

Performance of Multistep Finite Control Set Model Predictive Control for Power Electronics

Tobias Geyer, *Senior Member, IEEE* and Daniel E. Quevedo, *Senior Member, IEEE*

Abstract

The performance of direct model predictive control (MPC) with reference tracking and long prediction horizons is evaluated through simulations, using the current control problem of a variable speed drive system with a voltage source inverter as an illustrative example. A modified sphere decoding algorithm is used to efficiently solve the optimization problem underlying MPC for long horizons. For a horizon of five and a three-level inverter, for example, the computational burden is reduced by four orders of magnitude, compared to the standard exhaustive search approach. This work illustrates the performance gains that are achievable by using prediction horizons larger than one. Specifically, for long prediction horizons and a low switching frequency, the total harmonic distortion of the current is significantly lower than for space vector modulation, making direct MPC with long horizons an attractive and computationally viable control scheme.

Index Terms

Model predictive control, finite control set, sphere decoding, branch and bound, quantization, power electronics, drive systems

I. INTRODUCTION

The optimization problem underlying direct (also called *finite control set*) model predictive control (MPC) with reference tracking is typically solved by enumerating all possible solutions [2]. Since the number of possible solutions increases exponentially as a function of the length of the prediction horizon, when using enumeration, enlarging the prediction horizon entails an exponential increase in the computation time. For direct MPC with reference tracking, this combinatorial explosion has to date, in effect, limited the length of implementable prediction horizons to one [3], [4]. Conversely, solving the optimization problem of direct MPC with long prediction horizons in an *efficient* manner has been hitherto an unresolved problem.

A solution approach to this problem is proposed in [5], which adopts the notion of sphere decoding [6] and tailors it to the problem at hand. As a result, the optimization problem underlying MPC problem with long prediction horizons, such as ten, can be solved on average as quickly as the horizon of one case using enumeration.

T. Geyer is with ABB Corporate Research, Baden-Dättwil, Switzerland; e-mail: t.geyer@ieee.org

Daniel E. Quevedo is with the University of Newcastle, Australia; e-mail: daniel.quevedo@newcastle.edu.au

A preliminary version of this manuscript was presented at the IEEE Energy Conversion Congress and Exposition (ECCE) 2013 in Denver, CO, USA, see [1]

In this paper we start by summarizing the problem formulation and the key results obtained in [5]. Using a three-level neutral point clamped (NPC) voltage source inverter driving an induction machine as a case study, we present a detailed performance evaluation that highlights essential features of direct MPC with horizons significantly longer than one by using the optimization algorithm presented in [5]. At steady-state operating conditions, the key performance metrics are the total harmonic distortion (THD) of the current and the device switching frequency. Several comparisons with two commonly used modulation schemes are performed, i.e. space vector modulation (SVM) and optimized pulse patterns (OPPs) [7], [8]. Interestingly, in some cases MPC with long horizons has the potential to achieve a performance similar to that of OPPs, even at *steady state*. During transients, MPC with long horizons provides a transient response time as short as MPC with short horizons, often outperforming classic control arrangements such as field oriented control.

II. CONTROL PROBLEM AND SOLUTION METHOD

This section recapitulates the main subject matter of [5], by summarizing the control problem, model predictive control formulation and the efficient solution method based on sphere decoding.

A. Control Problem

In the stationary $\alpha\beta$ coordinate system, we consider a discrete-time linear power electronic system modelled as per

$$\mathbf{x}(k+1) = \mathbf{A}\mathbf{x}(k) + \mathbf{B}\mathbf{u}_{\alpha\beta}(k) \quad (1a)$$

$$\mathbf{y}(k) = \mathbf{C}\mathbf{x}(k) \quad (1b)$$

with system matrices \mathbf{A} , \mathbf{B} and \mathbf{C} , and $k \in \mathbb{N}$. We use $\mathbf{u}_{\alpha\beta} = \mathbf{P}\mathbf{u}$ with

$$\mathbf{P} = \frac{2}{3} \begin{bmatrix} 1 & -\frac{1}{2} & -\frac{1}{2} \\ 0 & \frac{\sqrt{3}}{2} & -\frac{\sqrt{3}}{2} \end{bmatrix} \quad (2)$$

to translate the three-phase switch positions $\mathbf{u} \triangleq [u_a \ u_b \ u_c]^T$ into the orthogonal coordinate system. Note that $\mathbf{u} \in \mathcal{U}^3$ is integer-valued. For a three-level inverter, for example, we have $\mathcal{U} = \{-1, 0, 1\}$. For a drive system with an induction machine, we choose as state vector $\mathbf{x} \triangleq [i_{s\alpha} \ i_{s\beta} \ \psi_{r\alpha} \ \psi_{r\beta}]^T$, where $\mathbf{i}_s \triangleq [i_{s\alpha} \ i_{s\beta}]^T$ denotes the stator current, and $\psi_{r\alpha}$ and $\psi_{r\beta}$ are the rotor flux linkages. The system output is chosen as $\mathbf{y} = \mathbf{i}_s$.

The control problem is to regulate the stator current along its reference \mathbf{i}_s^* , by manipulating the three-phase switch position \mathbf{u} . The switching effort, i.e., the switching frequency or the switching losses is to be kept small.

B. Model Predictive Control Formulation

The cost function

$$J = \sum_{\ell=k}^{k+N-1} \|\mathbf{i}_{e,\alpha\beta}(\ell+1)\|_2^2 + \lambda_u \|\Delta\mathbf{u}(\ell)\|_2^2, \quad (3)$$

is a suitable choice to penalize the predicted evolution of the current errors and the control effort over the prediction horizon of N steps, where

$$\mathbf{i}_{e,\alpha\beta}(\ell+1) \triangleq \mathbf{i}_{s,\alpha\beta}^*(\ell+1) - \mathbf{i}_{s,\alpha\beta}(\ell+1) \quad (4a)$$

$$\Delta\mathbf{u}(\ell) \triangleq \mathbf{u}(\ell) - \mathbf{u}(\ell-1). \quad (4b)$$

The scalar parameter $\lambda_u \geq 0$ is a tuning parameter that adjusts the trade-off between the tracking accuracy and the switching effort.

Introducing the switching sequence $\mathbf{U}(k) = [\mathbf{u}^T(k) \dots \mathbf{u}^T(k+N-1)]^T$, the optimization problem underlying direct MPC with current reference tracking is

$$\mathbf{U}_{\text{opt}}(k) = \arg \min_{\mathbf{U}(k)} J \quad (5a)$$

$$\text{subj. to } \mathbf{U}(k) \in \mathbb{U} \quad (5b)$$

$$\|\Delta \mathbf{u}(\ell)\|_{\infty} \leq 1, \forall \ell = k, \dots, k+N-1, \quad (5c)$$

where $\mathbb{U} = \mathcal{U}^{3N}$. The constraint (5c) is imposed since switching is only possible by one step up or down in each phase.

C. An Efficient Method for Calculating the Optimal Switch Positions

Using algebraic manipulations, the optimization problem (5) can be rewritten as

$$\mathbf{U}_{\text{opt}}(k) = \arg \min_{\mathbf{U}(k)} \|\bar{\mathbf{U}}_{\text{unc}}(k) - \mathbf{H}\mathbf{U}(k)\|_2^2 \quad (6a)$$

$$\text{subj. to } (5b) \text{ and } (5c). \quad (6b)$$

\mathbf{H} turns out to be an invertible lower-triangular matrix, see [5]. We use $\bar{\mathbf{U}}_{\text{unc}}(k) \triangleq \mathbf{H}\mathbf{U}_{\text{unc}}(k)$, where $\mathbf{U}_{\text{unc}}(k)$ is the switching sequence obtained from minimizing (5) *without* constraints, i.e., with $\mathbb{U} = \mathbb{R}^{3N}$ and ignoring (5c).

The MPC optimization problem in its form (6) is a (truncated) *integer least-squares* problem. As shown in [5], the sphere decoding algorithm [6], [9] can be adapted to solve (6). The algorithm iteratively considers candidate sequences $\mathbf{U} \in \mathbb{U}$ within a sphere of radius $\rho(k) > 0$ centered in $\bar{\mathbf{U}}_{\text{unc}}(k)$, i.e., $\|\bar{\mathbf{U}}_{\text{unc}}(k) - \mathbf{H}\mathbf{U}\|_2 \leq \rho(k)$, and which satisfy the switching constraint (5c). Since \mathbf{H} is triangular, finding candidate sequences is computationally simple, in the sense that at each step only a one-dimension problem needs to be solved. For more details, the reader is referred to [5] and [6].

III. FRAMEWORK FOR PERFORMANCE EVALUATION

A. Case Study

We consider a NPC voltage source inverter connected with a medium-voltage induction machine and a constant mechanical load, as shown in Fig. 1. A 3.3 kV and 50 Hz squirrel-cage induction machine rated at 2 MVA with a total leakage inductance of 0.25 pu is used as an example of a typical medium-voltage induction machine. The dc-link voltage is $V_{\text{dc}} = 5.2$ kV and assumed to be constant. The potential of the neutral point N is assumed to be fixed. The detailed parameters of the machine and inverter are summarized in Table I. The per unit (pu) system is established using the base quantities $V_B = \sqrt{2/3}V_{\text{rat}} = 2694$ V, $I_B = \sqrt{2}I_{\text{rat}} = 503.5$ A and $f_B = f_{\text{rat}} = 50$ Hz, with V_{rat} , I_{rat} and f_{rat} referring to the rated voltage, current and frequency, respectively.

All simulations are performed in Matlab, using an idealized setup with the semiconductors switching instantaneously. As such, second order effects such as deadtimes, controller computation delays, measurement noise, observer errors, saturation of the machine's magnetic material, parameter variations, etc. are neglected, resulting

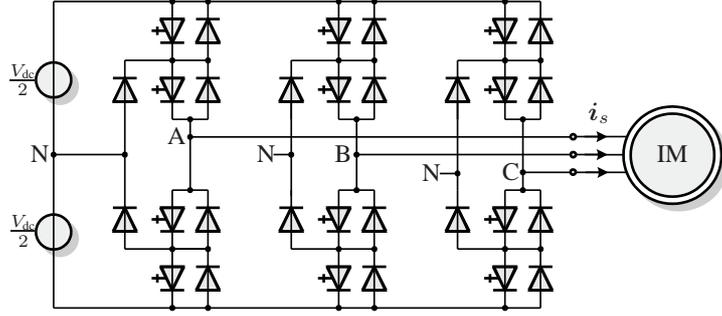


Fig. 1: Three-level three-phase neutral point clamped voltage source inverter driving an induction motor with a fixed neutral point potential

in a switched linear model for the drive system. The significance of such simulations is underlined by the very close match between previous simulations and experimental results using a similar model. The simulation results in [10] predicted the experimental results in [11] accurately to within a few percent. Throughout this paper, if not otherwise stated, all simulations were done at nominal speed and rated torque, implying a fundamental frequency of 50 Hz and rated currents. All results are shown in the pu system.

B. Modulation Methods used for Benchmarking

To evaluate the steady-state performance of multi-step optimal direct MPC, we benchmark this strategy with SVM and OPPs. The SVM gating signals are obtained by using a three-level regular sampled PWM with two triangular carrier signals, which are in phase (phase disposition). By adding to the reference voltage an appropriate common mode voltage, which is of the min/max plus modulo type, the modulator resembles a SVM, as shown in [12]. Synchronous modulation is used, i.e., the carrier frequency is an integer multiple of the fundamental frequency.

The OPPs were calculated offline for pulse numbers (ratio between the switching frequency and the fundamental frequency) of up to 15. The switching angles were computed by minimizing the squared differential-mode voltage harmonics divided by the order of the harmonic. For an inductive load such as a machine, this approach is effectively equivalent to minimizing the current THD [8].

Induction motor	Voltage	3300 V	r_s	0.0108 pu
	Current	356 A	r_r	0.0091 pu
	Real power	1.587 MW	x_{ls}	0.1493 pu
	Apparent power	2.035 MVA	x_{lr}	0.1104 pu
	Frequency	50 Hz	x_m	2.3489 pu
	Rotational speed	596 rpm		
Inverter			V_{dc}	1.930 pu
			x_c	11.769 pu

TABLE I: Rated values (left) and parameters (right) of the drive

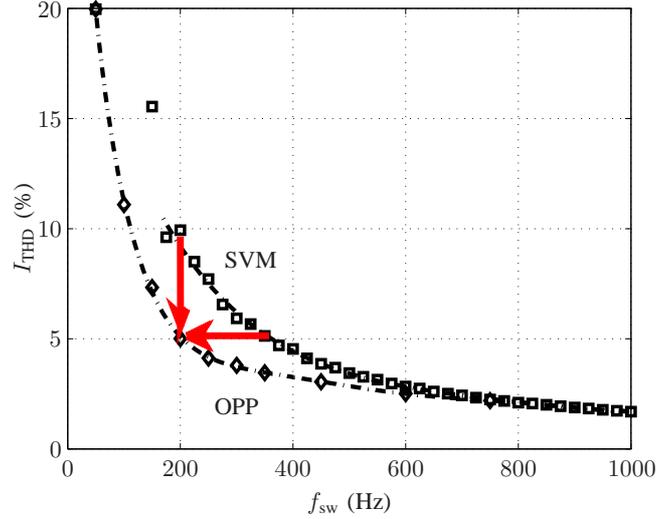


Fig. 2: Trade-off between the current THD I_{THD} and the switching frequency f_{sw} for synchronous space vector modulation (SVM) and optimized pulse patterns (OPPs)

C. Performance Criteria

The key criteria related to the control performance are the device switching frequency f_{sw} and the current THD, I_{THD} . In addition, we will also investigate the *a posteriori closed-loop* cost. It is obtained after the simulations, by evaluating the cost function J from (3) over all simulated time-steps and dividing it by the total number of time-steps k_{tot} :

$$J_{\text{cl}} = \frac{1}{k_{\text{tot}}} \sum_{\ell=0}^{k_{\text{tot}}-1} \|\mathbf{i}_{e,\alpha\beta}(\ell+1)\|_2^2 + \lambda_u \|\Delta \mathbf{u}(\ell)\|_2^2. \quad (7)$$

In summary, the closed-loop cost (7) captures the squared RMS current error plus the weighted averaged and squared switching effort over the closed-loop simulation.

D. Trade-Off between Current THD and Switching Frequency

Unavoidably, with switching power converters, there is a trade-off between the current THD I_{THD} and the switching frequency f_{sw} . It is convenient to plot these two quantities along two orthogonal axes. Figure 2 illustrates this performance trade-off for SVM and OPPs. In the figure, each square corresponds to a unique simulation with synchronous SVM. The squares are approximated using a polynomial, indicated by the dash-dotted line. Accordingly, the diamonds correspond to steady-state simulations with OPPs.

The switching frequency range between 200 and 350 Hz is of particular importance for medium-voltage power converters. As can be seen in Fig. 2, in the given range, there is scope for a significant reduction of the current THD, while maintaining the same switching frequency. For example, at $f_{\text{sw}} = 200$ Hz the current THD can be almost halved, when replacing SVM by OPPs. Conversely, the switching frequency can be drastically reduced for the same current THD. For $I_{\text{THD}} = 5\%$, for example, the switching frequency can be lowered from 350 to 200 Hz, when adopting OPPs instead of SVM. This is a reduction of 42%. Both examples are indicated by red arrows shown in Fig. 2.

At high switching frequencies, however, the performance benefit of OPPs compared to SVM tends to be marginal. For $f_{sw} > 600$ Hz and pulse numbers greater than 12, the difference is very small. Moreover, the optimization process to compute OPPs with very high pulse numbers becomes computationally demanding and the memory space required to store such patterns is significant. As a result, we have computed OPPs only up to pulse number 15, or equivalently, up to a switching frequency of 750 Hz.

With the above as a background and recalling that OPPs exhibit, to a large extent, optimal steady-state behaviour, we will quantify in the sequel the *relative* merits of MPC by normalizing the current THD to the one obtained by OPPs. Specifically, we introduce

$$\delta_{\text{THD}} = \frac{I_{\text{THD}} - I_{\text{THD,OPP}}}{I_{\text{THD,OPP}}}, \quad (8)$$

which is the *relative* current THD degradation, normalized to the current THD of OPPs and given in percent. The normalization is done with regard to the polynomial approximation of the OPPs shown in Fig. 2. For switching frequencies beyond 750 Hz, SVM is used as a baseline, since OPPs were computed only up to this frequency.

IV. STEADY-STATE PERFORMANCE

In this section, the performance of direct MPC with long prediction horizons is investigated at steady-state operating conditions, using the three-level inverter with an induction machine as a case study. We use the modified sphere decoding algorithm described in [5] to solve the optimization problem. To ensure that the drive system has settled at steady-state operating conditions, the system is first simulated over several fundamental periods without recording the results.

A. Comparison at 250 Hz Switching Frequency

Consider direct MPC with the horizon of $N = 1$, sampling interval $T_s = 125 \mu\text{s}$ and cost function (3) with the weighting factor $\lambda_u = 8.4 \cdot 10^{-3}$. This results in an average device switching frequency of $f_{sw} = 250$ Hz, which is typical for medium-voltage applications, and a current THD of $I_{\text{THD}} = 5.96\%$.

Fig. 3(a) illustrates three-phase stator current waveforms along with their (dash-dotted) references over one fundamental period. The colors blue, green and red correspond to the phases a , b and c , respectively. The evolution of the stator current is simulated with a time resolution of $25 \mu\text{s}$, based on which the spectrum of the stator current is computed with a Fourier transformation. The resulting current spectrum is shown in Fig. 3(b)

Control scheme	Control settings	$I_{s,\text{THD}}$ [%]	$T_{e,\text{THD}}$ [%]	f_{sw} [Hz]	P_{sw} [kW]
MPC	$N = 1, \lambda_u = 8.4 \cdot 10^{-3}$	5.96	4.65	250	6.05
MPC	$N = 10, \lambda_u = 8.3 \cdot 10^{-3}$	5.05	4.03	254	6.04
SVM	$f_c = 450$ Hz	7.71	5.35	250	5.73
OPP	$d = 5$	4.12	3.40	252	5.88

TABLE II: Comparison of direct MPC with SVM and an OPP in terms of the current THD $I_{s,\text{THD}}$, torque THD $T_{e,\text{THD}}$, switching frequency f_{sw} and switching losses P_{sw} . The penalty λ_u , carrier frequency f_c and pulse number d are chosen such that a switching frequency of approximately 250 Hz results

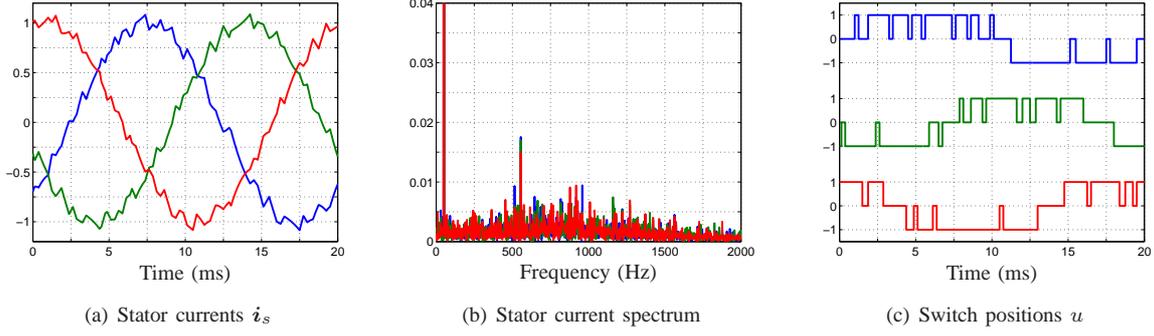


Fig. 3: Simulated waveforms for direct MPC with the horizon of $N = 1$, sampling interval $T_s = 125 \mu\text{s}$ and weighting $\lambda_u = 8.4 \cdot 10^{-3}$, at full speed and rated torque. The switching frequency is approximately 250 Hz

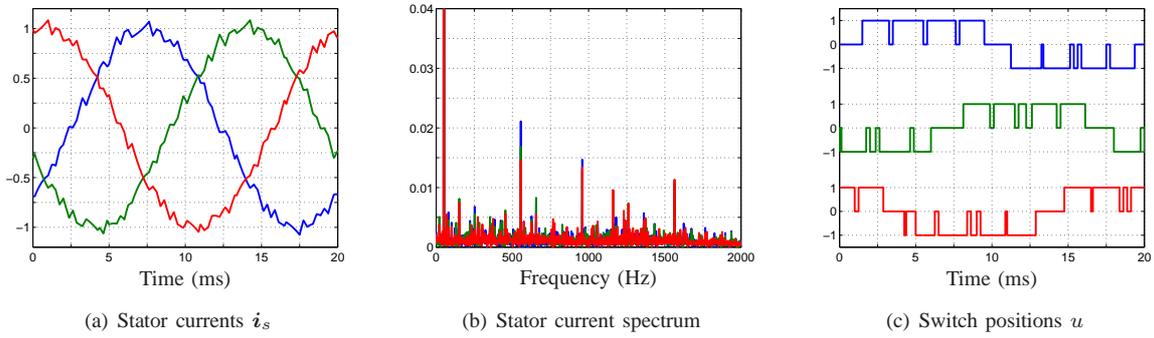


Fig. 4: Simulated waveforms for direct MPC with the horizon of $N = 10$, sampling interval $T_s = 125 \mu\text{s}$ and weighting $\lambda_u = 8.3 \cdot 10^{-3}$. The operating point and the switching frequency are the same as in Fig. 3

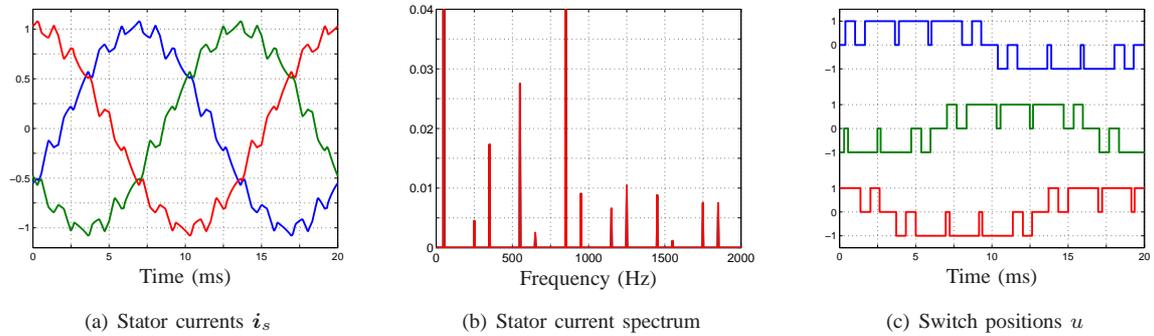


Fig. 5: Simulated waveforms for SVM with the equivalent carrier frequency $f_c = 450 \text{ Hz}$. The operating point and the switching frequency are the same as in Fig. 3

and the three-phase switching sequence is depicted in Fig. 3(c). For direct MPC, unlike PWM, a repetitive switching pattern is not enforced. As a result, the current spectrum is predominately flat without characteristic harmonics, despite a pronounced 11th harmonic.

Extending the prediction horizon to $N = 10$ reduces the current THD by about one percentage point, as stated in Table II. This first result indicates that long prediction horizons do indeed reduce the current THD, in this case by about 15%. The corresponding waveforms for the $N = 10$ case are shown in Fig. 4. It can be seen that the long horizon leads to a certain degree of repetitiveness in the switching pattern. Accordingly,

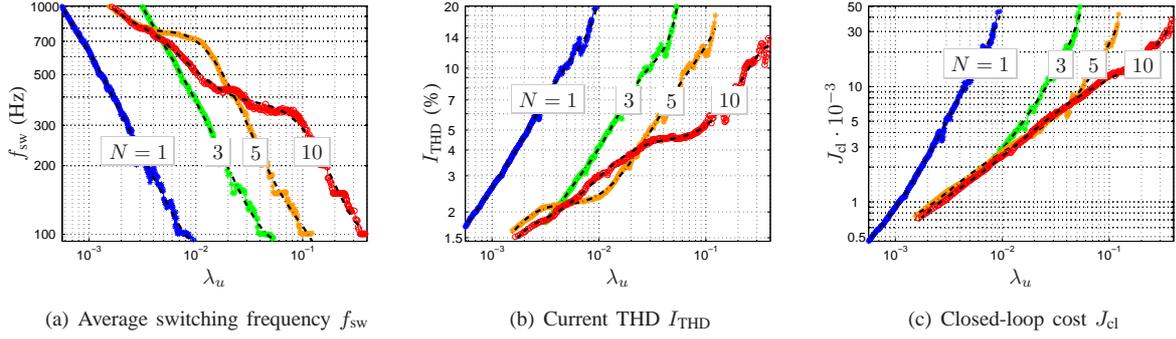


Fig. 6: Key performance criteria of MPC for the prediction horizons $N = 1, 3, 5, 10$ and sampling interval $T_s = 25 \mu s$. The switching frequency, current THD and closed-loop cost are shown as a function of the tuning parameter λ_u , using a double logarithmic scaling. The individual simulations are indicated using dots, their overall trend is approximated using dash-dotted polynomials

non-triplen odd-order harmonics are clearly identifiable in the current spectrum, such as the 11th, 13th and 19th harmonics. Indeed, the degree of repetitiveness in the switching pattern and, correspondingly, the magnitude of the discrete harmonics in the current spectrum are remarkable. It can be observed that, in this case, long prediction horizons foster a discrete current spectrum, by concentrating the harmonic power in harmonics of odd order. An analysis shows that the same applies to the triplen (common mode) voltage harmonics. The shift of some of the harmonic ripple power into common mode harmonics is one of the reasons, why direct MPC with long prediction horizons leads, in general, to a lower current THD than the horizon of one case. Moreover, longer horizons result in a shift of some of the differential-model voltage harmonics from the low-frequency range to higher frequencies, resulting in a lower current THD.

To facilitate a comparison with SVM, the corresponding waveforms of SVM are shown in Fig. 5. The equivalent carrier frequency $f_c = 450$ Hz results in the same switching frequency as for MPC, i.e. $f_{sw} = 250$ Hz. The current THD is at 7.71% significantly higher than with direct MPC, see Table II. As expected, due to the symmetry and repetitiveness of the switching pattern, SVM features a discrete current spectrum with distinctive harmonics at non-triplen and odd multiples of the fundamental frequency. Note that the 17th current harmonic has an amplitude of 0.066.

On the other hand, for the same switching frequency and the pulse number $d = 5$, an OPP leads to a current THD of 4.12%, which is approximately one percentage point lower than for direct MPC with $N = 10$.

B. Closed-Loop Cost

Next, the influence of λ_u on the switching frequency, the current THD and the closed-loop cost in (7) is investigated. For each of the horizons $N = 1, 3, 5$ and 10 and for more than 1000 different values of λ_u , ranging between 0 and 0.5, steady-state simulations were run. Focusing on switching frequencies between 100 Hz and 1 kHz, and current THDs below 20%, the results are shown in Fig. 6, using a double logarithmic scale. Each simulation corresponds to a single data point. Polynomial functions are overlaid, which approximate the individual data points. Figs. 6(a) and 6(b) suggest that, for small prediction horizons, the relationship between λ_u and the performance variables is approximately linear in the double logarithmic scale; for larger values of N , the relationship is more complicated, but still monotonic.

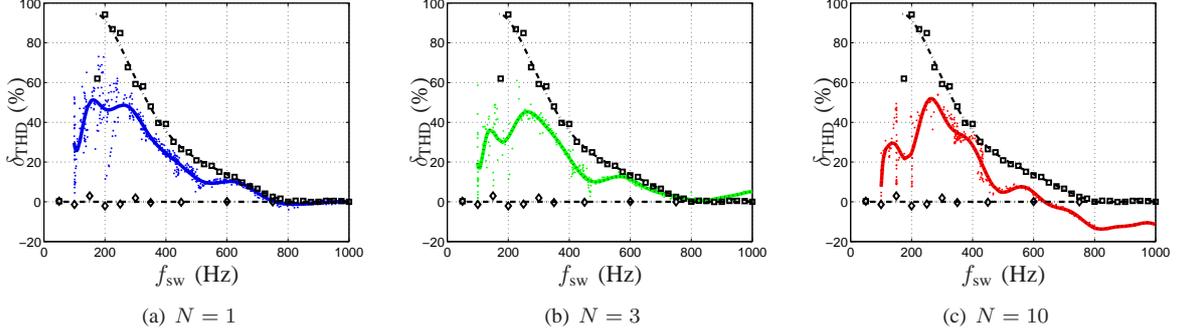


Fig. 7: Trade-off between the *relative* current THD and the switching frequency for MPC with the prediction horizons $N = 1, 3, 10$ and sampling interval $T_s = 25 \mu s$

When extending the horizon for a given λ_u , the switching frequency is increased, while the current THD is significantly reduced. This waterbed effect makes it difficult to assess from Figs. 6(a) and 6(b) the benefit long prediction horizons might have on these two key performance metrics. A more suitable measure is the a posteriori closed-loop cost, see (7), which is illustrated in Fig. 6(c). As the prediction horizon is increased the cost is clearly reduced, suggesting the use of horizons larger than one. For example, with $\lambda_u = 0.01$ and $N = 1$, we have $J_{cl} \approx 0.05$, whereas with the horizon of $N = 3$, the closed-loop cost can be reduced to $J_{cl} \approx 0.003$. This is a reduction by a factor of 17! We also note that, for this value of λ_u , the achieved a posteriori closed-loop cost is almost optimal. The benefit of long horizons on the current THD and the switching frequency is investigated in the subsequent sections, see also Figs. 7–11. For guidelines to tune λ_u please also refer to [13].

C. Relative Current THD for Sampling Interval $T_s = 25 \mu s$

Fig. 7 shows the *relative* current THDs of SVM and of MPC, as defined in (8). In this figure, the simulations referring to SVM are indicated by squares, those of OPPs are indicated with diamonds. Using the sampling interval $T_s = 25 \mu s$, hundreds of individual simulations of MPC with prediction horizons $N = 1, 3, 5$ and 10 were performed, using different weights λ_u . Specifically, λ_u was varied between 0 and 1. Each simulation corresponds to a dot in the figure. The individual simulation results were approximated by polynomials in a least-squares sense, shown in Fig. 7 as colored and dotted lines. The trend lines for the different prediction horizons are summarized in Fig. 9(a).

It can be clearly seen that using MPC with horizons larger than one reduces the current THD. In fact, for high switching frequencies above 600 Hz, the horizon of one case resembles the performance of SVM. Increasing the horizon to 10 steps reduces the current THD by about 15% compared to SVM. Interestingly enough, the OPP with pulse number $d = 15$ (and $f_{sw} = 750$ Hz) is outperformed by MPC with $N = 10$, indicating that the global minimum was not obtained when computing this particular OPP. In absolute terms, however, the differences are small, being in the range of a fraction of a percent (in terms of the absolute current THD).

For low switching frequencies between 100 and 250 Hz, the performance results are somewhat scattered. The trend lines suggest that around $f_{sw} = 200$ Hz the current THD can be reduced by about 30%, when increasing the prediction horizon from $N = 1$ to 5. Longer horizons do not appear to carry any additional performance benefit. Interestingly, for long horizons such as $N = 10$ and low switching frequencies, the switching frequency

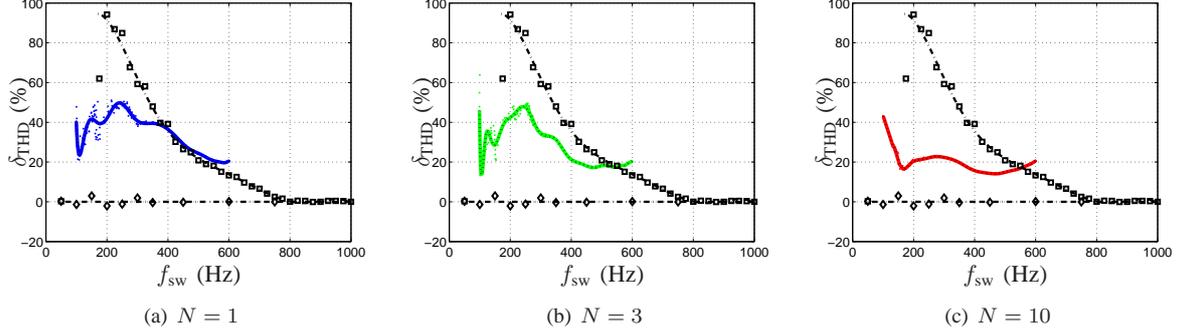


Fig. 8: Trade-off between the *relative* current THD and the switching frequency for MPC with the prediction horizons $N = 1, 3, 10$ and sampling interval $T_s = 125 \mu\text{s}$

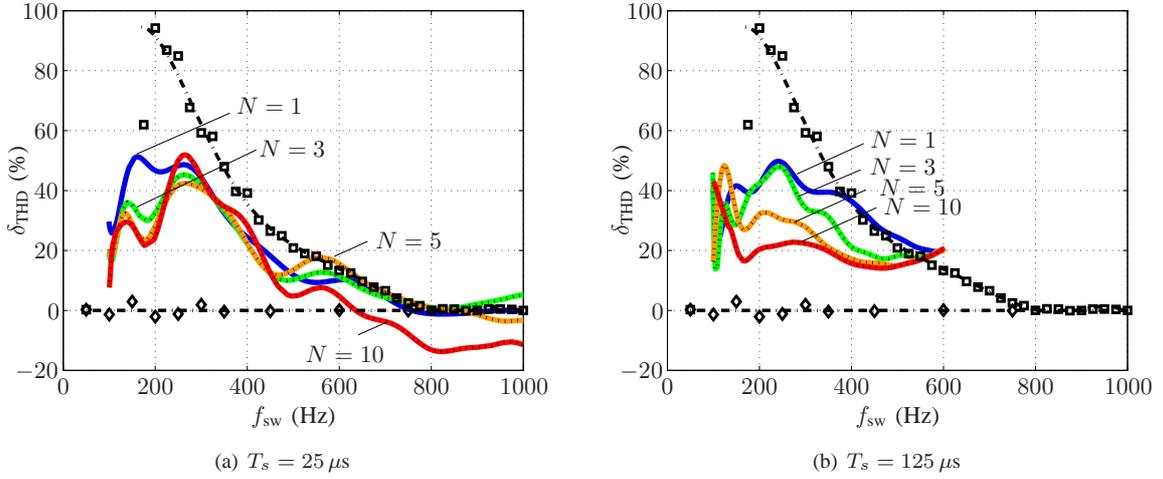


Fig. 9: Trade-off between the *relative* current THD and the switching frequency for MPC with the prediction horizons $N = 1, 3, 5, 10$ and sampling interval $T_s = 25 \mu\text{s}$ and $125 \mu\text{s}$, respectively

appears to lock into integer multiples of the fundamental frequency. This is apparent for $f_{\text{sw}} = 100, 150$ and 200 Hz . For these switching frequencies and for particular choices of λ_u , MPC almost reproduces the steady-state performance of OPPs, in terms of the current THD achieved for a given switching frequency.

D. Relative Current THD for Sampling Interval $T_s = 125 \mu\text{s}$

The simulations shown in the previous section are repeated here for a five times longer sampling interval, i.e., $T_s = 125 \mu\text{s}$. Fig. 8 shows the resulting trade-off relations, analog to those in Fig. 7. The summary plot is provided in Fig. 9(b). As in the $T_s = 25 \mu\text{s}$ case, longer prediction horizons improve the performance of MPC by lowering the current THD for a given switching frequency. This becomes particularly evident for switching frequencies between 150 and 450 Hz . In this range, MPC with $N = 10$ exhibits a relative current THD that is approximately 20% lower than that for the $N = 1$ case. When comparing Figs. 9(a) and 9(b), we note that in addition to the weight λ_u , the choice of sampling interval has a significant impact on the resulting closed-loop performance. This somewhat complicates the tuning procedure of direct MPC.

All trade-off curves converge at $f_{\text{sw}} = 600 \text{ Hz}$, which corresponds to a zero penalty on the switching effort,

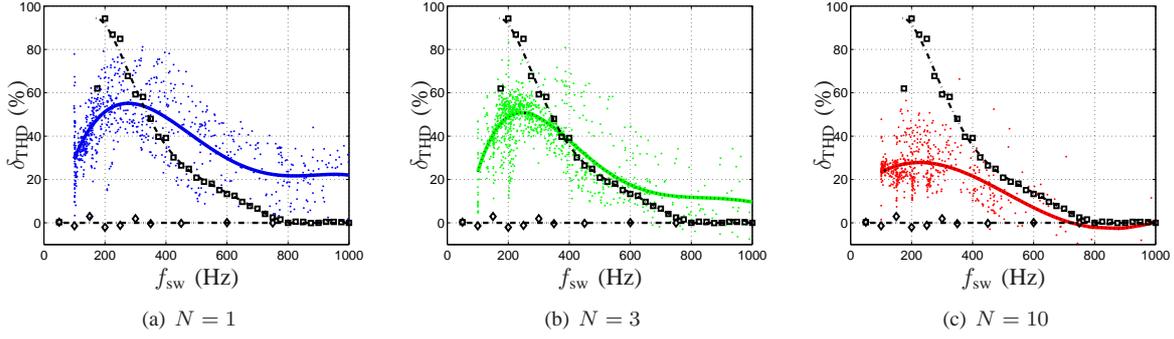


Fig. 10: Trade-off between the *relative* current THD and the switching frequency for MPC with the prediction horizons $N = 1, 3, 10$, using Monte Carlo simulations

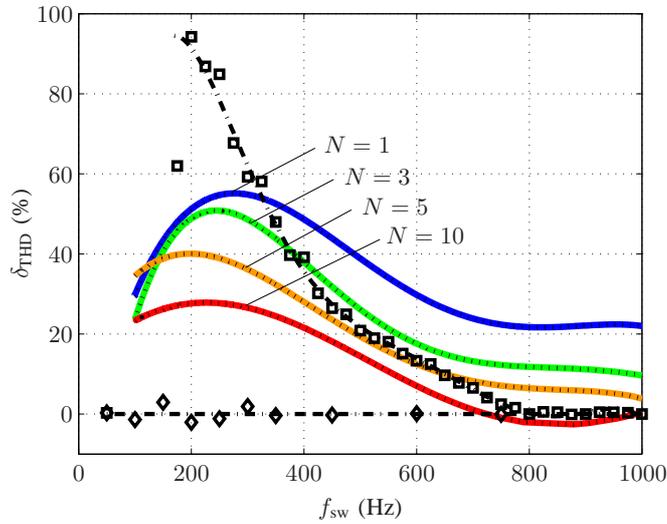


Fig. 11: Trade-off between the *relative* current THD and the switching frequency for MPC with the prediction horizons $N = 1, 3, 5, 10$, using Monte Carlo simulations

i.e., $\lambda_u = 0$. When not penalizing the switching transitions and thus only penalizing the predicted deviation of the current from its sinusoidal reference, MPC turns into a deadbeat controller. Here, the current loop effectively constitutes two first-order systems (one in the α - and another one in the β -axis) and the length of the prediction horizon ceases to have an impact on the performance of MPC. In this situation, MPC with $N = 1$ yields the same control action as MPC with $N > 1$, compare this also to the results in [14], [15].

E. Relative Current THD for Monte Carlo Simulations

We have seen that, in addition to the tuning parameter λ_u and the horizon N , the sampling interval T_s has a profound influence on the MPC performance. The reason for this is that the MPC cost function in (3) evaluates system predictions over a prediction horizon of length NT_s in time.

To derive results that take into account a variety of sampling intervals, we carried out Monte Carlo simulations with random sampling intervals and random weights. Specifically, the sampling interval was randomly chosen from the interval $T_s \in [5, 200] \mu\text{s}$, and the weight was chosen from $\lambda_u \in [0, 5]$. Moreover, the initial conditions of the drive system are random, including random initial stator currents and rotor fluxes for the induction

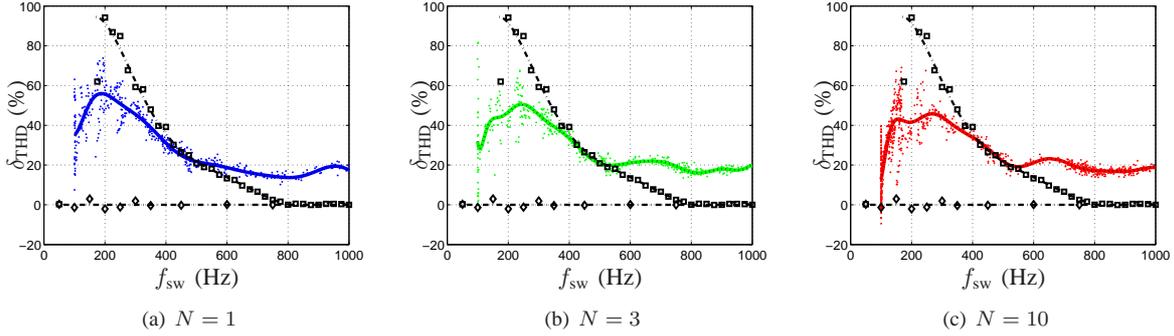


Fig. 12: Trade-off between the *relative* current THD and the switching frequency for *suboptimal* MPC with the prediction horizons $N = 1, 3, 10$ and sampling interval $T_s = 25 \mu\text{s}$, when rounding the unconstrained solution

machine, and random initial switch positions for the inverter. As previously, to ensure the simulations to be captured at steady-state, pre-simulations were run that were not recorded.

MPC with prediction horizons $N \in \{1, 3, 5, 10\}$ was considered and approximately 10^4 simulations were performed. The results are depicted in Fig. 10, where each datapoint corresponds to one closed-loop simulation. It is convenient to determine the trends, by fitting polynomials to the data points using a least-squares approach. The resulting curves are shown as dashed lines in Fig. 10, and they are summarized in Fig. 11.

It can be clearly observed that as the prediction horizon is extended, the performance of MPC is improved by reducing the current THD for a given switching frequency. For switching frequencies above 300 Hz, MPC with horizon of $N = 1$ performs worse than SVM. When enlarging the horizon from one to three, the performance improvement is most significant, whereas the performance gains level off when further increasing N to five and ten. MPC with the horizon of $N = 10$, on average, always outperforms SVM and achieves a steady-state performance close to the one of OPPs. For high switching frequencies, such as $f_{sw} = 750 \text{ Hz}$, MPC might even outperform the OPP, as discussed in Sect. IV-C.

Throughout the entire range of switching frequencies considered, MPC with $N = 10$ reduces the relative current THD by more than 20%, when compared to the popular $N = 1$ case. For very low switching frequencies, the trade-off curves for different prediction horizons converge, getting close to the point of six-step operation, i.e., fundamental frequency modulation.

V. SUBOPTIMAL MPC VIA DIRECT ROUNDING

We had seen in [5] that, in general, direct rounding of the unconstrained solution provides only suboptimal solutions. However, in some cases, the generator matrix \mathbf{H} is almost diagonal and its basis vectors are almost orthogonal. This motivates the investigation of an approximative solution, based on trivial quantization (rounding). This approach yields suboptimal solutions, but is computationally very fast, since it requires only basic matrix manipulations. Sphere decoding and branching is not required.

Specifically, instead of invoking Algorithm 1 in [5], the (suboptimal) sequence of switch positions is obtained by rounding (quantizing) the unconstrained solution componentwise to the nearest integer in the set \mathcal{U} :

$$\mathbf{U}_{\text{sub}}(k) = \text{round}_{\mathcal{U}}(\mathbf{U}_{\text{unc}}(k)). \quad (9)$$

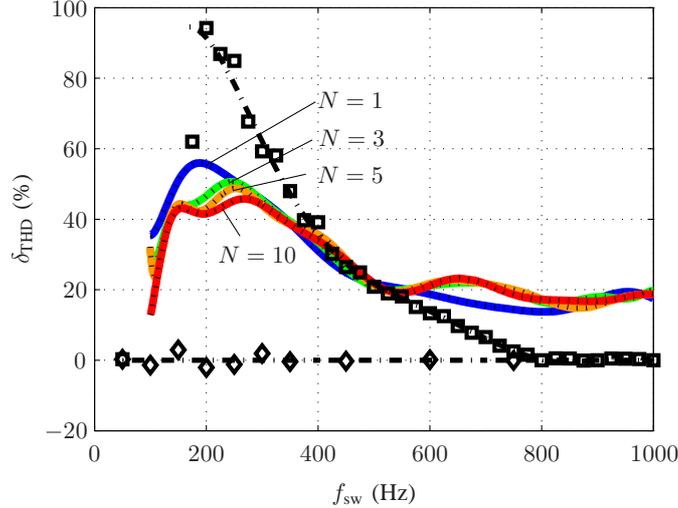


Fig. 13: Trade-off between the *relative* current THD and the switching frequency for *suboptimal* MPC with the prediction horizons $N = 1, 3, 5, 10$ and sampling interval $T_s = 25 \mu\text{s}$, when rounding the unconstrained solution

Recall that, as defined in Sect. II-C, $U_{\text{unc}}(k)$ denotes the unconstrained solution to the optimization problem (5) at time-step k .

The simulations in Sect. IV-C for (optimal) direct MPC were repeated for this suboptimal design. Using the sampling interval $T_s = 25 \mu\text{s}$, the resulting trade-off curves for suboptimal MPC are depicted in Fig. 12 for the horizons one, three and ten. The trendlines were fitted as before and are shown separately in Fig. 13.

For switching frequencies below 300 Hz, the weight λ_u is large and the diagonal terms dominate over the other terms in the generator matrix \mathbf{H} . As a result, the componentwise quantization in (9) yields solutions close to the optimal one. This can be seen, when comparing Figs. 9(a) and 13 with each other. For $N = 1$, suboptimal MPC exhibits a performance that is very similar to the one of optimal MPC. Longer horizons improve the performance of suboptimal MPC (9), but to a lesser degree than for the optimal case (6).

High switching frequencies are the result of small λ_u and generator matrices that are correspondingly less orthogonal. Using the trivial quantization in (9) for switching frequencies above 300 Hz leads to suboptimal solutions that are clearly inferior, with the $N = 1$ case being about 15% worse than the optimal solution. Moreover, extending the horizon appears to be of very little benefit, if at all. As λ_u is decreased and the switching frequency is increased, the relative deterioration of suboptimal MPC becomes more prominent. The absolute performance loss, in terms of current THD for a given switching frequency, is, however, very small.

VI. PERFORMANCE DURING TRANSIENTS

One of the major benefits of direct MPC is its very fast dynamical behavior during transients. Consider MPC with a horizon of one, the sampling interval $T_s = 25 \mu\text{s}$ and the weight $\lambda_u = 2.55 \cdot 10^{-3}$. At nominal speed, reference torque steps of magnitude one are imposed, see Fig. 14(a). The steps on the torque reference are translated into steps in the current reference, shown as dash-dotted lines in Fig. 14(b). The corresponding switching pattern is shown in Fig. 14(c), with the switching frequency being $f_{\text{sw}} = 252 \text{ Hz}$.

When switching from rated to zero torque, the voltage applied to the machine is momentarily inverted, leading to an extremely short settling time of 0.35 ms. On the other hand, the torque step from zero to one is with

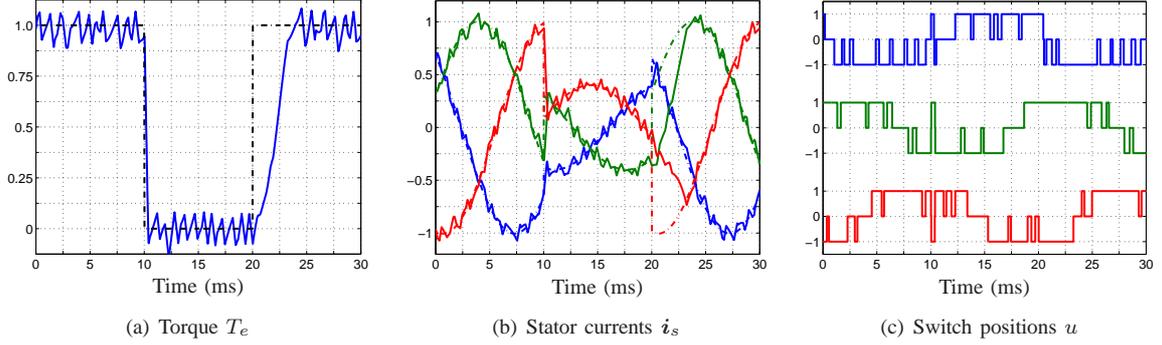


Fig. 14: Reference torque steps for MPC with the horizon of $N = 1$ at nominal speed

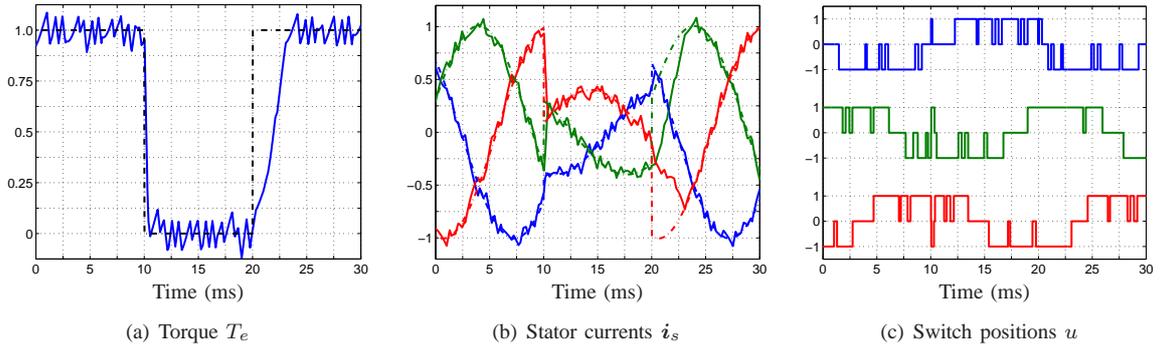


Fig. 15: Reference torque steps for MPC with the horizon of $N = 10$ at nominal speed

4 ms significantly slower. This is due to the small voltage margin available, which results from the machine operating at nominal speed. Nevertheless, as can be seen in Fig. 14(b), the currents are regulated as quickly as possible to their new references. Note that, due to the constraint (5c), switching between -1 and 1 is inhibited, and switching is performed via an intermediate zero switch position, which is applied for T_s .

Fig. 15 shows the corresponding step responses for MPC with horizon of ten. The settling times are nearly identical to the horizon of one case. The weight $\lambda_u = 120 \cdot 10^{-3}$ was chosen, which results in the same switching frequency as above, i.e. $f_{sw} = 250$ Hz.

When operating at 50% speed and applying the same torque steps as before, the torque settling times are 0.5 ms for the step down and 1.1 ms for the step up case, both for MPC with horizon of one and for horizon of ten. We conclude that during transients the dynamical performance of direct MPC is effectively limited only by the available voltage, regardless of the length of the prediction horizon. In particular, long horizons do not slow down the dynamic response of MPC.

VII. COMPUTATIONAL BURDEN

Next, we analyze the computational burden of the proposed modified sphere decoding algorithm and compare it with the one of the exhaustive search algorithm described in Sect. III-D in [5]. Through this section, the sampling interval $T_s = 25 \mu s$ is used. Different prediction horizons are investigated. The weight λ_u is chosen such that approximately the same switching frequency of $f_{sw} = 300$ Hz is obtained, regardless of the prediction horizon. As a measure of the computational burden, the number of switching sequences, which are investigated

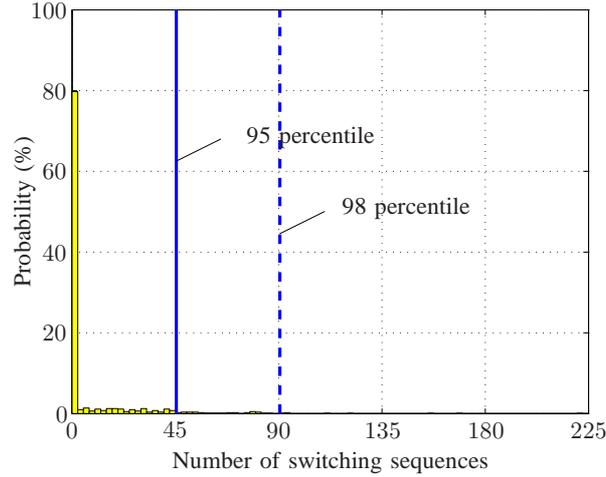


Fig. 16: Histogram of the number of switching sequences investigated by the modified sphere decoding algorithm, when considering the horizon of $N = 10$

by the algorithm at each time-step when computing the optimum, is considered. Over multiple fundamental periods, the average as well as the maximal number of sequences is monitored and summarized in Table III.

When using sphere decoding for the horizon of one case, the table shows that on average 1.18 switching sequences need to be considered by the algorithm at every time-step. The empirical upper bound on the number of sequences is five. This implies that, by choosing the initial radius of the sphere based on an educated guess (see (31) in [5]), the sphere is sufficiently tight. Specifically, in the vast majority of the cases, the sphere is perfectly tight, in the sense that out of all the admissible switching sequences only *one* is located within the sphere.

This is in stark contrast to exhaustive search. Here, depending on the optimal switch position obtained at the previous time-step, $\mathbf{u}_{\text{opt}}(k-1)$, and in accordance with the switching constraint, up to 18 sequences need to be investigated, with the average being 11.8. We conclude that for $N = 1$ and the three-level inverter at hand, sphere decoding is at least four times faster than exhaustive search; on average, it is 10 times faster. These numbers are reinforced by the fact that, for each switching sequence to be examined, the modified sphere decoding algorithm tends to require less computations than exhaustive search, as described in Sect. V-B in [5].

As the prediction horizon is increased, the computational burden associated with sphere decoding initially grows slowly, despite being exponential, whilst exhaustive search becomes computationally intractable for

Prediction horizon N	Sphere decoding		Exhaustive search	
	avg.	max.	avg.	max.
1	1.18	5	11.8	18
2	1.39	8	171	343
3	1.72	14	2350	4910
5	2.54	35	467'000	970'000
10	8.10	220		

TABLE III: Average and maximal number of switching sequences that need to be considered by the sphere decoding and exhaustive search algorithms to obtain the optimal result, depending on the length of the prediction horizon

horizons of five or more. Using sphere decoding, the optimization problem for direct MPC with long horizons such as $N = 10$ can be solved relatively quickly, with the upper bound on the number of switching sequences to be investigated being 220. Fig. 16 depicts the histogram of the average number of switching sequences, which need to be explored at each time-step, when using a horizon of ten steps. The histogram is highly concentrated at one and exhibits a long, yet very flat tail. It can be seen that, with sphere decoding, in 80% of the cases, the optimization problem can be solved by exploring only one switching sequence. The 95 and 98 percentiles are shown as straight and dashed lines, respectively, indicating that in 95% of the cases, less than 45 switching sequences need to be considered.

VIII. DISCUSSION AND CONCLUSIONS

In this final section, the proposed MPC algorithm, its performance during steady-state and transient operation, the choice of the cost function and its computational complexity are discussed and conclusions are provided, focusing on the three-level inverter drive system used as a case study.

A. Performance at Steady-State Operating Conditions

When assessing the steady-state performance of a current controller, the two key performance metrics are the current THD and the switching effort. Since the switching frequency is easy to quantify, the latter is usually used as a measure for the switching effort, rather than the switching losses, which might be more meaningful. OPPs are typically considered to yield the lowest achievable current THD for a given switching frequency, while SVM, particularly for low switching frequencies, entails a significantly higher current THD.

When tracking the current reference in MPC and directly setting the converter switch positions without the use of a modulator, a horizon of $N = 1$ is almost universally used [2]–[4]. Alas, the penalty on the switching effort is often omitted in the literature, resulting in a deadbeat control scheme. Such schemes are well-known to be highly sensitive to noise in the measurements and estimates. Adding a penalty on the switching effort not only reduces the switching frequency, but also lessens the sensitivity to such noise. By enlarging the prediction horizon, this sensitivity is further reduced, as shown for example in the related MPC formulation in [16].

For the low switching frequencies typically used in medium-voltage applications, the horizon of one case tends to improve on SVM, by reducing the current THD for a given switching frequency or vice versa. For higher switching frequencies, however, MPC with $N = 1$ performs similarly to SVM or worse. The use of long prediction horizons entails a significant reduction in the current THD. For a three-level inverter, for example, direct MPC with the horizon of $N = 10$ leads to a 20% reduction, when compared to the horizon of one case and the same switching frequency. Indeed, for long prediction horizons, the resulting steady-state performance in terms of current THD per switching frequency gets close to the one of OPPs. When considering multi-level inverters with a higher number of voltage levels, the benefit of long horizons is expected to be even more pronounced.

Not only the weight λ_u , but also the sampling interval T_s has a profound impact on the closed-loop characteristic of MPC. Even though this second degree of freedom complicates the tuning procedure, it can be exploited to one's advantage. Specifically, it is important to achieve a long prediction *interval* in time. If a low switching frequency is desired, it is beneficial to use a fairly long sampling interval such as $T_s = 125 \mu\text{s}$,

even though this reduces the granularity of switching. For high switching frequencies, a high granularity is important, requiring a high sampling frequency and thus a short sampling interval, such as $T_s = 25 \mu\text{s}$.

B. Performance during Transients

During transients, MPC achieves an excellent dynamic performance, similar to the one of deadbeat control, see also [2], [10], [17]. When applying torque steps, the settling time is limited in effect only by the available voltage. If required, MPC temporarily inverts the voltage applied to the load, in order to achieve as short a transient as possible. For the case investigated, the horizon length has no impact on the settling time, with long horizons resulting in the same transient performance as short ones.

Moreover, the transient performance of direct MPC is by far superior to the one typically achieved with OPPs, since traditionally it has only been possible to use OPPs in a modulator driven by a very slow control loop, see e.g. [18], albeit some recent improvements, see [16], [19], [20] and the references therein.

C. Cost Function

Horizons longer than one significantly reduce the closed-loop cost (7), when compared to the $N = 1$ case. For very long prediction horizons, however, when further increasing the horizon, the incremental cost reduction becomes very small and ceases at some point. This is a general characteristic of MPC, see e.g. [21]–[25], and can be seen in Fig. 6(c). The larger the weight on the switching effort, the later this levelling off occurs, indicating that long horizons are particularly beneficial when switching is expensive and the switching frequency is low.

The cost function consists of two terms. The first term is the RMS current error, which corresponds to the current THD, while the second term is represented by the squared switching effort. The latter is a direct measure of the switching frequency. Both terms are penalized in the cost function, and the trade-off between the two is adjusted by the weight λ_u . When increasing the length of the prediction horizon for a given λ_u , a drastic reduction of the closed-loop cost can be observed, but only minor reductions in the current THD and switching frequency are achieved. In particular, long horizons shift the trade-off point along the trade-off curve, while only marginally improving it.

As an alternative, in model predictive direct current control, this shift is avoided by fixing one of the two quantities [26]. More specifically, the width of the current bounds determines the current THD, and the cost function captures the switching effort, which is to be minimized [10], [27], [28]. Fixing one of the two performance metrics, while minimizing the other one, rather than aiming at minimizing both, merits further investigations. Apart from that, the effect of final state weighting is worth exploring, since it allows one to approximate infinite horizon problems, see [22], [29].

D. Control Objectives

It is conceivable to directly minimize the switching losses rather than the switching frequency, as proposed in [27] and [30]. To achieve this, one might replace the constant scalar weight λ_u by a time-varying and diagonal 3x3 matrix, with each term corresponding to a phase-specific weight. These weights are adjusted online according to the phase current. Specifically, the phases with high currents feature large weights, while

phases with low currents have accordingly a small weight. As a result, it is expected that the switching transitions are shifted from phases with high currents to phases with lower currents, thus reducing the average switching losses.

E. Computational Complexity

When considering multi-level converters with more than three levels, the computational complexity increases, in the worst case, exponentially with a large base. Nevertheless, our empirical results for the modified sphere decoding algorithm suggest that the average computational burden is effectively independent of the number of inverter levels, since the search for the optimal switching sequence is restricted to a sphere centered at the unconstrained solution. The size of the sphere is independent of the number of levels.

Therefore, this algorithm appears to be particularly suited to multi-level converter topologies with a very large number of levels. Even for a three-level converter, as shown in this paper, the modified sphere decoding algorithm provides significant computational savings, when compared to exhaustive search, which is commonly used in the power electronics community. Notably, even for the horizon of one case, an average reduction of the computational burden by one order of magnitude can be observed, making sphere decoding an attractive alternative to solve the optimization problem of direct MPC also in cases where long horizons are not strictly required.

ACKNOWLEDGMENT

This research was partially supported under Australian Research Council's Discovery Projects funding scheme (project number DP 110103074). The first author gratefully acknowledges a research grant from ABB Corporate Research, Switzerland, while being with the University of Auckland, New Zealand.

REFERENCES

- [1] T. Geyer and D.E. Quevedo. Multistep direct model predictive control for power electronics—Part 2: Analysis. In *Proc. IEEE Energy Convers. Congr. Expo.*, Denver, CO, USA, Sep. 2013.
- [2] 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.*, 55(12):4312–4324, Dec. 2008.
- [3] J. Rodríguez and P. Cortés. *Predictive control of power converters and electrical drives*. Wiley, 2012.
- [4] J. Rodríguez, M. P. Kazmierkowski, J. R. Espinoza, P. Zanchetta, H. Abu-Rub, H. A. Young, and C. A. Rojas. State of the art of finite control set model predictive control in power electronics. *IEEE Trans. Ind. Informatics*, 9(2):1003–1016, May 2013.
- [5] T. Geyer and D.E. Quevedo. Multistep finite control set model predictive control for power electronics. *IEEE Trans. Power Electron.*, 2014. in print.
- [6] B. Hassibi and H. Vikalo. On the sphere-decoding algorithm I. Expected complexity. *IEEE Trans. Sign. Process.*, 53(8):2806–2818, Aug. 2005.
- [7] H. S. Patel and R. G. Hoft. Generalized techniques of harmonic elimination and voltage control in thyristor inverters: Part I—Harmonic elimination. *IEEE Trans. Ind. Appl.*, 9(3):310–317, May/June. 1973.
- [8] G. S. Buja. Optimum output waveforms in PWM inverters. *IEEE Trans. Ind. Appl.*, 16(6):830–836, Nov./Dec. 1980.
- [9] U. Fincke and M. Pohst. Improved methods for calculating vectors of short length in a lattice, including a complexity analysis. *Math. Comput.*, 44(170):463–471, Apr. 1985.
- [10] T. Geyer, G. Papafotiou, and M. Morari. Model predictive direct torque control—Part I: Concept, algorithm and analysis. *IEEE Trans. Ind. Electron.*, 56(6):1894–1905, Jun. 2009.
- [11] G. Papafotiou, J. Kley, K. G. Papadopoulos, P. Böhren, and M. Morari. Model predictive direct torque control—Part II: Implementation and experimental evaluation. *IEEE Trans. Ind. Electron.*, 56(6):1906–1915, Jun. 2009.

- [12] B. P. McGrath, D. G. Holmes, and T. Lipo. Optimized space vector switching sequences for multilevel inverters. *IEEE Trans. Power Electron.*, 18(6):1293–1301, Nov. 2003.
- [13] P. Cortés, S. Kouro, B. La Rocca, R. Vargas, J. Rodríguez, J. I. León, S. Vazquez, and L. G. Franquelo. Guidelines for weighting factors design in model predictive control of power converters and drives. In *Proc. IEEE Int. Conf. Ind. Technol.*, 2009.
- [14] D. E. Quevedo, C. Müller, and G. C. Goodwin. Conditions for optimality of naïve quantized finite horizon control. *Int. J. Contr.*, 80(5):706–720, May 2007.
- [15] C. Müller, D. E. Quevedo, and G. C. Goodwin. How good is quantized model predictive control with horizon one? *IEEE Trans. Automat. Contr.*, 56(11):2623–2638, Nov. 2011.
- [16] T. Geyer, N. Oikonomou, G. Papafotiou, and F. Kieferndorf. Model predictive pulse pattern control. *IEEE Trans. Ind. Appl.*, 48(2):663–676, Mar./Apr. 2012.
- [17] 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*, 17(17):1572–1589, Nov. 2007.
- [18] J. Holtz and B. Beyer. Fast current trajectory tracking control based on synchronous optimal pulsewidth modulation. *IEEE Trans. Ind. Appl.*, 31(5):1110–1120, Sep./Oct. 1995.
- [19] J. Holtz and N. Oikonomou. Synchronous optimal pulsewidth modulation and stator flux trajectory control for medium-voltage drives. *IEEE Trans. Ind. Appl.*, 43(2):600–608, Mar./Apr. 2007.
- [20] N. Oikonomou, C. Gutscher, P. Karamanakos, F. Kieferndorf, and T. Geyer. Model predictive pulse pattern control for the five-level active neutral-point-clamped inverter. *IEEE Trans. Ind. Appl.*, 49(6):2583–2592, Dec. 2013.
- [21] D. E. Quevedo, H. Bölskei, and G. C. Goodwin. Quantization of filter bank frame expansions through moving horizon optimization. *IEEE Trans. Signal Processing*, 57(2):503–515, Feb. 2009.
- [22] D. E. Quevedo and G. C. Goodwin. Multistep optimal analog-to-digital conversion. *IEEE Trans. Circuits Syst. I*, 52(4):503–515, Mar. 2005.
- [23] D. E. Quevedo and G. C. Goodwin. Moving horizon design of discrete coefficient FIR filters. *IEEE Trans. Signal Processing*, 53(6):2262–2267, Jun. 2005.
- [24] D. E. Quevedo, G. C. Goodwin, and J. A. De Doná. Multistep detector for linear ISI-channels incorporating degrees of belief in past estimates. *IEEE Trans. Commun.*, 55(11):2092–2103, Nov. 2007.
- [25] L. Grüne and A. Rantzer. On the infinite horizon performance of receding horizon controllers. *IEEE Trans. Automat. Contr.*, 53(9):2100–2111, Oct. 2008.
- [26] T. Geyer. Model predictive direct current control: Formulation of the stator current bounds and the concept of the switching horizon. *IEEE Ind. Appl. Mag.*, 18(2):47–59, Mar./Apr. 2012.
- [27] T. Geyer. Generalized model predictive direct torque control: Long prediction horizons and minimization of switching losses. In *Proc. IEEE Conf. Decision Control*, pages 6799–6804, Shanghai, China, Dec. 2009.
- [28] T. Geyer. *Low Complexity Model Predictive Control in Power Electronics and Power Systems*. PhD thesis, Autom. Control Lab. ETH Zurich, 2005.
- [29] D. E. Quevedo, G. C. Goodwin, and J. A. De Doná. Finite constraint set receding horizon quadratic control. *Int. J. Robust Nonlin. Contr.*, 14(4):355–377, Mar. 2004.
- [30] S. Mastellone, G. Papafotiou, and E. Liakos. Model predictive direct torque control for MV drives with LC filters. In *Proc. Eur. Power Electron. Conf.*, pages 1–10, Barcelona, Spain, Sep. 2009.