2013 European Control Conference (ECC)
July 17-19, 2013, Zürich, Switzerland.

Reference Frame Re-alignment for Vector Control of the Brushless Doubly-Fed Machine

Alexander Broekhof, Richard McMahon and Jan Maciejowski
Department of Engineering, University of Cambridge, Cambridge, UK, CB2 1PZ, Email: ab940@cam.ac.uk

***Abstract*— Vector control provides stability and performance when applied to the brushless doubly-fed machine, however cross-coupling effects can arise between inputs and outputs. To address these effects, a procedure is proposed to both visualize and minimize the cross-coupling by means of steady-state mapping and a re-alignment of the dq reference frame. With this method implemented, gain-response tests show improved decoupling across the operating region.**

I. Introduction

Despite the application of numerous control methods to the brushless doubly-fed machine (BDFM), including scalar control [1], direct torque control [2], [3], feedback linearization and phase-angle control [4], vector control (VC) has emerged as a dominant strategy. Development of VC was conducted in [5]–[8] and has been successfully implemented in a number of test-rigs, including in a 250 kW generator [9].

The success of VC can be attributed to its ability to stabilize the machine while also tracking demands in active and reactive power (Q_1). These two quantities are of importance to one of the BDFM's main applications, wind power generation, where active power is used to regulate the torque output, and therefore the rotor angular velocity (ω_r) of the generator. Additionally, the ability to control reactive power is crucial for ensuring a favorable power factor.

Typically, VC is developed so that, in the dq reference frame, the converter-connected stator winding voltage phasor components v_{2d} and v_{2q} can be used to control Q_1 and ω_r, respectively. However, the dual stator windings of the BDFM result in additional complexities in implementing VC, specifically cross-coupling between the control inputs and outputs. These relationships are expressed in table I. For

TABLE I
Input-output relationships

	Q_1	ω_r
v_{2d}	Direct-coupling	Cross-coupling
v_{2q}	Cross-coupling	Direct-coupling

example, if a step change in ω_r is demanded, the controller would apply an appropriate change to v_{2q}. However, due to the non-diagonal elements in table I, this change in v_{2q} will also affect Q_1.

In [6], [8], the cross-coupling is addressed through the use of feed-forward signals, but the stability of such a solution is not discussed. In [7], the cross-coupling effects are treated as disturbances, which places additional demands on the closed-loop controller.

In this paper, the cross-coupling effects are analyzed directly through input-output mapping. A parameter-independent decoupling method is then given using a reference frame re-alignment. This method is demonstrated in simulation, with experimental results provided to confirm its practical applicability.

II. Vector control for the BDFM

The BDFM (fig. 1) is comprised of dual stator windings with unequal pole-pair numbers to prevent cross-coupling [10]. The primary winding, typically attached directly to the grid, is referred to as the power winding (PW) and has p_1 pole-pairs. The secondary winding, fed through a partially-rated voltage converter, is known as the control winding (CW) and has p_2 pole-pairs. A specially designed rotor, typically consisting of $p_1 + p_2$ sets of nested loops, is able to couple with the flux waveforms of both stator windings, thereby generating torque.

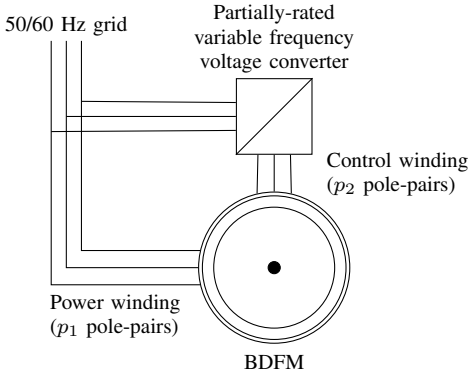

Fig. 1. Overview of the BDFM

Vector control is applied to the BDFM in a similar manner as is used for the standard induction machine (IM). By expressing the vector-form model equations in a particular reference frame using a dq transform, input-output relationships can be simplified. Loops can then be closed around these input-output relationships through the use of any number of control methods, with proportional-integral (PI) controllers being a popular choice.

A. Reference frames

With the two pole-pair numbers of the BDFM, angles can be expressed in terms of either p_1 or p_2 pole-pair

distributions. Quantities in the BDFM can then be expressed in one of four reference frames:

1) $\alpha\beta_1$: p_1 pole-pair reference frame fixed to the PW stator winding
2) $\alpha\beta_2$: p_2 pole-pair reference frame fixed to the CW stator winding. Any mechanical displacement between the PW and CW windings is represented by γ (radians)
3) $\alpha\beta_{r_1}$: p_1 pole-pair reference frame fixed to the rotor, which is located at θ_r mechanical radians
4) $\alpha\beta_{r_2}$: p_2 pole-pair reference frame, also fixed to the rotor

These four reference frames are expressed geometrically in fig. 2.

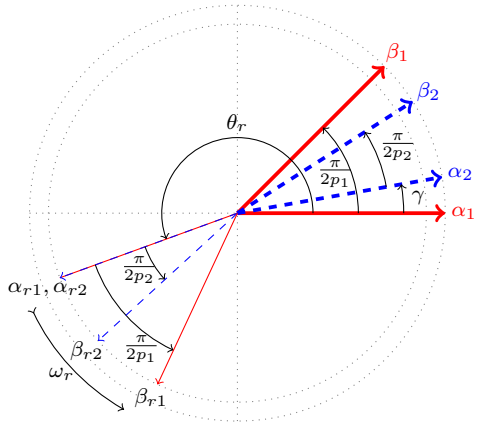

Fig. 2. Geometric representation of the four BDFM reference frames. All angles are in mechanical radians. Diagram is accurate for a BDFM with $p_1 = 2, p_2 = 4$

In [11], a method is given for expressing all values in terms of a single reference frame, which in this case is $\alpha\beta_1$. By analyzing fig. 2, a number of relationships are immediately apparent:

$$\vec{x}_{\alpha\beta_2} = e^{jp_2(\theta_r-\gamma)}\vec{x}_{\alpha\beta_{r_2}} \tag{1}$$

$$\vec{x}_{\alpha\beta_1} = e^{jp_1\theta_r}\vec{x}_{\alpha\beta_{r_1}} \tag{2}$$

where $\vec{x}$ represents an arbitrary phasor expressed in the subscripted reference frame. The final relationship needed to enable transformation between reference frames is:

$$\vec{x}_{\alpha\beta_{r_2}} = \vec{x}^*_{\alpha\beta_{r_1}} \tag{3}$$

where * represents the complex conjugate. This relationship emerges from the indirect cross-coupling requirement of the two stators, the full derivation of which can be found in [11]. Combining eqs. (1) to (3) gives:

$$\vec{x}_{\alpha\beta_1} = e^{j((p_1+p_2)\theta_r-p_2\gamma)}\vec{x}^*_{\alpha\beta_2} \tag{4}$$

With eqs. (1) to (4) it is possible to express any quantity in terms of the $\alpha\beta_1$ reference frame. The phasor can be rotated by any arbitrary angle θ_a (expressed in terms of a p_1 pole-pair distribution) using the dq transform:

$$\vec{x}_{dq} = \vec{x}_{\alpha\beta_1}e^{-j\theta_a} \tag{5}$$

If the dq reference frame is aligned with the PW flux phasor, $\vec{\psi}_1$, then the overall transformation from the CW frame to the dq frame is:

$$\vec{x}_{dq} = e^{j((p_1+p_2)\theta_r-p_2\gamma-\angle\vec{\psi}_1)}\vec{x}^*_{\alpha\beta_2} \tag{6}$$

with the inverse transformation given by:

$$\vec{x}_{\alpha\beta_2} = e^{j((p_1+p_2)\theta_r-p_2\gamma-\angle\vec{\psi}_1)}\vec{x}^*_{dq} \tag{7}$$

Recall that the BDFM operates in synchronism when the rotor angular velocity satisfies the following equality [10]:

$$\omega_r = \frac{\omega_1+\omega_2}{p_1+p_2} \tag{8}$$

where ω_2 is the angular velocity of the CW voltage phasor. It can then be shown that when a constant dq phasor is applied to the CW, the argument of eq. (7) ensures that ω_2 satisfies eq. (8). This property is desirable, as the ability to use constant inputs simplifies the control design.

Assuming that the PW is connected directly to the grid and that grid angular velocity, ω_1, is constant, then the flux waveform produced by the PW will also have angular velocity ω_1, i.e. $\frac{\mathrm{d}\angle\vec{\psi}_1}{\mathrm{d}t} = \omega_1$. As p_1, p_2 and γ are constant, differentiating the argument of eq. (7) produces:

$$\tfrac{\mathrm{d}}{\mathrm{d}t}((p_1+p_2)\theta_r - p_2\gamma - \angle\vec{\psi}_1) = (p_1+p_2)\omega_r - \omega_1 \tag{9}$$

Therefore, with the input of a constant dq-frame CW voltage phasor, eq. (7) always produces a physical CW voltage which satisfies eq. (8).

B. PW flux vector control

With these transformations in place, a simple, unified-reference-frame model of the BDFM is given by the following equations [11]:

$$\vec{v}_1 = R_1\vec{i}_1 + s\vec{\Psi}_1 + j\omega_1\vec{\Psi}_1 \tag{10}$$

$$\vec{\Psi}_1 = L_1\vec{i}_1 + L_{1r}\vec{i}_r \tag{11}$$

$$\vec{v}_2 = R_2\vec{i}_2 + s\vec{\Psi}_2 + j(\omega_1 - (p_1+p_2)\omega_r)\vec{\Psi}_2 \tag{12}$$

$$\vec{\Psi}_1 = L_2\vec{i}_2 + L_{2r}\vec{i}_r \tag{13}$$

$$0 = R_r\vec{i}_r + s\vec{\Psi}_r + j(\omega_1 - p_1\omega_r)\vec{\Psi}_r \tag{14}$$

$$\vec{\Psi}_1 = L_{1r}\vec{i}_1 + L_{2r}\vec{i}_2 + L_r\vec{i}_r \tag{15}$$

where L_x is self-inductance, L_{xr} is stator-rotor mutual inductance, R is resistance, $\vec{i}$ is current, subscripts $1, 2, r$ represent PW, CW and rotor quantities and for a generic quantity $\vec{x}$:

$$\vec{x} = x_d + jx_q \tag{16}$$

In [7], [8], using eqs. (10) to (15), transfer functions relating v_{2d}, v_{2q} and ω_r, Q_1 are derived:

$$\begin{bmatrix} Q_1 \\ \omega_r \end{bmatrix} = \begin{bmatrix} G_{1,1}(s) & G_{1,2}(s) \\ G_{2,1}(s) & G_{2,2}(s) \end{bmatrix} \begin{bmatrix} v_{2d} \\ v_{2q} \end{bmatrix} \tag{17}$$

Employing a number of simplifying assumptions to alleviate the complexity of the derivation, it is further shown that $G_{1,1}(s)$ and $G_{2,2}(s)$ are the dominant transfer functions. However the cross-coupling terms, $G_{2,1}(s)$ and $G_{1,2}(s)$,

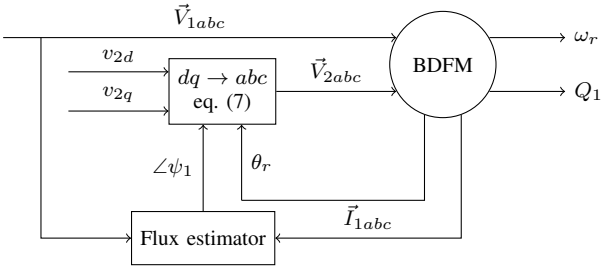

Fig. 3. Schematic of test setup used to produce steady-state maps with V_{2dq} inputs

have been shown to detract from the overall effectiveness of the strategy [12].

Instead of attempting to analyze the transfer functions, a simpler and more accurate approach is to directly obtain the gains from the BDFM, either through simulation or experimentation. The effects of these cross-coupling terms can be observed by creating a map from inputs v_{2d} and v_{2q} to outputs Q_1 and ω_r. This is accomplished by applying a range of constant v_{2d} and v_{2q} values to the BDFM and observing the resulting steady-state reactive powers and rotor angular velocities. This map then gives an indication of the DC gains in the d and q directions.

Simulations were carried out using the coupled-circuit model presented in [10] with parameters based on the BDFM prototype detailed in table II (There are too many parameters required for this model to be able to reproduce them concisely in this paper). Contour plots of these values are shown in figs. 4a and 4b.

For v_{2d} and v_{2q} to be considered 'decoupled', their axes must both be orthogonal to the contour lines of the output they are meant to control. This also demands that the contours themselves be orthogonal. The polarity of the dq axes with relation to the slope of the maps is not important, as this can always be corrected by reversing signs.

The input-output maps have been super-imposed in fig. 5 to demonstrate the relative orthogonality of the contours. There exist regions in fig. 5, particularly around the origin, where the contours are indeed orthogonal. To ensure decoupled control in this region, the dq axes can simply be re-aligned so that they are orthogonal to the appropriate contours.

While the Q_1 contours (fig. 4b) display a good level of uniformity, the 'fanning' pattern in fig. 4a means that as the operating point moves away from the origin, the v_{2q} axis will no longer be orthogonal to the ω_r contours. Therefore, a compromise must be made when re-aligning the axes so that they are maximally orthogonal, on average, across the operating range.

Applying this re-alignment is equivalent to multiplying the input vector by a rotation matrix:

$$R = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \tag{18}$$

Equation (18) can be considered as a decoupling network

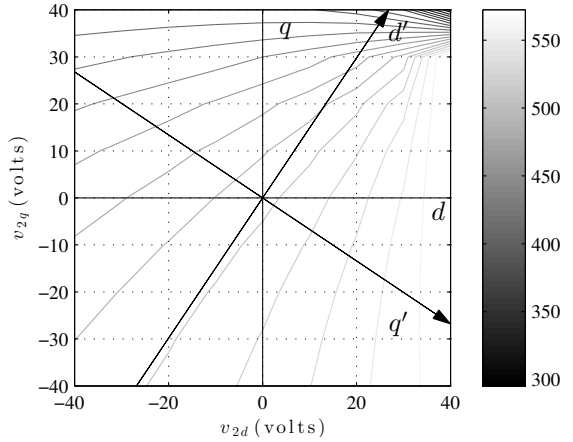

(a) Rotor angular velocity (RPM)

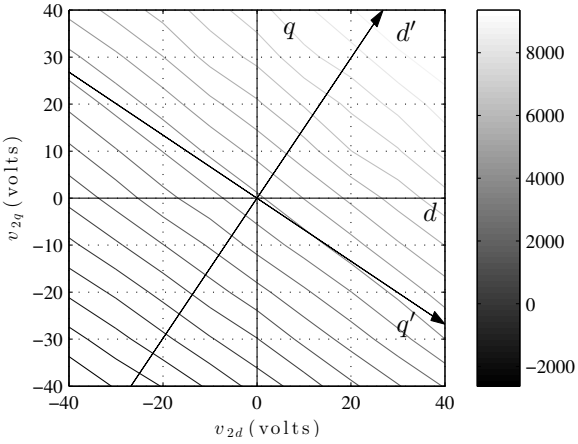

(b) PW reactive power (var)

Fig. 4. Steady-state maps with zero load torque (simulated). Arrow lines are re-aligned dq axes.

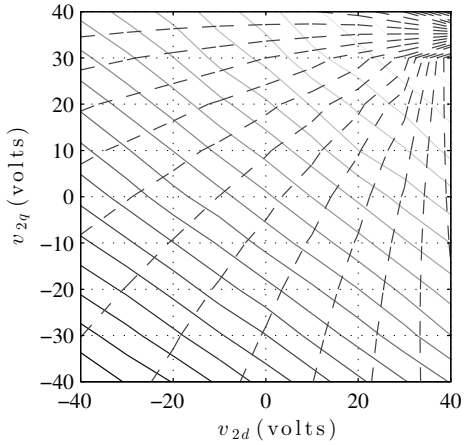

Fig. 5. Superimposed input-output maps. Dashed contours represent rotor angular velocity and solid contours represent PW reactive power

with constant gain transfer functions. However, since the dq transform is essential for control of the BDFM, the decoupling method will be discussed in terms of the reference frame re-alignment.

III. Re-alignment of the dq axes to maximize relative orthogonality

The functions from $v_{2dq} \to Q_1, \omega_r$ are defined as:

$$Q_1 = f(v_{2d}, v_{2q}) \tag{19}$$

$$\omega_r = g(v_{2d}, v_{2q}) \tag{20}$$

Assuming that f and g are differentiable on a bounded operating region, V_{2dq}, with area A then the gradient in that region is defined as:

$$\nabla f = \frac{\partial f}{\partial v_{2d}} \begin{bmatrix} 1 \\ 0 \end{bmatrix} + \frac{\partial f}{\partial v_{2q}} \begin{bmatrix} 0 \\ 1 \end{bmatrix} \tag{21}$$

∇g is defined similarly. The re-aligned axes are defined by the orthogonal unit vectors

$$\hat{i} = \begin{bmatrix} \cos(\vartheta) \\ \sin(\vartheta) \end{bmatrix}, \hat{j} = \begin{bmatrix} -\sin(\vartheta) \\ \cos(\vartheta) \end{bmatrix} \tag{22}$$

Orthogonality between two vectors is defined to be a zero inner product, that is:

$$\hat{x} \cdot \hat{y} = x_d y_d + x_q y_q = 0 \tag{23}$$

The objective is then to find an angle, ϑ, such that the average inner products $\nabla f \cdot \hat{i}$ and $\nabla g \cdot \hat{j}$ over the region is minimised. As the magnitude of the gradient vectors is not relevant, the inner products are normalised by dividing by these magnitudes. This objective can be expressed analytically as:

$$\min_{\vartheta \in [0,2\pi]} \frac{1}{A} \iint\limits_{V_{2dq}} \frac{\nabla f \cdot \hat{i}}{\|\nabla f\|} \, \mathrm{d}v_{2d} \, \mathrm{d}v_{2q} + \frac{1}{A} \iint\limits_{V_{2dq}} \frac{\nabla g \cdot \hat{j}}{\|\nabla g\|} \, \mathrm{d}v_{2d} \, \mathrm{d}v_{2q} \tag{24}$$

Using eq. (23), eq. (24) is expanded to :

$$\min_{\vartheta \in [0,2\pi]} \frac{1}{A} \iint\limits_{V_{2dq}} \frac{\nabla f_d \cos(\vartheta) + \nabla f_q \sin(\vartheta)}{\|\nabla f\|} + \frac{-\nabla g_d \sin(\vartheta) + \nabla g_q \cos(\vartheta)}{\|\nabla g\|} \, \mathrm{d}v_{2d} \, \mathrm{d}v_{2q} \tag{25}$$

As $\frac{x_d}{\|\hat{x}\|} = \cos(\angle \hat{x})$ and $\frac{x_q}{\|\hat{x}\|} = \sin(\angle \hat{x})$, eq. (25) is equivalent to:

$$\min_{\vartheta \in [0,2\pi]} \frac{\cos(\vartheta)}{A} \iint\limits_{V_{2dq}} \cos(\angle \nabla f) + \sin(\angle \nabla g) \, \mathrm{d}v_{2d} \, \mathrm{d}v_{2q} + \frac{\sin(\vartheta)}{A} \iint\limits_{V_{2dq}} \sin(\angle \nabla f) - \cos(\angle \nabla g) \, \mathrm{d}v_{2d} \, \mathrm{d}v_{2q} \tag{26}$$

Evaluating the integrals and dividing by A reduces the integrands to their mean values, denoted by overlines. Expressing eq. (26) in vector form produces:

$$\min_{\vartheta \in [0,2\pi]} \begin{bmatrix} \cos(\vartheta) & \sin(\vartheta) \end{bmatrix} \left\{ \begin{bmatrix} \overline{\cos(\angle \nabla f)} \\ \overline{\sin(\angle \nabla f)} \end{bmatrix} + \begin{bmatrix} \overline{\sin(\angle \nabla g)} \\ \overline{-\cos(\angle \nabla g)} \end{bmatrix} \right\} \tag{27}$$

The second term in eq. (27) is sum of the average ∇f angle and the average ∇g angle rotated by $-\frac{\pi}{2}$. Therefore, this expression can be minimized by setting $\hat{\vartheta}$ so that the first term is orthogonal to the second.

Though closed formed expressions do not exist for f and g, eq. (27) can easily accommodate simulated or experimental data using numerical gradients. This is calculated by finding the slopes between adjacent points in the d and q directions at each sampled point in the desired operating region.

Additionally, due to its simple form, eq. (27) can be adapted for a wide range of uses. For example, if one output is more important than the other, it can be made 'more orthogonal' to the rotated axis by applying a weighting function. Likewise, if certain areas of the operating region need to be emphasized, say for example if the machine rarely deviates substantially from the origin, then a probability distribution can be used to calculate the gradient angle averages.

IV. Results

Equation (27) is used to re-align the v_{2dq} axes in fig. 4 and are plotted on the same figure. It should be noted that, though the Q_1 contours are nearly parallel, the re-aligned axes do not follow their slope. This is because the average relative orthogonality between the contours is affected by the curvature of the ω_r map.

A step response is also used to observe the effects of the reference frame re-alignment,. In both the standard and re-aligned reference frames, unit steps are applied individually to v_{2d} and v_{2q} and the gains are measured at the ω_r and Q_1 outputs. This process is repeated on a grid of $\vec{v}_{2dq}$ operating points, producing fig. 6.

Were the input-output relationships strongly decoupled, the gain plots would exhibit large separation between the gain surfaces of the individual inputs across the entire operating region. For example, since it is desired that v_{2q} controls ω_r, its gain plot (white surface) should be 'above' that of v_{2d} (gray surface) for the entirety of the operating region.

However, in figs. 6a and 6c the desired gain separations are either negligible or only exist of portions of the operating region, indicating poor decoupling of the input-output relationships. This inconsistency reduces the effectiveness of any control method, as input-output assumptions no longer hold.

As seen in figs. 6b and 6d, the desired gain characteristics are increased significantly with re-alignment. In fig. 6b the gain from v_{2q} to ω_r is greater than that of v_{2d} for the majority of the operating region. Similarly, in fig. 6d, re-alignment has improved the gain from v_{2d} to Q_1, where previously (fig. 6c) the gains from v_{2d} and v_{2q} are nearly the same.

A. Experimental results

To confirm the applicability of this method to the physical system, input-output maps (fig. 7) were created using a prototype 180-frame BDFM, detailed in table II.

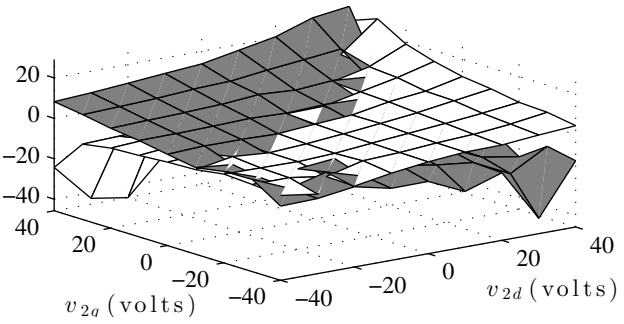

(a) ω_r (RPM)

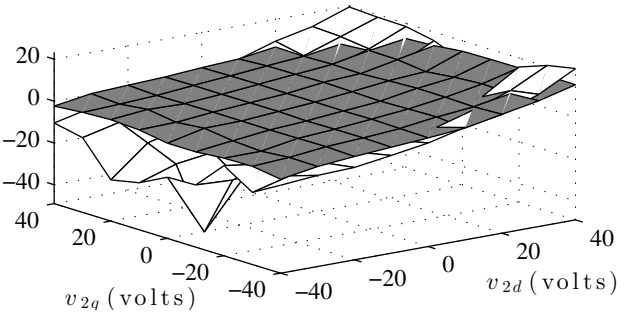

(b) ω_r re-aligned (RPM)

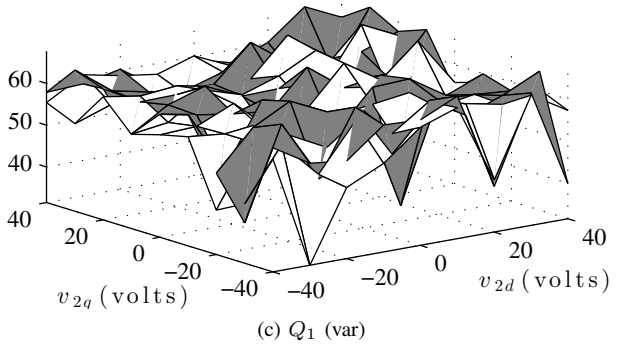

(c) Q_1 (var)

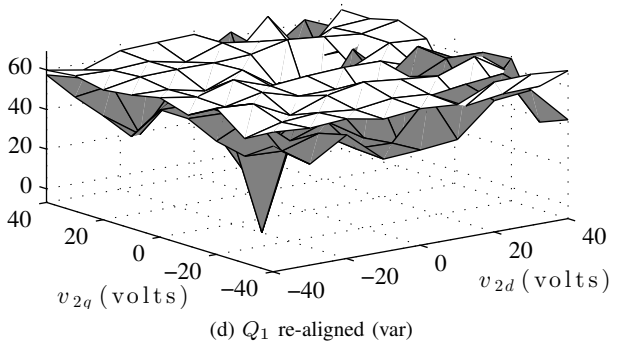

(d) Q_1 re-aligned (var)

Fig. 6. Gains from unit step inputs (dB). White: v_{2d}, Gray: v_{2q}

TABLE II
BDFM PROTOTYPE SPECIFICATIONS

Frame size	D180
p_1	2
p_2	4
Stator slots	48
Rotor slots	36
PW rated voltage	240 V (at 50 Hz)
CW rated voltage	240 V (at 50 Hz)
PW rated current	7 A
CW rated current	7 A
Rated torque	100 Nm
Rotor design	Nested-loop: 6, 3-loop nests

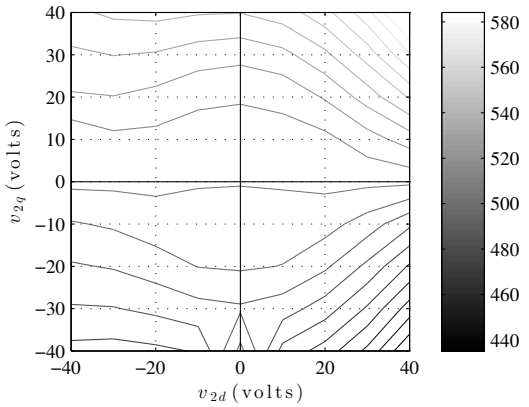

(a) Rotor angular velocity (RPM)

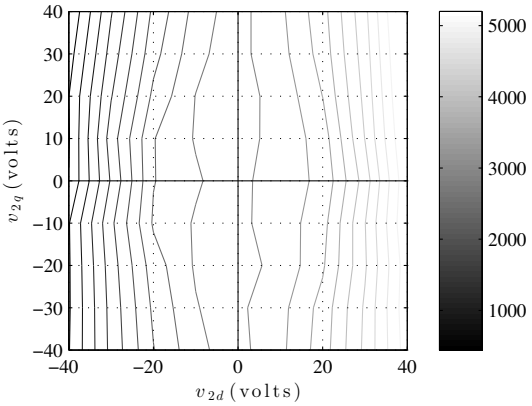

(b) PW reactive power (var)

Fig. 7. Experimental steady-state maps with zero load torque

In fig. 7 the v_{2dq} axes have been re-aligned to maximize relative orthogonality using the proposed method. In comparing the experimental results of fig. 7 with the simulated fig. 4, a noticeable 'bulge' occurs around the origin in the experimental case. This indicates that the curvature is flattened and BDFM exhibits lower gain behavior in this region. This behavior is likely due to issues with the inverter producing low-amplitude voltage waveforms.

However, the simulated and experimental maps are otherwise in good agreement. Figure 7a exhibits the 'fanning' which is present in simulation, while the contours in fig. 7b are predominantly parallel. Most importantly, the relative orthogonality between fig. 7a and fig. 7b is similar to that of fig. 4a and fig. 4b. This indicates that, with proper reference frame alignment, Q_1 and ω_r can be decoupled to a good extent.

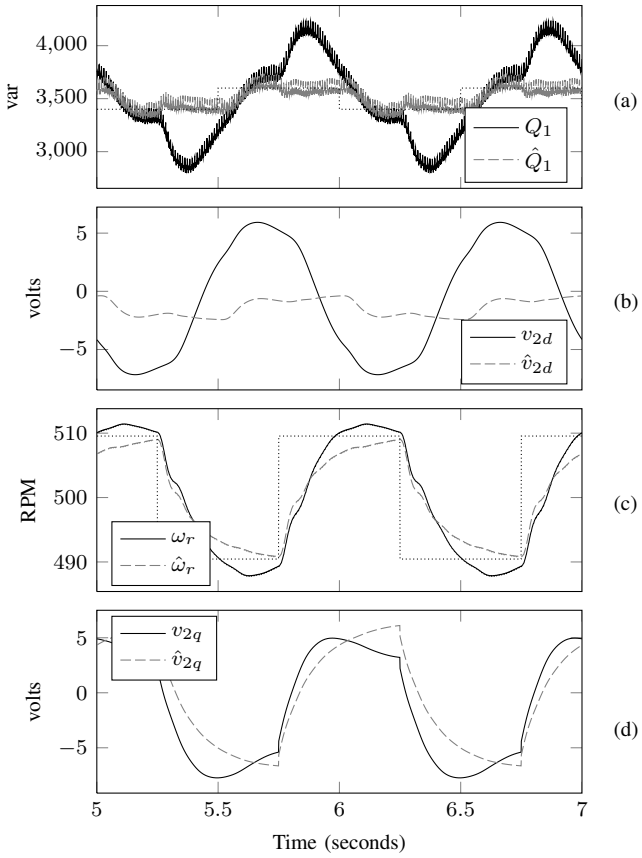

Fig. 8. Effects of reference frame re-alignment on PI control of Q_1, ω_r, where $\hat{x}$ indicates re-aligned quantity. Gains for v_{2d} are $K_P = 0.001, K_I = 0.1$. Gains for v_{2q} are $K_P = 0.5, K_I = 50$.

B. Control implementation

To ensure the efficacy and stability of the re-alignment, simulations are conducted using PI control of Q_1, ω_r. The controllers are given square wave references with period of 1 second and are 90^o out of phase. The results are seen in fig. 8.

Without re-alignment, it is clear that Q_1 is strongly coupled to both v_{2d} and v_{2q}, which leads to poor tracking of the reference signals and excessive control action. After re-alignment, the cross-coupling is drastically reduced, leading to improved performance on both outputs.

V. Conclusion

To ensure acceptable performance of VC applied to the BDFM, the cross-coupling effects must be not only understood, but also minimized to the furthest possible extent. In the derivation of the reference frame re-alignment presented in this paper, both demands have been addressed. The cross-coupling is easily visualizable in steady-state maps figs. 4 and 7, while the actual re-alignment maximizes the gain separation across the operating range.

This procedure is important in practical implementation as an observer will not always reliably determine the dq reference frame angle, due to low-speed integration errors, parameter estimation variance and other sources of error. Therefore, performing a reference frame re-alignment not only ensures optimal performance from the controller, but also precludes any potential observer errors.

However, while ensuring correct alignment of the reference frame alignment is important, during standard operation of the BDFM alignment will be affected by parameter variance, such as changes in resistance due to temperature fluctuations. Therefore, a topic of future research is to develop an adaptive controller to perform online re-alignment of the reference frame.

References

[1] D. Zhou, R. Spee, and A. K. Wallace, "Laboratory control implementations for doubly-fed machines," in *Proceedings of the IECON '93 International Conference on Industrial Electronics, Control, and Instrumentation*, vol. 2. IECON, November 1993, pp. 1181–1186.

[2] I. Sarasola, J. Poza, M. Rodriguez, and G. Abad, "Predictive direct torque control for brushless doubly fed machine with reduced torque ripple at constant switching frequency," in *IEEE International Symposium on Industrial Electronics*, June 2007, pp. 1074–1079.

[3] W. R. Brassfield, R. Spee, and T. G. Habetler, "Direct torque control for brushless doubly-fed machines," *IEEE Transactions on Industry Applications*, vol. 32, no. 5, p. 1098, September 1996.

[4] P. C. Roberts, T. Flack, J. Maciejowski, and R. McMahon, "Two stabilising control strategies for the brushless doubly-fed machine (bdfm)," in *International Conference on Power Electronics, Machines and Drives*, no. 487, June 2002, pp. 341–346.

[5] D. Zhou, R. Spee, G. Alexander, and A. Wallace, "A simplified method for dynamic control of brushless doubly-fed machines," in *Industrial Electronics, Control, and Instrumentation, 1996., Proceedings of the 1996 IEEE IECON 22nd International Conference on*, vol. 2, aug 1996, pp. 946 –951 vol.2.

[6] K. Protsenko and X. Dewei, "Modeling and control of brushless doubly-fed induction generators in wind energy applications," *IEEE Transactions on Power Electronics*, vol. 23, no. 3, pp. 1191–1197, May 2008.

[7] S. Shao, E. Abdi, F. Barati, and R. McMahon, "Stator-flux-oriented vector control for brushless doubly fed induction generator," *IEEE Transactions on Industrial Electronics*, vol. 56, no. 10, October 2009.

[8] J. Poza, E. Oyarbide, I. Sarasola, and M. Rodriguez, "Vector control design and experimental evaluation for the brushless doubly fed machine," *IET Electric Power Applications*, vol. 3, no. 4, pp. 247–256, July 2009.

[9] R. A. McMahon, E. Abdi, P. D. Malliband, S. Shao, M. E. Mathekga, and P. J. Tavner, "Design and testing of a 250 kw medium-speed brushless dfig," in *Power Electronics, Machines and Drives (PEMD 2012), 6th IET International Conference on*, march 2012, pp. 1 –6.

[10] P. C. Roberts, "A study of brushless doubly-fed (induction) machines," Ph.D. dissertation, University of Cambridge, September 2004.

[11] J. Poza, E. Oyarbide, D. Roye, and M. Rodriguez, "Unified reference frame dq model of the brushless doubly fed machine," in *IET Proceedings Electric Power Applications*, vol. 153, no. 5. IET, September 2006.

[12] A. Broekhof, "Analysis of stator-flux-oriented vector control for the brushless doubly-fed machine," Master's thesis, University of Cambridge, 2012.