction of f
I	25 % or 60 Hz	42 % or 120 Hz
II	23 % or 66 Hz	42 % or 132 Hz
III	23 % or 91 Hz	49 % or 195 Hz

TABLE III: Reduction of the average switching frequency f of MPDTC with $N = 2$ relative to the standard DTC for the three case studies

It is interesting to note that the largest performance improvement is achieved around a modulation index of 0.5. Here, at the transition between the inner and the outer hexagon in the tree-level inverter, the "density" of voltage vectors and thus the degrees of freedom is at its maximum. This allows the model predictive DTC scheme to choose among several different switching strategies the one that meets the given performance objective best. On the other hand, it seems that the standard DTC scheme struggles to take advantage of this.

The figures for Case Studies II and III are very similar and thus omitted here. As shown in Table III, the average and the maximal reductions are similar for all three case studies indicating that the proposed control scheme works equally well for drives with very different characteristics and ratings.

C. Average Switching Frequency for Horizon $N = 1$

For the case of MPDTC with $N = 1$ as introduced in Section V-B, we consider only Case Study I as an example. For Case Study 1, the average reduction of the switching frequency over all grid points is given by 16 % (or 40 Hz), while the maximum improvement amounts to 38 % (or 91 Hz). The bounds on the torque, stator flux and neutral point potential are almost perfectly respected over the whole range of operating points. The corresponding figures are omitted here due to space limitations. The interested reader is referred to [24, Section 7.4.4].

Compared to the performance results for $N = 2$, the performance improvement for $N = 1$ is smaller by one third. However, at least in the absence of noise, the MPDTC scheme with $N = 1$ keeps the controlled variables tightly within their bounds, whereas small violations of the bounds are tolerated for $N > 1$.

VIII. Conclusions

In this paper, we have presented the new drive control concept *model predictive DTC* that is based on an internal controller model, switching sequences comprising multiple time-steps, a controller objective function, and an optimization stage. The latter minimizes the objective function with respect to the internal model dynamic, the bounds on the torque, stator flux and neutral point potential, and the admissible switch transitions. This optimization is performed over a multiple step prediction horizon. The control scheme is available in two forms – with a switching horizon larger than one and

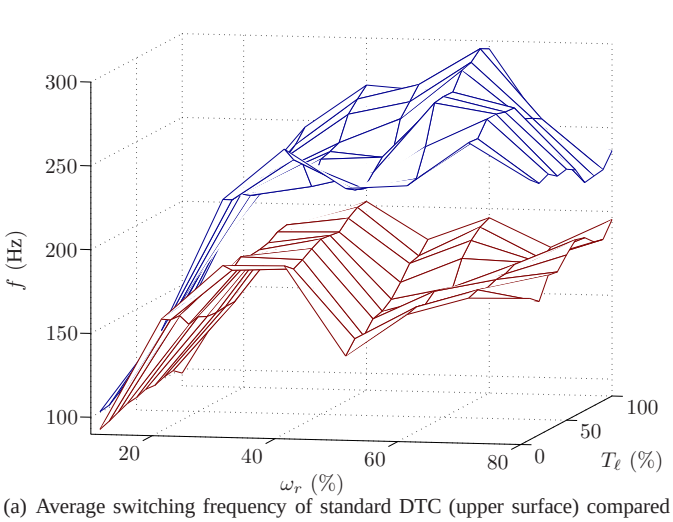

(a) Average switching frequency of standard DTC (upper surface) compared with MPDTC (lower surface)

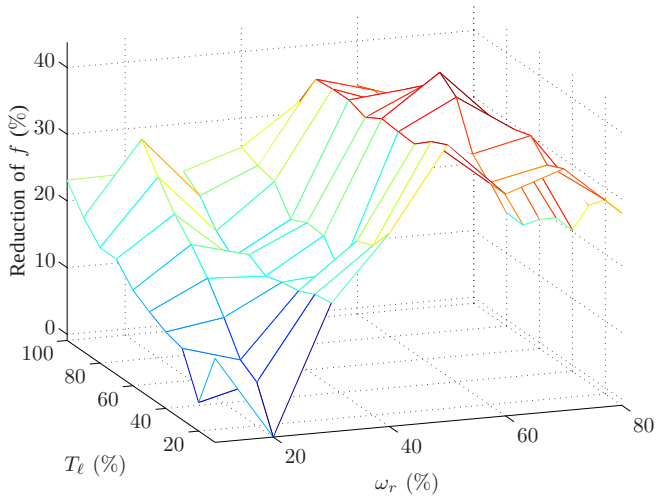

(b) Relative (percentage-wise) reduction of the average switching frequency

Fig. 8: Comparison of the average switching frequency f of standard DTC with respect to MPDTC with $N = 2$ over the grid of operating points (speed and torque)

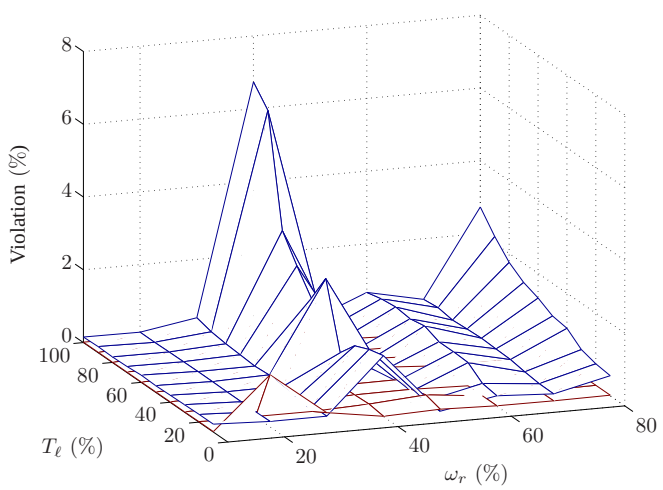

(a) Torque violation of the bounds

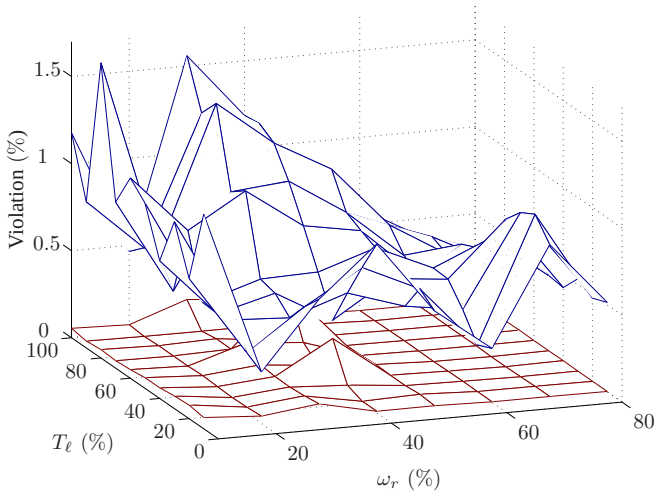

(b) Stator flux violation of the bounds

Fig. 9: Percentage-wise root mean squared violation of the torque and the stator flux (hysteresis) bounds for standard DTC (upper surface) compared with MPDTC with $N = 2$ (lower surface) over the grid of operating points (speed and torque)

with a horizon equal to one. For illustration purposes, we have focused in this paper on a three-level voltage source inverter driving a squirrel-cage induction motor.

Compared to the state of the art in drive control, the proposed control scheme offers two major advantages. (i) Performance: MPDTC inherits the very fast dynamic torque response of standard DTC that is inherent in hysteresis-based control concepts. By penalizing the short-term switching frequency in the objective function, the average switching frequency is reduced – compared to ABB's ACS 6000 scheme – by up to 50 %. In average over the whole range of operating points, the switching frequency is reduced by a quarter. At the same time, the imposed bounds on the torque, flux and neutral point potential are more strictly respected. In the second part of this paper [28], these simulation results are accurately confirmed by experimental results. Moreover, the performance improvement is practically independent from the rating and the specific drive characteristics.

(ii) Flexibility and simplicity: The MPDTC scheme is highly flexible. It is straightforward to incorporate additional or different performance and control objectives by simply modifying the cost function. Most importantly, the controller can be directly applied to a large class of three-phase AC drives – since the controller is based on an internal prediction model, only this internal model needs to be adapted. This adaptation can be done on-line as a parameter adaption to account for a varying stator resistance for example, or it can be done off-line as a model structure change to make the controller applicable to another drive with a different inverter topology and/or a different electrical machine. The fact that the

internal model is based on first principle dynamic equations simplifies the controller design. Except for the width of the hysteresis bounds, tuning parameters (like controller gains in field oriented control) are not required.

However, these advantages impose a significant computational burden inherent to model predictive control. It is our opinion that the computational burden associated with the algorithms is close to the lowest possible that is achievable for a model predictive control approach. Moreover, ABB's successful implementation of the MPDTC approach shown in the second part of this paper [28] proves that such algorithms can be run on the already existing hardware in the time scales required.

APPENDIX

The matrices and vectors of the discrete-time drive model (14) are given as follows.

$$A = \begin{bmatrix} -r_s\frac{x_{rr}}{D} & 0 & r_s\frac{x_m}{D} & 0 \\ 0 & -r_s\frac{x_{rr}}{D} & 0 & r_s\frac{x_m}{D} \\ r_r\frac{x_m}{D} & 0 & -r_r\frac{x_{ss}}{D} & -\omega_r \\ 0 & r_r\frac{x_m}{D} & \omega_r & -r_r\frac{x_{ss}}{D} \end{bmatrix}, \quad (18)$$

$$B_1 = \frac{V_{dc}}{2}\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} P\,, \quad (19)$$

$$B_2(x(k)) = x^T(k)\frac{1}{2x_C}\begin{bmatrix} \frac{x_{rr}}{D} & 0 & 0 \\ 0 & \frac{x_{rr}}{D} & 0 \\ -\frac{x_m}{D} & 0 & 0 \\ 0 & -\frac{x_m}{D} & 0 \\ 0 & 0 & 0 \end{bmatrix} P^{-T} \quad (20)$$

and

$$g(k) = \begin{bmatrix} \frac{x_m}{D}(x_2(k)x_3(k) - x_4(k)x_1(k)) \\ \sqrt{x_1^2(k) + x_2^2(k)} \\ x_5(k) \end{bmatrix}, \quad (21)$$

where x_i denotes the ith component of the vector x.

The percentage-wise root mean squared violation of the torque over the discrete-time axis $k = 1, 2, \ldots, k_{max}$ is defined as

$$100\%\sqrt{\frac{1}{k_{max}}\sum_{k=1}^{k_{max}}(\epsilon_T(k))^2} \quad (22)$$

with

$$\epsilon_T(k) = \begin{cases} T_e(k) - T_{max} & \text{if } T_e(k) > T_{max} \\ T_{min} - T_e(k) & \text{if } T_e(k) < T_{min} \\ 0 & \text{else}\,, \end{cases} \quad (23)$$

where $T_e(k)$ and the bounds are given in p.u.. The violations for the stator flux and the neutral point potential are defined accordingly.

ACKNOWLEDGMENT

The authors would like to thank Christian Stulz, Pieder Jörg and Petri Schroderus of ABB ATDD, Turgi, Switzerland, and Andreas Poncet of ABB Corporate Research, Baden-Dättwil, Switzerland, for their continuous advice.

REFERENCES

[1] I. Takahashi and T. Noguchi, "A new quick response and high efficiency control strategy for the induction motor," *IEEE Trans. Ind. Applicat.*, vol. 22, no. 2, pp. 820–827, Sep./Oct. 1986.

[2] I. Takahashi and Y. Ohmori, "High-performance direct torque control of an induction motor," *IEEE Trans. Ind. Applicat.*, vol. 25, no. 2, pp. 257–264, Mar./Apr. 1989.

[3] P. Pohjalainen, P. Tiitinen, and J. Lulu, "The next generation motor control method - direct torque control, DTC," in *Proc. European Power Electron. Chapter Symp.*, vol. 1, Lausanne, Switzerland, 1994, pp. 115–120.

[4] G. S. Buja and M. P. Kazmierkowski, "Direct torque control of PWM inverter-fed AC motors – a survey," *IEEE Trans. Ind. Electron.*, vol. 51, no. 4, pp. 744–757, Aug. 2004.

[5] J. Maciejowski, *Predictive Control*. Prentice Hall, 2002.

[6] ABB Asea Brown Boveri Ltd, "Product webpage of ACS 6000," online document, www.abb.com/motors&drives.

[7] G. Papafotiou, T. Geyer, and M. Morari, "A hybrid model predictive control approach to the direct torque control problem of induction motors (invited paper)," *Int. J. of Robust Nonlinear Control*, vol. 17, no. 17, pp. 1572–1589, Nov. 2007.

[8] T. Geyer and G. Papafotiou, "Direct torque control for induction motor drives: A model predictive control approach based on feasibility," in *Hybrid Systems: Computation and Control*, ser. LNCS, M. Morari and L. Thiele, Eds. Springer, 2005, vol. 3414, pp. 274–290.

[9] R. Kennel and A. Linder, "Predictive control of inverter supplied electrical drives," in *Proc. 31st IEEE Power Electron. Specialists Conf.*, vol. 2, Galway, Ireland, 2000, pp. 761–766.

[10] P. Cortés, M. P. Kazmierkowski, R. M. Kennel, D. E. Quevedo, and J. Rodríguez, "Predictive control in power electronics and drives," *IEEE Trans. Ind. Electron.*, vol. 55, no. 12, pp. 4312–4324, Dec. 2008.

[11] M. Pacas and J. Weber, "Predictive direct torque control for the PM synchronous machine," *IEEE Trans. Ind. Electron.*, vol. 52, no. 5, pp. 1350–1356, Oct. 2005.

[12] G. Abad, M. A. Rodriguez, and J. Poza, "Predictive direct torque control of the doubly fed induction machine with reduced torque and flux ripples at low constant switching frequency," in *Proc. Annual Conf. on IEEE Ind. Electron.*, 2006, pp. 1000–1005.

[13] P. Correa, M. Pacas, and J. Rodríguez, "Predictive torque control for inverter-fed induction machine," *IEEE Trans. Ind. Electron.*, vol. 54, no. 2, pp. 1073–1079, Apr. 2007.

[14] R. Morales-Caporal and M. Pacas, "A predictive torque control for the synchronous reluctance machine taking into account the magnetic cross saturation," *IEEE Trans. Ind. Electron.*, vol. 54, no. 2, pp. 1161–1167, Apr. 2007.

[15] M. Nemec, D. Nedeljković, and V. Ambrožič, "Predictive torque control of induction machines using immediate flux control," *IEEE Trans. Ind. Electron.*, vol. 54, no. 4, pp. 2009–2017, Aug. 2007.

[16] R. Kennel, A. Linder, and M. Linke, "Generalized predictive control (GPC) – ready for use in drive applications?" in *Proc. 32nd IEEE Power Electron. Specialists Conf.*, vol. 4, Vancouver, Canada, 2001, pp. 1839–1844.

[17] G. Escobar, A. M. Stanković, E. Galvan, J. M. Carrasco, and R. Ortega, "A family of switching control strategies for the reduction of torque ripple in DTC," *IEEE Trans. Contr. Syst. Technol.*, vol. 11, no. 6, pp. 933–939, Sep. 2003.

[18] J. Rodríguez, J. Pontt, C. Silva, P. Cortés, U. Amman, and S. Rees, "Predictive direct torque control of an induction machine," in *Proc. 11th IEEE Power Electron. and Motion Control Conf.*, Riga, Latvia, 2004.

[19] J. M. Retif, X. Lin-Shi, A. M. Llor, and F. Morand, "New hybrid direct-torque control for a winding rotor synchronous machine," in *Proc. 35st IEEE Power Electron. Specialists Conf.*, vol. 1, Aachen, Germany, 2004, pp. 1438–1442.

[20] A. Linder and R. Kennel, "Model predictive control for electrical drives," in *Proc. 36st IEEE Power Electron. Specialists Conf.*, Recife, Brasil, 2005, pp. 1793–1799.

[21] ——, "Direct model predictive control – a new direct predictive control strategy for electrical drives," in *Proc. European Conf. on Power Electron. and Applicat.*, Dresden, Germany, 2005.

[22] J. Rodríguez, J. Pontt, C. A. Silva, P. Correa, P. Lezana, P. Cortés, and U. Ammann, "Predictive current control of a voltage source inverter," *IEEE Trans. Ind. Electron.*, vol. 54, no. 1, pp. 495–503, Feb. 2007.

[23] R. Vargas, P. Cortés, U. Ammann, J. Rodríguez, and J. Pontt, "Predictive control of a three-phase neutral-point-clamped inverter," *IEEE Trans. Ind. Electron.*, vol. 54, no. 5, pp. 2697–2705, Oct. 2007.

[24] T. Geyer, "Low complexity model predictive control in power electronics and power systems," Dr. sc. tech. thesis, Automatic Control Laboratory ETH Zurich, 2005.
[25] H. du Toit Mouton, "Natural balancing of three-level neutral-point-clamped PWM inverters," *IEEE Trans. Ind. Electron.*, vol. 49, no. 5, pp. 1017–1025, Oct. 2002.
[26] N. Celanovic and D. Boroyevich, "A comprehensive study of neutral-point voltage balancing problem in three-level neutral-point-clamped voltage source PWM inverters," *IEEE Trans. Power Electron.*, vol. 15, no. 2, pp. 242–249, Mar. 2000.
[27] P. C. Krause, *Analysis of Electric Machinery*. McGraw-Hill, NY, 1986.
[28] G. Papafotiou, J. Kley, K. Papadopoulos, P. Bohren, and M. Morari, "Model predictive direct torque control - part II: Implementation and experimental evaluation," *IEEE Trans. Power Electron.*, special issue on Predictive Control of Power Electronics Drives, 2009.

Tobias Geyer (M'08) received his Degree (Dipl.-Ing. ETH) in 2000 and his Ph.D. in 2005 both in Electrical Engineering from ETH Zurich, Switzerland. From 2006 to 2008, he was with GE's Global Research Center in Munich, Germany, where he was a project leader working on control and modulation schemes for large electrical drives. In 2008 he joined the Power Electronics Group of the University of Auckland, New Zealand, as a Research Fellow. Dr. Geyer's research interest is at the intersection of the fields of control, optimization and power electronics, which includes model predictive control, hybrid systems, electrical drives and dc-dc conversion.

Georgios Papafotiou (M'03) received the Diploma and Ph.D. degrees, both in Electrical Engineering, from the Aristotle University of Thessaloniki, Thessaloniki, Greece in 1997 and 2002, respectively. In 2006 he joined the ABB Corporate Research Center, Dättwil, Switzerland, where he is currently working on the development of modern control and estimation methods for power electronics applications. From 2003 to 2006 he was with the Automatic Control Laboratory, ETH Zurich, Zurich, Switzerland, where he worked on MPC for hybrid systems, with a special focus on dc-dc converters and induction motor drives.

Manfred Morari (F'05) was appointed head of the Automatic Control Laboratory at ETH Zurich in 1994. Before that he was the McCollum-Corcoran Professor of Chemical Engineering and Executive Officer for Control and Dynamical Systems at the California Institute of Technology. He obtained the diploma from ETH Zurich and the Ph.D. from the University of Minnesota, both in chemical engineering. His interests are in hybrid systems and the control of biomedical systems. In recognition of his research contributions, he received numerous awards, among them the Donald P. Eckman Award and the John R. Ragazzini Award of the Automatic Control Council, the Allan P. Colburn Award and the Professional Progress Award of the AIChE, the Curtis W. McGraw Research Award of the ASEE, Doctor Honoris Causa from Babes-Bolyai University, Fellow of IEEE and IFAC, the IEEE Control Systems Field Award, and was elected to the National Academy of Engineering (U.S.). Manfred Morari has held appointments with Exxon and ICI plc and serves on the technical advisory boards of several major corporations.