

A Doubly-Fed Permanent Magnet Generator for Wind Turbines

by

Andrew Joseph Thomas

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2004

© 2004 Andrew Joseph Thomas. All rights reserved.

The author hereby grants to MIT permission to reproduce and distribute publicly
paper and electronic copies of this thesis document in whole or in part.

Author


Department of Electrical Engineering and Computer Science
May 25, 2004

Certified by

James L. Kirtley, Jr.
Professor of Electrical Engineering
Thesis Supervisor

Accepted by

Arthur C. Smith
Chairman, Department Committee on Graduate Students



BARKER

A Doubly-Fed Permanent Magnet Generator for Wind Turbines

by

Andrew Joseph Thomas

Submitted to the Department of Electrical Engineering and Computer Science
on May 25, 2004, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

Optimum extraction of energy from a wind turbine requires that turbine speed vary with wind speed. Existing solutions to produce constant-frequency electrical output under wind-speed variations are undesirable due to complexity, cost, inefficiency or reliability issues. We propose a novel variation of a doubly-fed induction generator which aims to improve power density and simplify construction. Our design is a doubly-fed, dual-rotor, axial-flux, permanent-magnet machine. Progress in the construction of a prototype is described. Analysis of the steady state and dynamic behavior of the machine is detailed, and a control algorithm developed therefrom.

Thesis Supervisor: James L. Kirtley, Jr.

Title: Professor of Electrical Engineering

Acknowledgements

I would like to thank Professor Kirtley for his assistance, advice, and supervision over the course of the project. Funding for the project was generously provided by a grant from the National Renewable Energy Laboratory. I would also like to thank fellow graduate student Shiv Reddy for help with the design and construction of the prototype. The assistance of Colin Wu and Saagar Gupta in various construction operations has been critical to the progress of the project. Steve Englebretson, Jack Holloway, Riad Wahby, Mariano Alvira, Rob Cox, and Tim Abbott provided extra sets of hands when occasionally required for construction. Finally, I would like to thank Kate Baker, Cat Miller, and Reid Barton for their patience while editing this document.

Contents

1	Introduction	11
1.1	Prior Work	11
1.1.1	Induction Generators	12
1.1.2	Synchronous Generators	12
1.1.3	Doubly-fed Generators	13
1.2	Our Solution	14
1.3	Overview	16
2	Construction	17
2.1	Structural Elements	17
2.1.1	Power Rotor Plates	18
2.1.2	Magnet Rotors	19
2.1.3	End Plates	20
2.1.4	Housing	20
2.2	Magnetic Cores	21
2.3	Windings	23
2.3.1	Stator	24
2.3.2	Rotor	25
2.4	Other Parts	28
2.4.1	Shafts	29
2.4.2	Slip-Rings	31
2.4.3	Ventilation	32
2.4.4	Mounting and Testbed	33
2.5	Assembly	34

2.5.1	Power Rotors	34
2.5.2	Stator Mounting	35
2.5.3	Final Assembly	36
3	Analysis and Modeling	39
3.1	Steady-State Behavior	39
3.2	Dynamics	41
3.2.1	Power Rotor	42
3.2.2	Stator	45
3.2.3	Overall Behavior	46
3.2.4	Model Simplification	47
4	Control	49
4.1	Algorithm	49
4.2	Electronics	50
4.2.1	Sensors	52
4.2.2	Controller	53
4.2.3	Power Converter	54
5	Conclusion	59
5.1	Completed Tasks	59
5.2	Further Work	60
5.3	Summary	61
A	Structural Part Drawings	63
B	Machine Photos	71

List of Figures

1-1	Machine Cutaway View	15
2-1	Shape of Permanent Magnets	19
2-2	Assembled “Skeletal” Housing	22
2-3	Stator B-H Characteristics for 100 V _{rms} 60 Hz Sinusoidal Drive	25
2-4	Schematic Representation of Rotor End-Turn Configuration	27
2-5	Histogram of Rotor Winding Inductances, First Pass	28
2-6	Histogram of Rotor Winding inductances, Second Pass	29
2-7	Driveshaft Plan View	30
2-8	Form of Stator Support Blocks	36
4-1	Machine Controller Block Diagram	51
4-2	Bi-directional 3-phase AC/DC/AC converter	55
4-3	Sub-synchronous Rotor Drive Cricuit	55
4-4	Hysteresis Phase-Current Control Circuit	56
A-1	Power Rotor Plate Drawing	64
A-2	Magnet Rotor Plate Drawing	65
A-3	Endplate Drawing	66
A-4	Original Housing Drawing, Page 1	67
A-5	Original Housing Drawing, Page 2	68
A-6	New Housing Mounting Ring Drawing	69
A-7	New Housing Stringer Drawing	70
B-1	Completed Stator Core	71
B-2	Rotor Coil Winding Apparatus	72

B-3	Wound and Partially Fiberglassed Stator	73
B-4	Assembled Power Rotor, without Windings	74
B-5	Power Rotor with Windings	75

Chapter 1

Introduction

The world increasingly relies on electrical power to function, in areas ranging from transportation to home comfort. Most of the electricity in the world, and in the United States in particular, is generated from ever-dwindling reserves of petroleum. This presents a large problem both in the increasing scarcity of fuel and in the environmental effects of its use. Hence, investigation into renewable energy sources seems prudent in both financial and environmental terms.

Wind power is one such renewable source which shows fair promise of becoming economically viable in the near future. This will be particularly true if it can be made to more efficiently capture energy, thereby reducing cost for a given capacity. The intention of this thesis work, and the project of which it is a part, is to develop a novel generator design to that end. We also aim to produce and test a prototype to validate our analysis. The generator is to be a doubly-fed permanent magnet machine, constructed based on the design work in [1].

1.1 Prior Work

Extraction of the maximum energy attainable from the wind for a given turbine requires a turbine speed which is proportional to the speed of the wind. Thus, if one is to extract energy with maximum efficiency, one must allow the turbine to change speeds as the wind changes speeds.

1.1.1 Induction Generators

Many installed generators ignore the problem completely, running an induction generator at essentially constant speed. Although in principle an induction generator could operate over a large range of speeds (all super-synchronous) by allowing large slip, this would require that the rotor resistance be large. However, large rotor resistance translates to much dissipation in the rotor, and consequent inefficiency.

1.1.2 Synchronous Generators

Synchronous generators do not have the rotor resistance problem that induction machines do. They also allow more control of the reactive power generation. However, synchronous generators have an output electrical frequency proportional to the shaft speed, so some additional consideration is needed before interfacing the generator to a power grid.

To some extent, the problem that a given turbine has an optimal speed which changes with wind speed can be solved by varying the turbine geometry with wind speed. Ideally, doing so allows the optimal speed for the turbine to remain always at the synchronous speed of the generator even as the wind speed varies. Although work is progressing in this area, variable-geometry turbines are largely restricted to low-power applications, due to the immense stresses to which such a mechanism would be subjected in a large turbine (50 m in diameter, perhaps). Better results can still be obtained if the turbine speed is allowed to vary with wind speed.

One possible solution is to use a continuously variable transmission (CVT) between the turbine shaft and the generator shaft, and to use a synchronous generator. A CVT, appropriately controlled, allows the ratio of the two shaft speeds to be varied continuously over some range. Unfortunately, CVTs are mechanically complex and, in this application, are subject to substantial wear. Consequently, the devices are prone to mechanical failure and potentially have problems with low control-loop bandwidth.

The other obvious solution is to use a synchronous generator, but allow the shaft speed of the generator to vary with wind speed, as described in [2]. Obviously, the output electrical frequency then varies with the wind speed, and a power electronics package (generally a back-to-back PWM converter for low power installations, or a cycloconverter for high power) converts the frequency to match the constant line frequency. Such a scheme also

offers complete control of real and reactive power components. The main drawback to such an indirect grid connection is the high cost of power electronics, which for large installations may have to handle 1 MVA or more. Furthermore, the power converter must be carefully designed to be highly efficient, or the advantage of increased extraction efficiency will be offset by losses in the converter. Naturally, more efficient converters are yet more expensive. Finally, the power quality of the output needs to be considered. Requiring that the converter also be low-distortion adds further to the cost.

1.1.3 Doubly-fed Generators

The doubly-fed induction generator (DFIG) offers a convenient electromechanical solution, which should be reasonably cheap and not prone to wear problems. The operation of a simple DFIG is analyzed in [3]. The essential principle of operation is as follows: by injecting an AC current into the field windings of the rotor, the flux wave in the machine can be made to move with respect to the rotor. Seen another way, in the reference frame of the rotor, the flux wave will then rotate with a speed proportional to the electrical frequency of the AC current injected. From the reference frame of the stator, the apparent motion of the flux wave is then due to the frequency of the AC rotor current added to the mechanical motion of the rotor. Thus, the stator electrical frequency can be controlled independently of shaft speed by varying the frequency of the AC current. Specifically, then, with an appropriate control loop, the stator electrical frequency can be kept constant over variations in shaft speed. The power level of the rotor excitation will typically be an order of magnitude less than the rating of the machine, so the cost of the necessary electronics is similarly reduced over that required by an indirect grid connection scheme.

The DFIG has a number of other useful properties which also recommend it as a solution. Firstly, if the power electronics used to drive the rotor of the machine are designed correctly, perhaps using a back-to-back PWM inverter scheme as in [4], the system has complete control over output power (both real and reactive), and may operate in both sub- and super-synchronous regimes. Super-synchronous operation is particularly attractive, as it allows power to be extracted from the rotor as well, effectively increasing the capacity of the machine in as much as capacity is limited by armature heating. Secondly, it has been shown (e.g. in [5]) that DFIG wind turbines are more stable under grid fluctuations than standard induction generators. As wind power takes a larger share of the generation

load, stability will become a bigger issue, because local instabilities can potentially disable distribution on a very large scale.

1.2 Our Solution

Our proposed solution builds upon the apparent advantages of the DFIG by adding permanent magnets to the design. Modern permanent magnet materials offer savings in efficiency and space when used appropriately in electric machines [6]. Additionally, the use of permanent magnets simplifies construction, as they allow for much larger air gaps, which loosens the manufacturing tolerance. Furthermore, the magnetic cores in the machine can be simplified by eliminating the winding slots. Both of these considerations also greatly simplify the cooling of the windings in the machine.

Adding permanent magnets to a DFIG presents engineering challenges. The first is the question of how the permanent magnets, which inherently provide DC excitation, can be used in a generator in which the excitation needs to be AC. This is solved by placing the magnets in a separate, freely spinning rotor between the wound rotor and the stator. Essentially, the magnet rotor is dragged along at the stator electrical frequency and serves to drive flux across the air gap of the machine.

Having decided upon such a dual-rotor approach, manufacturability becomes a major concern. A standard radial-flux design would require the installation of the magnets in a free-spinning cylindrical shell between the rotor and stator. Such a shell would be difficult to design to be sufficiently structural to support the required loading, and even more difficult to manufacture and assemble. Our solution thus uses axial flux.

The specific geometry we have chosen is derived from the TORUS geometry introduced in [7]. The TORUS design consists of a toroidal stator with a Gramme-Ring winding, interacting axially with a pair of rotors (one on each side), each of which consist of an annulus of magnets, arranged in alternating polarity. Our modification is to allow the magnet rotors to spin independently of the shaft and to add two additional rotors outside of the magnet rotors. These rotors are also toroidal in form, but are face-wound because they have only one active surface, and the conventional end-turns are shorter. A cut-away view of our machine prototype, with its original housing design, is shown in Figure 1-1.




Figure 1-1: Machine Cutaway View

1.3 Overview

The material of this thesis is divided into two principal sections. The first of these sections concerns the construction of the prototype, which has been the major thrust of work so far in the project. Chapter 2 details the work which has thus far been completed on the prototype and discusses that which remains to be done.

Once the prototype is constructed, some means will be required to control it to generate power in an optimum fashion. Before a control scheme can be developed, the behavior of the machine must be analyzed, and a dynamic model constructed. Chapter 3 derives a steady-state and dynamic model of the machine. Chapter 4 then considers the control objectives, and how the machine may be controlled to meet those objectives.

Chapter 2

Construction

The primary objective of this project has been the testing of a prototype generator to determine its suitability as a wind-energy solution. Thus, the majority of the work invested so far has focused on the construction of such a prototype. At this time, the prototype generator has not yet been completed. This section describes the work which has been completed and suitable designs for aspects which are unfinished.

Section 2.1 considers the design and implementation of the structural elements of the machine: the housing and the rotor support plates. Electromechanical aspects of the design are discussed in Section 2.2 and Section 2.3. The remaining miscellany of parts are discussed in Section 2.4, and assembly concerns are addressed in Section 2.5.

2.1 Structural Elements

Design of the structural elements of the machine was largely accomplished in [1]. Some modification of the plans therein was necessary to account for changes in the electrical parameters and dimensions of the machine.

Parts were redimensioned in metric units to simplify engineering aspects of the design. The electrical and structural engineering of the machine is conducted almost entirely in metric units. Thus, to reduce the number of unit conversions and associated chances of error, the parts are dimensioned in millimeters.

Mechanical drawings of the parts were produced with SolidWorks so that the parts could be manufactured by a machine shop. These drawings are reproduced in Appendix A. The initial batch of parts as specified were produced by the MIT Central Machine Shop.

Specific design issues for each part are reviewed in the following sections. Most notably, the machine housing specified in the original drawings was determined to be much too small given the overhang of the power rotor windings and the leads on the stator. The housing of the machine was thus redesigned in a much cheaper form which would accommodate the machine wiring. This redesign is discussed in Section 2.1.4.

2.1.1 Power Rotor Plates

The design objective of the power rotor plates is to support the core and windings associated with the electrically-excited rotor. In the assembled configuration (shown in Figure 1-1), the power rotor plates are the outermost of the electrically-interacting structures. The power rotors thus do not have the inherent balancing of static attractive forces that the magnet rotors and stator do, as they interact on only one face. Thus, the power rotor plates will have to support the entire attractive force of the permanent magnets acting on the rotor core. As calculated in [1], this attractive force requires an aluminum disc of approximately 45 mm thickness to prevent deflection from substantially impinging on the designed air-gap of 1 mm.

The drawing of the power rotor plates as manufactured is shown in Figure A-1. In the United States, aluminum plate is only easily available in fractional-inch sizes, so the actual plate thickness is 1.75 in. instead of 45 mm.

The design includes an array of holes drilled through the plate to allow for ventilation flow axially through the center of the machine to the active air-gaps. Attachment to the driveshaft is accommodated with a keyed hole at the center of the disc. The keyway is designed (as per [1]) to support the maximum expected torque in the machine, although the dimensions have been changed to fit the nearest convenient metric size. The concern in [1] of using ANSI standard shafting is not important, as the shaft must be custom-manufactured anyway (see Section 2.4.1).

At the time of manufacture, it was intended that the power rotor core attach to the power rotor plate with a series of U-bolts. This attachment scheme is accommodated with two annuli of smaller holes through which the bolts were to pass. As discussed in Section 2.5, the U-bolt scheme was abandoned in favor of directly bolting the core to the plate, which involved drilling new holes not shown in the part drawing.

2.1.2 Magnet Rotors

The magnet rotors serve to support the permanent magnets between the power rotor and the stator. It was determined that the optimal combination of manufacturability and function would be obtained with magnets shaped as 22 mm thick sectors of an annulus, with inner radius of 270 mm, outer radius 350 mm, and spanning 12° , as shown in Figure 2-1.




Figure 2-1: Shape of Permanent Magnets

In the assembled state, the attractive forces acting on each magnet are approximately balanced between the power rotor core and the stator core, so the magnet rotors are not required to support substantial axial loading. As described in Section 2.5, the unbalanced forces encountered during assembly will be supported by other means. Thus, the magnet rotor plates are aluminum discs 22 mm thick, with windows slightly larger than the magnets machined into the active region of the machine. The magnets are then retained in the magnet rotor by the combined action of a thin (1 mm thick) aluminum faceplate on each side and epoxy filling the remaining space in the windows.

The part drawing from which the magnet rotor plates were machined is shown in Figure A-2. As in the power rotor plates, an array of holes allow for axial ventilation to flow through the magnet rotor to reach the innermost air-gaps. Due to a mistake in manufacturing, twenty evenly spaced holes were drilled instead of the specified eighteen, but the impact on airflow and structural strength will be negligible.

The magnet rotors must spin independently of the other elements of the machine. They are designed to fit a bearing in the center which bears on the main driveshaft. The two magnet rotors should, however, rotate at identical speeds, so it is reasonable to connect them together. Connecting the magnet rotors with a hollow free-wheeling shaft simplifies the problem of axially constraining the magnet rotors. Thus, six holes around the periphery of the bearing mount allow bolting to such a hollow shaft and serve the additional function of securing a locator plate to constrain the bearing.

2.1.3 End Plates

The end plates serve a number of ancillary functions in the machine. Primarily, they support the weight of the driveshaft and constrain it against axial motion with respect to the housing. To this end, each has a hole in which to mount a bearing for the driveshaft and bolt holes at the outer radius to attach to the housing, as shown in the part drawing, Figure A-3.

The end plates additionally serve as mounting points for the slip-ring brushes and ventilation system (see Section 2.4.2 and Section 2.4.3, respectively). At the time the end plates were designed, the details of the slip-ring and ventilation systems were not known, so no provision for them was included in the part drawing. The expectation was that it would be simple to modify the end plates to accommodate the mounting requirements of the slip-rings and ventilation system.

2.1.4 Housing

The housing of the machine serves to hold the stator in place and provides a means of attaching the machine to the test-bed. Thus, the housing bears the entire torque-load of the machine, but at a large radius, so that the required material thickness is fairly small. The housing must transfer the weight of the machine to the test-bed mounting points, and provide attachment points for the end plates.

Original Housing

The original design of the housing was manufactured according to the part drawing of Figure A-4. The housing design has slots to allow airflow to escape radially from the air-gaps of the machine. Twenty mounting holes are provided to allow the stator to be bolted

to the housing. Twenty larger holes interspaced with the stator-mounting holes allow the electrical connections to the stator windings to be brought outside of the case. The end plates mount to ten holes in flanges at each end of the housing.

Revised Housing

Unfortunately, two problems arose after the housing was machined which necessitated a redesign of the housing. First, the wire from which the stator windings were constructed was quite stiff. The stiffness of the wire made it extremely difficult for the wires to be bent to fit out through the provided holes. This problem was compounded by the fact that the actual stator winding form took up more space than was anticipated at the time of the housing design, leaving very little space in which to maneuver the wire ends.

Even more disastrously, it was found that the design of the rotor windings (as described in Section 2.3.2) required about a centimeter more radial space than was available in the housing. Because the design of the windings was dictated by geometric concerns alone, available housing space was not considered, and the interference between windings and housing was not discovered until after the windings had been manufactured.

Given these problems, it was determined that the simplest solution was to redesign the housing of the machine. Budget requirements dictated that the new housing be much simpler than the old design. Thus a “skeletal” design was developed. This design consists simply of a rolled-and-welded aluminum ring to hold the stator, and ten aluminum stringers running from the ring to the end plates. The design of the ring is shown in Figure A-6, and that of the stringers in Figure A-7. The assembled housing, including end plates, is shown in Figure 2-2. The cross-sectional area of the stringers is similar to that of the original housing (they are much thicker than the housing wall), so ability to bear the torque loading will not be compromised.

2.2 Magnetic Cores

The overall machine design contains three magnetic cores: two rotors and one stator. These cores serve to transport flux between poles in the machine, closing the flux loops without adding any appreciable reluctance to the magnetic circuit. The cores are made from silicon electrical steel, which has high permeability and relatively low conductivity. The particular




Figure 2-2: Assembled “Skeletal” Housing

steel used is M-19, a medium-grade non-oriented steel, which was determined to be sufficient for our purposes.

The use of permanent magnets in the machine greatly simplifies construction of the cores. The magnets serve to drastically reduce the effect of the air-gap of the machine on its function. Consequently, the coils in the machine may be wound in the air-gap. Winding slots are not needed, as the magnets already constrain the flux to flow across the air-gap.

To prevent the formation of large eddy currents in the cores, which would substantially decrease machine efficiency, the cores have a laminated construction, with thin sheets of steel separated by a layer of insulating coating. The core is tape-wound, which is to say that the required toroidal form is constructed from a single very long strip of steel wound up in a long coil from the required inner radius to the required outer radius. The tape-wound form makes the core construction relatively simple.

The required axial thickness of the cores, and consequently the width of the steel strip, is determined by the expected maximum flux density which the cores must carry. This flux density must be less than the saturation flux density of the steel. For our machine, a rotor thickness of 16.5 mm is sufficient, while the stator, which must carry flux from both rotors, must be 31 mm thick. Thickness of the steel strip is determined so as to be much less than the skin depth in M-19 steel at the electrical frequency of concern: $t \ll \delta = \sqrt{\frac{1}{\sigma\mu\pi f}}$. The

selected 29 ga. steel has a thickness of about 0.35 mm, which is satisfactory. The steel was ordered with an AISI C-5 insulating coating to ensure insulation between layers.

To reduce the cost from processing overhead, the steel strip was ordered only in 16.5 mm width. Four essentially identical cores were then wound up from the strip, using about 400 m of strip per core. As the strip was shipped in several coils, the strip was spliced together, once per core, during the winding process. The resulting slight discontinuities should have negligible effect on the electrical properties of the cores. After winding, the cores were soaked in a low-viscosity epoxy (West System 105/206) to hold the laminations together.

Due to the relatively uneven nature of the winding form used, the cores were neither perfectly round nor perfectly flat. Furthermore, the nature of the imperfections varied somewhat from core to core. The two most similar cores were selected to be combined to form the stator core. All three core assemblies were then wrapped in 2-inch-wide 9 oz. fiberglass tape, and epoxied again. This has the additional benefit of further ensuring the integrity of the cores, and protecting their surfaces.

Figure B-1 shows a photograph of the cores after the fiberglass wrapping had been applied. Assembly of the cores with the windings and the structural elements of the machine is discussed in Section 2.5.

2.3 Windings

The windings provide the means for electrical interaction between the fields in the machine and the outside world. As such, they are perhaps the most critical element to the operation of the machine. Ignoring for the moment the effects of the permanent magnets, windings on the rotor are responsible for the electrical poles which generate the rotating flux wave in the machine. To ensure a moderately sinusoidal overall flux distribution, three phases of excitation are used on the rotors. Thus, for our 20-pole, dual-rotor machine, there are a total of 120 rotor windings. Stator windings are then responsible for linking the rotating flux wave generated on the rotor, and thereby producing voltage at the output terminals of the machine. The stator is a 3-phase system for similar reasons, and thus there are a total of sixty stator windings per active face.

2.3.1 Stator

The stator windings are responsible for the majority of the power handling capability of the machine. Consequently, in our design, they are constructed from 7×1.25 mm flat litz wire (SF2 transpose). Litz wire is used to overcome skin-effect-induced resistive losses in the windings due to the 60 Hz alternation of the current in the windings. The total copper cross-sectional area (860 mm^2) was chosen to be sufficient to carry the target output power of about 10 kW. The rectangular cross-section of the wire simplifies the winding operation.

The stator interacts with a rotor on both of its axial faces,¹ so the stator windings must cover both axial faces of the stator. This is most simply accomplished with a Gramme-Ring winding, in which the windings wrap all the way around the rectangular cross-section of the toroidal stator. Current flowing in a given direction in some stator winding will thus flow radially inward along one axial face of the stator, axially across the inner radial face, radially outward along the second axial face, and axially back across the outer radial face. This has the advantage of creating both face windings with a single relatively short pair of end-turns.

Each of the individual stator windings is composed of 10 turns of wire in 2 layers. The end-turns were formed around half-pieces of 1-inch diameter acrylic rod, to keep the bend radius high at the corners of the stator cross-section. After winding, the coils were clamped in place and epoxied to each other and the core. The whole structure was wrapped in fiberglass tape and epoxied again to protect the windings. Figure B-3 shows the structure with the fiberglass wrapping partially complete.

Measurements

Before the stator was wrapped in fiberglass, its electrical characteristics were tested. This was done by applying a 60 Hz sinusoidal voltage across one phase of the winding. This voltage was produced with a variable transformer (variac) from one phase of a 3-phase supply. The current and voltage waveforms were then measured using an oscilloscope with appropriate probes. This test was carried out for 4 voltage amplitudes: 75, 100, 150, and $200 V_{\text{rms}}$.

¹Where faces are specified as axial, radial, or circumferential, the direction refers to the direction of a normal vector to the surface.




Figure 2-3: Stator B-H Characteristics for 100 V_{rms} 60 Hz Sinusoidal Drive

An approximation of the hysteresis loop of the core can be obtained by plotting applied volt-seconds (the integral of the input voltage) against terminal current. The resulting curve for 100 V_{rms} is shown in Figure 2-3. The field intensity in the core is then approximately $H = Ni$, and the flux density $B \approx \frac{\int v dt}{NA_c}$, where A_c is the cross-sectional area of the core. Note that to the extent that the winding resistance is significant $v > \frac{d\lambda}{dt}$ and thus we have overestimated the flux density in the core. A plot of the approximate hysteresis loop for the 100 V_{rms} excitation is shown in Figure 2-3.

The curve of Figure 2-3 allows determination of the saturation field and flux of the core, along with the remanent magnetization and the coercivity of the material. We can read $B_{\text{sat}} \approx 0.5$ T, which is fairly low. This corresponds to a useful drive limit of about 1 Vs. The field at B_{sat} is about $H_{\text{sat}} \approx 150$ A. Finally, the hysteresis parameters are $H_c \approx 100$ A, and $B_r \approx 0.4$ T.

2.3.2 Rotor

The required power level of the rotor windings is much lower: less than 1 kW for our expected range of slip. The rotor windings are thus constructed of standard 16 AWG magnet wire. The electrical frequency at the rotor is the slip frequency, which will be much lower than 60 Hz, so skin effect is not important in the rotor windings.

In contrast with the stator, which interacts on both axial faces, the rotor interacts on only one axial face, and consequently only needs windings on that face. Although it would be possible to use a Gramme-Ring winding pattern on the rotors as well, more wire is used by a Gramme-Ring winding than conventional end-turns. Conventional end-turns thus offer lower resistance and consequently higher efficiency. Additionally, securement of the rotor core to the power rotor plate will be stronger if there is no wire between them. With conventional end-turns, all of the wire remains in the plane of the rotor's axial face. Current in a given winding runs radially inward at one pole, crosses to an adjacent pole at the inner radius, flows radially outward at that pole, then crosses back to the first pole at the outer radius. The difference between the two winding schemes is visible in the assembled core photos of Figure B-3 and Figure B-5.

Although the conventional end-turn pattern improves efficiency, it introduces a geometric problem: how to physically arrange the end-turns of different phases. Our machine uses a non-consequent-pole design, which means that in any given bundle of wires traversing the core face (what would be a slot if our machine had slots), there are two windings, one above and one beneath. If the windings are to neatly lie flat, the end-turns traversing from bundle to bundle in phase A must pass four other end turns corresponding to the two bundles from phases B and C lying between the phase A bundles. Conceptually, this can be solved if the end turns from A to A' pass above the end-turns which come from the bottom windings in the intervening B and C bundles, and below the end-turns which come from the top windings.

If the windings are to be symmetric, this suggests a hexagonal footprint, so that the nearly-parallel top- or bottom-winding end-turns do not interfere with each other. A schematic representation of the winding plan is shown in Figure 2-4.

Two layers of 12 turns of 16 AWG wire were determined to appropriately fill the axial face of the core. The geometry problem associated with finding appropriate angles and side-lengths for the hexagon produces bend angles at the inner and outer ends of the active legs of the winding of 35° and 65° respectively. Two forms were constructed having a hexagonal center spacer of appropriate angle and side length, and a height of about 0.11 in. to accommodate two 16 AWG wires side by side. By splitting the form in half and displacing the halves vertically with respect to each other by about 0.11 in. the winding was constrained to have the correct vertical displacement between its two active legs.




Figure 2-4: Schematic Representation of Rotor End-Turn Configuration

The form was affixed to an aluminum mandrel to allow it to be turned by a lathe, and the wire fed by a combination tensioner and electronic turn counter. Forms could then be tightly and fairly uniformly wound with the required 24 turns in about 30 seconds. The top half of the form was divided into thirds. Each third was then removed one at a time, and the underlying winding epoxied with an epoxy gel. Piecewise gluing was required to prevent springback in the winding from destroying the coil, as it would if the entire top half of the form were removed at once. Unfortunately, the overhead associated with preparing a coil for winding and gluing it after winding was sufficiently long that only two forms could be run in parallel, and the process was quite slow as a result. Figure B-2 is a photograph of the winding apparatus.

Measurements

The 120 rotor windings were tested on an impedance analyzer in 2 batches of sixty. The objective of this testing was to ensure that none contained a shorted turn, and that all had approximately the correct number of turns. A shorted winding turn would be disastrous, but coils having 23 or 25 turns would not appreciably affect the operation of the machine, especially if appropriately balanced around the machine. From the inductances of the coils, it was possible to detect shorted coils and coils with too few turns. One coil was found to

have a shorted turn, evidenced by an inductance too low by half an order of magnitude. A histogram of the inductances of the coils after replacement of that coil is shown in Figure 2-5.

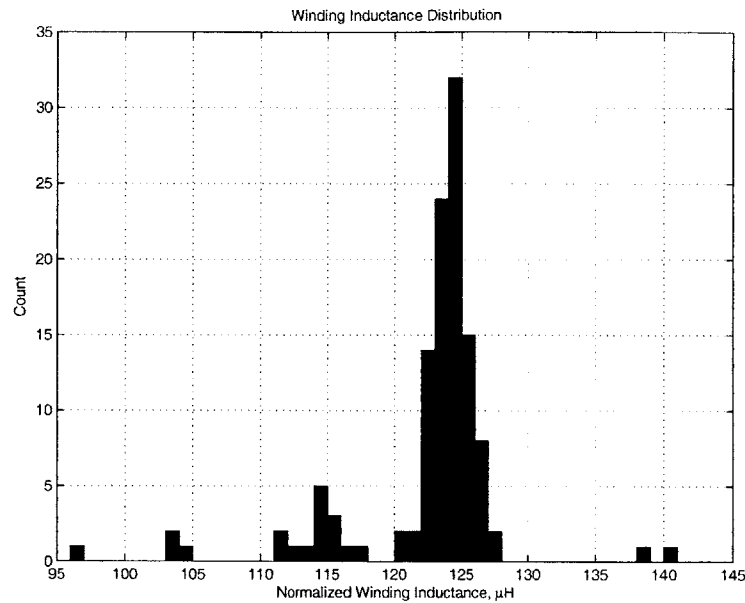


Figure 2-5: Histogram of Rotor Winding Inductances, First Pass

In Figure 2-5, the coils clearly clump into groups whose centers correspond well to the expected inductances of coils with numbers of turns ranging from 21 to 25. Four coils can be seen to have fewer than 23 turns, and these coils were also replaced and the replacements tested. The final inductance distribution is as shown in Figure 2-6. In this distribution, there are 14 coils with 23 turns, and 2 with 25. If the coils with 25 turns are matched pairwise with 23-turn coils, the effects of the two will essentially cancel out, leaving 12 coils short a turn to distribute about the machine. Conveniently enough, these 12 coils can be evenly distributed among the 3 phases and physically about the machine by placing them at locations 1, 11, 21 ... 101, 111, where the 60 coils of the first rotor are numbered sequentially around the rotor from 1 to 60, and those of the second rotor from 61 to 120.

2.4 Other Parts

A few other miscellaneous parts are necessary to the function of the machine, none of which have been fabricated as of this writing. This set of parts includes the shafts, slip-rings, ventilation system, and mounting hardware. This section briefly discusses the requirements

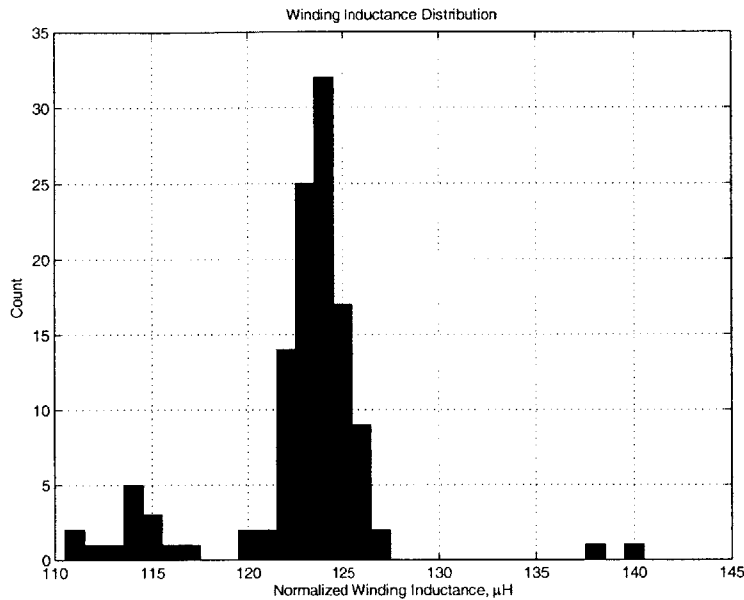


Figure 2-6: Histogram of Rotor Winding inductances, Second Pass

of each part and outlines possible designs. Fabrication of these parts, assembly, and testing of the machine are the bulk of the further work described in Section 5.2.

2.4.1 Shafts

The 2 shafts in the machine cannot be produced until final dimensions of the machine are known. Due to the inaccuracy of some of the details of the assembly process, these dimensions are not yet known. The specifications of the shafts and their dimensions as a function of the as-of-yet-unknown machine dimensions are examined in turn.

Driveshaft

The driveshaft serves to transmit torque from the external prime mover to the power rotors of the machine. To sustain 10 kW at a nominal synchronous speed of 360 rpm requires a torque capability of about 265 N·m. Given such a substantial load, the shaft must be made from a reasonably high-grade steel alloy. Converting the analysis of [1] to metric units indicates that the minimum shaft diameter should be 30 mm. The shaft must also bear the substantial axial load associated with supporting the power rotors against the attractive force of the permanent magnets. This may be accomplished by a shoulder on

the shaft stepping up to a 45 mm diameter. Appropriate bearings for these shaft diameters have already been purchased (to measure for fit into the endplates and magnet rotors). One more shoulder on the shaft is required to locate the permanent magnet rotors, but the anticipated loading is quite small, so the next easily available steel rod diameter—probably 2 in.—should be used. Overall form of the driveshaft is then as shown in Figure 2-7.




Figure 2-7: Driveshaft Plan View

The distance between the inner shaft shoulders should be set as given in (2.1), and between the outer shoulders as in (2.2).

$$d_i = t_s + 2g + (t_{mr1} - t_f1) + (t_{mr2} - t_f2) - 2t_b \quad (2.1)$$

$$d_o = t_s + 4g + t_{mr1} + t_{pc1} + t_{mr2} + t_{pc2} \quad (2.2)$$

In (2.1), t_s is the total stator thickness including windings, g is the desired operating air-gap, t_{mr} is the thickness of the magnet rotor, including faceplates, t_f is the thickness of a magnet-rotor faceplate, and t_b is the axial length of the magnet-rotor bearings. In (2.2), t_{pc} is the thickness of the power rotor core, measured from the winding face to the surface of the power rotor plate, in an assembled state. Numerical subscripts are used to indicate the fact that it should not be assumed that thicknesses are identical for nominally symmetrical parts. The distance $t_{mr} - t_f$ should be almost exactly 23 mm for both magnet rotors, and $t_b = 19$ mm.

The keyways in the shaft should be machined 5 mm wide, 2.5 mm deep, and 45–50 mm long to accommodate standard square keys which interface properly with the power rotor plates.

Free-Wheeling Shaft

The free-wheeling shaft mostly serves to help locate the magnet rotors with respect to the shaft, and thereby the rest of the machine. As such, it is not expected to bear substantial axial or torque load. The free-wheeling shaft must have an internal diameter sufficient to be well clear of the outer diameter of the driveshaft. The outer diameter of the shaft should be at least 4.75 in. \approx 120 mm to allow mounting to the magnet rotors. Mounting to the magnet rotors will be most simply accomplished if 6 holes are drilled and tapped to fit an M6x1.0 screw in the ends of the shaft. In principle, there is no reason that the function of the free-wheeling shaft cannot be accomplished with 6 standoffs with threaded holes in their ends. Using 6 standoffs is attractive for cost reasons, but care must be taken that the standoff lengths are very well matched to avoid skewing the rotor plates.

Regardless of the nature of the free-wheeling shaft, its length should be as given in (2.3).

$$l = t_s + 2g \tag{2.3}$$

Included in the free-wheeling shaft assembly should be locator plates for the magnet-rotor bearings. These plates should be annular in form, with an inner diameter of 80 mm, an outer diameter of at least 120 mm, and thickness on the order of 2 mm, with holes drilled to align with the six mounting holes in the magnet rotors. These plates then lie against the outer faces of the magnet rotor plates and are held in place by the same bolts which secure the free-wheeling shaft.

2.4.2 Slip-Rings

The slip-ring assembly connects the rotor windings to an external source of excitation current for the rotor windings. This is accomplished with a set of conductive rings mounted to the shaft and connected to the rotor windings. Spring-loaded carbon brushes are affixed to the case and are held against the conductive rings, making a rotating contact to the rotor windings. The number of rings in the set of slip-rings determines the variability of the rotor connection topology. The most general configuration requires 12 rings, one for each end of each phase on each rotor. Such a scheme would allow the the rotors to be configured in either a Δ - or Y-connection, and for the 2 rotors to be connected in series or parallel.

Unfortunately, it is somewhat difficult to run wires from the rotors to the outside of the machine. Thus, if the most logical scheme of locating all of the slip-rings on the outside at one end of the machine is followed, wires from the far rotor will have to travel under 3 bearings to reach the slip-rings, and as such it is desirable to limit the number of wires needed.

Fixing the configuration of the far rotor as a buried-neutral Y-connection and bringing out all 6 ends from the near rotor seems a reasonable compromise. This allows the minimum number of wires (3) to transit under the magnet-rotor bearings, but still allows some variation in the overall configuration. This scheme is especially attractive if the slip-rings are located on the end of the machine opposite the prime mover, as the driveshaft is not required to transmit any torque at the point that the 9 wires must pass under a bearing. The shaft may therefore be drilled out (“gun-bored”) from the slip-ring end, and the wires run along the center of the shaft to the slip rings. Such a scheme is not practical for the far rotor for two reasons. Firstly, the required length of drill to reach the far rotor is prohibitive. Secondly, gun-boring the shaft might easily compromise its ability to bear the required torque load, given the weakening effect of the keyways.

It is worth noting that by having two distinct sets of slip-rings, one for each rotor, and affixing these slip-rings to the inside of the endplates instead of the outside, it is possible to avoid passing the rotor wires under *any* bearings. It may be infeasible to use such a slip-ring design due to the additional space required inside the machine, but the benefits are sufficiently attractive to merit careful consideration.

2.4.3 Ventilation

Some amount of ventilation is required to keep the temperature of the cores and wiring of the machine within reasonable limits; a maximum temperature of 85° C should allow all of the insulations and epoxies to operate within specified limits. With the original housing design described in Section 2.1.4, the intended method of cooling was to mount a blower to each endplate. Air would then be forced axially through the center of the machine to the air-gaps, where it would flow radially outward through the air-gaps and then through ventilation slots in the housing. Airflow around the exterior of the power rotor plates would be prevented by a simple plastic rotating seal against the outside of the power rotor plates.

With the new skeletal housing design, this scheme is more difficult to enact. However, an extension of the plastic-rotating-seal concept can likely be used. An annulus of thin flexible plastic (preferably polyethylene sheet) mounted to the inside of the endplates such that it will touch the outer surface of the rotor plates may be used. The thin film would then be held against the outer surface of the rotor plates by the positive pressure created by the blower. Although the resulting seal would not be very good, all that is required is that leaking through the seal be a higher-resistance path than flowing out through the air-gaps.

As per the thermal analysis of [1], an airflow of $0.05 \text{ m}^3/\text{s}$ per air-gap will cause a temperature rise of approximately 20 K. Given the number of approximations required to obtain such an estimate, a factor of three below the desired maximum temperature rise seems wise. Thus, if two blowers are used (one in each endplate), each should have a rating of approximately $0.1 \text{ m}^3/\text{s} \approx 200 \text{ CFM}$.

2.4.4 Mounting and Testbed

The machine must be mounted to a test-bed so that its performance can be analyzed and the validity of the models of Section 3 ascertained. Our intended test facility is a large dynamometer which was donated to the lab by the Ford Motor Company. The dynamometer is rated at 20 horsepower over a range of 575–4000 rpm. Interfacing with the dynamometer has two aspects: mounting and torque transmission.

Mounting

The housing of the machine must be secured to the test facility. The currently-intended method of attachment is to attach 3 aluminum “feet” to the ring in which the stator is mounted. These feet would bolt to the outer radial face of the ring and be of appropriate length to support the machine in a level state. These feet would then bolt to an aluminum block whose height is determined by the requirements of the transmission between the dynamometer and the machine. The block would then bolt to the dynamometer baseplate.

Torque Transmission

The minimum specified speed of the dynamometer is 575 rpm, and its shaft is closer to the baseplate than the radius of our machine. Thus, some form of transmission is required between the machine and the dynamometer.

The transmission should likely be a toothed belt and two sheaves in a 2:1 diameter ratio. The diameter difference will allow the machine to run at its nominal synchronous speed of 360 rpm while the dynamometer runs at 720 rpm. The offset between the two sheaves can also be used to allow the machine shaft to be considerably higher than the dynamometer shaft. It is likely that a tensioning idler will be required to get the necessary torque transmission between the two. Details of the transmission depend very heavily on the availability of sheaves and belts appropriate to the connection, so the design is not pursued further here.

2.5 Assembly

Assembly of the machine requires careful consideration due to the large static and dynamic axial forces in the machine. At the time of this writing, relatively little assembly beyond what has already been described is completed. The individual components have largely been constructed, but they have yet to be interfaced.

2.5.1 Power Rotors

Both of the power rotor cores have been secured to their respective power rotor plates, and one of the sets of windings has been affixed to its core. As mentioned in Section 2.1.1, the original mounting scheme for the power rotor cores required constraining the cores with U-bolts over the top of the core. Difficulty in obtaining appropriate U-bolts made this scheme less attractive. Furthermore, the nature of the rotor windings left little available space on the axial surface of the core upon which the U-bolts might bear.

As specified in [1], the U-bolt attachment scheme would result in a larger-than-necessary structural safety margin for the attachment. Thus, it was possible to revise the attachment mechanism to simultaneously resolve both concerns with the U-bolts. The solution chosen bolts the core directly to the power rotor plate with standard socket-head cap screws. To prevent interference between the heads of the bolts and the windings, the holes in the core are counterbored.

In this application, it is not feasible to use flat-head screws and countersunk holes due to the laminar construction of the core. The glue-joint between laminations in the core is fairly weak, even though the overall structure is reinforced with fiberglass. Thus, there is substantial concern that the conical head undersides of flat-head screws would force the

laminations apart under heavy load, and cause disintegration of the core. The flat underside of socket-head cap screws and attendant flat bottom of appropriately counterbored holes do not provide any means of transferring lateral force to the core, and therefore do not engender this concern.

The attachment used 18 evenly spaced 65 mm M6x1.0 grade 12.9 socket-head cap screws, all radially centered in the core. The high grade steel of the screws is necessary, given the expected loading. The screws were all tightened to a torque of 16.9 N m, approximately the maximum recommended tightening torque. This gives an expected bolt preload of 15.6 kN, which should be sufficient that the joint does not come out of compression under any expected loading. If the joint were to leave the compressive regime, the comparatively small but not insignificant alternating force load on the bolts during operation could easily lead to fatigue failure of the bolts, which would be disastrous.

With the cores appropriately secured to the plates, all that remains to be done in assembly of the power rotors is attachment of the windings. At present, this has been accomplished for one of the rotors. Windings were laid out on the surface of the core, in a distribution as described in Section 2.3.2. The coils were then wired and soldered together with all of the coils in a given phase in series, appropriately alternating in sense. After soldering and insulating the joints, continuity of each phase was tested, and DC isolation between phases was tested by applying ± 400 V between each pair of phases.

Following the electrical tests, a relatively small amount of epoxy was spread over the top of the windings, securing them weakly to each other and the core, and the whole structure was clamped with a large aluminum sheet above to flatten the windings into place. Further attachment was then accomplished with a layer of fiberglass tape spread over the exposed windings to protect and secure them. The fiberglass tape was liberally soaked in epoxy, then similarly clamped.

2.5.2 Stator Mounting

The structure mounting the stator to the housing need sustain relatively little torque, transverse (gravitational), or axial load in operation. However, during assembly, the attachment will have to support the full attractive force between the stator core and the permanent magnet rotors. It is intended that this load be borne by L-shaped spacer blocks which separate the stator from the mounting ring of the new housing design. These L-shaped blocks,

shaped as shown in Figure 2-8, would lie around the periphery of the stator between the sets of three windings corresponding to poles.




Figure 2-8: Form of Stator Support Blocks

Space of about 0.75 in. has been left between the phase A winding of each pole and the phase C winding of the adjacent pole. By alternating the axial direction of the spacer blocks, the stator can be effectively constrained from motion in the axial direction by ten blocks each way. Constraint from rotational or transverse motion is trivially satisfied as long as the blocks are somewhat securely attached to the core and the ring.

2.5.3 Final Assembly

The majority of the remaining physical difficulty of creating the machine will lie in the final assembly of the various components. During assembly, the full attractive force of the permanent magnets will be at issue several times. Worse yet, at several stages in assembly, components whose attractive forces are roughly balanced in the assembled machine will have to bear the full attractive force in one direction.

Furthermore, some consideration will be required to allow careful positioning of elements which, due to the magnets, will be attracted to each other with a force on the order of 50 kN. Thus, a screw-jack must be constructed to lower the elements together in a controlled fashion. The simplest method of construction would be to use a stiff, thick plate of diameter larger than the machine diameter. Three locations around the edge of the plate would then have standoffs aligned to bear on the ends of the stringers in the housing. Necessarily,

these standoffs would not be perfectly symmetrically placed, as there are ten stringers in the housing structure.

The plate would also have three symmetrically placed holes at a smaller radius, through which large threaded rods (perhaps 0.75 in. diameter) could pass. Affixed to the ends of the threaded rods would be hooks, to which steel cables could be attached. On the top side of the plate, a nut and washer on each rod would bear on the surface of the plate. Thus, by loosening the nuts, the threaded rods, bearing the load of an element under assembly, could be extended in a controlled manner.

To lower one of the rotors into place, steel cables would be run from each hook, through the ventilation holes in the plate, and back to the hook. Thus, after the rotor was in place, the cable could be pulled back through the holes without substantial difficulty. If thin plastic spacers are inserted into the air-gaps before assembly, they would then support the attractive loading after the cable and screw-jack were removed. Once the machine is completely assembled, the plastic spacers would be pulled out.

A likely appropriate order of assembly would be as follows.

1. Stator and housing (not including end plates) assembled and stood on one end (axial direction vertical).
2. Magnet rotor A lowered into place over stator.
3. Free-wheeling shaft attached to magnet rotor A.
4. Power rotor A lowered into place over magnet rotor A , taking care that the two are well aligned so that the driveshaft fits both.
5. Endplate A secured to housing, aligning for shaft fit again.
6. Assembly flipped over.
7. Driveshaft inserted.
8. Magnet rotor B lowered over driveshaft.
9. Free-wheeling shaft attached to magnet rotor B.
10. Power rotor B lowered over driveshaft.
11. Endplate B secured to housing.
12. Testbed mounting hardware attached.
13. Slip-rings, ventilation, and wiring assembled.

Chapter 3

Analysis and Modeling

If results from measurement of the prototype are to be usefully understood, a model of the machine must first be developed which creates a context for the measurements. Moreover, the nature of the measurements performed will be dictated to some extent by the model of the machine. Finally, to obtain the desired end-behavior, a control circuit must be built. Controller design is considered in Chapter 4. However, it is not possible to design a controller, or even adequately describe the aims of one, without first developing a dynamic model of the machine.

It is therefore clear that some effort must be expended in development of models of the machine, and that is thus the aim of this chapter. A general understanding of the operation of the machine and its steady-state behavior are developed in Section 3.1. Based on these concepts, Section 3.2 then constructs a full-blown tenth-order model of the system dynamics and considers some simplifying assumptions which may be appropriate depending on the desired results.

3.1 Steady-State Behavior

The analysis of our machine topology is simplified by the fact that direct interaction between the rotor and stator windings is very weak in comparison to interaction with the magnet rotor. A quick verification of this assumption can be made by comparison of the peak flux-densities due to the magnets and currents in the stator, measured at the rotor.

The magnets in the machine are neodymium-iron-boron (NdFeB) material, and as such have a remanent flux density $B_r \approx 1.3$ T. For an air gap which is relatively small compared

to the thickness and lateral dimension of the magnets, the flux density at the other side of the gap can be approximated as given in (3.1), where t_m is the magnet thickness, and g is the air-gap thickness.

$$B \approx B_r \frac{t_m}{g + t_m} \approx 0.9 \text{ T} \quad (3.1)$$

Peak flux density due to stator winding currents can be estimated as by assuming the gap to be relatively small compared to the pole spacing and applying standard approximations. Firstly, from the point of view of excitation from the stator, the entire structure of the magnet rotor has permeability of μ_0 , so the effective air gap is the entire distance from the stator iron to the power rotor iron. Thus, the gap is not in fact small, so our calculation will therefore err on the conservative side and overestimate the flux density. Given the relative permeabilities of the core and the air gap, we can assume that the field intensity H is only non-zero in the air-gap. For a small gap, the field in the gap will be constant and directed strictly across the gap. Thus Ampere's law for a contour enclosing one stator coil and traversing straight across the gap gives $\oint H \cdot d\vec{l} = 2gH_g = N_s I_s$. Solving for H_g , we can then determine B_g , using a conservative stator heating limit of $N_s I_s < 1000 \text{ A}$.

$$B \leq \frac{\mu_0 N_s I_s}{2g} \leq 15 \text{ mT} \quad (3.2)$$

Comparison of (3.1) and (3.2) shows that the stator currents have more than an order of magnitude less effect on the power rotor than do the magnets. Thus, we are justified in the approximation that each set of windings interact only with the magnet rotor.

In the reference frame of the power rotor, the behavior of the magnet-rotor is equivalent to that of a permanent-magnet synchronous motor. That equivalent motor is then loaded by a torque derived from the interaction of the stator and the permanent magnets. The motor runs in synchronicity with the electrical frequency of the power rotor. Thus, in the reference frame of the stator, the motion of the magnet rotor is that synchronous rotational speed plus the speed of rotation of the rotor reference frame, which is the shaft speed. The interaction between the magnet rotor and the stator is equivalent to a permanent magnet synchronous generator in which the shaft speed is the same as the magnet rotor speed.

Basic operation of the machine can then be explained more rigorously. If the mechanical shaft speed is Ω_p , and the electrical frequency of excitation in the rotor is ω_r , then the mechanical rotation speed of the magnet rotor in the power rotor reference frame is $\frac{\omega_r}{p}$. In the stator reference frame, the mechanical speed of the magnet rotor is $\frac{\omega_r}{p} + \Omega_p$. Thus, the output electrical frequency is $\omega_s = \omega_r + p\Omega_p$. As expected, the output electrical frequency is the electrical-equivalent shaft speed plus the excitation frequency.

Assuming perfect energy conversion efficiency, we can also easily derive the rotor and shaft input power as a function of the total stator power, P_s . If a given power leaves the stator, the back-torque exerted on the magnet rotor must be $T_b = \frac{P_s}{\omega_r + p\Omega_p}$. Torque balance on the magnet rotor then requires that the power rotor exert an exactly opposite torque. In the reference frame of the rotor, then, the power transmitted is $T_b\omega_r$. Substitution gives (3.3), in which we have defined the slip $s = \frac{\omega_r}{\omega_s}$ in an analogous manner to the usual definition for induction machines.

$$P_r = \frac{P_s}{1 + \frac{p\Omega_p}{\omega_r}} = sP_s \quad (3.3)$$

3.2 Dynamics

The dynamic model of the machine is also simplified by separation into two separate machines, as in the steady-state model of Section 3.1. Specifically, we can treat the dynamics of the interaction between power rotor and magnet rotor with a set of differential equations almost completely independent from those describing the interaction between magnet rotor and stator.

These equations describe, in the case of the rotor-rotor interaction, the dynamics of a rotating synchronous motor. For the rotor-stator interaction, the equations describe a standard synchronous generator. The two sets of equations are interconnected only through the external torque loading of the motor, which is equal to the shaft torque of the generator. The operation of the two sub-machines will be derived separately, borrowing from analysis in [9], and then a unified model created by combining the two sub-models.

3.2.1 Power Rotor

The interface between the power rotor and the magnet rotor is, as mentioned, similar to the internal action of a standard permanent magnet synchronous motor. First, the operation of the machine with a stationary stator will be derived. The effects of the rotation of the stator (due to rotation of the machine driveshaft) will then be added to the model.

To simplify the state equations, the problem is first transformed into a coordinate frame which rotates synchronously with the electrical excitation. This transformation is effected by use of the Parks Transform (also known as the dq0 transform). Effectively, use of the Parks Transform changes the basis of the analysis from the standard phase-variables to quantities which are *direct*, *quadrature*, or *zero-sequence*.

Precise definition of these terms and the transformation itself are given in [6] and [9]. For our purposes, the machine will operate in a balanced condition, which is equivalent to no zero-sequence behavior. A direct-axis quantity is then one which is aligned with the axis of the transform, and a quadrature-axis quantity is perpendicular to the axis of the transformation. In as much as windings and currents which are physically located 90 electrical degrees apart do not interact, the new basis is orthogonal, and currents in direct-axis windings produce only direct-axis flux.

Within this framework, we can describe the structure of the machine as having a number of direct and quadrature axis windings which generate and link flux, producing mechanical torque and back-voltage. The inverse transformation can be applied to the results if it is then desired to extract the phase-variables. Reverse transformation is an important step, as the phase-variables are the ones which are measurable and controllable from the outside system.

The direct axis effectively has three windings. In reality, only one of these windings corresponds to real windings in the machine, the direct axis stator winding. In a wound-field synchronous motor, there would be a direct-axis field winding on the stator. In our machine, the field winding is replaced with permanent magnets, which behave like a field winding with a constant current. Finally, we include the effects of distributed conduction in the structure of the rotor with a third “damper” winding, whose ends are shorted. This winding simulates the effect of time-varying flux at the rotor which produces eddy currents in the retaining plate and magnet, rejecting some of the flux on a short timescale. The

excitation from the permanent magnets lies only on the direct axis, so the quadrature axis has only two windings, one corresponding to the stator windings, and one to the damper winding.

We can then write state equations relating the flux linked by each winding to the other fluxes in the system, the winding currents, and the motion of the machine. The form of the resulting differential equations depends on the nature of the connection to the external system. In particular, the stator windings may be driven by either a voltage source or a current source. If the machine is current-source driven, the stator fluxes λ_d and λ_q (direct and quadrature, respectively) are set by the current. For our purposes this is highly desirable, as it corresponds to reducing the order of the system by two, which will improve control behavior, as described in Chapter 4.

We define a number of quantities for use in our model. The significance of the variables used are as given in Table 3.1.

Table 3.1: Synchronous Motor Model Variables

λ_d	Flux linked by the direct-axis stator winding
λ_q	Flux linked by the quadrature-axis stator winding
λ_{kd}	Flux linked by the direct-axis damper winding
λ_{kq}	Flux linked by the quadrature-axis damper winding
λ_0	Flux produced by permanent magnets
i_d, i_q, i_{kd}, i_{kq}	Currents in the windings
R_d, R_q, R_{kd}, R_{kq}	Winding resistances
T_m	Mechanical torque loading on the machine
T_e	Electrical torque produced by the machine
J_m	Rotational inertia of the magnet rotor
ω	Rotational velocity of the magnet rotor
ω_0	Electrical angular velocity and angular velocity of the transform
δ	Angular offset between the transform and the magnet rotor
$\theta = \omega_0 t + \delta$	Angle of the transform
$p = 10$	Number of pole-pairs in the machine
L_d, L_q, L_{kd}, L_{kq}	Winding self-inductances
M_d, M_q	Mutual inductance between stator and damper windings

The resulting state equations are then as given in (3.4) through (3.7).

$$\frac{d\lambda_{kd}}{dt} = -R_{kd}i_{kd} \quad (3.4)$$

$$\frac{d\lambda_{kq}}{dt} = -R_{kq}i_{kq} \quad (3.5)$$

$$\frac{d\omega}{dt} = \frac{1}{J_m}(\mathcal{T}_m + \mathcal{T}_e) \quad (3.6)$$

$$\frac{d\delta}{dt} = \omega - \omega_0 \quad (3.7)$$

The auxiliary equations are as given in (3.8) through (3.12).

$$\lambda_d = L_d i_d + M_d i_{kd} + \lambda_0 \quad (3.8)$$

$$\lambda_q = L_q i_q + M_q i_{kq} \quad (3.9)$$

$$\mathcal{T}_e = \frac{3p}{2}(\lambda_d i_q - \lambda_q i_d) \quad (3.10)$$

$$i_{kd} = \frac{\lambda_{kd} - M_d i_d - \lambda_0}{L_{kd}} \quad (3.11)$$

$$i_{kq} = \frac{\lambda_{kq} - M_q i_q}{L_{kq}} \quad (3.12)$$

The behavior above, as expressed in (3.4) through (3.12), represents the situation in which the machine shaft is held still and the magnet rotor is moved only by application of motor action from the power rotor. The equations must therefore now be expanded to include the effects of driveshaft speed on the behavior of the magnet rotor.

The motion of the power rotor is derived from a torque input on the machine shaft as given in (3.13).

$$\frac{d\omega_p}{dt} = \frac{1}{J_p}(\mathcal{T}_{shaft} - \mathcal{T}_e) \quad (3.13)$$

The simplest way to encapsulate this effect on the machine dynamics is to examine the effects of this motion in the reference frame of the rotor. Acceleration of the rotor reference frame produces an apparent opposite acceleration of the magnet rotor. Thus, staying within the rotor reference frame, we can simply replace (3.6) with (3.14).

$$\frac{d\omega}{dt} = \frac{3p}{2J_m}(\mathcal{T}_m + \mathcal{T}_e) - \frac{1}{J_p}(\mathcal{T}_{shaft} - \mathcal{T}_e) \quad (3.14)$$

Note that ω no longer represents the absolute rotational velocity of the magnet rotor. The actual velocity in a static reference frame is given by $\omega_m = \omega + \omega_p$.

3.2.2 Stator

The model of the interaction between magnet rotor and stator is quite similar to that between the power rotor and the magnet rotor. Obviously, the stator is stationary, and thus does not require the correction term of (3.14). Additionally, the stator loading differs from the power rotor excitation. Specifically, we expect that in operation as a generator, the stator will be connected directly to the 3-phase mains voltage, which is reasonably approximated by an infinitely stiff voltage source. The net result is that the stator model has 2 additional state variables due to the flux linkage of the direct- and quadrature-axis stator windings.

Thus, for the voltage-source-loaded generator, our state equations are as given by (3.15) through (3.20).

$$\frac{d\lambda_d}{dt} = v_d + \frac{\omega}{\omega_0}\lambda_q - R_d i_d \quad (3.15)$$

$$\frac{d\lambda_q}{dt} = v_q - \frac{\omega}{\omega_0}\lambda_d - R_q i_q \quad (3.16)$$

$$\frac{d\lambda_{kd}}{dt} = -R_{kd} i_{kd} \quad (3.17)$$

$$\frac{d\lambda_{kq}}{dt} = -R_{kq} i_{kq} \quad (3.18)$$

$$\frac{d\omega}{dt} = \frac{1}{J_m}(T_m + T_e) \quad (3.19)$$

$$\frac{d\delta}{dt} = \omega - \omega_0 \quad (3.20)$$

The auxiliary equations are then as given in (3.21) through (3.24).

$$\begin{bmatrix} i_d \\ i_{kd} \end{bmatrix} = \begin{bmatrix} L_d & M_d \\ M_d & L_{kd} \end{bmatrix}^{-1} \begin{bmatrix} \lambda_d - \lambda_0 \\ \lambda_{kd} - \lambda_0 \end{bmatrix} \quad (3.21)$$

$$\begin{bmatrix} i_q \\ i_{kq} \end{bmatrix} = \begin{bmatrix} L_q & M_q \\ M_q & L_{kq} \end{bmatrix}^{-1} \begin{bmatrix} \lambda_q \\ \lambda_{kq} \end{bmatrix} \quad (3.22)$$

$$T_e = \frac{3p}{2}(\lambda_d i_q - \lambda_q i_d) \quad (3.23)$$

$$(3.24)$$

3.2.3 Overall Behavior

An overall model of the behavior of the machine requires combining the models derived in Section 3.2.1 and Section 3.2.2. In the derivation that follows, quantities which come originally from the rotor-rotor interaction model will be denoted with an additional subscript R and those from the rotor-stator model with an S . For example, the flux linked by the direct-axis damper winding in the rotor-rotor interaction will be denoted λ_{kdR} , while the identical linkage in the rotor-stator model will be λ_{kdS} .

Interaction between the two models takes place through the torques exerted upon and speed of the magnet rotor. The magnet rotor must move the same way in both static reference frames, so $\omega_{mR} = \omega_S$. The torque of each sub-machine on the magnet rotor has already been accounted for in the other sub-machines model by the inclusion of \mathcal{T}_m . The “external” load torque in one model is the electrical torque in the other model, which is to say, $\mathcal{T}_{mR} = \mathcal{T}_{eS}$ and $\mathcal{T}_{mS} = \mathcal{T}_{eR}$. We are then in a position to combine the two models, with the resulting state model shown in (3.25) through (3.34), and auxiliary equations in (3.35) through (3.42).

$$\frac{d\lambda_{kdR}}{dt} = -R_{kd}i_{kdR} \quad (3.25)$$

$$\frac{d\lambda_{kqR}}{dt} = -R_{kq}i_{kqR} \quad (3.26)$$

$$\frac{d\omega_R}{dt} = \frac{1}{J_m}(\mathcal{T}_{eR} + \mathcal{T}_{eS}) - \frac{1}{J_p}(\mathcal{T}_{shaft} - \mathcal{T}_{eR}) \quad (3.27)$$

$$\frac{d\delta_R}{dt} = \omega_R - \omega_{0R} \quad (3.28)$$

$$\frac{d\lambda_{dS}}{dt} = v_{dS} + \frac{\omega_S}{\omega_{0S}}\lambda_{qS} - R_{dS}i_{dS} \quad (3.29)$$

$$\frac{d\lambda_{qS}}{dt} = v_{qS} - \frac{\omega_S}{\omega_{0S}}\lambda_{dS} - R_{qS}i_{qS} \quad (3.30)$$

$$\frac{d\lambda_{kdS}}{dt} = -R_{kd}i_{kdS} \quad (3.31)$$

$$\frac{d\lambda_{kqS}}{dt} = -R_{kq}i_{kqS} \quad (3.32)$$

$$\frac{d\omega_S}{dt} = \frac{1}{J_m}(\mathcal{T}_{eR} + \mathcal{T}_{eS}) \quad (3.33)$$

$$\frac{d\delta_S}{dt} = \omega_S - \omega_{0S} \quad (3.34)$$

$$\lambda_{dR} = L_{dR}i_{dR} + M_{dR}i_{kdR} + \lambda_0 \quad (3.35)$$

$$\lambda_{qR} = L_{qR}i_{qR} + M_{qR}i_{kqR} \quad (3.36)$$

$$\mathcal{T}_{eR} = \frac{3p}{2}(\lambda_{dR}i_{qR} - \lambda_{qR}i_{dR}) \quad (3.37)$$

$$i_{kdR} = \frac{\lambda_{kdR} - M_{dR}i_{dR} - \lambda_0}{L_{kdR}} \quad (3.38)$$

$$i_{kqR} = \frac{\lambda_{kqR} - M_{qR}i_{qR}}{L_{kqR}} \quad (3.39)$$

$$\begin{bmatrix} i_{dS} \\ i_{kdS} \end{bmatrix} = \begin{bmatrix} L_{dS} & M_{dS} \\ M_{dS} & L_{kdS} \end{bmatrix}^{-1} \begin{bmatrix} \lambda_{dS} - \lambda_0 \\ \lambda_{kdS} - \lambda_0 \end{bmatrix} \quad (3.40)$$

$$\begin{bmatrix} i_{qS} \\ i_{kqS} \end{bmatrix} = \begin{bmatrix} L_{qS} & M_{qS} \\ M_{qS} & L_{kqS} \end{bmatrix}^{-1} \begin{bmatrix} \lambda_{qS} \\ \lambda_{kqS} \end{bmatrix} \quad (3.41)$$

$$\mathcal{T}_{eS} = \frac{3p}{2}(\lambda_{dS}i_{qS} - \lambda_{qS}i_{dS}) \quad (3.42)$$

These equations represent a non-linear tenth-order system which, to the extent that our approximations are valid, encapsulates the entire dynamics of the system at all timescales. Note that a number of quantities, such as R_{kd} , have been assumed to be identical for the rotor and stator.

3.2.4 Model Simplification

The model of (3.25) through (3.42) is quite complicated, and as such presents problems for simulation and control applications. Firstly, the system is non-linear. While this is not necessarily a problem for simulation, it prevents the use of standard dynamic analysis techniques. Thus, while it is possible to simulate the model dynamics for any particular set of parameters, it is much more difficult to develop an understanding of how variations in those parameters affect the results without performing a very large number of simulations. Furthermore the non-linearity presents a thorny problem for standard linear control algorithms. Linearization of the system may be accomplished as usual with Taylor-series approximation of the non-linear terms. It is not clear that carrying out the linearization at this time is useful to the control algorithm of Section 4.1. Linearization of the equations will inevitably require more terms to describe the model, and it is therefore probably not helpful to understanding the machine dynamics. Thus, linearization is not carried out. If it is needed for later analysis, it is easily produced.

Secondly, the system has ten state variables, with a wide variety of timescales. This means that any simulation attempted will be quite slow. The high order of the system

inherently slows down the simulation because more state variables must be calculated with each iteration. Additionally, accurate simulation of short-timescale events will require a very small simulation timestep, but the overall simulation time-range will have to be fairly long to observe the final settling behavior of the long-timescale states. Obviously, the high order of the system also complicates control, but here the variation in timescales is advantageous. As long as the desired control does not regulate the short-timescale dynamics, they can largely be ignored in the control design.

Order Reduction

Reduction of the order of our model can be accomplished if we are willing to sacrifice accuracy on a short timescale. The intended purpose of this section is to develop a model appropriate for developing a controller. Our controller will attempt to phase-lock the output currents from the stator to the line voltage. Such a control goal means that control bandwidth cannot be more than an order of magnitude less than the AC line frequency. Thus, we are safe in discarding a number of relatively fast time-scale effects for purposes of control modeling.

The most obvious effects to discard, as described in [9], are the dynamics associated with the stator time-constants. This is accomplished by setting the left hand side of (3.29) and (3.30) equal to zero. The equations then become auxiliary, with the consequence that the order of the system is reduced by two.

In many permanent magnet motor applications, there is relatively little rotor conductivity, so the machines behave as though they have no damper winding. This would, in our case, allow the removal of four more equations: those which describe the evolution of the λ_k s. Given the construction of the magnet rotors described in Section 2.1.2, which contains a substantial amount of aluminum supporting the magnets, the validity of assuming little conductivity is dubious at best. It may still be the case that the transient behavior of the damper windings is fast enough to disregard. Such determination is best made empirically with a functional prototype.

Chapter 4

Control

With the models of Section 3.2.3 in hand, we can approach the problem of controlling the machine to behave as we desire. We must examine both the control objective for the machine and the means of accomplishing that objective. Section 4.1 describes a reasonable objective and lays out a high-level algorithm which should allow the system to reach that objective. Details of the control algorithm, which hinge on determination of model parameters, are left to be determined later. Section 4.2 considers the hardware implementation of the controller described in Section 4.1 using a microcontroller. Also considered are the peripheral electronics required to interface the microcontroller with the machine.

4.1 Algorithm

Given that the stator of the machine is connected directly to the 3-phase mains voltage, the objective of the machine controller is two-fold. The power delivered to the mains should be purely real, therefore the controller should ensure that the stator output current is in phase with the mains voltage. Secondly, the controller should control the machine shaft speed so as to extract an optimum amount of power from the prime mover.

The goal of the control algorithm can thus be concisely described as varying the frequency, amplitude, and phase of the current injected into the rotor so as to produce stator currents which are phase-locked with the mains voltage, while attempting to maintain a given shaft speed. This is perhaps best accomplished by breaking down the overall control problem into several sub-problems. Working backwards from the mains voltage, we can determine the appropriate orientation of fields between the stator and magnet rotor which

will produce the desired output current. The location of the fields is determined by the position of the magnet rotor, so we can then determine a correct location of the magnet rotor which corresponds to the desired fields.

We must then construct a position controller for the magnet rotor. This controller will determine both how to arrange the fields between the power rotor and the magnet rotor to position the magnet rotor correctly and how to inject currents into the rotor to produce those fields. Rotation of the machine shaft would then correspond to a disturbance torque on the magnet rotor which the controller must reject. Alternately, the driveshaft motion can be accounted for by feed-forward correction, which would lessen the requirements on the rotor controller.

An overall block diagram of the system is shown in Figure 4-1. A few points bear special notice in the diagram. The effects of the prime mover have been included in this diagram, as it is the characteristics of the prime mover which determine the relationship between shaft speed and shaft torque. As discussed in Section 4.2.1, it is most convenient to sense the magnet rotor angle relative to the driveshaft angle. This quantity, θ_{mm} , can then be added to the measured driveshaft angle, θ_{pm} , to obtain the absolute angle of the magnet rotor. Control of shaft speed is accomplished with a slow outer loop which adds an adjustment torque (\mathcal{T}_{adj}) to the torque commanded by the position controller.

The implementations of the blocks shown are fairly complicated in a number of cases. Discussion of the details of these implementations is beyond the scope of this thesis. It is expected that most of the control functionality will be implemented in a microcontroller, which allows the fairly complex numerical operations involved (e.g. trigonometric functions) to be easily implemented. At this point, the author takes the electrical engineer's privilege and declares the controller implementation SMoP (a Small Matter of Programming).

4.2 Electronics

The electronics used to implement the control algorithm of Section 4.1 comprise three parts: sensors to measure appropriate variables, control hardware which implements the algorithm, and a power converter to drive the rotor windings.

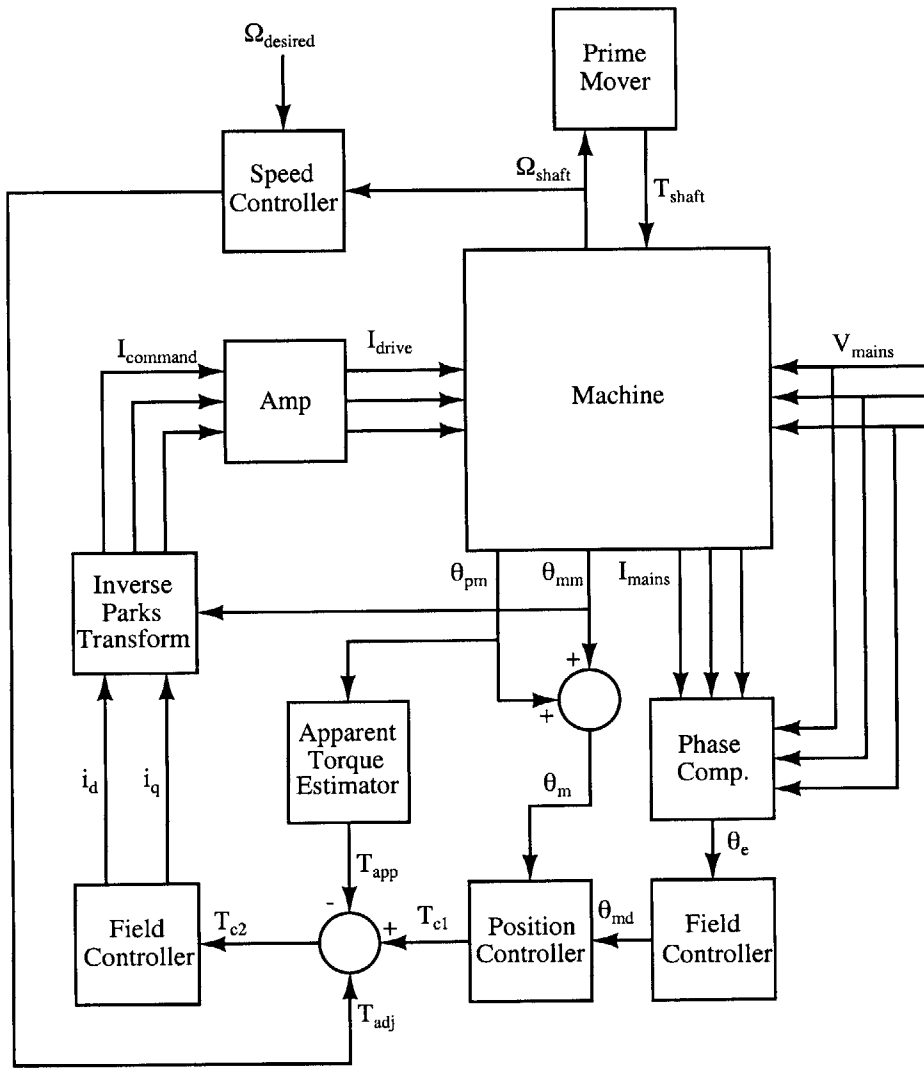


Figure 4-1: Machine Controller Block Diagram

4.2.1 Sensors

As indicated in Figure 4-1, the machine controller must sense a number of quantities to perform its function. Specifically, these quantities are the mains voltages, the stator output currents, the magnet rotor angle, and the driveshaft angle (Ω_{shaft} can be constructed by differentiating the shaft angle).

Voltage Sensors

Voltage sensing is relatively trivial, as the control variables external to the microprocessor, will be represented by voltages anyway. The two concerns that arise are scaling and isolation. The mains voltage will be on the order of several hundred volts (depending on the nature of the circuit to which the machine is connected, the exact value will change somewhat). Such a large voltage is not compatible with the components of the controller, so some scaling is required. It may also be the case that there is common-mode voltage between the mains neutral and the controller ground. Thus, some form of isolation may be desirable.

These two problems can be simultaneously solved with a standard isolation amplifier. The isolation amplifier (for example Analog Devices AD215AY) provides all three terminals of an input op-amp, and then galvanically isolates the output of the op-amp from the output of the package. Thus, by connecting the input op-amp to produce a small inverting gain (perhaps 0.01), we obtain an output voltage which is isolated and of reasonable magnitude.

Current Sensors

Sensing the stator currents is also not particularly difficult. We are guaranteed that the stator currents will be alternating at roughly 60 Hz, so simple current transformers will suffice for our measurement purposes. To get a voltage output from the transformer, we simply need to connect an appropriate resistance across each transformer secondary, with one side grounded. If the transformer turns ratio is $1:n$, then we should use a terminating resistance of $R_{term} \approx \frac{n}{20}$, which will produce 5 V output for a 100 A stator current. The current transformer already gives us isolation, so no further circuitry is required.

Angle Sensors

The controller described in Section 4.1 requires measurement of the angle of the driveshaft and the magnet rotor. Determination of the driveshaft angle merely entails connection of

an encoder of some sort to the driveshaft of the machine. For reasons of accuracy and simplicity of interface with the microcontroller, it is probably desirable that the angle be extracted with an optical encoder.

Mechanical concerns make determination of the magnet rotor angle substantially more complicated. The free-wheeling shaft is not easily accessible, and it is not therefore possible to mount a standard encoder to it. One possible solution is to mount a number of small permanent magnets to the inside of the freewheeling shaft. These magnets would be arranged in two bands around the circumference, with polarity alternating within each band, and the two bands offset by a quarter of the spacing within each band. Two coils mounted to the driveshaft—one aligned with each band—will then produce two periodic voltage waveforms which are in quadrature with each other. The frequency of each waveform will represent the relative speed of the free-wheeling shaft. The direction of motion can be ascertained by observing the relationship between the two waveforms. If band A leads band B, the motion is in one direction, and if band B leads band A, the motion is in the opposite direction.

4.2.2 Controller

The heart of the control electronics is a microcontroller which implements most of the functional blocks of Figure 4-1. Use of a microcontroller is more or less required, because many of the operations required in the control algorithm are messy (at best) to construct with analog circuitry.

Interfacing between the microcontroller and sensors will require A/D conversion. A total of 8 channels are required to digitize all of the sensed variables: 6 for voltage and current sense, and 2 to decode the magnet rotor angle sensor. It is assumed that the optical encoder used to determine driveshaft angle has a digital interface of some kind. Note that the sense variables are bipolar, so either the A/D converter used must accept bipolar input or additional level-shifting circuitry will be required.

Controlling the power converter requires 3 channels of D/A output. The required output frequency is quite low, so D/A conversion can be done with PWM on 3 digital I/O pins. Filtering the PWM with an RC lowpass should be adequate. The RC corner frequency can be as low as 60 Hz, so the ripple will be negligible.

As described, the control scheme does not require a particularly fast microcontroller. The maximum bandwidth of the controller should be considerably less than 60 Hz, so all

that is required is that the microcontroller be able to do the required calculations to produce an output command in much less than $\frac{1}{120}$ s. Using a slow microprocessor which executes one million instructions cycles per second, and allowing 1000 cycles for computation (which is *very* conservative) gives better than a factor of 8 above the Nyquist rate for the maximum conceivable bandwidth.

4.2.3 Power Converter

The power level required for the rotor excitation is fairly large. If we intend to operate over a 10% range of slip, then the rotor drive must handle at least 1 kW. Clearly, at such a power level, the only viable option is some kind of switching converter. To allow operation with both positive and negative slip (sub- and super-synchronous rotor speed), the converter must be able to manage net power flow into and out of the rotor.

It is also desirable that the power converter operate directly from the 3-phase mains voltage. This description implies that the required power converter is a bi-directional 3-phase frequency converter. There are essentially, two standard AC/AC converter topologies, cycloconversion and AC/DC/AC.

The principle behind cycloconversion is as follows. Suppose we are given an input 3-phase supply and 3 references which correspond to an output waveform. Between the input terminals and output terminals, we arrange a matrix of switches which allow connection of each of the output lines to any of the input lines. We then control the switches so that at any instant in time, each output is connected to whichever input is closest to the reference for that output. Cycloconversion is typically used at high power levels because the required switching frequency is on the order of the output frequency, so slow devices with high power-handling capability, such as SCRs, are usable. Unfortunately, the output has very substantial distortion, as the output voltage may be easily be in error by as much as 50% of the input amplitude at any time. Furthermore, it is somewhat difficult to implement current-mode output.

Fortunately, our desired power-handling capability lies easily within the range of commercially-available MOSFETs, so a high-frequency PWM AC/DC/AC scheme is viable. With such a converter, power conversion consists of rectification to produce a DC bus, followed by inversion to produce the desired AC output. Bidirectional power flow can be accomplished with a circuit as in Figure 4-2, which represents a general-purpose converter.




Figure 4-2: Bi-directional 3-phase AC/DC/AC converter

This circuit consists of two controlled 3-phase bridges back to back. If it is desired to transfer power from V_1 to V_2 , the switches on the left are controlled as a PWM rectifier. To the extent that the DC bus capacitor is large, and bus variation at the frequency of V_1 is acceptable, we can extract power from V_1 with unity power factor, which is a desirable feature. The switches on the right are controlled to produce the appropriate outputs at V_2 . Another feature of this scheme is the relatively ease of controlling the inversion switches to produce a controlled current at the outputs instead of a controlled voltage.

For purposes of demonstration with our machine, we may reasonably restrict operation to the sub-synchronous regime, and the converter design is thereby somewhat simplified. Specifically, the input PWM rectifier may be replaced with a simple diode rectifier. Additionally, the power rotor of the machine is an inductive load, so the output filter inductors of Figure 4-2 are not required. A simplified circuit which will then suffice is shown in Figure 4-3.




Figure 4-3: Sub-synchronous Rotor Drive Circuit

Due to the simple passive rectifier, this circuit does *not* draw power from the input with unity power factor, but this is not really a concern for our testing purposes. Input power factor can also be improved by reducing the size of the bus capacitor, or possibly eliminating it entirely. The output bridge will be controlled to produce currents which will probably be a balanced 3-phase set in steady state operation, so the overall current drawn

from the DC bus is approximately constant. This corresponds to an input power factor of better than 90%.

Converter Control

There is then an issue of how to control the output switches to produce the desired behavior. This problem is simplified if we assume that the machine will have no net neutral current, either because it is Δ -connected, or because it uses a buried-neutral Y-connection as suggested in Section 2.4.2. We must then measure the currents I_A , I_B , and I_C . This measurement should likely be performed with Hall-effect sensors to allow operation of the drive down to DC excitation, equivalent to operation of the machine up to synchronous speed.

The simplest control scheme is to then construct a hysteresis controller for each pair of switches in the output bridge. In principle, this involves turning on the top switch in the pair until the measured current exceeds the reference plus some hysteresis voltage. The top switch is then turned off and the bottom switch turned on until the current falls below the reference minus the hysteresis voltage. Such a control scheme results in a switching frequency which varies depending on the bus voltage and the commanded current, but this is not of concern, as the switching frequency (and its harmonics) will be well above all of the time constants in the system.

This control is relatively simply accomplished with the circuit of Figure 4-4. The circuit describes control of one of the phases; the other two are identical. Width of the hysteresis band and the transconductance of the amplifier (gain from reference voltage to output current) are set by the resistor values.




Figure 4-4: Hysteresis Phase-Current Control Circuit

Control of the power converter could also be accomplished with a similar algorithm in the microcontroller. However, obtaining high switching frequencies on the order of 100 kHz would require a substantial amount of processor power, and would likely introduce sizable errors due to Nyquist limitations.

Chapter 5

Conclusion

The objective of our project is the production and testing of a prototype generator to evaluate the feasibility of our proposed permanent-magnet DFIG. This work is well under way on a number of fronts, but some further construction and analysis will be required before the project ends. Section 5.1 recapitulates the design and construction work which has already been completed. Work remaining is described in Section 5.2. Finally, the status of the project is summarized in Section 5.3.

5.1 Completed Tasks

At the time of this writing, the majority of the project has been completed. Almost all of the design work has been accomplished at this point, and much of the construction is finished or in progress.

Fabrication of the structural elements of the machine is complete, with the exception of the new housing structure. The magnetic cores have been constructed, as have their windings. Performance of the cores and windings have been tested to a limited extent. Rotor cores have been attached to the rotor support plates. Final attachment of the windings to the cores is in progress.

The remainder of the construction has been almost entirely planned out. Designs of the new housing parts have been specified, and fabrication according to these designs has begun. Layout of the slip-rings has been considered, and two viable plans described. The decision of which to use is to be left until more is known about the exact available space inside the machine. The design of the machine shafts has been specified completely in terms of the

machine dimensions, which are to be measured as soon as construction of the rotors and housing is finished. Ventilation and mounting concerns have been addressed in a general fashion. More detailed analysis is left until the machine is assembled. Finally, a plan has been developed to allow assembly of the machine to proceed as simply as possible.

Some effort has also been expended to develop models of the machine. The behavior in steady-state has been shown to be similar to that of a permanent-magnet synchronous motor driving a permanent-magnet synchronous generator. The steady-state description has been extended to model dynamic conditions, with a fairly complete state-space model of the machine as the result. As the model is rather complicated, some consideration has been given to ways in which the model may be simplified for control applications.

Finally, the design of a controller for the machine has been considered. The controller specified aims to deliver maximum real power to a 3-phase voltage source attached to the stator windings by controlling excitation and (indirectly) shaft speed. An overall controller structure which accomplishes these aims has been determined, and the hardware required for the structure specified.

5.2 Further Work

A number of tasks remain before the project is completed. These tasks consist of implementation of plans which have been made, including, in some cases, specification of design details for those plans.

A few fabrication operations will be required before the machine is assembled, to produce the machine housing, shafts, stator mounting structure, and slip-rings. Also before final assembly, the power rotors must be completed, and the stator mounting structure attached. Once those operations are complete, the prototype machine may be assembled according to the plans which have been described.

In the area of modeling, further analysis should be done to more fully specify the operating characteristics of the machine. The equivalent of capability, compounding, and vee-curves should be produced. Simulation of the machine dynamics as specified by the state-space model presented should be performed. Such simulation results may then be first sanity-checked, and then compared to measurements of the basic open-loop behavior

of the machine. This will then allow both verification of the model accuracy and determination of appropriate model parameters.

Once the basic validity of the model has been established, a controller may be built as has been described. Construction of the controller will require selection of sensors and components as recommended in the controller design specification. Finally, before the controller is functional, appropriate program code must be written for the microcontroller. The code required has been only loosely described. An overview, in block-diagram form, of the behavior of the code has been provided, but nothing further. It is thus expected that the majority of the remaining intellectual work (as opposed to physical work of assembly) will be devoted to implementation of the controller functional description in the microcontroller code.

5.3 Summary

This thesis has described the design and implementation of a prototype generator that we hope will have desirable properties for wind-energy applications. Design decisions for parts completed and as-of-yet unfinished have been discussed. Additionally, a substantial amount of the analysis of the machine has been explicated, and a control system designed.

Although the work remaining before the project is completed is not trivial, it is reasonable to expect that it may be accomplished in the remaining time. It is our hope that in so doing we will produce proof of a useful new machine geometry. Development of a cost-effective, reliable, efficient wind-turbine generator could have substantial influence on our future energy production, hopefully leading to more sustainable generation solutions.

Appendix A

Structural Part Drawings

This appendix contains mechanical drawings of the structural parts. These drawings were produced with the intention that the parts be fabricated from the drawings. As the drawings were intended to occupy a full page, the figures of this appendix do not begin until the following page.

Figure A-1: Power Rotor Plate Drawing




Figure A-2: Magnet Rotor Plate Drawing




Figure A-4: Original Housing Drawing, Page 1





Figure A-6: New Housing Mounting Ring Drawing






Figure A-7: New Housing Stringer Drawing

Appendix B

Machine Photos




Figure B-1: Completed Stator Core




Figure B-2: Rotor Coil Winding Apparatus




Figure B-3: Wound and Partially Fiberglassed Stator




Figure B-4: Assembled Power Rotor, without Windings



Figure B-5: Power Rotor with Windings

Bibliography

- [1] A. J. Thomas, “Structural and Thermal Design of a Dual-Rotor, Constant-Frequency, Variable-Speed Generator,” B.S. thesis, MIT, 2003.
- [2] Z. Chen and E. Spooner, “Grid interface options for variable-speed. permanent-magnet generators,” IEE Proceedings — Electric Power Applications, vol. 145, pp. 273–283, Jul. 1998.
- [3] M. S. Vicatos and J. A. Tegopoulos, “Steady state analysis of a doubly-fed induction generator under synchronous operation,” IEEE Transactions on Energy Conversion, vol. 4, pp. 495–501, Sep. 1989.
- [4] R. Pena, J. C. Clare and G. M. Asher, “A doubly fed induction generator using back-to-back PWM converters supplying an isolated load from a variable speed wind turbine,” IEE Proceedings — Electric Power Applications, vol. 143, pp. 380-387, Sep. 1996.
- [5] L. Holdsworth, X. G. Wu, J. B. Ekanayake and N. Jenkins, “Comparison of fixed speed and doubly-fed wind turbines during power system disturbances,” IEE Proceedings — Generation, Transmission, and Distribution, vol. 150, pp. 343–352, May 2003.
- [6] A. E. Fitzgerald, C. Kingsley, S. D. Umans, Electric Machinery. New York: McGraw-Hill, 2003 (6th ed.).
- [7] E. Spooner and B. J. Chalmers, “‘TORUS’: a slotless, toroidal-stator, permanent-magnet generator,” IEE Proceedings — Electric Power Applications, vol. 139, pp. 497–506, Nov. 1992.

[8] J. L. Kirtley, Jr., "Doubly Fed, PM Excited Axial Field Generator." MIT Laboratory for Electromagnetic and Electronic Systems, Cambridge, MA. Report to National Renewable Energy Laboratory, 8 Dec. 2002.

[9] J. L. Kirtley, Jr., Class Notes for 6.685 Electric Machines, MIT 2003