Mixed-Domain Performance Model of the Piezoelectric Traveling-Wave Motor and the Development of a Two-Sided Device

by

Timothy Scott Glenn

S.B., Aeronautics and Astronautics
Massachusetts Institute of Technology, 1993

S.M., Aeronautics and Astronautics
Massachusetts Institute of Technology, 1995

SUBMITTED TO THE DEPARTMENT OF AERONAUTICS AND ASTRONAUTICS
IN PARTIAL FULFILLMENT OF THE DEGREE OF

DOCTORATE OF PHILOSOPHY
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2002

Signature of Author . . . .
Department of Aeronautics and Astronautics
May 16, 2002

Certified by . . . .
Nesbitt W. Hagood IV, Ph.D., Principal Research Engineer
Thesis Committee Chair

Certified by . . . .
Marthinus C. van Schoor, Ph.D., Lecturer
Thesis Committee Member

Certified by . . . .
David W. Miller, Professor of Aeronautics and Astronautics
Thesis Committee Member

Accepted by . . . .
Wallace E. Vander Velde, Professor of Aeronautics and Astronautics
Chair, Committee on Graduate Students

Mixed-Domain Performance Model of the Piezoelectric Traveling-Wave Motor and the Development of a Two-Sided Device

by

TIMOTHY S. GLENN

Submitted to the Department of Aeronautics and Astronautics
on May 16, 2002 in Partial Fulfillment of the
Requirements for the Degree of Doctorate of Philosophy
at the Massachusetts Institute of Technology

ABSTRACT

In recognition of the growing consideration of piezoelectric traveling-wave motors as suitable replacements for small-scale electromagnetic motors, the present work addresses two parallel objectives: (1) to develop an advanced modeling approach for the accurate prediction of steady-state performance of piezoelectric traveling-wave motors, and (2) to improve upon the typical piezoelectric traveling-wave motor configuration by investigating the novel implementation of two-sided operation.

Firstly, a performance model of the piezoelectric traveling-wave motor has been developed that accounts for loss at the material level through the integration of complex material constants. In effect, all model parameters can be attributed to physically relevant properties, and, as demonstrated by the success of several experimental correlation studies, motor performance can be predicted independently of the measured characteristics of an existing device. Other key features of the model include a dynamic rotor model, a hysteretic stick-slip friction contact model, and generalization to accommodate non-ideal traveling-wave excitation. Critical to the cohesion of the model, a mixed-domain formulation combines the ease and accuracy of deriving the nonlinear contact forces in the time domain and the efficiency of equilibrating the modal forces in the frequency domain.

Secondly, a two-sided motor configuration has been developed that is theoretically capable of doubling the torque and power output of the typical configuration with little or no increase in size or mass. Controlled experimental testing has been performed concurrently on essentially identical one-sided and two-sided prototypes, and the results verify a dramatic performance improvement due to the implementation of two-sided operation.

Thesis Supervisor:
Dr. Nesbitt W. Hagood
Department of Aeronautics and Astronautics

ACKNOWLEDGMENTS

TABLE OF CONTENTS

LIST OF FIGURES

11

LIST OF TABLES

Chapter 1

INTRODUCTION

The advancements in piezoelectric motor research in recent years have led to widespread recognition of their advantages and to growing commercialization. However, bringing a product to market often follows a more frantic timeline than the research behind the technology, so prototype iteration of the component technologies is often too costly. Product manufacturers must then rely on experienced engineers and packaged modeling tools for the efficient and timely design of components to meet their application's requirements. In recognition of the growing market for piezoelectric actuators, the work presented in this thesis is intended to provide the basis for such a modeling tool by developing a comprehensive model founded entirely on physical material properties and geometry. The model is tailored specifically to the dynamics of the piezoelectric traveling-wave motor, but the core approach is applicable to most other friction-driven piezoelectric motors. By adequately capturing the fundamental dynamics and material behavior, it is claimed that accurate prediction of motor performance is possible without having first realized a physical prototype. To put this in context, previously published models have all relied on the use of adjustment parameters or modal damping measurements in order to achieve accurate correlation. In other words, they all required measurements from an existing device.

Appealing to the need for a fully predictive simulation tool, a fairly generalized traveling-wave motor model has been developed that accounts for loss at the material level through the application of complex material constants. The approach incorporates key elements

from several previously published models, but the introduction of the lossy material properties sets this approach apart from the others. In effect, all model parameters can be specified by physically relevant properties, and, in theory, performance can be predicted accurately without the need to first assemble and test a working prototype. The use of complex numbers has necessitated harmonic analysis of the modal dynamics in the frequency domain, but to accurately capture the hysteretic stick-slip behavior of the frictional interface, a time-domain analysis of the contact dynamics is preferred. Therefore, a mixed-domain solution procedure has been implemented that concurrently supports the ease of simulating the nonlinear contact forces in the time domain and the efficiency of equilibrating the modal forces in the frequency domain. In doing so, it has been possible to integrate seamlessly many of the more advanced techniques from the previously published models while keeping the model as general as possible.

To demonstrate the validity and effectiveness of the proposed model, the analytical efforts have naturally been carried out in parallel with an equally substantial experimental endeavor. In addition to characterizing all the necessary complex material properties, several motors have been tested extensively to provide a controlled set of measurements for model correlation. Two of the motors are in-house prototypes and have been developed for the investigation of two-sided operation. This novel implementation of the traveling-wave motor concept aims to improve the torque and power density compared to existing configurations. Lastly, considerable focus has been placed on the importance of good rotor design by conducting a parametric rotor study for each of the tested motors. The rotor studies are directed at showing conclusively that the dynamics of the rotor are just as important as those of the stator in determining motor performance. While each of these experiments have been designed primarily to support the proposed model, the simultaneous efforts to improve the performance potential of traveling-wave motors also merit their own places as contributions to the field.

1.1 Motivation

The initial phase of this program was funded by the Jet Propulsion Laboratory (JPL) in support of a subtask to the NASA Mars Micro-Lander Dexterous Manipulator project. With a strong emphasis on mass savings, the direction of the subtask was to investigate alternative actuator technologies for use in the joints of the robotic manipulator arm. Until then, the tried-and-true practice in servo robotics applications had been the use of electromagnetic motors, but questions were raised as to whether newer technologies were available that could remedy some of their drawbacks. Specifically, though electromagnetic motors exhibit high power density and require relatively simple drive electronics, they also necessitate the added mass and complexity of a gearbox to transform the available power to higher torque at lower speed and a brake to supply holding torque. Furthermore, gearing introduces undesirable backlash, and the combined length and aspect ratio of the assembly can be awkward to integrate into applications such as a robotic arm. Direct drive would be very desirable but would require an alternative actuator technology with considerably different torque-speed characteristics than an electromagnetic motor.

1.1.1 Experimental Research

Piezoelectric ultrasonic motors were immediately identified as an attractive option for the manipulator application. Though the technology was relatively young at the time and had never been tested in an extraterrestrial environment, the power delivery of a piezoelectric motor was recognized as being ideally suited for intermittent applications requiring high torque at low speed. Power densities and efficiencies had been achieved which were comparable to those of geared DC motors (see Tables 1.1 & 1.3 later is Section 1.3.2), and the compact shape factors of certain implementations seemed more compatible to joint actuation. Indeed, a number of the features of the piezoelectric motor were considered desirable:

- Power is produced in the form of high torque at low speed, thus eliminating the need for speed reducing gears.
- Owing to a friction-based drive principle, even higher holding torque is maintained when unpowered, thus eliminating the need for a separate brake.

- In contrast with electromagnetic motors, peak efficiency is achieved in the upper torque band, i.e. near stall.

- Negligible backlash aids in precise positioning.

- Low inertia of the rotating components contributes to quick transient response (start/stop characteristics).

- Simple construction and flexibility in the selection of vibration modes promotes flexibility in choosing the motor shape, e.g. hollow for cable transport.

- Ultrasonic drive frequencies and low rotational speeds result in quiet operation and low transmission of vibration.

- Electromagnetic fields do not interfere with operation, nor are they directly produced.

Though it was not of particular interest to the immediate task at hand, the simple structure of most piezoelectric motors is also conducive to miniaturization.

The decision to pursue piezoelectric motor research was not without doubt, however, because certain disadvantages can be associated with their use:

- A high-voltage, high-frequency power source is required.

- Life span is limited by the wear of the frictional interface material.

- Increasing scale can be difficult because of the use of vibration modes.

Furthermore, the technology had never been proven at the cryogenic temperatures that would be experienced during the mission. Nevertheless, the advantages sparked enough interest to warrant an in-depth investigation into the viability of using piezoelectric actuators in the Micro-Lander Dexterous Manipulator project.

At the time, the commercial availability of piezoelectric motors was greatly limited. In fact, nothing was available that could meet the direct-drive torque requirements of the manipulator arm. The directive agreed upon by this laboratory, therefore, was to improve the direct-drive torque capability of a piezoelectric motor while keeping mass to a minimum. Attempting first to meet requirements of the wrist joint, a stall torque of 200N-cm was identified as the primary goal.

With the need for bidirectional operation and the desire for a more suitable aspect ratio, attention was focused specifically on the disk-type piezoelectric traveling-wave motor. Compared with other types of piezoelectric motors, the traveling-wave motor is inherently geared toward high torque density and low speed. Its flat shape acts to increase the length of the lever arm of the friction forces, thus resulting in greater torque output and slower rotational speed than more cylindrical implementations. Plus, its pancake-like aspect ratio is ideal for joint actuation as it can be affixed directly in the axis of rotation without protruding as much as a long actuator. A handful of standing-wave motor designs [Kumada, 1985; Sashida, 1993] that offered higher power density and efficiency than traveling-wave motors were available at the time, but their operation was limited to one direction. The rotation of the traveling-wave motor, on the other hand, is bidirectional. By inverting the signal to one of its two electrical inputs, the progression of the traveling-wave is reversed. Furthermore, the idea of the traveling-wave motor had been conceived to remedy the disappointing life span of standing-wave piezoelectric motors. Whereas wear is high in a standing-wave motor due to intermittent contact, contact in a traveling-wave motor is continuous, greatly reducing wear and extending its serviceable life.

Rather than simply scale the existing motor technology for the application at hand, however, the pressure to minimize actuator mass motivated the need for a significantly novel approach to the piezoelectric traveling-wave motor concept. In the typical traveling-wave motor, the motive force is derived from the elliptical displacement trajectories on one side of a flexurally vibrating ring or disk. The configuration and assembly are elegant, but an untapped source of motive force was identified in the vibrations of the opposing surface. Specifically, it was theorized that the torque and power output of a traveling-wave motor could be doubled, without significantly increasing its mass, by contacting both sides of the vibrating ring or disk. Consequently, two-sided operation was conceived for investigation in this thesis work, and experimental research commenced with the development of the first two-sided traveling-wave motor prototype.

Early progress was promising, and the program continued at JPL with performance and survivability testing of piezoelectric traveling-wave motors at cryogenic temperatures and vacuum [Bar-Cohen, 1998]. But before the experimental development of the two-sided motor was fully carried out to fruition, the subtask funding drew to a close. Nevertheless, with strong internal interest to see the two-sided motor succeed, the developmental research was continued after separation from JPL, later expanding to include a study of the importance of the rotor dynamics. Along the way, it was recognized that there was a need for a better analytical model of the traveling-wave motor. Existing models hadn't included the dynamics of the rotor, nor had they anticipated the implementation of two-sided operation with two rotors. Along with a practical interest in simulating non-ideal operation, developing a new model soon became the primary focus of this thesis research. The experimental work, therefore, would not only be recognized for its own merits, but it would also serve as a controlled testbed to which the new model could be applied for validation.

1.1.2 Analytical Research

As just suggested, the decision to pursue advanced analytical modeling of the traveling-wave motor was motivated by the lack of completeness and generality in existing models (see Section 1.3.3). At the time, the rotor had always been modeled as being rigid even though experiments had shown that it, too, supported a traveling wave during operation. Furthermore, although the compliance of the contact interface had spurred the development of hysteretic stick-slip contact models, no one had yet successfully integrated these advanced contact models into a full motor model. Instead, existing models still relied on simplistic contact mechanisms that assumed pure slip through Coulomb friction. Other shortcomings included the use of lumped modal damping coefficients and even fixed, pre-determined modal amplitudes.

How was it possible, then, that any of these models were published with satisfactory correlation to experimental data? The acceptable truth at the time was that each of the models relied on some combination of model adjustment factors to achieve apparent correlation. In

other words, the simulated data would be 'fit' to the experimental data by tweaking one or more unknown parameters such as the interface stiffness or damping coefficients. This approach, while allowing small extrapolations from an existing design, does not permit accurate prediction of the absolute performance of an unrealized motor. Nevertheless, earlier models were certainly capable of demonstrating relative performance trends and, therefore, might have been used for improving upon a current design.

Understandably, there was need for a better model that did not depend on post-production measurements to achieve accurate correlation. As a design tool for commercialization, it would at least compliment any experimental development efforts and, if implemented correctly, could supplant the need for iterative prototyping. After all, parametric studies can be carried out in much less time via simulation than through experimentation. If the model were made computationally efficient, automated design optimization would even be possible, e.g. via a technique such as a genetic algorithm. As an academic tool, an accurate model would provide the means for answering certain questions that are difficult or too costly to resolve through experimental study. For example, it would be useful to understand why the efficiency of traveling-wave motors has been limited to about 50% in laboratories.

It was decided, therefore, to pursue the advancement of piezoelectric traveling-wave motor modeling in the hopes of eliminating the need for model adjustment factors. The new model would include rotor dynamics and the hysteretic friction contact model. It would also eliminate lumped modal damping coefficients by attributing loss to the materials directly. Furthermore, there was experimental evidence that supported the need to account for non-ideal wave generation in the model. Specifically, each successive prototype exhibited some degree of mode splitting due to asymmetry that impeded the generation of an ideal traveling wave. It was not understood at the time how much mode splitting was affecting performance, if at all, so the desire to simulate "realistic" performance motivated the need to keep the model general in terms of the independently excited standing-wave modes. Finally, the new model would be capable of simulating the two-sided motor and, in theory, would sup-

port the experimental data in showing the performance boost of two-sided operation over one-sided operation.

1.2 Objectives

Beyond the scope of the initial JPL subtask, therefore, the work accomplished in this thesis was defined by the following two main objectives:

1. to improve the potential of the piezoelectric traveling-wave motor by investigating novel two-sided operation, and
2. to develop an advanced modeling approach for the accurate prediction of piezoelectric traveling-wave motor performance.

The specific approach taken to achieve these goals will be described in detail following a discussion of the early history of the traveling-wave motor, its working principle, and a review of the relevant work accomplished by others.

1.3 Background

The concept of extracting rectified motion from piezoelectric-induced vibrations dates at least as far back as 1948 when Williams and Brown filed a U.S. patent describing the general principle behind a piezoelectric motor [Williams, 1948]. Unfortunately, due to the lack of suitable piezoelectric materials and the corresponding power electronics at the time, they were prevented from realizing their idea. It wasn't until decades later with the development of high-power piezoceramics that successful implementations began to appear worldwide. The earliest experimental realization of a piezoelectric motor is generally attributed to Barth, also from the U.S., who in 1973, successfully rectified the ultrasonic vibrations of longitudinal vibrators by obliquely impacting the periphery of a circular rotor [Barth, 1973]. The general configuration of Barth's invention, crude by today's standards, is depicted in Figure 1.1. Numerous others soon followed with slight variations on this basic principle, most of the work being done in the Soviet Union [Vishnevsky, 1975; Gromakovskii, 1978; Vasiliev, 1979].
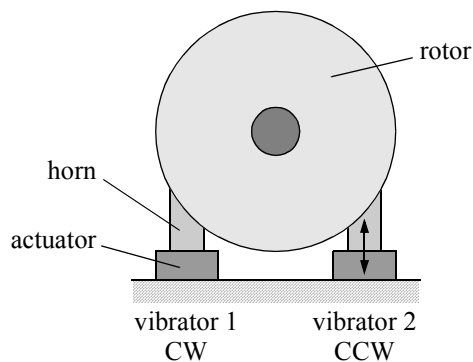
Figure 1.1 Principle of Barth's motor. Only one vibrator is driven at a time.
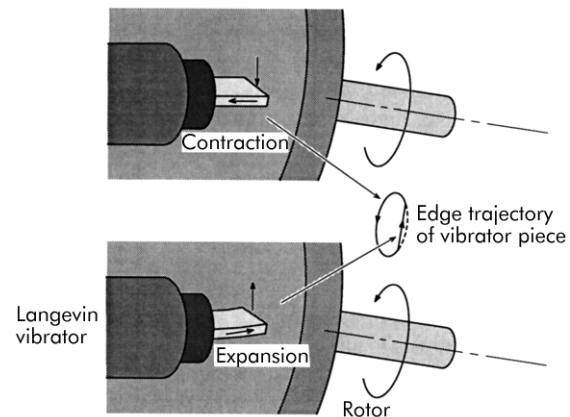
Figure 1.2 Principle of Sashida's wedge-type motor.

The most noteworthy implementation using longitudinal vibrators came from Japan in the early 80s when Sashida developed the wedge-type motor [Sashida, 1980; Sashida, 1982]. The principle of this far superior design in depicted in Figure 1.2. By amplifying the vibration amplitude with the use of a bolt-tightened Langevin vibrator [Sashida, 1993] and orienting the impact angle closer to perpendicular, sliding at the interface was considerably reduced, thus greatly improving efficiency. In fact, a maximum efficiency of 87% was achieved with a single vibrator under carefully controlled conditions. A lower, yet still impressive, efficiency of 60% was achieved with a prototype motor with 36 vibrating elements.

Regardless of its exceptional performance, however, the wedge-type motor still had one major drawback: it had a remarkably short life span due to rapid wear of the contact surfaces. In light of this, Sashida promptly directed his efforts away from the violent impact mechanism of the wedge-type motor and, by 1982, had completed a working prototype of his most acclaimed contribution to the field: the piezoelectric traveling-wave motor [Sashida, 1983]. The repeated impact of discrete vibrator tips was essentially supplanted in favor of continuous contact at the peaks of a traveling flexural wave. In doing so, wear was greatly reduced and life span extended. Today, through careful selection of low-wear materials, the life spans of traveling-wave motors are comparable to those of typical brushed DC

motors. Unfortunately, the efficiency of traveling-wave motors have generally been limited to 50% in laboratory studies. Nevertheless, the improved reliability of Sashida's traveling-wave motor concept spiked interest in the search for potential applications.

In early 1983, Canon, Inc., began fundamental development of the traveling-wave motor after having consulted with Sashida at Shinsei Industries, the company Sashida had founded a decade and a half earlier and where he had gained the invaluable experience in precision lapping that led to his success in ultrasonic motor development. Four years later, in 1987, Canon introduced to the world its new line of autofocus camera lenses driven by ring-type, traveling-wave ultrasonic motors [Hosoe, n/a]. This is probably the most widely known application of traveling-wave motors in today's market. With work being done by such large consumer goods manufacturers as Matsushita (Panasonic) [Panasonic, n/a], Toyota [Kawai, 1995], and Daimler-Chrysler [Schreiner, 2000], the future of the traveling-wave motor looks very promising.

1.3.1 Working Principle

The speed and torque of a piezoelectric rotary traveling-wave motor are derived from frictional contact forces acting at the interface of two bodies, one of which has a traveling flexural wave excited about its circumference by piezoceramics. The displacements at the surface of the excited body follow elliptical trajectories, the velocity of which are horizontal, or in-plane, at the peaks of the wave (see Figure 1.3). The active structure is most often called the stator, the passive structure, the rotor. By preloading the rotor into contact with the wave peaks of the stator, the ultrasonic vibrations are rectified via the frictional motive force into continuous, rotary stoke of the rotor. As with most piezoelectric ultrasonic motors, performance is maximized by driving the system near resonance.

Figure 1.3 specifically illustrates the working principle behind the novel two-sided configuration where both surfaces of a symmetric stator are utilized to increase torque output. Note that although the passive rotors are shown rigid in the illustration for simplicity, they
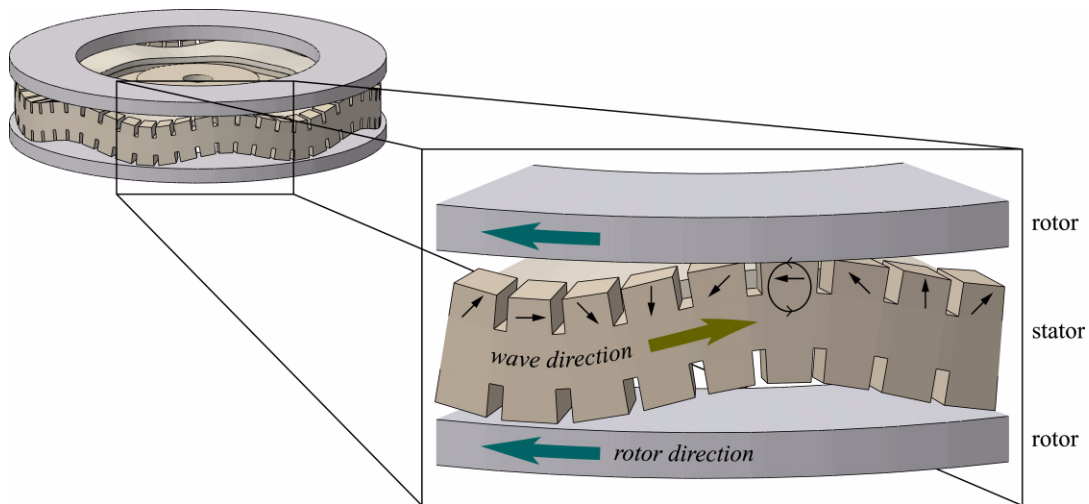
Figure 1.3   Principle of motion of the traveling-wave motor showing instantaneous tooth tip veloc-
ities (arrows) and elliptical surface displacement trajectory (not to scale). A symmetric,
two-sided configuration is depicted.

do flex in response to the spatially periodic contact forces. In fact, their dynamic behavior
plays a very important role in the motor's performance.

The rendered cutaway of the experimental one-sided prototype in Figure 1.4 helps to illus-
trate the fundamental components and assembly of a typical piezoelectric, rotary traveling-
wave motor. At the heart is the stator, a comb-toothed structure to which are bonded one or
more piezoceramic actuators for exciting a traveling wave about its circumference. The
teeth provide a means for increasing the thickness of the stator for greater amplification of
the transverse surface displacements without significantly increasing flexural stiffness. In
contact with the stator is the rotor. A thin ring of polymer is typically bonded to the rotor
at the contact interface as a necessary means of providing local compliance. This helps to
avoid contact instability and reduce wear when the motor is driven at large amplitudes. A
Belleville disk spring is shown here as the means for preloading the rotor against the stator.
Within limits, the available torque is related to this axial preload force and the coefficient
of friction at the contact interface.

As illustrated in Figure 1.5(a), excitation of the traveling wave is typically realized by the
phased superposition of two standing waves. This is easily achieved with the piezoelectric
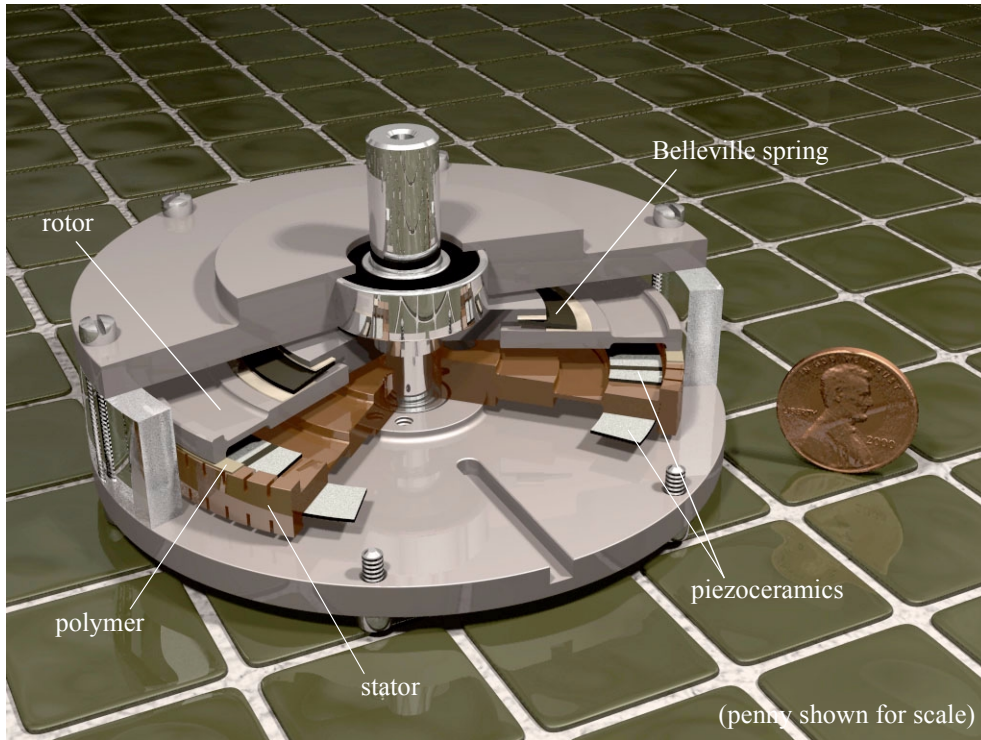
Figure 1.4 Rendered cutaway of the one-sided prototype, an example of a $B_{0,n}$ disk-type motor with symmetric stator.

ring or disk motor configuration because the circular symmetry exhibits degeneracy of the circumferential modes. Two piezoceramic arrays, bonded to the stator substrate and identified in the figure as A and B, are etched and poled so as to spatially coincide with the desired pair of orthogonal degenerate modes of circumferential wavelength λ. Specifically, each array consists of segments of width $\lambda/2$ and alternating polarization, and the two arrays are offset from one another by $\lambda/4$. Figure 1.5(b) shows the typical configuration of the piezoceramic arrays in a ring or disk motor, whether they share the real estate of a single piezoceramic washer or comprise two full rings. Mathematically, how a traveling wave is derived from orthogonal standing waves can be explained by the trigonometric identity

$$\cos N\theta \cos \omega t + \sin N\theta \sin \omega t \ = \ \cos(N\theta - \omega t), \tag{1.1}$$

where $N = 2\pi/\lambda$ describes the wave number or, rather, the number of wavelengths around the circumference of the stator. The two arrays are driven temporally out of phase, i.e. in
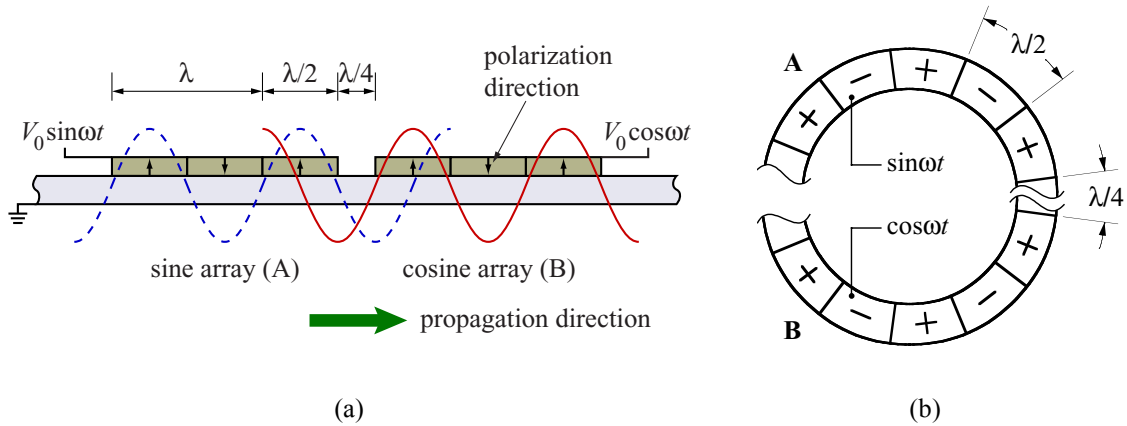
Figure 1.5   (a) Conceptual illustration of a traveling wave excited by piezoceramic arrays via orthogonal mode superposition. (b) Typical configuration of the arrays in a ring- or disk-type motor.

this example by $\cos\omega t$ and $\sin\omega t$, respectively, and a traveling wave results. The principle of motion for a linear motor is very much the same, but the traveling wave is generally realized through more difficult means.

1.3.2 Performance Review

A great variety of traveling-wave motors have been realized in the two decades since Sashida first proposed the concept. The performance metrics of several of these are compared in Table 1.1. Some, such as those by Canon, Shinsei, and Matsushita, have been available commercially for several years either as stand-alone products or as actuators built into other consumer goods. The models investigated in this work include the Shinsei USR60 and the two in-house prototypes by Glenn. Accounting for scale, the three most relevant performance metrics being used for comparison are the normalized values of efficiency (ratio of mechanical power out to electrical power in), torque density (stall torque per unit mass), and power density (output power per unit mass).

To make a fair comparison, the motors have been categorized by stator configuration, the two most common being the ring-type motor and the $B_{0,n}$ disk-type motor. Examples of these are shown in Figure 1.6(a) and Figure 1.4, respectively. In accordance with common labeling of the bending modes of an annular disk, the term $B_{0,n}$ refers to this class of disk

TABLE 1.1 Performance Comparison of Piezoelectric Rotary Traveling-wave Ultrasonic Motors

	Maker	Model or Description	Stator Dia. (mm)	Mass (g)	Stall Torque (N-cm)	No-load Speed (rpm)	Max. Power (W)	Peak Eff. (%)	Torque Density (N-cm/kg)	Power Density (W/kg)
Ring	Flynn[a]		8	0.26	0.054	1750	0.03	n/a	210	108
	Matsushita[b]	R36	36	12	2.1	270	2.1	45	175	19.2
	Canon[c]	UA80	60	235	49	150	3.1	n/a	209	13.2
	Canon[d]	EF 300/F2.8 L autofocus lens	77	45	16	40	0.26	35	349	5.8
$B_{0,n}$ Disk	Canon[c]	UA40	<40	55	7.8	500	1.9	n/a	142	33.6
		UA60	60	195	49	150	3.1	n/a	251	15.8
	Shinsei[e]	USR60	60	240	50	145	3.3	24	208	13.9
	Toyota[f]	symm., 1 rotor	67	n/a	207	200	24	39	—	—
	Glenn[e]	symm., 1 rotor	76	420	120	97	7.5	29	286	17.8
		symm., 2 rotors	76	350	160	116	12.2	27	457	34.8
	Daimler-Chrysler[g]	robot wrist actuator	90	480	620	240	50	20	1290	104
$B_{1,n}$ Disk	Matsushita[h]		10	3.5	0.25	1200	0.04	20	70	11.4
	Matsushita[i]	USM-17D	17	6	0.78	700	0.23	n/a	131	38.3
		USM-40D	40	70	7.8	800	2.6	45	112	37.1
		USM-80D	80	400	59	500	12.3	n/a	147	30.8
Rod Actuators	Petit[j]	symm., 2 rotors	20	40	11	250	1.3	15	275	32.5
	Petit[k]	symm., 1 rotor	45	240	40	115	1.9	15	167	7.9
		symm., 2 rotors	45	280	80	120	4.0	25	286	14.3
	Guyomar[l]	"butterfly"	60	n/a	250	90	10	n/a	—	—

a. [Flynn, 1995]

b. [Inaba, 1987]

c. Canon website, *http://www.canon.co.jp/Motor/ua-e/*, last viewed June 1998

d. [Hosoe, n/a]

e. As tested in-house and presented in Chapter 4

f. [Kawai, 1995]

g. [Schreiner, 2000]

h. [Kawasaki, 1993]

i. [Panasonic, n/a]

j. [Petit, 2001]

k. [Petit, 1998]

l. [Guyomar, 1996]

(a)   (b)   (c)
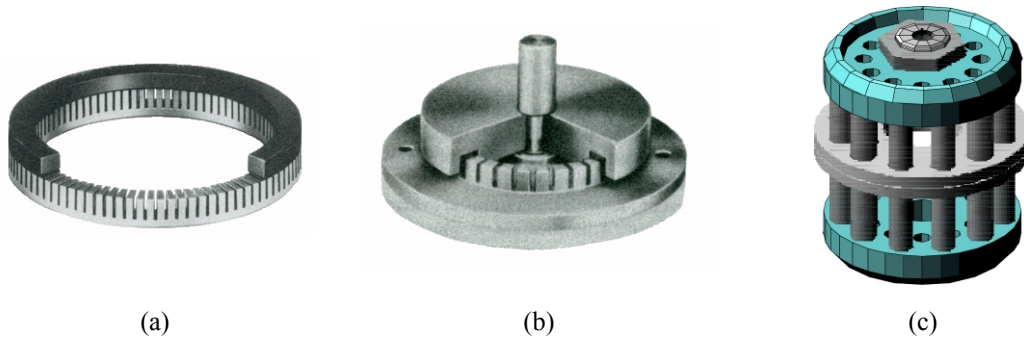
Figure 1.6   Types of traveling-wave motors in addition to the $B_{0,n}$ disk-type motor: (a) ring-type (Matsushita), (b) $B_{1,n}$ disk-type (Matsushita), (c) stator with piezo rod actuators (Petit).

motors as being configured for the excitation of stator bending modes with n nodal diameters and zero nodal rings. All of the motors investigated in this work are of this class. Canon developed a series of ring-type motors for use in their autofocus lens application because the hollow structure integrates well with the internal lens system. However, a ring-type stator must be supported by felt or some other soft backing and so is generally not well suited to accommodate the large preload needed to produce high torque. The $B_{0,n}$ disk-type motor, on the other hand, overcomes this disadvantage by supporting the ring by a thin flange connected to an inner mounting hub. Without significantly affecting the flexural dynamics of the chief outer ring, the flange can react substantially higher preload than felt. It also provides the necessary torsional rigidity for efficiently reacting the torque load to the fixed stator hub. Consequently, the disk configuration is a much better choice when torque output weighs more heavily than shape factor.

In contrast to the $B_{0,n}$ disk-type motor, the motor in Figure 1.6(b) by Matsushita is representative of the class of $B_{1,n}$ disk-type motors which operate via bending modes with one nodal ring. Though this configuration never quite reached the popularity of the $B_{0,n}$ motor, it brings with it a very purposed advantage over the latter. The problem with $B_{0,n}$ modes is that the circumferential and radial curvatures oppose each other, thus limiting the effective coupling that a bonded piezo wafer has to this type of mode shape. In the $B_{1,n}$ motor, on the other hand, the piezo actuators are placed inside the nodal ring where the two curvatures are in phase. Through greater electromechanical coupling, this configuration is considered to

be capable of higher efficiency and mechanical output power. The drawback to the design, however, is that the teeth are moved inward to the radius of maximum modal amplitude. In doing so, torque output is greatly reduced while rotational speeds are increased. The benefit of improved actuator coupling is questionable when the loss of effective "gearing" is considered.

The last type of traveling-wave motor included in the comparison is a relatively recent introduction. Whereas each of the previously described motor types use bonded piezo wafers to excite bending modes in the stator, motors of this not-yet-classified type of traveling-wave motor use prestressed piezo rod actuators. Though not necessary, a stator assembled with rod actuators is more naturally configured in a symmetric fashion. An example of such a symmetric stator is shown in Figure 1.6(c). The construction of a stator with rod actuators may seem more complex, but the inventors suggest that the use of rods over wafers should reduce cost and, through improved piezo stack technologies, should greatly reduce the operating voltage. As seen in Table 1.1, realized performance is already on par with some of the more traditional types of traveling-wave motors. So as to take full advantage of the symmetric stator design, Petit *et al.* have even gone as far as testing their stators with two rotors, similar to the proposed two-sided implementation of the disk motor presented in this work. Their results clearly support the benefits of two-sided operation, i.e. twice the torque and output power with little added mass.

As far as achieving high torque and power is concerned, two prototypes from the literature stand out. The first is that of Kawai *et al.* from Toyota, who specifically aimed to improve the power throughput of the conventional $B_{0,n}$ disk-type motor by increasing the volume of active material. Coincidentally, they arrived at a motor configuration that was essentially identical in construction to that of the one-sided prototype shown in Figure 1.4. By bonding piezoelectric elements to both sides of a symmetric stator, they claimed it was possible to increase the electrical input power and, thus, the vibration amplitude and motive force. As seen in Table 1.1, the symmetric stator approach proved worthy with a 70-mm diameter motor producing a peak output power of 24W and a maximum torque of 207N-cm, all while

achieving a peak efficiency of 39%. Their work was published soon after the onset of the in-house development of the two-sided motor and further supported the decision to pursue a symmetric stator assembly. Strangely, however, Kawai *et al*. did not propose the possibility of two-sided operation even though their stator had teeth on both sides as well.

More impressive is the recent work by Schreiner *et al*. at DaimlerChrysler. The outstanding performance of their conventional prototype is in a class of its own and truly demonstrates what can be achieved when every aspect of the traveling-wave motor is thoroughly investigated. Specific areas of experimental optimization included the axial preload, the friction material, the piezoceramic material, and bonding. Furthermore, as the axial preload and friction material stiffness were increased for greater torque output, measures were taken to damp the parasitic oscillations of the stator-rotor contact. In doing so, the limit of operation with respect to drive frequency was shifted closer to resonance, and the motor could be driven at significantly greater amplitude and input power. Setting a new bar in traveling-wave motor performance, Schreiner *et al*. developed a 90-mm diameter prototype that was ultimately capable of producing a mechanical power of 50W with a maximum torque of 620N-cm. Efficiency, however, was somewhat limited at 20%.

A curious note on the work by Schreiner *et al*. pertains to the stator mode-splitting problem that will be discussed throughout this thesis. They understood that optimal motor behavior first required that an ideal traveling wave be generated in the stator and that doing so necessitated perfect stator symmetry. However, it appears that they, too, battled stator asymmetry. In reference to the results that were being presented, they admitted that an ideal traveling wave was achieved by "a proper selection of stators with perfect symmetry." In other words, a number of stators had to be manufactured in order to yield at least one that behaved appropriately.

A reasonable degree of traveling-wave imperfection was also observed and commented on by Petit *et al*. and Guyomar *et al*. with regard to their rod-actuator motors. Interestingly, Guyomar *et al*. suggested that an advantage of using rod actuators is that individual pre-

stressing can be used to counter the effects of stator asymmetry. Whatever the cause, it has been demonstrated that stator asymmetry leads to non-ideal wave behavior and to less-than-optimal performance. And even though a prototype might be manufactured with near-perfect symmetry, it is unlikely that the same would be true for commercial mass-production. Therefore, to better understand the effect stator asymmetry has on motor performance, non-ideal wave generation has been accounted for in the proposed model.

Trends with Scale

Except for a few outliers in Table 1.1, particularly those with outstanding performance, a number of trends are readily perceived in the performance metrics as a function of scale. Of course, trends with scale are most valid within a particular classification of motor, even more so for scaling by the same manufacturer. With that in mind, it is noticed that torque density is relatively constant across scale whereas power density goes down with increasing dimension. With mass logically scaling as the cube of dimension, this would indicate that torque output does as well. A simple analysis helps to explain these observed trends with scale and provides a quick tool for sizing a motor to meet one's needs.

Consider a typical motor with the general dimension, D, e.g. its diameter. To a good approximation, the maximum achievable stall torque of a (one-sided) rotary motor is equal to the product of the dynamic coefficient of friction, μ, the axial preload force, $f_{preload}$, and the radius of contact, r_c:

$$\tau_{stall} \approx \mu f_{preload} r_c . \tag{1.2}$$

The appropriate contact pressure is assumed to be independent of size, so the axial preload force is generally proportional to the contact area, i.e. the square of the dimension. Needless to say, the radius of contact scales proportionally with the dimension. By this argument, the maximum achievable stall torque scales with the dimension cubed, or

$$\tau_{stall} \propto D^3 . \tag{1.3}$$

The rotational speed is proportional to the product of vibration amplitude and drive frequency divided by the radius of contact. As the amplitude generally scales linearly with the dimension, and resonant frequencies scale inversely with dimension, the no-load speed should scale inversely with the dimension:

$$\omega_{no\text{-}load} \propto \frac{1}{D}. \tag{1.4}$$

Mechanical output power is the product of torque and speed, so putting (1.3) and (1.4) together,

$$P_{out} \propto D^2. \tag{1.5}$$

As expected, torque and power increase dramatically with scale. However, greater absolute performance is achieved at the cost of increased mass which grows by the dimension cubed. Consequently, all other things being equal, the normalized metric of torque density is presumed constant across scale. Power density, however, scales with the inverse of dimension. Consequently, as with many types of actuators, the power density of ultrasonic motors can be improved through miniaturization. The relationships between the motor characteristics and dimension are summarized in Table 1.2.

TABLE 1.2   Scaling: Metrics vs. Dimension

Property	Symbol	Scale
dimension	D	–
stall torque	τ_{stall}	D^3
no-load speed	$\omega_{no\text{-}load}$	D^{-1}
output power	P_{out}	D^2
preload	$f_{preload}$	D^2
mass	m	D^3
torque density	$\tau_{stall}\,/\,m$	constant
power density	$P_{out}\,/\,m$	D^{-1}

Compared to Electromagnetic Motors

To put the performance metrics of the traveling-wave motors in Table 1.1 in perspective, they should be compared with those of traditional electromagnetic motors. Therefore, a representative cross-section of comparably-sized permanent magnet (brushed) and brushless DC motors has been compiled from manufacturer product catalogs. The performance metrics of these motors are summarized in Table 1.3. As electromagnetic motors produce relatively low torque at high speed, most applications require that the output first be geared down so as to produce higher torque output at more manageable speeds. Thus, to make this comparison more valid, each of the motors has also been listed in one or more relevant gearhead configurations.

Without gearing, DC motors can be remarkably efficient and have very respectable power densities. Unfortunately, gears not only contribute additional mass, but they also introduce considerable loss and sometimes undesirable backlash. The mass alone greatly reduces the power density, and it is reduced even farther by the inefficiency of the gears to transform the power. Consequently, the efficiency of a geared motor is notably less than that of the motor alone, and its power density is significantly so.

Looking first at the permanent magnet DC motors that have been geared to produced torque densities on the order of those in Table 1.1, it is seen that traveling-wave motors compare reasonably well with respect to efficiency and power density. The permanent magnet DC motors may have a slight edge on paper, but the effect of backlash should be considered as well. Brushless DC motors, however, are generally more efficient than their brushed counterparts and often produce greater power per unit mass. Consequently, even with gearing, the brushless DC motors in Table 1.3 clearly outperform the average traveling-wave motor when compared solely on the merit of peak metrics.

Peak performance metrics, however, don't tell the whole story because where the peaks occur can be just as important as the absolute values themselves. To examine this, typical performance curves for electromagnetic (EM) motors are plotted in Figure 1.7 against those

TABLE 1.3 Performance of Comparably Sized Electromagnetic (EM) Motors

Type	Maker	Model	Mass (g)	Stall Torque (N-cm)	No-load Speed (rpm)	Max. Power (W)	Peak Eff. (%)	Torque Density (N-cm/kg)	Power Density (W/kg)
Permanent Magnet DC Motors	MicroMo	2338-024S[a]	70	1.8	7600	3.58	64	25	51
		2338-024S-30/1 *66:1 gear ratio*	241	81	115	2.44	45	337	10.1
		2338-024S-30/1 *134:1 gear ratio*	273	141	57	2.10	39	518	7.7
	MicroMo	GNM 2145[a]	350	10	7260	19.0	50	29	54
		GNM 2145-G5 *30:1 gear ratio*	550	242	242	15.3	40	440	28
	Pittman	8712[b]	127	3.6	7730	7.28	52	28	57
		GM8712 *30.9:1 gear ratio*	202	73	250	4.76	34	361	24
		GM8712 *95.9:1 gear ratio*	206	202	81	4.28	31	980	21
	Pittman	9413[b]	255	11	5600	16.1	66	43	63
		GM9413 *38.3:1 gear ratio*	432	308	146	11.8	48	712	27
	Barber-Colman	EYQM-33300-51[c]	150	2.7	6520	4.61	50	18	31
		EYQF-33300-640 *36.2:1 gear ratio*	210	55	180	2.59	36	262	12
		EYQF-33300-650 *114:1 gear ratio*	210	151	57	2.25	32	721	11
Brushless DC Motors	MicroMo	3153-024BRE[a]	157	3.6	5000	4.45	58	23	30
		3153-024BRE-30/1 *43:1 gear ratio*	328	108	116	3.28	41	330	10.1
	MicroMo	3056-024B[a]	190	9.3	8200	20.0	73	49	105
		3056-024B-30/1 *14:1 gear ratio*	329	104	586	16.0	58	317	48
		3056-024B-30/1 *43:1 gear ratio*	361	280	191	14.0	51	775	39
	Pittman	34X1[b]	170	9.4	7150	17.6	75	55	103
		GM34X1 *19.5:1 gear ratio*	241	159	367	15.3	65	661	63
		GM34X1 *31:1 gear ratio*	245	234	231	14.2	60	956	58

a. MicroMo Electronics, Inc., website data sheets, *http://www.micromo.com/*

b. Pittman Co., website data sheets, *http://www.pittmannet.com/*

c. Servo Systems, Co., 1996–1997 Product Catalog.
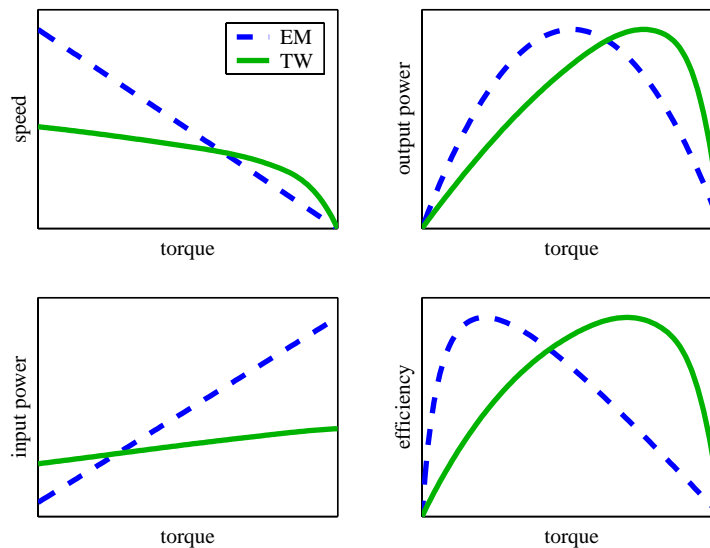
Figure 1.7 Typical shapes of the performance curves for a (geared) electromagnetic (EM) motor and a traveling-wave (TW) motor. Scales are set such that the two have the same stall torque, peak output power, and peak efficiency.

for traveling-wave (TW) motors. Whereas the speed of an electromagnetic motor decreases linearly from no load to stall, the speed of a traveling-wave motor is more constant throughout its torque range. Consequently, both the output power and efficiency of a traveling-wave motor peak at relatively high load. Conversely, although the peak efficiency of an electromagnetic motor can be quite high, it is achieved at relatively low load. This is because the electrical power draw of an electromagnetic motor is nearly linearly proportional to the torque load. At low load, where the efficiency is maximized, the electrical input power is rather small. The traveling-wave motor, on the other hand, which achieves peak performance metrics at high load, does so at the cost of drawing high electrical input power even at no load. Clearly, the power delivery of the two classes of motors are very different.

In light of the comparisons between Tables 1.1 & 1.3 where only peak steady-state performance metrics were considered, one would conclude that traveling-wave motors don't offer any benefits as direct replacements for traditional electromagnetic motors. However, it's when the other advantages of traveling-wave motors come into play such as configurable

shape factor and quick response time that they merit consideration for particular applications. Furthermore, as demonstrated by Figure 1.7, one should also consider how power is delivered over the full torque range, not just compare peak values, when drawing conclusions as to which type of motor would best meet the performance requirements of a particular application.

1.3.3 Model Review

A comprehensive literature review of various modeling approaches was presented by Wallaschek [Wallaschek, 1998]. Even to this date, the most noteworthy advancements in traveling-wave motor modeling can be found in that review and are summarized here. As each modeling endeavor is discussed, its significance is related in terms of how the following key model elements were handled:

- stator vibration   (prescribed or force dependent)
- piezoceramic forcing (neglected or electromechanically modeled)
- contact model   (elastic or viscoelastic, pure slip or stick-slip)
- rotor dynamics   (rigid or flexible)
- solution procedure

One of the earliest theoretical analyses of a traveling-wave motor was that of Maeno *et al*. [Maeno, 1990, 1992] for the ring-type motor used in Canon's autofocus lens. In this analysis, the finite-element method was used to calculate the vibration mode of the stator and then to simulate the elastic stator/rotor contact. Though the rotor itself was modeled as rigid, both normal and tangential deformations of a flexible flange on the rotor were taken into account so as to capture the stick-slip behavior at the contact interface. Note that the flange in Canon's design serves the same purpose as the compliant polymer in most other designs. Experimental measurements of the interface displacements were shown to be in good agreement with the numerical results. Through arbitrary selection of the friction coefficients, correlation of the speed-torque characteristics was also demonstrated. The analysis, however, did not model the electromechanical coupling of the stator nor the mutual interaction of the contact forces with the stator dynamics. Instead, the amplitude of the sta-

tor vibration mode was prescribed in the simulation to correspond with the value measured experimentally. Furthermore, although the contact dynamics of the rotor were carefully modeled, its modal dynamics were neglected. Nevertheless, the work by Maeno *et al.* is noteworthy as it was the first to identify the stick-slip regions of the traveling-wave motor. Owing to the numerical intensity of the finite-element method, each test case was reported as taking an hour to process on a supercomputer. For this reason, few others have pursued a dynamic finite-element approach for simulating stator/rotor contact.

Following a more analytical approach, Cao and Wallaschek [Cao, 1995a,b] presented a viscoelastic contact model that not only included normal and tangential stiffness of the compliant interface but also accounted for dissipation in either direction. Like Maeno *et al.*, however, the stator motion was still given as a kinematic constraint. Shortly thereafter, Schmidt *et al.* [Schmidt, 1996] adopted the same viscoelastic foundation model for the description of the contact layer and formulated a motor model in which the stator motion was dependent on the contact forces. Nevertheless, focus was still restricted to the contact problem, so the bending waves in the stator were generated by applying an arbitrary distributed force rather than by modeling the electromechanical forces of the piezoceramics. A rigid rotor was assumed, and the coupled motor model was solved iteratively using a Fourier series representation for the stator deformations. Stick-slip regions resulted that were similar to those observed by Maeno *et al.*, but more importantly, the introduction of dissipation in the deformable contact layer brought about an asymmetric profile to the normal contact forces.

The first comprehensive traveling-wave motor model was developed by Hagood and McFarland [Hagood, 1995]. Rather than focusing specifically on the contact mechanics, Hagood and McFarland set out to simulate the overall motor system behavior, including electromechanical forcing of the stator and full coupling of the stator dynamics with the contact forces. In doing so, the response of the system was appropriately tied to the voltage inputs of the piezoceramic actuators. However, to accommodate the complexity of the fully coupled motor model, the contact model was greatly simplified. The normal stresses were

determined from an elastic foundation model, and pure slip was assumed such that the tangential stresses were given by Coulomb's friction law. As usual, only rigid-body displacements were assumed for the rotor. The transient analysis was formulated purely in the time domain, and a number of parametric studies were carried out to show the effects of such variables as preload and tooth height.

Recognizing the significance and shortcomings of Hagood and McFarland's work, Hagedorn *et al*. [Hagedorn, 1998] adopted the model approach and incorporated dynamic flexibility of the rotor. Furthermore, to improve the efficiency of the model as a design tool, steady-state solutions to the equations of motion were formulated in the frequency domain, greatly reducing computation time. The contact model, however, was left unchanged, and a lumped damping coefficient still had to be determined arbitrarily by fitting the simulated speed-torque characteristics to experimental measurements. Additional model parameters including the electromechanical coupling factor and coefficient of friction were extracted more directly from the measurements. Upon fitting, excellent agreement was obtained with the speed-torque curves for the case of a rotor with finite flexibility but not for the case of a rigidly modeled rotor. The numerical results clearly demonstrated the importance of including rotor flexibility, though as stated by the authors, the model was still only suited for qualitative investigations.

In each of the approaches discussed thus far, only ideal traveling waves were considered. Motivated by the approaches of Hagood and McFarland and Cao and Wallaschek, Maas *et al*. [Maas, 1995a,b] extended the elastic contact model with pure slip to accommodate the general case of a non-ideal traveling wave. The effects of varying the relative amplitudes and phases of the constitutive standing waves were investigated. As intended by Maas *et al*., the work proved that a perfect traveling wave is, indeed, the optimal mode for operation.

1.4 Approach

At the onset of this modeling endeavor, the state of the art in traveling-wave motor modeling was best represented by the comprehensive approach of Hagood and McFarland and the vis-

coelastic contact model of Schmidt *et al.* Coincidentally, the strength of the latter was the weakness of the former and vice versa. Therefore, a combination of the two was seen as the logical jumping-off point for achieving the modeling objective set forth in Section 1.2. There were still factors that needed to be addressed, however, before the combined model would be suitable as a design tool. First, several experimental studies [Guyomar, 1996; Kawai, 1995] had begun to report on the unmistakable influence of rotor vibration on motor performance, thus invalidating the rigid rotor assumption. Furthermore, accurate performance prediction would require a better damping model, i.e. one that could be attributed directly to the material properties. Lastly, for the model to be useful, its formulation should be conducive to computational efficiency and allow for real-world imperfections.

As a result, the modeling efforts of this thesis were defined by the following approach:

- adopt the general approach of Hagood and McFarland, particularly the electromechanical stator model
- integrate a viscoelastic polymer model and derive a finite-time hysteretic stick-slip contact algorithm
- extend the rotor model to include dynamic flexibility, i.e. modal degrees of freedom
- introduce complex material properties to account for system losses
- formulate the model in steady state by developing a mixed-domain formulation to handle the specific analysis requirements of each model element
- generalize the model to support simulation of a non-ideal traveling wave

In support of the modeling endeavor, the experimental studies in this thesis were designed to have twofold purpose. As the primary experimental objective, much of the laboratory effort was focused on the development of the novel two-sided motor. Along the way, however, various aspects of the research were tailored to provide a meaningful base of experimental data for thorough validation of the newly proposed model. The various elements of this dual-purpose experimental approach included the following specific objectives:

- develop an effective two-sided traveling-wave motor configuration
- construct both one-sided and two-sided prototypes from essentially identical components for a direct comparison study

- manufacture an array of rotors for each motor for a parametric study on the importance of rotor design

- characterize the modal quality of each stator for the express purpose of accurately simulating their non-ideal behavior

- characterize all the necessary complex material constants for the new loss model

- conduct performance testing of the various stator/rotor combinations over a broad range of operating conditions

This last objective was in response to the typical practice of reporting model-experiment correlation at only a single set of operating conditions. However, to prove this model capable of predicting motor performance, model-experiment correlation would need to be demonstrated for varying parametric conditions such as drive frequency and rotor stiffness. To this end, the final and most important task of this thesis would be to implement the model and simulate various case studies representing a broad range of operating conditions and geometric parameters. The chosen studies included directly comparing one-sided and two-sided operation, observing the effects of rotor stiffness, and quantifying performance loss due to non-ideal behavior. The validity of the model would ultimately be determined by its ability not only to capture relative performance trends but also to predict absolute performance metrics.

1.5 Outline

The development of the generalized motor model is presented analytically in Chapter 2. After an introduction to the assumed material models such as the use of complex material properties, the nonlinear equations of motion are formulated in modal coordinates using the Rayleigh-Ritz energy method then linearized by harmonic balance for representation in the frequency domain. However, limited by having to calculate the hysteretic-friction forces in the time domain, a mixed-domain solution procedure is proposed for computational efficiency.

Chapter 3 lays the foundation for the extensive experimental work that was carried out not only to demonstrate the potential of two-sided operation but also to provide a controlled performance sample for model validation. The candidate motors are introduced first, then each of their stators are compared in light of various qualification studies. Material characterization studies are then presented for the purpose of supplying the model with the appropriate material properties in complex form. The performance test apparatus are described last, including the drive electronics and the performance measurement testbed.

The results of experimental performance testing are presented in Chapter 4. In addition to directly comparing the behavior of the one-sided and two-sided prototypes, particular focus is given to demonstrating the sensitivity of motor performance on the design of the rotor. Additional parametric studies include the effects of drive voltage and excitation frequency.

With the motors fully characterized and a broad experimental data set clearly defined, the model is revisited in Chapter 5 in the context of numerical implementation and simulation. Specific details on building the model are worked through including how to handle stator asymmetry and two-sided operation. Computational details are also discussed with regard to model reduction and the selected nonlinear solver. Ultimately, the validity of the model is put to the test by presenting simulations of several of the experimental case studies. Upon comparing the calculated performance metrics to the actual measurements, the ability and limitations of the model to accurately predict motor performance without the need for adjustment factors are discussed.

Finally, conclusions are presented in Chapter 6 along with contributions and recommendations for further work.

Chapter 2

MODELING

Building upon the modeling endeavor by Hagood and McFarland [Hagood, 1995], the model presented here features five key advancements: (1) addition of the flexural rotor dynamics, (2) use of lossy (complex) material properties, (3) integration of a viscoelastic, hysteretic stick-slip friction contact model, (4) implementation of a mixed-domain solution technique, and (5) generalization to the non-ideal traveling-wave case. These contributions are intended to extend the work by Hagood and McFarland and, as believed by the author of this work, constitute the elements needed to take the model to completion for use as a design tool.

The fact that the rotor dynamics greatly affect motor performance is well known at this point, but until recently, the rotor has generally been modeled as rigid. A decisive paper regarding the importance of rotor flexibility is that published by Hagedorn [Hagedorn, 1998] (see Section 1.3.3). Adopting the approach of Hagood and McFarland, Hagedorn appended modal rotor dynamics to the model and demonstrated through fitting of the simulated speed-torque characteristics to measured data that good agreement was possible only for the case where rotor flexibility was included. However, fitting the data relied on the adjustment of an arbitrary damping coefficient and was illustrated for only a single curve of speed-torque data. Data was not provided to confirm agreement with the simulated input power and efficiency. Furthermore, without simultaneously fitting data measured at various operating conditions, the accuracy of the simple contact model was not addressed.

Though better than previous models employing numerous adjustable parameters, still one model parameter remained in Hagedorn's formulation that was not explicitly related to physical material properties or geometry. To remedy this, complex material properties [Sherrit, 1998] have been employed here as a unified approach to attributing motor losses to the inherent material damping present in the metallic stator and rotor substrates, the polymer interface layer, and the piezoceramic actuator. The piezo material in particular exhibits not only elastic losses but also piezoelectric and dielectric losses. The lossiness of each of these material properties can be represented by means of complex constants, thus avoiding the need for lumped modal damping coefficients.

Just as important as rotor dynamics is the contact model. Although areas of contact between the stator and rotor surfaces are known to exhibit regions of stick and slip due to tangential deformation of the compliant polymer lining, previous contact models have often been limited to the case of pure slip in order to facilitate the model solution [Wallaschek, 1998]. In this simplification, the normal deformations of the polymer lining are considered, but the tangential deformations are neglected. However, capturing the stick-slip behavior is an important factor in the accuracy of performance computations since the tangential deformations of the viscoelastic polymer contribute to the losses in the piezo motor. Thus, tangential deformations have been taken into account in the proposed model derivation.

Implementing both complex material constants and a stick-slip contact model has prompted a novel solution procedure. The stick-slip contact model is inherently hysteretic and necessitates simulation in the time domain. The simplicity of the rest of the motor model, however, is conducive to steady-state solution in the frequency domain, and, in fact, the complex representation is founded on the assumption of harmonic analysis in the frequency domain. Preserving as best as possible the computational benefits of the frequency domain, a mixed-domain solution procedure has been implemented that iteratively switches between domains while searching for the nonlinear model solution.

Lastly, realized motors are not perfect and can suffer lower-than-expected performance due to either geometric or material asymmetries. Possible geometric imperfections include insufficient machining tolerances, uneven electrode etch patterns, and piezo array misalignment. Asymmetry can also be the result of non-isotropic material properties associated with metal extrusion, nonuniform piezoceramic composition, or inconsistent poling. Either the standing-wave modes aren't excited exactly 90° apart, or the perfect circular symmetry is lost, and the degenerate eigenmodes separate in frequency and preferentially orient themselves in a frame not necessarily aligned with the piezo arrays. In either case, the result is the inability to excite a pure traveling wave. Hence, the model approach presented here has not been limited to the ideal traveling-wave case. Instead, the equations have been generalized to allow performance prediction of under non-ideal traveling-wave conditions. In addition to accounting for asymmetries, this approach permits simulation of intentional non-ideal operation. For instance, motor speed can effectively be controlled by adjusting the relative phase of the input voltages. This results in a less-than-pure traveling wave while maintaining consistent wave amplitudes and smooth operation and is arguably the preferred method of speed control in the traveling-wave motor.

2.1 Material Model

Before formulating the full dynamic model, various elements will be introduced. These include complex material property representation, a viscoelastic polymer interface model, and the generalized electromechanically coupled constitutive relations. The next few sections provide the material models that will be used in deriving the equations of motion and calculating the interface forces in Section 2.2

2.1.1 Complex Material Properties

Acknowledging the shortcomings of the modal damping model has motivated the use of the hysteretic material damping model and complex material property representation to account for the internal losses in the motor system. In essence, the stress-strain response of a linear material under harmonic loading generally exhibits hysteresis that is (relatively)

independent of frequency. This hysteresis is often observed as a constant phase lag in the elastic response. Thus, for harmonic analysis of an elastic material, the phase lag can be accounted for by the use of a complex elastic constant. For piezoelectric materials, phase lag is not only observed in the elastic response but also in the piezoelectric and dielectric responses. The use of complex material constants to represent the losses in a piezoelectric material subject to an AC field was demonstrated decades ago by Holland [Holland, 1967].

Notion of Complex Constants

Examining the energy dissipated per cycle under harmonic loading, e.g. as done in the derivation by Inman [Inman, 1994], offers some insight into the use of complex constants to represent internal material damping. Consider the equation of motion for a viscously damped mass-spring system:

$$m\ddot{x} + c\dot{x} + kx = F_0 e^{i\omega_{dr}t}, \tag{2.1}$$

which has an assumed solution of the form $x(t) = Xe^{i\omega_{dr}t}$. (When using the complex exponential to express the harmonic forcing and response of a physical system, it is implicitly assumed that only the real parts are retained.) By integrating the area of the elliptical force-displacement hysteresis loop, the energy dissipated per cycle is given by

$$\Delta E_{viscous} = \pi c\omega_{dr}X^2. \tag{2.2}$$

Notice that the energy loss is dependent on the drive frequency. However, the area defined by the stress-strain hysteresis loop of many structural materials is found to be independent of frequency and proportional to the stiffness of the material. In this case, the energy dissipated per cycle is assumed to be represented by

$$\Delta E_{material} = \pi k\beta X^2, \tag{2.3}$$

where the constant β (sometimes denoted by η) is defined as the hysteretic damping constant. By equating ΔE in each of the two cases, an equivalent viscous damping coefficient

can be defined for the material damping case which is proportional to stiffness and inversely proportional to frequency:

$$c_{eq} = \frac{k\beta}{\omega_{dr}}. \tag{2.4}$$

The analogy allows a mass-spring system with material damping to be written as

$$m\ddot{x} + \frac{k\beta}{\omega_{dr}}\dot{x} + kx = F_0 e^{i\omega_{dr}t}. \tag{2.5}$$

Replacing the velocity term only with the appropriate exponential response, the lossy term can be lumped with the stiffness term to form a complex stiffness:

$$m\ddot{x} + k(1 + i\beta)x = F_0 e^{i\omega_{dr}t}. \tag{2.6}$$

It is important to note that the analogy assumes harmonic loading, so the use of complex constants is restricted to frequency-domain, or steady-state, analyses.

Application of Complex Constants

In general, the angle of a complex material constant represents the frequency-independent phase lag in the dynamic response attributable to the lossy property, and β is equivalent to the tangent of the angle. Thus, β is commonly referred to as $\tan\delta$, or the loss tangent. For resonance analyses, the loss tangent is often replaced by the quality factor, Q, where $Q = 1/\tan\delta$. For the purposes of the model presented here, the necessary complex elastic, piezo-electric, and dielectric constants have been defined as

$$\begin{aligned}
s_{11}^{E} &= s_{11}^{E}{}'(1 - i/Q_m) \\
d_{31} &= d_{31}{}'(1 - i/Q_p) \\
\varepsilon_{33}^{T} &= \varepsilon_{33}^{T}{}'(1 - i\tan\delta),
\end{aligned} \tag{2.7}$$

where Q_m is the mechanical quality factor, Q_p is the piezoelectric quality factor, and $\tan\delta$ is the dielectric loss tangent.

In practice, the losses of piezoceramic materials, though relatively independent of frequency, can be highly dependent on field level and temperature, typically increasing as drive field and operating temperature increase. Unfortunately, piezo properties are most often published only for low field at room temperature and are of little use for modeling high-power actuators. To properly capture the losses in these materials, it is then very important to measure the complex constants at the appropriate operating conditions. The same goes for all the materials used in a high-power system, and since all materials exhibit some frequency-dependent behavior over a large enough frequency band, it is best to gather measurements close to the appropriate frequency as well.

A very thorough discussion on the use of complex piezoelectric material constants is presented by Sherrit [Sherrit, 1998]. In his paper, he not only introduces the topic both macroscopically and microscopically but also provides meaningful interpretations of the complex coefficients and cautionary notes on their use, especially regarding dispersion and attempts to transform the complex material models to the time domain.

2.1.2 Viscoelastic Polymer Interface

Researchers in the field have learned that the stator-rotor interface must incorporate some form of compliance mechanism to transmit the oscillation energy of the stator to the rotor effectively. Most often this is accomplished by bonding a thin polymer sheet to the contact surface of the metallic rotor. The local compliance provided by the softer material ensures continuous contact with the many wave crests. Theoretical performance suffers, however, manifested by lower speeds, less available torque, and greater losses as a result of the expanded contact area, the increased shear compliance, and the polymer's viscoelastic nature, respectively. Nevertheless, realistic performance is made possible by this necessary evil. Accordingly, accurate prediction of motor performance requires a detailed model of the interface dynamics.

Contact is where the nonlinearities enter the motor model, so the contact dynamics must be formulated in the time domain. This precludes the use of complex property representation
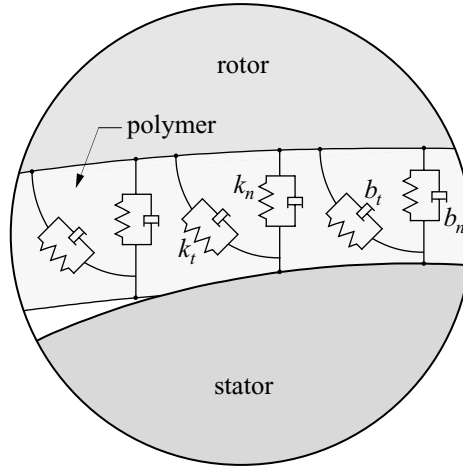
Figure 2.1   Viscoelastic polymer interface model.

in the interface model. Instead, a traditional viscoelastic representation is used. Accounting for both compressive and shear compliance, the polymer layer is modeled by distributed normal and tangential springs and dampers, as seen in Figure 2.1. This is identical to the viscoelastic polymer model used by Cao and Wallaschek [Cao, 1995a,b] and Schmidt [Schmidt, 1996] as described earlier in Section 1.3.3.

Utilizing (2.4), the spring and damper values can be related to the material's complex elongation and shear moduli, $E = E' + iE''$ and $G = G' + iG''$, by

$$k_n = \frac{\alpha E'}{t}, \quad b_n = \frac{\alpha E''}{\omega t}, \quad k_t = \frac{G'}{t}, \quad b_t = \frac{G''}{\omega t} \quad \text{where} \quad \alpha = \frac{1}{1-2\nu} \cdot \frac{1-\nu}{1+\nu}. \quad (2.8)$$

The coefficient α has been introduced here under the assumption of zero in-plane strain. For a typical Poisson's ratio, ν, of around 0.33, $\alpha \approx 1.5$, equating to a 50% increase in normal stiffness due to the in-plane incompressibility assumption. Note that mass has been ignored in the interface model as the internal dynamics of the polymer layer are assumed quasistatic relative to the drive frequency (at least one order of magnitude difference in frequency). Also note that although Figure 2.1 illustrates one-dimensional line contact, the interface model has been implemented to accommodate area contact because the contact profile can, indeed, vary significantly across the thin, radial dimension.

It is important to understand that while the contact forces are periodic with drive frequency, ω, they are not necessarily harmonic and cannot be solved for in the frequency domain. Nevertheless, though the viscous damping analogy presented in the previous section assumes harmonic forcing, the analogy is used here as an approximation, with heed to Sherrit's cautionary notes on dispersion.

2.1.3 Constitutive Relations

The stator and rotor are modeled as thin, laminated plates. Plane stress conditions are assumed, and electric field in the piezoceramic wafer is assumed nonzero in the thickness direction only. As such, the stress, strain, electric field, and electric displacement vectors are reduced, respectively, as follows:

$$\mathbf{T} = \begin{bmatrix} T_1 \\ T_2 \\ T_6 \end{bmatrix} = \begin{bmatrix} \sigma_r \\ \sigma_\theta \\ \tau_{r\theta} \end{bmatrix}, \quad \mathbf{S} = \begin{bmatrix} S_1 \\ S_2 \\ S_6 \end{bmatrix} = \begin{bmatrix} \varepsilon_r \\ \varepsilon_\theta \\ \gamma_{r\theta} \end{bmatrix}, \quad \mathbf{E} = E_3, \quad \mathbf{D} = D_3. \qquad (2.9)$$

The constitutive relations for the electromechanically coupled material are represented using the standard convention [IEEE Std 176-1987],

$$\begin{bmatrix} \mathbf{D} \\ \mathbf{T} \end{bmatrix} = \begin{bmatrix} \varepsilon_{33}^S & \mathbf{e} \\ -\mathbf{e}^\mathsf{T} & \mathbf{c}^E \end{bmatrix} \begin{bmatrix} \mathbf{E} \\ \mathbf{S} \end{bmatrix}, \qquad (2.10)$$

where superscripts $(\cdot)^S$ and $(\cdot)^E$ signify that the values are calculated at constant strain and electric field, respectively. The plane stress stiffness matrix, $\mathbf{c}^E$, should be calculated as the inverse of the plane stress compliance matrix, $\mathbf{s}^E$, reduced from the full transversely isotropic 6x6 system:

$$\mathbf{c}^E = (\mathbf{s}^E)^{-1} \quad \text{where} \quad \mathbf{s}^E = \begin{bmatrix} s_{11}^E & s_{12}^E & 0 \\ s_{12}^E & s_{11}^E & 0 \\ 0 & 0 & 2(s_{11}^E - s_{12}^E) \end{bmatrix}. \qquad (2.11)$$

The piezoelectric and dielectric constants required for (2.10) are related to more commonly used constants by the relationships

$$\mathbf{e} = \mathbf{dc}^E \quad \text{and} \quad \varepsilon_{33}^S = \varepsilon_{33}^T - \mathbf{ed}^\mathsf{T} \quad \text{where} \quad \mathbf{d} = \begin{bmatrix} d_{31} & d_{31} & 0 \end{bmatrix}. \tag{2.12}$$

Within the isotropic structural materials, the constitutive relationship reduces to $\mathbf{T} = \mathbf{cS}$, where $s_{11} = 1/E$ and $s_{12} = -\nu/E$ may be substituted in (2.11) to calculate $\mathbf{c}$.

2.2 Motor Model

The model approach begins with formulating the elastic modal dynamics of the stator and rotor using the Rayleigh-Ritz energy method of assumed modes and transforming the equations to frequency domain. The rigid body dynamics of the rotor, assumed constant in steady state, are described using simple Newtonian dynamics and relate force and torque equilibrium between the external loads and the reaction forces at the interface. The nonlinear, state-dependent interface forces that couple the stator and rotor dynamics are then derived in the time domain via viscoelastic contact and stick-slip friction models. Finally, a mixed-domain, iterative-search solution procedure is implemented for steady-state performance calculation.

2.2.1 Modal Dynamics

To derive the elastic modal dynamics of the motor system, an approach based on variational principles is used. The analysis begins with the Generalized Hamilton's Principle for electromechanical systems [Hagood, 1990] (ignoring the magnetic energy term):

$$\delta \int_{t_1}^{t_2} (T - U_m + U_e)\,dt + \int_{t_1}^{t_2} (\delta W_m + \delta W_e)\,dt = 0. \tag{2.13}$$

The kinetic, elastic, and electrical energy terms are

$$T = \frac{1}{2}\int_V \rho\,\dot{\mathbf{u}}^\mathsf{T}\dot{\mathbf{u}}\,dV, \quad U_m = \frac{1}{2}\int_V \mathbf{S}^\mathsf{T}\mathbf{T}\,dV, \quad \text{and} \quad U_e = \frac{1}{2}\int_V \mathbf{E}^\mathsf{T}\mathbf{D}\,dV. \tag{2.14}$$

The virtual mechanical and electrical work terms,

$$\delta W_m = \int_{S_c} \delta \mathbf{u}^\mathsf{T} \mathbf{f} dS_c \quad \text{and} \quad \delta W_e = \int_{S_{el}} \delta \mathbf{v}^\mathsf{T} \mathbf{q} dS_{el}, \tag{2.15}$$

are due to the surface stresses, $\mathbf{f}$, at the contact interface, S_c, and the charge density, $\mathbf{q}$, on the free electrode surface, S_{el}.

Strain-Displacement Relations

The strain-displacement relations are derived under Kirchhoff's assumption for thin plates which states that any point's displacement is a function of the plate's centerline displacements and centerline slopes and the offset of the point from the centerline. In cylindrical coordinates, the displacement assumptions are

$$u = u_0 - z\frac{\partial w_0}{\partial r}, \quad v = v_0 - z\frac{1}{r}\frac{\partial w_0}{\partial \theta}, \quad \text{and} \quad w = w_0, \tag{2.16}$$

where u, v, and w correspond to the r, θ, and z directions, respectively. Defining the displacement operator matrix, $\mathbf{L}_u$, the displacement vector, $\mathbf{u}$, is expressed as a function of the centerline displacement vector, $\mathbf{u}_0$:

$$\mathbf{u} = \begin{bmatrix} u \\ v \\ w \end{bmatrix} = \mathbf{L}_u \mathbf{u}_0 \quad \text{where} \quad \mathbf{L}_u = \begin{bmatrix} 1 & 0 & -z\dfrac{\partial}{\partial r} \\ 0 & 1 & -z\dfrac{1}{r}\dfrac{\partial}{\partial \theta} \\ 0 & 0 & 1 \end{bmatrix}. \tag{2.17}$$

The strain-displacement relations in cylindrical coordinates are

$$\mathbf{S} = \begin{bmatrix} S_1 \\ S_2 \\ S_6 \end{bmatrix} = \begin{bmatrix} \dfrac{\partial u}{\partial r} \\ \dfrac{u}{r} + \dfrac{1}{r}\dfrac{\partial v}{\partial \theta} \\ \dfrac{1}{r}\dfrac{\partial u}{\partial \theta} + \dfrac{\partial v}{\partial r} - \dfrac{v}{r} \end{bmatrix}. \tag{2.18}$$

Defining the strain operator matrix, $\mathbf{L}_S$, relating the strains to the centerline displacements, the strain vector can then be rewritten as

$$\mathbf{S} = \mathbf{S}_0 + z\kappa = \mathbf{L}_S \mathbf{u}_0, \tag{2.19}$$

where

$$\mathbf{S}_0 = \begin{bmatrix} \dfrac{\partial u_0}{\partial r} \\[3ex] \dfrac{u_0}{r} + \dfrac{1}{r}\dfrac{\partial v_0}{\partial \theta} \\[3ex] \dfrac{1}{r}\dfrac{\partial u_0}{\partial \theta} + \dfrac{\partial v_0}{\partial r} - \dfrac{v_0}{r} \end{bmatrix}, \quad \kappa = \begin{bmatrix} -\dfrac{\partial^2 w_0}{\partial r^2} \\[3ex] -\dfrac{1}{r}\dfrac{\partial w_0}{\partial r} - \dfrac{1}{r^2}\dfrac{\partial^2 w_0}{\partial \theta^2} \\[3ex] -\dfrac{2}{r}\dfrac{\partial^2 w_0}{\partial r\partial \theta} + \dfrac{2}{r^2}\dfrac{\partial w_0}{\partial \theta} \end{bmatrix}, \tag{2.20}$$

and

$$\mathbf{L}_S = \begin{bmatrix} \left(\dfrac{\partial}{\partial r}\right) & 0 & \left(-z\dfrac{\partial^2}{\partial r^2}\right) \\[3ex] \left(\dfrac{1}{r}\right) & \left(\dfrac{1}{r}\dfrac{\partial}{\partial \theta}\right) & \left(-\dfrac{z}{r}\dfrac{\partial}{\partial r} - \dfrac{z}{r^2}\dfrac{\partial^2}{\partial \theta^2}\right) \\[3ex] \left(\dfrac{1}{r}\dfrac{\partial}{\partial \theta}\right) & \left(\dfrac{\partial}{\partial r} - \dfrac{1}{r}\right) & \left(-\dfrac{2z}{r}\dfrac{\partial^2}{\partial r\partial \theta} + \dfrac{2z}{r^2}\dfrac{\partial}{\partial \theta}\right) \end{bmatrix}. \tag{2.21}$$

Assumed Shape Functions

Employing the Rayleigh-Ritz approach, shape functions are assumed for the midplane displacements, allowing separation of the spatial and temporal variables:

$$\mathbf{u}_0 = \begin{bmatrix} u_0 \\ v_0 \\ w_0 \end{bmatrix} = \begin{bmatrix} \phi_u & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \phi_v & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \phi_w \end{bmatrix} \begin{bmatrix} \mathbf{p}_u \\ \mathbf{p}_v \\ \mathbf{p}_w \end{bmatrix} = \Phi\mathbf{p}. \tag{2.22}$$

In vector form, the extensional and flexural shape functions are

$$\phi_u = \left[\phi_u^1(r, \theta) \cdots \phi_u^{n_u}(r, \theta) \right]$$

$$\phi_v = \left[\phi_v^1(r, \theta) \cdots \phi_v^{n_v}(r, \theta) \right] \tag{2.23}$$

$$\phi_w = \left[\phi_w^1(r, \theta) \cdots \phi_w^{n_w}(r, \theta) \right],$$

and the generalized coordinates are

$$\mathbf{p}_u = \begin{bmatrix} p_u^1(t) \\ \vdots \\ p_u^{n_u}(t) \end{bmatrix}, \quad \mathbf{p}_v = \begin{bmatrix} p_v^1(t) \\ \vdots \\ p_v^{n_v}(t) \end{bmatrix}, \quad \mathbf{p}_w = \begin{bmatrix} p_w^1(t) \\ \vdots \\ p_w^{n_w}(t) \end{bmatrix}. \tag{2.24}$$

By replacing $\mathbf{u}_0$ in (2.17) with the shape function expansion in (2.22), the displacement vector, $\mathbf{u}$, is expressed as a linear summation of the assumed deflection shapes, $\mathbf{N}_u$:

$$\mathbf{u} = \mathbf{N}_u \mathbf{p} \quad \text{where} \quad \mathbf{N}_u = \mathbf{L}_u \Phi. \tag{2.25}$$

Similarly, substituting (2.22) into (2.19) transforms the strain vector, $\mathbf{S}$, into a linear summation of the assumed strain shapes, $\mathbf{N}_S$:

$$\mathbf{S} = \mathbf{N}_S \mathbf{p} \quad \text{where} \quad \mathbf{N}_S = \mathbf{L}_S \Phi. \tag{2.26}$$

Shape functions are also defined for the electric potential boundary conditions. Rather than model the alternating polarization of neighboring segments in the piezoceramic ring, it is simpler to simulate the polarization by specifying an alternating electric potential boundary condition at the free surface of a uniformly polarized ceramic. The two approaches produce the same result. The surface potentials for arrays A and B are then related to the applied voltages and the potential shape functions by

$$\varphi_{surface} = \begin{bmatrix} \varphi_A \\ \varphi_B \end{bmatrix} = \begin{bmatrix} \phi_{\varphi_A} & 0 \\ 0 & \phi_{\varphi_B} \end{bmatrix} \begin{bmatrix} v_A(t) \\ v_B(t) \end{bmatrix} = \Phi_\varphi \mathbf{v}, \tag{2.27}$$

where $\phi_{\phi_A}(r, \theta)$ and $\phi_{\phi_B}(r, \theta)$ are discontinuous functions taking the value 1 for positively poled sectors, -1 for negatively poled sectors, and zero elsewhere. The potential at the bonded surface is grounded to the metallic substrate.

The corresponding electric field is nonzero in the thickness direction only, so $\mathbf{E} = -\partial\phi/\partial z$. Assuming a constant electric field through the piezoceramic thickness, t_p,

$$\mathbf{E} = \mathbf{L}_E \phi_{surface} \quad \text{where} \quad \mathbf{L}_E = -1/t_p. \tag{2.28}$$

Hagedorn [Hagedorn, 1998] argued that the field must actually depend linearly on z in order to satisfy Maxwell's equations since the strain is linear in z. The linearity effectively changes the piezoceramic stiffness, but the difference is small. Thus, the variation is neglected here for simplicity. Substituting (2.27) into (2.28), $\mathbf{E}$ is given by the electric field shape functions, $\mathbf{N}_E$, and the voltage degrees of freedom:

$$\mathbf{E} = \mathbf{N}_E \mathbf{v} \quad \text{where} \quad \mathbf{N}_E = \mathbf{L}_E \Phi_\phi. \tag{2.29}$$

Equations of Motion

Substituting expressions (2.14) and (2.15) into (2.13) and replacing $\mathbf{T}$ and $\mathbf{D}$ with the constitutive relations in (2.10), the variational principle describing an electromechanically coupled structure, e.g. the stator, is given by:

$$\int_{t_1}^{t_2} \left[\int_V (\delta\dot{\mathbf{u}}^\mathsf{T} \rho\dot{\mathbf{u}} - \delta\mathbf{S}^\mathsf{T}\mathbf{c}\mathbf{S} + \delta\mathbf{S}^\mathsf{T}\mathbf{e}^\mathsf{T}\mathbf{E} + \delta\mathbf{E}^\mathsf{T}\mathbf{e}\mathbf{S} + \delta\mathbf{E}^\mathsf{T}\varepsilon_{33}^S \mathbf{E}) dV \right] dt +$$

$$\int_{t_1}^{t_2} \left[\int_{S_c} \delta\mathbf{u}^\mathsf{T}\mathbf{f} dS_c + \int_{S_{el}} \delta\mathbf{v}^\mathsf{T}\mathbf{q} dS_{el} \right] dt = 0. \tag{2.30}$$

Substituting (2.25), (2.26), and (2.29) for displacement, strain, and electric field, integrating the kinetic energy term by parts, and allowing arbitrary variations of $\mathbf{p}$ and $\mathbf{v}$, two matrix equations of motion are obtained. For the stator, these equations are

$$\mathbf{M}_s \ddot{\mathbf{p}}_s + \mathbf{K}_s \mathbf{p}_s = \Theta\mathbf{v} + \mathbf{f}_{sc} \tag{2.31}$$

$$\Theta^{\mathsf{T}}\mathbf{p}_s + \mathbf{C}_p\mathbf{v} = \mathbf{q}, \tag{2.32}$$

where the mass and stiffness matrices are defined as

$$\mathbf{M}_s = \int_{V_s} \mathbf{N}_u^{(s)\mathsf{T}}\rho(r,\theta,z)\mathbf{N}_u^{(s)}dV_s \quad \text{and} \quad \mathbf{K}_s = \int_{V_s} \mathbf{N}_S^{(s)\mathsf{T}}\mathbf{c}(r,\theta,z)\mathbf{N}_S^{(s)}dV_s, \tag{2.33}$$

and the electromechanical coupling and capacitance matrices are defined as

$$\Theta = \int_{V_p} \mathbf{N}_S^{(s)\mathsf{T}}\mathbf{e}^{\mathsf{T}}\mathbf{N}_E dV_p \quad \text{and} \quad \mathbf{C}_p = \int_{V_p} \mathbf{N}_E^{\mathsf{T}}\varepsilon_{33}^S\mathbf{N}_E dV_p. \tag{2.34}$$

The superscript $(\cdot)^{(s)}$ denotes the use of assumed stator shape functions. The mass and stiffness integrations are carried out over the full stator volume, V_s, whereas terms associated with the electrical domain need only be evaluated over the piezoceramic volume, V_p. Though not needed to calculate the mechanical response of the stator, (2.32) is used to determine current, and subsequently, the electrical input power delivered to the piezoceramics later in Section 2.2.6.

In the case of the rotor, the constitutive relations are uncoupled, and electrical terms are inconsequential. This results in a single matrix equation describing the motion of the rotor:

$$\mathbf{M}_r\ddot{\mathbf{p}}_r + \mathbf{K}_r\mathbf{p}_r = \mathbf{f}_{rc}, \tag{2.35}$$

where

$$\mathbf{M}_r = \int_{V_r} \mathbf{N}_u^{(r)\mathsf{T}}\rho(r,\theta,z)\mathbf{N}_u^{(r)}dV_r \quad \text{and} \quad \mathbf{K}_r = \int_{V_r} \mathbf{N}_S^{(r)\mathsf{T}}\mathbf{c}(r,\theta,z)\mathbf{N}_S^{(r)}dV_r. \tag{2.36}$$

The superscript $(\cdot)^{(r)}$ denotes the use of assumed rotor shape functions. Volume V_r is the full rotor volume, inclusive of both the metal substrate and polymer layer.

The spatial interface forces, $\mathbf{f}(r, \theta, t)$, acting on the stator are equal and opposite those acting on the rotor. Expressions for these are formulated later in Section 2.2.3, but for the moment,

consider $\mathbf{f} = \mathbf{f}_c$ defined as positive acting on the rotor and $\mathbf{f} = -\mathbf{f}_c$ acting on the stator. This nonlinear interface force vector, $\mathbf{f}_c$, is of the form

$$\mathbf{f}_c = \begin{bmatrix} 0 \\ f_v \\ f_w \end{bmatrix}, \tag{2.37}$$

where f_w is the normal contact force per unit area reacting the externally applied axial pre-load, and f_v is the circumferential frictional shear force per unit area reacting the torque load. Assuming radial deformations and their attributable losses to be negligible, the radial shear force is ignored. Hence, the modal interface forces in (2.31) and (2.35) acting on the generalized degrees of freedom are derived from (2.30) as

$$\mathbf{f}_{sc}(t) = \int\int_{S_c} \mathbf{N}_u^{(s)\mathsf{T}}(-\mathbf{f}_c(r, \theta, t)) r\,dr\,d\theta \quad \text{and} \quad \mathbf{f}_{rc}(t) = \int\int_{S_c} \mathbf{N}_u^{(r)\mathsf{T}} \mathbf{f}_c(r, \theta, t) r\,dr\,d\theta. \tag{2.38}$$

2.2.2 Rigid Body Dynamics

The rigid body degrees of freedom of the motor system consist of the relative rotor height, h_r, and the rotor's angular velocity, ω_r. It is through formulation of the rigid body dynamics that force and torque equilibrium are enforced between those forces applied externally to the rotor and those reacting internally at the contact interface. The equations of motion describing the rigid body degrees of freedom are

$$M_r \ddot{h}_r = f_{int} - f_{preload} \tag{2.39}$$

$$I_r \dot{\omega}_r + \gamma_\omega \omega_r = \tau_{int} - \tau_{load}. \tag{2.40}$$

The externally applied axial preload and torque load, $f_{preload}$ and τ_{load}, are constant operating condition parameters and are each defined in the negative direction for intuitive convenience, i.e. as to oppose positive motion. f_{int} and τ_{int} are the modal interface forces associated with the rigid body modes $\phi_w = 1$ and $\phi_v = r$. Thus, the rigid body interface forces are calculated similarly to the flexural mode contact forces in (2.38):

$$f_{int}(t) = \int\limits_{S_c} \!\! \int f_w(r, \theta, t) r \, dr \, d\theta \quad \text{and} \quad \tau_{int}(t) = \int\limits_{S_c} \!\! \int f_v(r, \theta, t) r^2 \, dr \, d\theta . \quad (2.41)$$

The coefficient γ_ω in (2.40) represents rotational damping of the rotor and could be specified to account for bearing friction. However, physical damping of the rigid-body mode is assumed negligible at the slow rotational speeds of the traveling-wave motor. Instead, the coefficient has proven useful as a numerical restoring term to improve convergence efficiency of the solution procedure. In effect, the term maintains finite sensitivity of (2.40) to ω_r even for unrealistic solution iterations on ω_r where the calculated τ_{int} saturates due to uniform slip of the contact interface. Its value is chosen such that the effective viscous damping term is negligible and does not noticeably effect steady-state performance simulations (<0.1% of peak metrics).

Lastly, a stiffness term could be added to (2.39) to account for the preload spring. However, as this term has little or no effect in steady state, an ideal preload is assumed, as is the case for a properly selected Belleville washer.

2.2.3 Interface Forces

The Kirchhoff assumptions defined by (2.16) are used to describe the displacements of the stator and rotor at the contact interface due to global flexural and extensional deformations. The interface model, in turn, describes the local deformations, v_c and w_c, of the compliant polymer lining relative to the global displacements, v and w, that are needed to evaluate the contact normal and friction forces. An essential simplifying assumption is that the distributed interface forces depend only on these local deformations at the contact points. Thus, utilizing the distributed spring and damper model in Figure 2.1, the interface forces are related as functions of the local deformations via the viscoelastic constants defined in (2.8), or $f_v(r, \theta, t) = f(v_c, \dot{v}_c)$ and $f_w(r, \theta, t) = f(w_c, \dot{w}_c)$.

The through-thickness compression, w_c, of the polymer is related to the overlap, Δw, of the normal stator and rotor deformation profiles by

$$w_c(r, \theta, t) = \begin{cases} \Delta w & \Delta w > 0 \\ 0 & \Delta w \leq 0 \end{cases} \quad \text{where} \quad \Delta w = w_s - (w_r + h_r). \quad (2.42)$$

As such, the contact normal force, f_w, can be expressed as the sum of the elastic and viscous reaction forces given the following conditions:

$$f_w(r, \theta, t) = \begin{cases} k_n w_c + c_n \dot{w}_c & k_n w_c + c_n \dot{w}_c > 0 \\ 0 & k_n w_c + c_n \dot{w}_c \leq 0. \end{cases} \quad (2.43)$$

The conditions in the force calculation limit the contact normal force from exhibiting tension in the presence of the viscoelastic term.

The hysteretic friction force, f_v, is inherently less straight forward to calculate than the contact normal force, f_w, because the conditions of stick and slip require knowledge of the state history. Following the Coulomb friction law with a single (dynamic) coefficient of friction, the magnitude of the friction force is limited by the product of the normal contact force and the coefficient of friction, μ:

$$f_{v_{max}}(r, \theta, t) = \mu f_w(r, \theta, t). \quad (2.44)$$

Slipping of the contact surfaces occurs in regions where $f_{v_{max}}$ is exceeded, resulting in hysteresis of the friction forces and shear deformations. To capture the hysteretic response, it is necessary to integrate slip over time, thereby keeping a cumulative record of the relative motion between surface points during periods of slip. To this end, the slip variable, Δv, is introduced. The shear deformation, v_c, of the polymer can then be expressed as the difference in the relative stator and rotor displacements at the contact surface (including angular rotation of the rotor) and this "slip record":

$$v_c = v_s - (v_r + \alpha_r r) - \Delta v. \quad (2.45)$$

Following the viscoelastic polymer model, the friction force is related to v_c by

$$f_v(r, \theta, t) = k_t v_c + c_n \dot{v}_c. \quad (2.46)$$

Substituting (2.45) into (2.46), the friction force at each contact point (r,θ) at time index k is given by

$$f_v^k = k_T(v_s^k - (v_r^k + \alpha_r^k r) - \Delta v^k) + c_T(\dot{v}_s^k - (\dot{v}_r^k + \omega_r^k r) - \Delta \dot{v}^k).$$ (2.47)

The tangential force cannot exceed the maximum allowable friction force, or $|f_v| \leq f_{v_{max}}$, for the force is limited by slipping at the interface. Stating the obvious, while the contact sticks, the slip variable remains unchanged, giving $\Delta v^k = \Delta v^{k-1}$ and $\Delta \dot{v}^k = 0$. Thus, defining

$$f_{v_{test}}^k = (k_T(v_s^k - (v_r^k + \alpha_r^k r) - \Delta v^{k-1}) + c_T(\dot{v}_s^k - (\dot{v}_r^k + \omega_r^k r))),$$ (2.48)

the test criterion for stick is to verify that

$$\left| f_{v_{test}}^k \right| \leq f_{v_{max}}^k.$$ (2.49)

For contact points where (2.49) holds true, v_c^k is calculated directly from (2.45) using the prior value for the slip variable:

$$v_c^k = v_s^k - (v_r^k + \alpha_r^k r) - \Delta v^{k-1}.$$ (2.50)

The time derivative of the shear deformation, $\dot{v}_c^k$, is then expanded using the reverse difference equation

$$\dot{v}_c^k = \frac{v_c^k - v_c^{k-1}}{\Delta t}.$$ (2.51)

Conversely, where (2.49) is not true, slip occurs, and the friction force is limited by (2.44):

$$f_v^k = \text{sign}(f_{v_{test}}^k) \cdot f_{v_{max}}^k.$$ (2.52)

Substituting (2.51) and (2.52) into (2.46), v_c^k during slip is instead calculated using the update equation

$$v_c^k = \frac{1}{c_T + k_T \Delta t}(c_T v_c^{k-1} + \Delta t \,\text{sign}(f_{v_{test}}^k) \cdot f_{v_{max}}^k).$$ (2.53)

With v_c^k appropriately determined from either (2.50) or (2.53) and $\dot{v}_c^k$ from (2.51), the friction force can then be computed using

$$f_v^k = k_T v_c^k + c_T \dot{v}_c^k. \tag{2.54}$$

Lastly, in preparation for the next time step, the present slip record is determined by revisiting (2.45):

$$\Delta v^k = v_s^k - (v_r^k + \alpha_r^k r) - v_c^k. \tag{2.55}$$

2.2.4 Uncoupled Modal Equations

The matrices in equations (2.31) and (2.35) are fully populated, coupling the generalized degrees of freedom. Alternatively, eigenvalue decomposition is used to transform the generalized coordinates, $\mathbf{p}$, to uncoupled modal coordinates, $\mathbf{z}$, by

$$\mathbf{p}_s = \mathbf{V}_s \mathbf{z}_s \quad \text{and} \quad \mathbf{p}_r = \mathbf{V}_r \mathbf{z}_r, \tag{2.56}$$

where $\mathbf{V}_s$ and $\mathbf{V}_r$ are the eigenvector matrices of the stator and rotor, respectively. At this point, only the eigenvectors corresponding to modes that resonate near the drive frequency are retained for the purpose of model reduction. For the stator, this means eliminating all but the two modes corresponding to the intended traveling wave. Modes that should be kept for the rotor model, on the other hand, depend on how the rotor dynamics relate to the stator dynamics. Typically, the rotor is designed to have similar eigenmodes in the vicinity of the stator resonance, so all but those two modes are dropped for the rotor.

For mass normalized eigenvectors,

$$\mathbf{V}_s^\mathsf{T} \mathbf{M}_s \mathbf{V}_s = \mathbf{I}, \quad \mathbf{V}_s^\mathsf{T} \mathbf{K}_s \mathbf{V}_s = \Lambda_s \quad \text{and} \quad \mathbf{V}_r^\mathsf{T} \mathbf{M}_r \mathbf{V}_r = \mathbf{I}, \quad \mathbf{V}_r^\mathsf{T} \mathbf{K}_r \mathbf{V}_r = \Lambda_r. \tag{2.57}$$

Thus, the modal equations of motion become

$$\ddot{\mathbf{z}}_s + \Lambda_s \mathbf{z}_s = \tilde{\Theta} \mathbf{v} + \tilde{\mathbf{f}}_{sc}(\mathbf{z}_s, \mathbf{z}_r, h_r, \omega_r) \tag{2.58}$$

$$\ddot{\mathbf{z}}_r + \Lambda_r \mathbf{z}_r = \tilde{\mathbf{f}}_{rc}(\mathbf{z}_s, \mathbf{z}_r, h_r, \omega_r), \tag{2.59}$$

where

$$\tilde{\Theta} = \mathbf{V}_s^\mathsf{T} \Theta, \quad \tilde{\mathbf{f}}_{sc} = \mathbf{V}_s^\mathsf{T} \mathbf{f}_{sc}, \quad \text{and} \quad \tilde{\mathbf{f}}_{rc} = \mathbf{V}_r^\mathsf{T} \mathbf{f}_{rc}. \tag{2.60}$$

Similarly, (2.32) transforms in modal coordinates to

$$\tilde{\Theta}^\mathsf{T} \mathbf{z}_s + \mathbf{C}_p \mathbf{v} = \mathbf{q}. \tag{2.61}$$

Note how the otherwise uncoupled, forced equations of motion in (2.58) and (2.59) remain coupled through the contact forcing terms.

2.2.5 Frequency Domain

In the general case of non-ideal traveling-wave excitation, the nonlinearities of the interface model introduce superharmonics in both the modal states and the forcing terms. The voltage signals themselves aren't necessarily pure sinusoids either and can contain harmonic content as a result of the drive electronics. Nevertheless, for voltage inputs that are periodic at the drive frequency, ω, the states and forces will also be periodic in steady state. Thus, for steady-state analysis, the voltage signals, states, and forces can all be expressed as Fourier series expansions.

In exponential form [Siebert, 1986], the Fourier series expansion for an arbitrary periodic function, $x(t)$, with period $T = 2\pi/\omega$ is given by

$$x(t) = \Re e\left(\sum_{n=0}^{\infty} \mathbf{X}^{(n)} e^{in\omega t} \right) \quad \text{where} \quad \mathbf{X}^{(n)} = \frac{1}{T} \int_{-T/2}^{T/2} x(t) e^{-in\omega t} dt. \tag{2.62}$$

As mentioned in Section 2.1.1, the real operator can be implicitly assumed when using the complex exponential representation and will be dropped accordingly in the following formulations. Expanding the voltage, modal state, and modal force vectors as Fourier series

in exponential form $(\mathbf{v}(t) \Leftrightarrow \mathbf{V}^{(n)}, \quad \mathbf{z}_s(t) \Leftrightarrow \mathbf{Z}_s^{(n)}, \quad \mathbf{z}_r(t) \Leftrightarrow \mathbf{Z}_r^{(n)}, \quad \tilde{\mathbf{f}}_{sc}(t) \Leftrightarrow \tilde{\mathbf{F}}_{sc}^{(n)},$
$\tilde{\mathbf{f}}_{rc}(t) \Leftrightarrow \tilde{\mathbf{F}}_{rc}^{(n)})$ and substituting into (2.58) and (2.59), the modal equations of motion become

$$\sum_{n=0}^{\infty} (\Lambda_s - n^2\omega^2\mathbf{I})\mathbf{Z}_s^{(n)} e^{in\omega t} = \tilde{\Theta} \sum_{n=0}^{\infty} \mathbf{V}^{(n)} e^{in\omega t} + \sum_{n=0}^{\infty} \tilde{\mathbf{F}}_{sc}^{(n)} e^{in\omega t} \tag{2.63}$$

$$\sum_{n=0}^{\infty} (\Lambda_r - n^2\omega^2\mathbf{I})\mathbf{Z}_r^{(n)} e^{in\omega t} = \sum_{n=0}^{\infty} \tilde{\mathbf{F}}_{rc}^{(n)} e^{in\omega t}. \tag{2.64}$$

Applying harmonic balance by matching like harmonics and dividing out $e^{in\omega t}$, one arrives at the following frequency domain expressions for the modal dynamics of the system:

$$\boxed{\begin{aligned} \mathbf{Z}_s^{(n)} &= (\Lambda_s - n^2\omega^2\mathbf{I})^{-1}(\tilde{\Theta}\mathbf{V}^{(n)} + \tilde{\mathbf{F}}_{sc}^{(n)}) \\ \mathbf{Z}_r^{(n)} &= (\Lambda_r - n^2\omega^2\mathbf{I})^{-1}\tilde{\mathbf{F}}_{rc}^{(n)} \end{aligned}} \tag{2.65}$$

Considering that the interfaces forces are state dependent and must be calculated in the time domain, the above expression cannot be solved explicitly for $\mathbf{Z}_s^{(n)}$ and $\mathbf{Z}_r^{(n)}$. Instead, the solution process, as described in Section 2.2.7, involves searching iteratively for the implicit modal state solutions.

Understandably, the ultrasonic harmonic components of the interface forces have negligible effect on the inertial rigid body degrees of freedom. The rotor height and speed in (2.39) and (2.40) can then be considered constant variables in steady state even in the presence of unsteady interface forces due to an impure traveling wave. Thus, force and torque equilibrium are given by

$$\boxed{\begin{aligned} F_{eq} &= \langle f_{int} \rangle - f_{preload} = 0 \\ T_{eq} &= \langle \tau_{int} \rangle - \tau_{load} - \gamma_\omega\omega_r = 0 \end{aligned}}, \tag{2.66}$$

where the interface forces are now time averaged, leaving only the influence of DC components. The time average, $\langle \cdot \rangle$, is equivalent to the constant Fourier term, i.e. $n = 0$.

2.2.6 Performance Metrics

The desired performance metrics of a motor in steady state include speed, output power, and efficiency. The model solution yields the rotor speed, and thus the output power, but further computation is required to determine efficiency. Efficiency for an electromechanical system is defined here as the ratio of time-averaged mechanical output power to time-averaged electrical input power:

$$\eta = \frac{\langle P_{out} \rangle}{\langle P_{in} \rangle}. \tag{2.67}$$

Output power is the product of the externally applied torque load and the angular rotor speed and is constant because speed is assumed constant in steady state:

$$\langle P_{out} \rangle = \tau_{load} \omega_r. \tag{2.68}$$

Input power, however, is periodic, and additional steps are required to extract the DC component. The instantaneous value of input power is simply the product of the voltage and the current delivered to the piezo arrays:

$$P_{in}(t) = \mathbf{v}(t)^\mathsf{T} \mathbf{i}(t). \tag{2.69}$$

Until now the real operator in (2.62) has been implicitly assumed in each of the Fourier series expansions. However, the power calculation is a nonlinear operation, i.e. the product of two time-dependent variables, so the transformation of (2.69) to the frequency domain requires explicit retention of the real operators. In other words, the series expansions of voltage and current must be given explicitly by

$$\mathbf{v}(t) = \Re e \left(\sum_{n=0}^{\infty} \mathbf{V}^{(n)} e^{in\omega t} \right) \quad \text{and} \quad \mathbf{i}(t) = \Re e \left(\sum_{n=0}^{\infty} \mathbf{I}^{(n)} e^{in\omega t} \right). \tag{2.70}$$

Upon substituting (2.70) into (2.69) and integrating over one period, the time-averaged input power is expressed as a function of the frequency-domain voltage and current coefficients:

$$\langle P_{in} \rangle = \frac{1}{T} \int_{-T/2}^{T/2} \mathbf{v}(t)^\mathsf{T} \mathbf{i}(t) dt$$

$$= \frac{1}{T} \int_{-T/2}^{T/2} \left\{ \Re e \left(\sum_{n=0}^{\infty} \mathbf{V}^{(n)\mathsf{T}} e^{in\omega t} \right) \Re e \left(\sum_{n=0}^{\infty} \mathbf{I}^{(n)} e^{in\omega t} \right) \right\} dt . \tag{2.71}$$

Expanding the product of the two real operators in (2.71) can be avoided, thankfully, by the direct application of Parseval's Theorem [Siebert, 1986]. Specifically, the theorem states that the average power of a periodic signal is equal to the sum of the powers in its harmonic components. As a result, the calculation of the input power in the frequency domain reduces to

$$\langle P_{in} \rangle = \frac{1}{2} \Re e \left(\sum_{n=0}^{\infty} \mathbf{V}^{(n)\mathsf{H}} \mathbf{I}^{(n)} \right), \tag{2.72}$$

where $(\cdot)^\mathsf{H}$ denotes the Hermitian (complex conjugate) transpose.

The voltage is specified, and the current can be calculated from the time derivative of the charge, $\mathbf{q}$, in (2.61):

$$\mathbf{i}(t) = \tilde{\Theta}^\mathsf{T} \dot{\mathbf{z}}_s + \mathbf{C}_p \dot{\mathbf{v}} . \tag{2.73}$$

Substituting the appropriate series expansions into (2.73) and applying harmonic balance, current in the frequency domain is given by

$$\mathbf{I}^{(n)} = in\omega \tilde{\Theta}^\mathsf{T} \mathbf{Z}_s^{(n)} + in\omega \mathbf{C}_p \mathbf{V}^{(n)} . \tag{2.74}$$

Finally, replacing $\mathbf{I}^{(n)}$ in (2.72) with the above expression, the input power can be expressed as a function of the voltages and the modal states:

$$\langle P_{in} \rangle = \frac{1}{2} \Re e \left(\sum_{n=-\infty}^{\infty} in\omega \mathbf{V}^{(n)\mathsf{H}} \tilde{\Theta}^\mathsf{T} \mathbf{Z}_s^{(n)} \right) + \frac{1}{2} \Re e \left(\sum_{n=-\infty}^{\infty} in\omega \mathbf{V}^{(n)\mathsf{H}} \mathbf{C}_p \mathbf{V}^{(n)} \right). \quad (2.75)$$

It is seen here that power is comprised of two contributions, one electromechanical and one representing the dielectric loss in the piezoceramic. The latter term is non-zero only when $\mathbf{C}_p$ is complex, i.e. only when complex material properties are used to attribute losses to inherent material damping. In fact, it has been determined that the dielectric loss accounts for a significant fraction of the total electrical power.

2.2.7 Mixed-domain Solution Procedure

To summarize the model thus far, the linear modal and rigid-body dynamics have been formulated in the frequency domain whereas the nonlinear interface forces have been formulated in the time domain. The hysteretic nonlinearities prohibit steady-state solution completely in the frequency domain, but it is useful to preserve as best as possible the computational benefits of frequency-domain analysis. To that end, a mixed-domain solution procedure is preferred that alternates between domains while iteratively searching for the states that equate the full set of nonlinear equations.

The concept of switching between domains to analyze the steady-state response of nonlinear systems is relatively new. In 1989, Cameron proposed the Alternating Frequency/Time (AFT) method [Cameron, 1989] as an alternative to popular methods such as the Harmonic Balance (HB) method [Nayfeh, 1979] and Incremental Harmonic Balance (IHB) method [Lau, 1982]. The latter two require that the nonlinear terms be approximated by series expansions so that the system can be solved entirely in the frequency domain. Expanding the nonlinear terms demands additional work specific to the nonlinearity and can be very difficult, if possible, especially for discontinuous functions such as those involving contact or friction.

The AFT method, on the other hand, retains the ease and accuracy of analyzing nonlinear functions in the time domain where the state history is readily available while taking advantage of performing certain linear operations efficiently in the frequency domain. It also remains more general than either the HB or IHB methods as it does not require case-specific expansion of the nonlinearities.

Various iterative search techniques are available for solving a set of nonlinear algebraic equations. At each iteration, the frequency-domain states are transformed to the time domain where the nonlinear forces can be evaluated as functions of the state trajectories. The time-domain forces are then transformed back to the frequency domain where the linear operations in the system equations are performed. Specific details regarding the iterative solver used in the model simulations in Chapter 5 are given later in Section 5.1.3, but an overview of the solution procedure is illustrated in Figure 2.2 and described as follows.

The complete motor problem is described in the frequency domain by (2.65) and (2.66). The equations are fully coupled by the nonlinear, state-dependent interface forces. To summarize:

$$
\begin{aligned}
\mathbf{Z}_s^{(n)} &= (\Lambda_s - n^2\omega^2\mathbf{I})^{-1}(\tilde{\Theta}\mathbf{V}^{(n)} + \tilde{\mathbf{F}}_{sc}^{(n)}(\mathbf{Z}_s^{(n)}, \mathbf{Z}_r^{(n)}, h_r, \omega_r)) \\
\mathbf{Z}_r^{(n)} &= (\Lambda_r - n^2\omega^2\mathbf{I})^{-1}\tilde{\mathbf{F}}_{rc}^{(n)}(\mathbf{Z}_s^{(n)}, \mathbf{Z}_r^{(n)}, h_r, \omega_r) \\
F_{eq} &= \langle f_{int}(\mathbf{Z}_s^{(n)}, \mathbf{Z}_r^{(n)}, h_r, \omega_r)\rangle - f_{preload} = 0 \\
T_{eq} &= \langle \tau_{int}(\mathbf{Z}_s^{(n)}, \mathbf{Z}_r^{(n)}, h_r, \omega_r)\rangle - \tau_{load} - \gamma_\omega\omega_r = 0
\end{aligned}
\tag{2.76}
$$

The external inputs (operating conditions) are $\mathbf{V}^{(n)}$, $f_{preload}$, and τ_{load}. Obtaining a solution involves searching iteratively for the modal states, $\mathbf{Z}_s^{(n)}$ and $\mathbf{Z}_r^{(n)}$, and the rigid-body states, h_r and ω_r, that simultaneously balance those equations. Given that the assumed frequency-domain solution is truncated to N harmonic terms, this equates to solving $4N+2$ equations of just as many unknowns.

The interface forces are evaluated in the time domain by the models described in Section 2.2.3. To do this, the frequency-domain states guess at each iteration must first be

frequency domain
(FD)

time domain
(TD)

```
Initial FD        Gauss-Newton      Transform complex       Expand TD states to
states guess  →   algorithm     →   FD states to TD     →   spatial displacements
                                    (two periods in time)   spanning contact area
```

IFFT

Calculate spatial
interface forces
by time marching

N

FFT

```
                      Y   Do all EOM        Transform modal        Reduce spatial
Solution      ←           equilibrate?  ←   forces to truncated ←  interface forces
                                            set of FD harmonics    to modal form
```

Figure 2.2   Iterative mixed-domain solution procedure.

inverse transformed over a time period appropriate to the evaluation of the forces. The contact forces are expressed as an explicit function of the state trajectories, so f_w can be evaluated over a single period. The friction forces, however, are expressed as a hysteretic function of the state trajectories. Fortunately, the "term of memory" of this hysteretic nonlinearity is less than two periods in time for an arbitrary initial value of the slip variable, thus allowing evaluation of f_v from a finite state time history. Essentially, the arbitrary initial condition integrates out by the end of the first period, and evaluation over the second period correctly traces the periodic friction forces. The interface forces are converted to modal form then transformed to the frequency domain for an equilibrium assessment of (2.76) by the iterative solver. Upon convergence of a solution, postprocessing by (2.75) provides the input power and, subsequently, the efficiency of the motor.

Chapter 3

EXPERIMENTAL PREPARATION

To demonstrate the extent to which the preceding model can accurately simulate motor performance, real motors must be tested under controlled conditions where every aspect of the geometry, the materials, and the operating conditions are known. To this end, a few different motors have been constructed and the materials fully characterized. Modes have been qualified and the degree of non-ideality quantified. Numerous measurements have been made so that every model variable is accountable. This chapter is dedicated to describing the motors and characterization done in preparation for the performance tests and model simulations presented in the following chapters.

3.1 Motors

For the purpose of distinction amid parametric rotor studies, each "motor" is identified here by the use of a particular stator. Specifically, four motors have been included in this investigation: the Shinsei USR60, a "mock" Shinsei USR60, a "one-sided" prototype, and a "two-sided" prototype. The three unique assemblies are shown in Figure 3.1. The Shinsei USR60 is a commercially produced traveling-wave motor available from Japan that is highly recognized in this field. It's inclusion helps substantiate the performance measurements and claims documented in this work; however, the piezoceramic material that Shinsei uses proved difficult to obtain and could not be characterized to the extent necessary for the model. Therefore, a mock-up of the Shinsei USR60 stator was fabricated from accessible
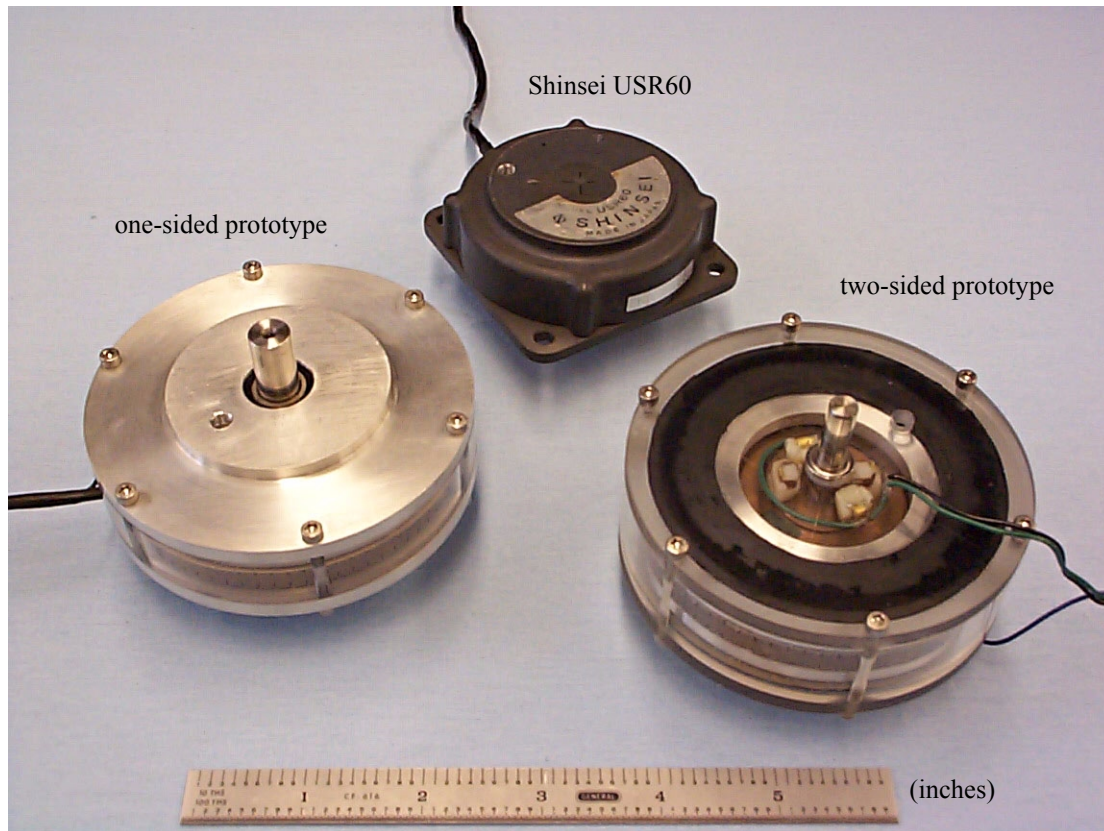
Figure 3.1   The three unique motor assemblies investigated in this work.

materials. This stator is called the mock Shinsei and shares the authentic Shinsei housing and parts for performance testing.

The one-sided and two-sided motors are in-house prototypes and are larger than the Shinsei motor. They each share the same basic symmetric stator geometry, but differ in the symmetry of operation. Like the Shinsei USR60, the one-sided motor generates torque via the frictional motive force created between one stator and one rotor. The two-sided motor, however, is a novel application of the traveling-wave motor concept in that it utilizes two rotors in contact with the opposing surfaces of a single symmetric stator. The motivation for this improvement was the potential to double the available torque and power output without significantly increasing mass.
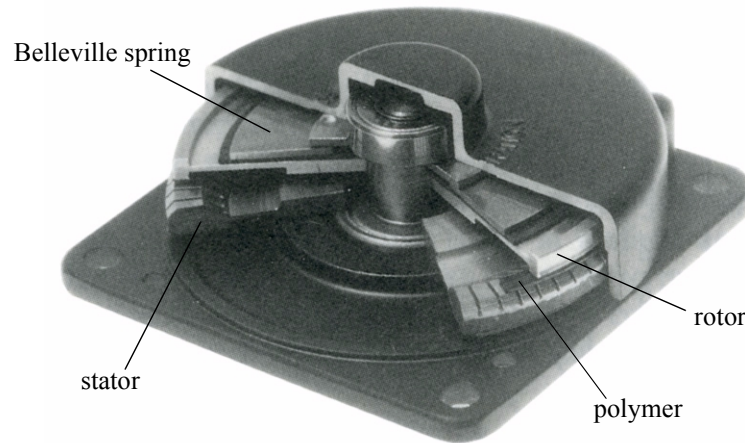
Figure 3.2   Cutaway of the Shinsei USR60 motor (older model shown).

To investigate the role of the rotor in determining motor performance, a series of rotors has been machined for each particular assembly. However, though the design of the rotor is arguably just as important as that of the stator in determining motor performance, its mention comes second to the description of each motor. Therefore, the rotor arrays will be described accordingly in Section 3.1.5 following the introduction of each of the four motors with regard to stator design and assembly.

3.1.1 Shinsei USR60

As indicated above, the Shinsei USR60 is a well-documented example of a piezoelectric traveling-wave motor as it is often a standard for comparison. Commercially produced by the Shinsei Corporation[1] in Tokyo, Japan, the USR60 with 60-mm diameter stator is the largest in a series of traveling-wave motors (TWMs), including the USR30 and USR45. It's construction is illustrated in Figure 3.2 and is typical of a disk-type traveling-wave motor.

Stator.   Shown in Figure 3.3(a), the stator substrate is phosphor bronze and has seventy-two teeth machined on one side for lateral displacement amplification. Often used in bells, phosphor bronze has a very high mechanical quality factor (low damping) and is an excellent choice for ultrasonic resonators. It also has excellent wear characteristics and is com-

1. Shinsei Corp., 2-1-8 Kasuya Setagaya-ku, Tokyo 157 Japan; phone: 81-3-3302-7677; fax: 81-3-3329-0066
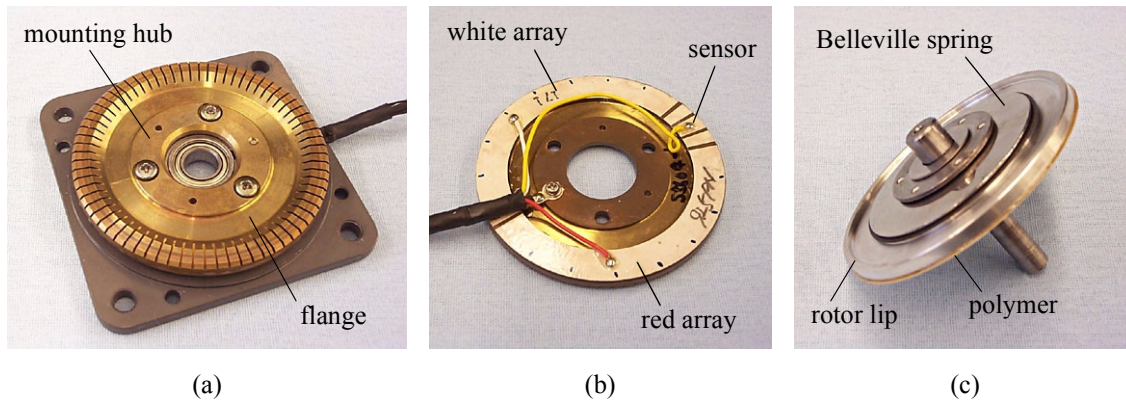
Figure 3.3 Components of the Shinsei USR60: (a) the mounted stator, (b) back side of the stator show-ing the piezo arrays, and (c) assemblage of the rotating parts.

monly used for mechanical bushings. A single piezoceramic ring, 0.5mm thick, is bonded to the bottom of the stator directly beneath the teeth, as seen in Figure 3.3(b). Chosen for its high power characteristics, the piezoceramic is a low-loss, hard lead-zirconate-titanate (PZT) composition known as NEPEC N-61. The particular composition is unique to Tokin Ceramics Corp. in Hyogo, Japan.

Each electrode array constitutes slightly less than half of the available piezoceramic area. They are spatially orthogonal and each consist of eight alternately poled sectors in order to excite a nine-wavelength traveling wave. Specifically, the USR60 utilizes the $B_{0,9}$ vibration mode, the bending mode of an annular disk with zero nodal rings and nine nodal diameters. A quarter-wavelength electrode sliver between the two arrays is used as a sensor electrode whereas the remaining three quarter-wavelength slivers are left unused.

The cross-section in Figure 3.2 shows that the radial geometry of the stator consists of a thick outer ring connected via a thin flange to a thick center hub. The flange provides sup-port to the outer ring and torsional rigidity for the transmission of torque but does not sig-nificantly contribute to nor affect the primary bending dynamics of the of the outer ring. Effectively a node of the $B_{0,n}$ modes, the center hub provides the means to bolt the stator to the lower housing cap. Another unique feature of the stator design is the trapezoidal radial cross-section of the teeth. This profile extends the height of the teeth without significantly

increasing mass or inertia while still efficiently transmitting the vibrational energy of the larger width piezoceramic to the protruding teeth.

Rotor. Shown in Figure 3.3(c), the aluminum rotor is simpler in construction and considerably lighter than the stator. Its radial geometry is similarly comprised of a prominent, thick outer ring, but the transition from thin flange to thick inner hub is interrupted by an intermediate ridge where the preload is applied. Bonded to the lower surface of the outer "lip" is a 0.2-mm thick polymer film. The polymer Shinsei chose to use is an Ekonol/PTFE blend manufactured by Sumitomo Chemical Co., Ltd., in Osaka, Japan, under the trade name Sumikasuper. Ekonol is a thermally stable, aromatic polyester. When combined with polytetrafluoroethylene (PTFE, or Teflon), it produces a composite material that has excellent temperature stability and wear resistance. Sumitomo also claims that the blend will not wear or damage soft metal counterparts. From the limited information provided by a Shinsei representative, the specific composition is thought to be Sumikasuper S300, a blend of 30% Ekonol in 70% PTFE.

Assembly. The rotor is pressed to the stator by preloading a Belleville washer, or conical disk spring, against the aforementioned rotor ridge. The axial load is reacted by a flange on the stainless steel shaft and through the large radial ball bearing in the upper housing cap. A smaller bearing centers the shaft in the lower housing cap just after it passes through the stator. The preload is applied by closing the housing and can be controlled by varying the thickness of shims placed between the shaft flange and the upper bearing. Because the housing reacts the axial preload near the centers of the upper and lower surfaces, it has been cast of sufficiently thick aluminum to avoid bending of the housing caps when assembled.

The advantages of using a disk spring are it's thin profile and nonlinear force-displacement relationship. With properly selected geometry, as Shinsei has done, the spring constant approaches zero under a predetermined load as the washer is pressed flat. Consequently, it responds like a much softer spring while occupying considerably less space. Better control of the preload is maintained than with a stiffer spring.

Slipping between the shaft flange and Belleville washer is prevented by six nubs on the flange that interlock with notches in the washer. A high-friction ring is instead used to prevent slipping between the Belleville washer and rotor. In Shinsei's earlier USR60 model in Figure 3.2, black rubber was used, but that was later replaced by a thin, lacquered metal ring that fits snugly around the rotor ridge.

Factory Specs. As specified in Shinsei's product literature, the nominal operating conditions and performance numbers for the USR60 are shown in Table 3.1 and Figure 3.4. Roughly calculated from the speed-torque curve, the maximum output of the USR60 is 3.5W around 45N-cm. The total mass is 240g. Torque density and power density are then calculated to be 258N-cm/kg and 14.6W/kg, respectively. Oddly, Shinsei does not provide information regarding input power or efficiency.

3.1.2 Mock Shinsei

Mock Shinsei refers to a stator that was constructed in-house with identical geometry and dimensions as the Shinsei USR60 but with a different piezoceramic composition. The reproduction endeavor was initiated for a number of reasons. Of primary interest, recognition of the USR60 in this field made it an excellent candidate as a testbed for model-experiment correlation. However, the NEPEC N-61 ceramic could not be obtained for

TABLE 3.1   Factory Specifications for the Shinsei USR60

Drive frequency	40kHz
Drive voltage	$100V_{rms}$
Preload	160N
No-load speed	100rpm
Maximum torque	62N-cm
Rated speed	90rpm
Rated torque	31N-cm
Rated power	3W
Limit of continuous operation	1000hr
Operating temperature	−10 to +50°C
Mass	240g
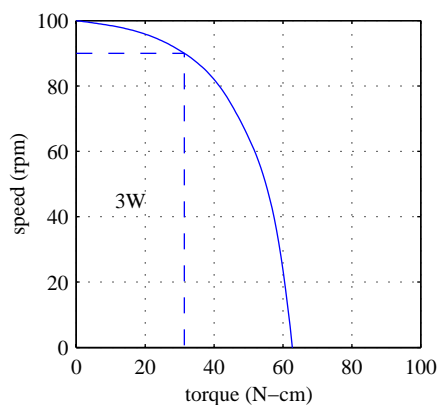
Figure 3.4   Nominal USR60 speed-torque characteristics.

characterization. Attempts were made to purchase ceramics from Tokin, but as the company appears to deal only in large volume for commercial applications, neither Tokin nor its distributors were able to provide the requested materials. Consequently, a different approach was taken: to reproduce the USR60 stator with similar, available materials.

As shown in Figure 3.5, the mock Shinsei stator is identical to the Shinsei USR60 stator in every dimensional detail, that is, to the tolerances of measurement and machining. Its substrate is made of free-cutting phosphor bronze (alloy 544), assumed to be the same phosphor bronze alloy used in the authentic stator. The one major difference between the two is the use of a piezoceramic composition designated PZT-4 (Navy Type I). Like NEPEC N-61, PZT-4 is a low-loss, hard ceramic used in high power applications. The two compositions are compared later in Section 3.3.1. Precut rings and various sample shapes were purchased from Staveley Sensors, Inc., in East Hartford, CT. Specifically, EDO Electro-Ceramic Products in Salt Lake City, UT, supplies Staveley with the bulk sintered ceramic, and Staveley performs the cutting, electroding, and poling (if requested). Additional preparation such as etching and repoling will be addressed in Section 3.1.6.

Another purpose of the mock Shinsei stator is that it provides a more controlled testbed for performance comparison with the one-sided and two-sided prototype motors. They have all been manufactured from the same materials, using the same techniques, under the same
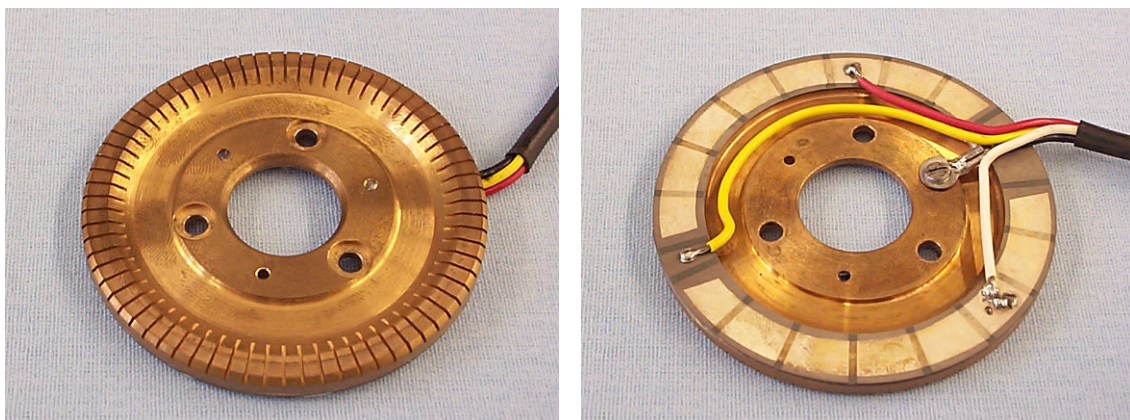
Figure 3.5 Mock Shinsei stator.

conditions. In effect, performance differences can be attributed more to geometric design and implementation of the traveling-wave motor principle and less to piezoceramic efficiency. On the flipside, manufacturing and testing of the mock Shinsei stator permits direct comparison with the commercial Shinsei USR60 stator with regard to manufacturing quality and material selection. Their differences will become evident in the Section 3.2.

3.1.3 One-sided Prototype

The first of the two in-house prototypes is the one-sided motor. The name "one-sided" is in reference to its conventional, asymmetric operation where only one side of the stator is utilized to propel a single rotor. For the most part, as seen by the cross-section in Figure 3.6 (also see the cutaway rendering in Figure 1.4), it's configuration and assembly are identical in concept to the Shinsei USR60.

The greatest inherent difference between the two implementations lies in the symmetry of the stator construction. In contrast to the asymmetric USR60 stator, the prototype stator integrates piezoceramic actuators and teeth on both sides. This symmetric stator design was specifically conceived for the purpose of two-sided operation, the ultimate objective driving the development of these prototypes, but found practical use in this standard implementation as well. The purpose of the one-sided assembly was to provide an intermediate testbed for studying the performance benefits of going from one-sided to two-sided operation under controlled conditions. Thus, both motors utilize nearly identical stators differing only by the means of support. The rotors are likewise nearly identical and will be discussed in Section 3.1.5.

Regardless of the symmetry of operation, the symmetric stator construction has the advantage of being sandwiched by two piezoceramic rings. This provides for more actuator authority because the effective moment of a bi-morph structure is greater than that of a mono-morph with the same total active thickness. Also, where voltage limits are a concern, a bi-morph configuration is considered advantageous because the total active volume can be doubled without having to increase the drive voltage. Coincidentally, Kawai *et al.* devel-
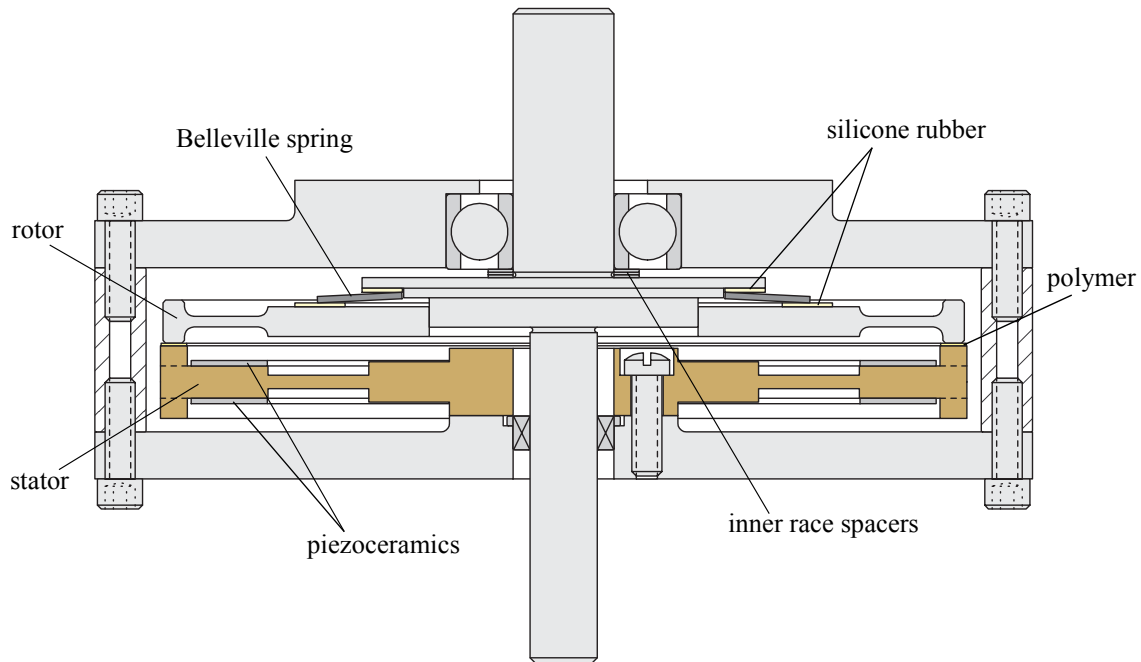
Figure 3.6 One-sided prototype motor assembly.

oped a similar sandwiched stator as a way to increase the electric input power to the piezo-ceramics and thus increase their motor's mechanical output [Kawai, 1995]. Aside from the improved stator design, however, their motor implementation was conventional in that it utilized only a single contact interface; they made no mention of the potential for two-sided operation.

Amidst its advantages, the symmetric stator does have its disadvantages. In the asymmetric construction, e.g. the USR60, the piezoceramic ring is bonded directly beneath the teeth, extending to the edge where the deflection is maximum and the greatest electromechanical coupling is achieved. Both the teeth and actuator exist where they are most effective. To create the bi-morph structure, however, the rings are moved inward of the teeth, away from the radius of maximum deflection. This reduces their authority to excite the desired mode. As a design choice to maximize torque output, the teeth are maintained at the outer radius at the cost of actuator placement and efficiency. Furthermore, the symmetric construction involves more machining and assembly, leading to greater production costs.

The prototype stator was designed with the intention of achieving at least 200N-cm of maximum torque under two-sided operation. Hypothetically, the available torque should be twice under two-sided operation than it is under one-sided operation for the same preload. Therefore, while the prototype stator was designed with a symmetric profile to enable the two-sided configuration, its diameter was sized appropriately to achieve at least 100N-cm in the conventional configuration. For comparison, the Shinsei USR60 manages a stall torque of 60N-cm with a 60-mm diameter stator under 160N (36lb) of preload. Assuming that Shinsei had already done some experimental optimization of the axial preload in the development of the USR60, an approach based on direct scaling was used in specifying the diameter and respective preload of the prototype stator.

As demonstrated by the analysis of scaling in Section 1.3.2, the torque output of a traveling-wave motor scales proportionally with the cube of the overall dimension. This scaling rule assumes that the axial preload is scaled by the square of the dimensional change, i.e. proportional to contact area, such that contact pressure is held constant. For example, by the assumptions of the scaling analysis, increasing the dimension by 26% should double the stall torque ($\times 1.26^3$) given that the force is scaled accordingly ($\times 1.26^2$). Therefore, to meet the desired torque goal with some room for error, the prototype stator was designed around a diameter requirement of 3" (76.2mm) and a predetermined preload of 58lb (258N).

As with the mock Shinsei, the stator substrate is made of free-cutting phosphor bronze for its excellent low loss and wear characteristics. Machined on either side are seventy-two rectangular teeth, a number which facilitates machining and electrode alignment. The piezoceramic rings are made of PZT-4 and are 0.020" (0.5mm) thick. Like the USR60, the prototype stator is designed to support a nine-wavelength traveling wave, but unlike the USR60, each piezo ring is etched into a full array of eighteen equally spaced sectors with alternating polarization. The two arrays, one on either side of the stator, are angularly offset by half a sector so as to excite orthogonal $B_{0,9}$ flexural modes.
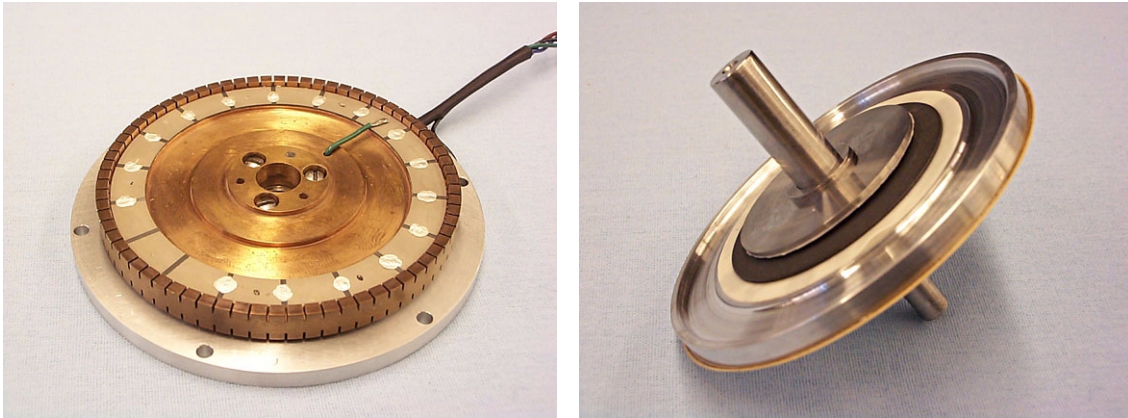
Figure 3.7 Components of the one-sided prototype motor: (a) mounted stator, and (b) assemblage of the rotating components.

In the one-sided configuration, only seventeen electrode sectors on each array are active so that one sector on the lower array can be used as a sensor during preliminary testing. To maintain actuator symmetry, a single sector on the upper array is left disconnected, as seen in Figure 3.7(a). The laminated stator is bolted at three locations through its center hub to the lower housing cap, and a small hole in the hub's outer ring provides access for the wire that supplies power to the upper actuator array.

In general, the rotating assembly shown in Figure 3.7(b) is very much like that found in Shinsei's earlier USR60 model that was illustrated in Figure 3.2. Specific details concerning the rotor design will be discussed in Section 3.1.5, but as with the USR60, the rotor's radial profile is defined by a distinct outer lip, a thin connecting flange, and a thick inner hub. The rotor is loaded by a Belleville spring located between the rotor's hub and a wide flange extending from the stainless steel shaft. The load is transmitted to the upper housing cap via a large, radial ball bearing that has been carefully chosen to handle the thrust load as well as provide radial support for the shaft. A much smaller, thin-walled ball bearing is used in the lower cap. Slipping between both the rotor and spring and the spring and shaft flange is prevented by thin rings of silicone rubber inserted at these junctions. The high coefficient of friction of the rubber guarantees that slip first occurs at the stator-rotor interface under high torque load. The axial preload can be adjusted by adding or subtracting inner race spacers (shims) between the shaft flange and large upper bearing.
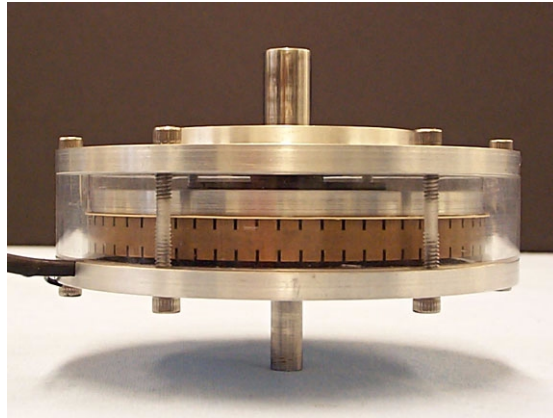
Figure 3.8 Fully assembled one-sided motor as seen through the transparent housing wall.

The upper and lower housing caps are machined from aluminum to resist bending under the axial preload whereas the cylindrical housing wall is made of tubular polycarbonate. As seen in Figure 3.8, the transparency of the housing wall provides visual access for inspection during testing and is invaluable for the purpose of demonstration. The one-sided prototype incorporates a through-shaft configuration similar to an encoder-ready version of the USR60. During testing, the torque load is applied to the quarter-inch diameter shaft protruding from the lower housing cap. The larger diameter end protruding from the upper housing cap, on the other hand, facilitates manual braking during demonstration. Total mass for the entire assembly is 420g.

As a final note regarding the design of the prototype stator, the multi-stepped radial contour of the center hub was utilized to manipulate the stator's modal characteristics. A coincident or nearby mode can couple with the desired mode and significantly hinder the ability to excite a traveling wave. Therefore, care must be taken during the profile design process to guarantee sufficient modal spacing around the desired mode. In the case of the prototype stator, the various dimensional variables of the center hub provided the necessary degrees of freedom to facilitate precise frequency placement of undesirable extensional and higher-order, flexural modes (e.g. $B_{1,n}$, $B_{2,n}$, etc.). The resonant frequencies of higher $B_{0,n}$ modes, on the other hand, are relatively insensitive to these dimensional changes because of the hub's location at a common node of these modes. Consequently, the process of unwanted

mode placement was decoupled from the general stator design process. Modal analysis was conducted using ANSYS finite-element analysis software. In the final stator design, careful dimensioning isolated the $B_{0,9}$ mode from all other modes by more than a couple kilohertz in either direction.

3.1.4 Two-sided Prototype

Last, but certainly not least, is the two-sided motor, a novel endeavor that was the ultimate focus of the experimental development efforts. At the onset of this research, the traveling-wave motor concept was already well understood and researched by many, but implementations were limited to the conventional asymmetric configuration where only one contact surface of the stator is utilized. It can easily be seen in Figure 1.3 that elliptical trajectories are generated by the traveling flexural wave on both sides of the stator, and, indeed, the transverse velocities at the opposing peaks are in the same direction. The unused dynamic surface presented a unique opportunity to improve on the traveling-wave motor concept by potentially doubling the available frictional motive force of a single stator. Thus, in an attempt to achieve 200N-cm of stall torque from a motor not too much larger than the Shinsei USR60, it was decided to investigate the novel progression to two-sided operation.

The concept of two-sided operation essentially involves pressing two identical rotors to the opposing surfaces of a symmetric stator. As one might expect, two-sided implementation presents a bit of a challenge as it requires a difference approach to a number of facets of the overall assembly. In an asymmetric, one-sided motor, the rotor is typically preloaded by the shaft flange, virtually sealing off the upper stator surface. That is why a small hole is necessary in the symmetric prototype stator to connect the upper piezo array for one-sided operation. If two opposing rotors were to be preloaded in a similar fashion, neither side of the stator would be accessible for hub mounting. Instead, an entirely different approach would have to be developed for mounting the stator ring at its outer periphery. Compounding the issue even further, electrical connections to the piezoceramics would likely have to be run between the slits in the teeth, or the piezo rings would need to be moved outward of the

teeth. Needless to say, it is desirable to maintain a hub-based (disk) geometry for the stator with the teeth at the outer radius so as to maximize torque output per overall motor dimension.

To overcome the accessibility problem, an inherently different configuration has been adopted for two-sided operation. Rather than being affixed to the housing, the active drive element (stator) has been press fit onto the shaft and subsequently given the role as rotor as well. Conversely, the passive aluminum washers (rotors) that would typically have rotated with the shaft are now stationary and are preloaded by Belleville springs in contact with the housing. A cross-sectional representation of this unconventional assembly is illustrated in Figure 3.9, followed by photographs of the shafted stator in Figure 3.10 and the fully assembled motor in Figure 3.11. Technically, the rotating drive element should be labeled as the rotor and the stationary washers as stators, but it will be less distracting to remain consistent with the accepted traveling-wave motor nomenclature. So, from here on, the active drive element will still be called the stator, and the passive washers will still be called rotors.
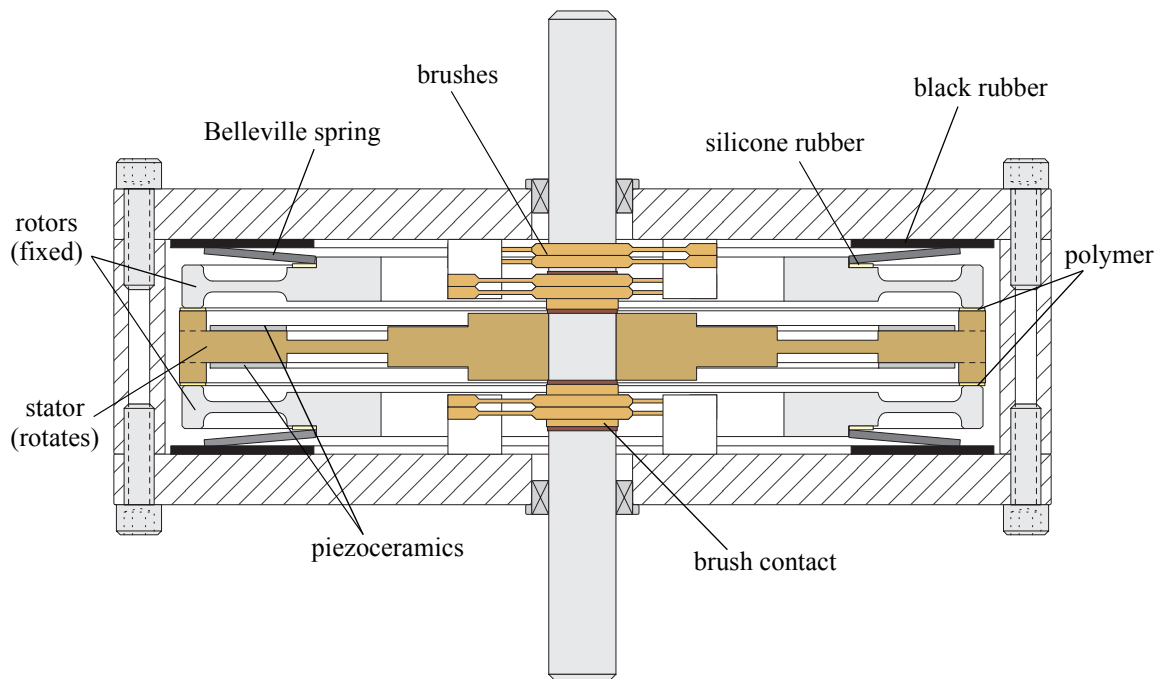
Figure 3.9 Two-sided prototype motor assembly.
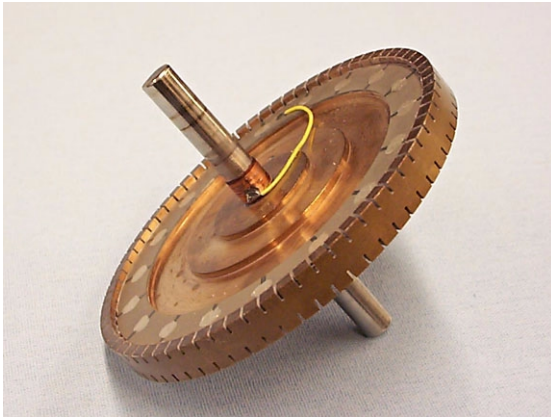
Figure 3.10 Shafted stator used in the two-sided motor.
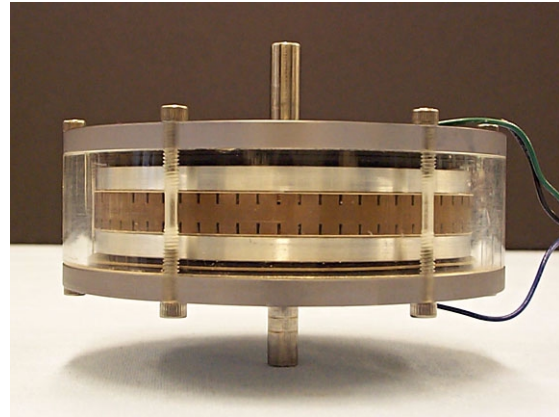
Figure 3.11 Fully assembled two-sided motor as seen through the transparent housing.

In accordance with such naming, the stator is essentially identical to the one described previously for the one-sided motor. The symmetric stator was, after all, primarily designed for use in the two-sided motor. The only difference in the two implementations is that, rather than having holes in the hub for mounting and an inner diameter large enough for the shaft to pass through, the stator hub has been precision bored to provide a nominal 0.0005" interference fit with the 1/4" stainless-steel shaft. As described later in Section 3.1.6, a hot-cold procedure was used to insert the shaft. At a radius of only 0.125", the stator-shaft interface might be feared as a weak link in the transmission of torque. Luckily, the interference fit held throughout motor testing. Nevertheless, future iterations would likely benefit from a larger diameter shaft or at least a shaft with a larger diameter waist.

One can quickly deduce that by placing the electromechanical stator in the rotating frame, brushes are necessary to deliver power to the piezoceramic arrays. This undoubtedly is the greatest disadvantage to having the stator rotate. Of greatest concern, the traveling-wave ultrasonic motor is electrically driven at high voltage and high frequency, neither of which are generally considered conducive to the use of simple brushes. Unlike with brushed DC motors, however, the rotational speeds are very low. The quality of signal transmission via a brush is dependent on a number of factors including the linear sliding speed between the brush and its contact. Therefore, to minimize the sliding speed of the brushes in the two-

sided motor assembly, the brush systems have been configured about the small-diameter shaft. At a radius of 1/8", a typical motor speed of 100rpm translates to only 1.3in/s. At such low sliding speeds, the brushes have proven to be fully capable of delivering clean, uninterrupted power to the piezoceramics.

Three electrical connections are required by the stator: the two high-voltage inputs and ground. On either side of the press-fit stator (see Figure 3.12(a)), a band of copper-coated Kapton is bonded around the shaft to act as the electrical contact for the appropriate input. A single short wire on each side then connects the copper bands with the actuator electrodes. The shaft itself is used as the ground contact. As seen in Figure 3.12(b,c), the brushes are essentially brass prongs and have been scavenged from cheap, off-the-shelf electromagnetic motors. At each contact location, four brushes are used, a pair on either side of the shaft. The redundancy ensures clean power transmission through every contact, especially when passing across the seams of the copper-coated Kapton bands. When fully assembled, the brush mechanisms make efficient use of the space inward of the rotors' inner diameters, as seen in Figure 3.12(d) through the transparent housing cap. Note that no sensor electrode is supported in the two-sided motor because running a second wire from one of the electrode arrays would have greatly complicated the brush contact implementation. Consequently, all eighteen sectors on either side are actuated in the two-sided motor.

Brush access on either side of the stator is made possible by using larger diameter Belleville springs and by reversing their orientation. In other words, the inner diameter of each Belleville spring, not the outer, now presses each rotor at approximately the same radius as in the one-sided motor. The axial preload is then reacted through the Belleville springs' outer diameters to the housing caps. The rotors are identical to those used in the one-sided motor except for a slightly larger inner diameter and a minor ridge on the inner hub that helps center the Belleville springs. Slipping between each of the interfaces is prevented by the insertion of black rubber and silicone rubber shims, and the axial preload is controlled by varying the total thickness of the rubber.

(a)   (b)

(c)   (d)
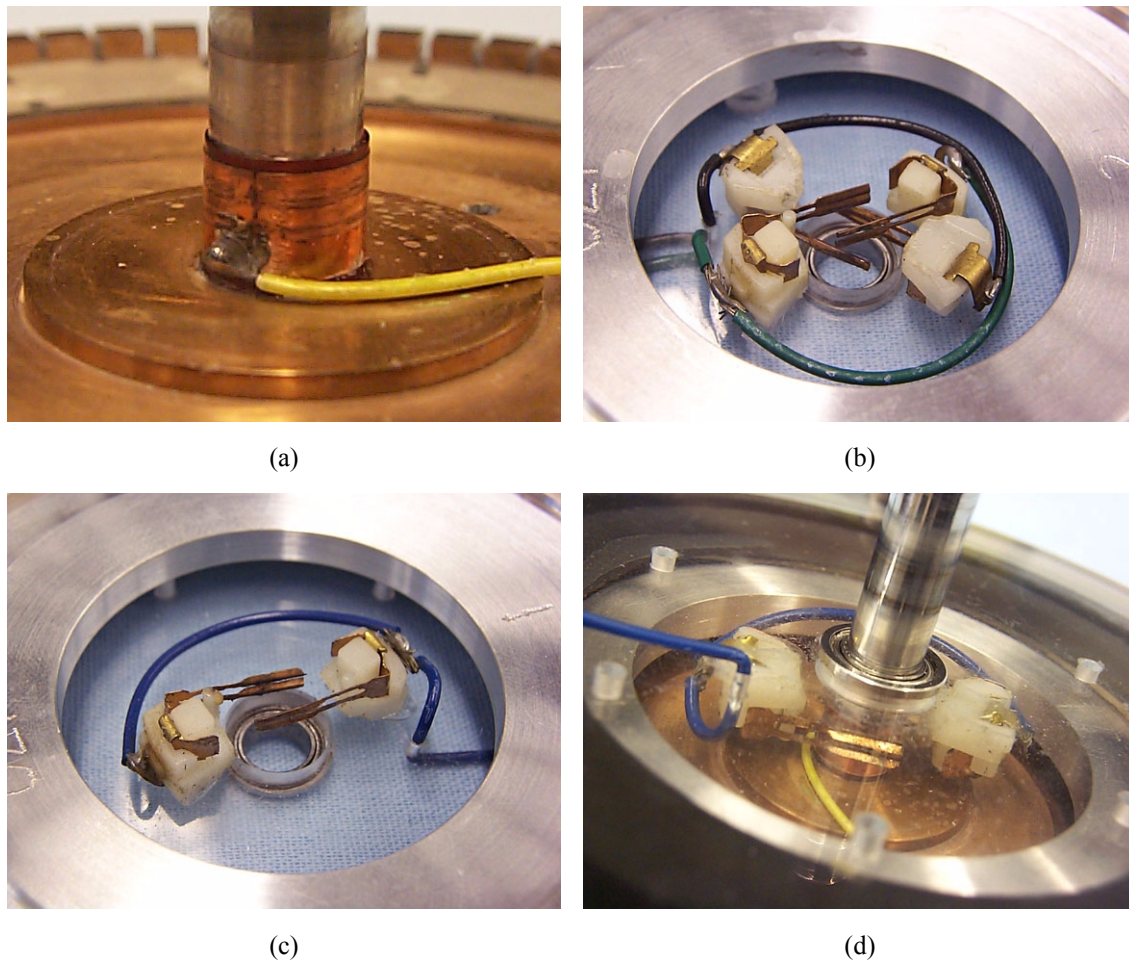
Figure 3.12   Brush system in the two-sided motor: (a) copper-coated Kapton brush contact, (b) upper brushes including ground set, (c) lower brushes, and (d) lower brushes in action as seen through the transparent housing.

As neither bearing is required to react thrust loads in the two-sided assembly, small, thin-walled ball bearings are used to center the shaft in both housing caps. Similarly, because the housing caps no longer react the preload near their centers but rather near the housing wall, their stiffness is no longer a critical issue. Thus, it was possible to construct the entire housing of transparent polycarbonate. In addition to providing visual access for monitoring the brushes and their contacts during testing, the clear housing caps proved to be invaluable during the manual process of aligning the flexible brushes.

The total mass of the two-sided motor is 350g, or about 17% lighter than the one-sided motor. Most of the mass savings can be attributed to the lighter-weight housing. Though

the second rotor and larger Belleville springs are sources of additional mass, these have been offset by the use of a lighter shaft and considerably smaller upper bearing. The significant difference in total mass was unintentional; rather, the difference is a direct reflection of the different load transfer paths in the two configurations. Understandably, some mass could have been shaved off the one-sided motor by sacrificing the stiffness of its load-bearing housing caps. Still, the 70g difference between the two motors is equivalent to the mass of one of the aluminum caps, so unless a different material were to be used, the one-sided motor would still weigh more than the two-sided motor.

3.1.5 Rotor Arrays

As a major focus of experimental study, a series of rotors have been machined for each motor in order to investigate the importance of rotor dynamics. Though commonly assumed rigid in older models, the rotor does indeed deform dynamically in response to the spatially periodic contact forces at the stator-rotor interface. More specifically, modes are excited that have the same circumferential periodicity as the excited stator modes. The vibration amplitudes of the rotor modes depend on the their resonant frequencies in relation to the stator resonance and the drive frequency. So, for each motor under investigation, the impact that a rotor has on performance is greatly defined by where the resonant frequencies of its $B_{m,9}$ modes lie in relation to the stator's $B_{0,9}$ modes.

The resonant frequency of the USR60 stator's $B_{0,9}$ degenerate mode pair is about 39.3kHz. This was obtained experimentally via the two methods outlined in Section 3.2. Excitation was naturally provided by the bonded piezoelectric actuators. The rotor, however, is a passive component. It has no built-in method of exciting modes, and direct contact with the rotor would change its modal characteristics. Any number of non-contacting techniques could easily be implemented to sense the rotor's vibrations, but the same does not go for exciting them. Therefore, rather than overcome the challenge of experimentally characterizing the passive rotor, numerical approximations have been made using finite element analysis.

Utilizing the tools provided in the ANSYS 5.7 Product Suite, an unconstrained, 2-D finite element model (FEM) of the USR60 rotor was created from axisymmetric-harmonic 8-node structural solid elements. Unlike standard axisymmetric elements, this element features the ability to model axisymmetric structures with circumferentially harmonic loading. When used in modal analysis, this translates to the ability to solve for modes with not only nodal radii but also nodal diameters. Thus, the element is well suited for this purpose. Furthermore, compared to a 3-D element model, an axisymmetric model provides an exact solution in the circumferential direction while the computational savings allow for greater refinement of the 2-D cross-sectional mesh. The result is that eigenvalue convergence is achieved with reasonable computational effort.

The initial, coarse mesh of the USR60 rotor FEM is shown in the top left of Figure 3.13, and the nodal solution of the $B_{0,9}$ mode is presented beneath. The lack of displacement boundary conditions reflects the assumed ideal load behavior of the nonlinear Belleville spring and its location of contact near the node. Beside these, the resonant frequency is shown to converge quickly for increasing mesh refinement. Specifically, the number of element divisions along each line have been scaled globally at each stage of mesh refinement by the multiplier shown on the x-axis. Not surprisingly, as a testament to the accuracy of the higher-order 8-node element and the simplicity of this particular mode, the so-called "coarse" mesh results in a frequency only about 0.1% higher than the converged value.

The resulting FEM prediction for the USR60 rotor's $B_{0,9}$ mode is 50.051kHz. Given a drive frequency of about 40kHz, the nearby location of the $B_{0,9}$ mode indicates that the rotor dynamics most certainly play a significant role. The vibration amplitude is likely very sensitive to dimensional changes in this transition region between quasistatic and resonant excitation. The hypothesis of the rotor study is that motor performance is largely determined by the dynamic response of the rotor. In other words, it is believed that performance can be optimized by appropriately "tuning" the rotor's resonant frequency in relation to that of the stator. No evidence could be found in the literature that anyone had previously related rotor design in this way, nor is it known how Shinsei arrived at this particular design. There-

| (a) | (b) |
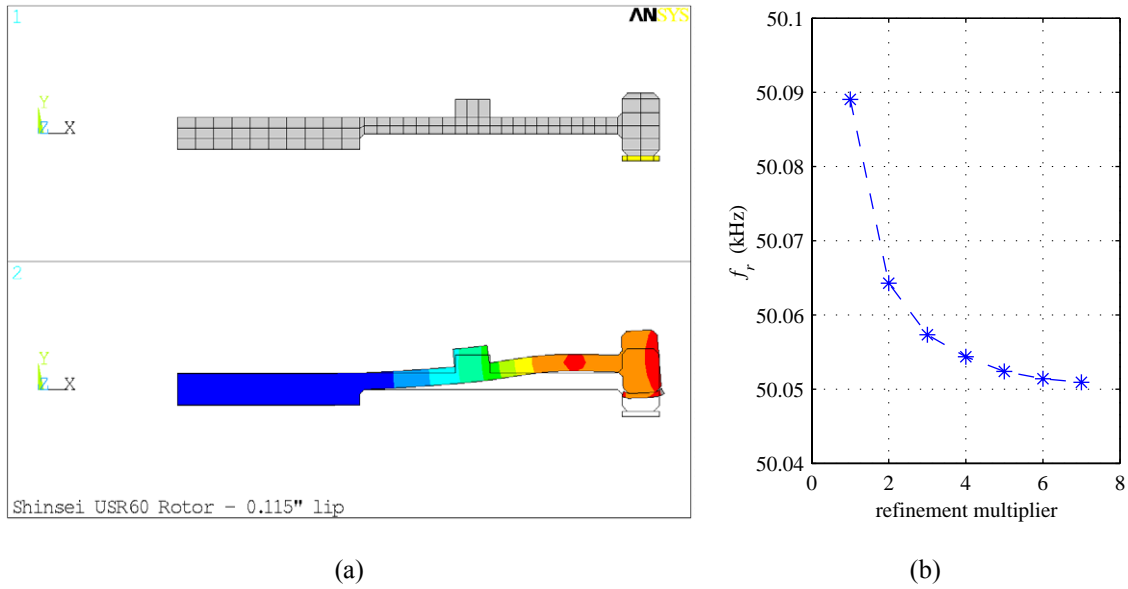
Figure 3.13 Axisymmetric-harmonic FEM modal analysis of the Shinsei USR60 rotor: (a) initial coarse element mesh and resulting contoured nodal solution of the $B_{0,9}$ mode, and (b) convergence of the $B_{0,9}$ resonant frequency vs. mesh refinement.

fore, there is no reason to believe that the nominal USR60 rotor has been optimally "tuned" for maximum performance.

Just as it was intended with the design of the disk-type stator, the $B_{0,n}$ bending modes of the rotor are predominantly influenced by the geometry of the thick outer ring. In the case of the rotor, this feature is generally thinner in the radial direction than though its thickness and takes on the appearance of a "lip". Similar to a beam in bending, the resonant frequency of the $B_{0,n}$ modes are particularly sensitive to the thickness of this lip but no so much to its width. Therefore, tuning of the rotor's resonant frequency for this study has been done by changing the thickness, or total height, of this outer lip. Actually, so as to maintain a consistent preload for varying rotor thickness without having to add or remove shim, only the upper lip height is modified while holding the lower lip height constant.

Using this approach, five rotors have been machined from 6061 aluminum for use in the USR60 motor: two thicker, two thinner, and one identical to the nominal Shinsei rotor. The array of lip heights is listed in Table 3.2 along with their respective frequency predictions as calculated from the ANSYS finite element model. Highlighted in the table, the nominal

TABLE 3.2 USR60 Rotor Array: FEM Predictions

Lip Height (in.)	$B_{0,9}$ Mode (kHz)
0.095	43.408
0.105	46.768
0.115	50.051
0.125	53.049
0.135	55.568

Shinsei design has a lip height of 0.115", not including polymer, and an estimated frequency of about 50.1kHz. The other rotor designs range from 0.095" to 0.135" in height and 43.4kHz to 55.6kHz in frequency, respectively.

The nominal rotor was made anew to remain consistent with the rest of the array, particularly in the use of consistent materials. As described in Section 3.1.1, Shinsei uses an Ekonol/PTFE blend for the interface polymer. A large, pretreated sample of this material was graciously supplied by Sumitomo Chemical for use in these rotor studies. The exact composition of the sample is 20% Ekonol in 80% PTFE, otherwise known as Sumikasuper S200. Unfortunately, though not verified by the Shinsei representative, the composition used by Shinsei is thought to be Sumikasuper S300 which has a higher fraction of Ekonol. The effective differences in the blends will be discussed later in Section 3.3.3. By including the nominal rotor in the manufactured lot, all rotors in the experimental array are made with the same materials via consistent techniques.

This same approach to making rotor arrays went into the design of rotors for the one-sided and two-sided motors except, of course, that there was no "nominal" rotor design from which to start. Instead, the Shinsei rotor was loosely scaled up in dimension to match the size and frequency range of the symmetric stator prototype. ANSYS was again used to calculate resonant frequencies. The finite element models of the resulting rotor configurations for the two prototypes are depicted in Figure 3.14. The two designs are essentially identical where it matters the most: the outer lip and the connecting flange. Their hubs differ, however, to accommodate the respective preloading methods and, in the case of the two-sided motor, a larger inner diameter for brush access.

(a)	(b)
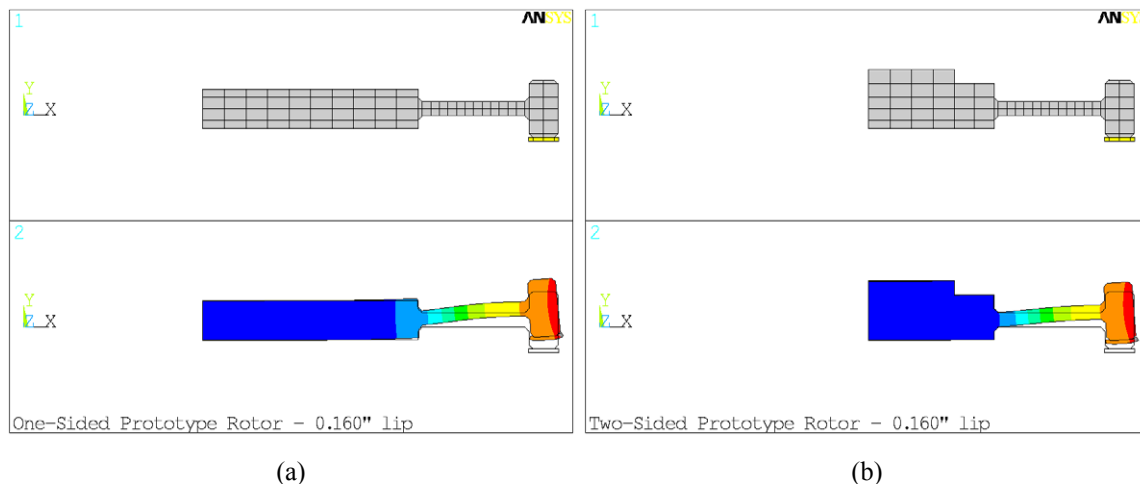
Figure 3.14 Axisymmetric-harmonic FEM modal analysis of the prototype rotors: (a) one-sided,   and (b) two-sided. Coarse mesh and nodal solution of the $B_{0,9}$ mode are shown for 0.160" lip rotors in both cases.

The $B_{0,9}$ mode of the symmetric stator prototype is around 29.6kHz, again measured experimentally using the techniques described in Section 3.2. Early experimental performance testing with the USR60 array indicated that the nominal design was a bit on the thin side, so the prototype arrays were scaled slightly stiffer in relation to the USR60 array. This resulted in the lip height selections and predicted frequencies in Table 3.3. Four rotors designs were made for each prototype motor with heights ranging from 0.150" to 0.180". Notice that the resonant frequencies are slightly different between the two designs because of the difference in hub geometry.

3.1.6 Manufacturing Details

Because of the exacting tolerances needed to build a functioning ultrasonic motor, all the parts that required machining were outsourced to professional shops. Outlined here, then, are the steps and practices followed during the preparation and assembly of the various motor components. These include etching, poling, bonding, lapping, and wiring.

Piezo preparation

The piezoceramic rings were supplied by Staveley Sensors in a uniformly poled state with electroless nickel electrodes plated on either side. In-house preparation included segment-

TABLE 3.3 Prototype Rotor Arrays: FEM Predictions

Lip Height (in.)	$B_{0,9}$ Mode (kHz)	
	One-sided	Two-sided
0.150	39.789	40.099
0.160	41.764	42.118
0.170	43.709	44.115
0.180	45.588	46.057

ing one electrode on each ring into the appropriate patterned array and then repoling the sectors in an alternating sequence. To minimize stress fracture of the fragile wafers, chemical etching was adopted over more violent abrasive or electric-arc techniques. Per Staveley's recommendation, a strong mixture of nitric and hydrochloric acid was needed to etch the newer nickel alloy, while simple flash tape was used to mask off the desirable electrode surfaces. As seen in Figure 3.15, etch lines between electrode sectors are approximately 1mm wide.

Repoling of the hard PZT-4 ceramic was done in a silicone oil bath at 3kV/mm and 100°C for one hour. Though the rings came uniformly poled from the supplier such that only half the sectors technically needed repoling, all sectors were subjected to this intensive schedule to guarantee equal poling strength in either direction. Simultaneous, multi-sector poling was facilitated with the aid of the crude yet effective paper clip technique illustrated in
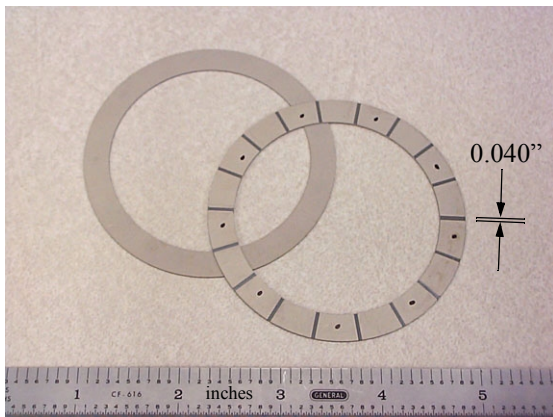
Figure 3.15 Virgin and etched piezoceramic rings for the symmetric prototype stator.
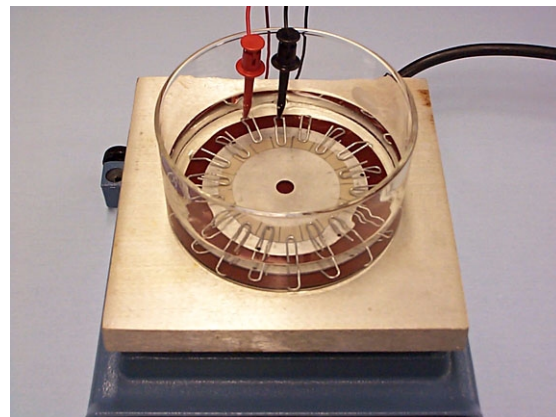
Figure 3.16 Method for simultaneous, multi-sector poling.

Figure 3.16. All sectors with like poling could be processed at once in this manner, thus requiring only two poling steps for each ring.

Stator preparation

Before bonding the piezoceramic rings, the stator teeth were lapped on a flat glass surface using grease-based silicon carbide compound. Edge effects, i.e. rounding, were minimized by rotating the part about its axis with a small DC motor. Starting with a 400-grit compound for coarse material removal, the process was repeated with increasingly finer compounds down to 1200-grit to produce the very flat, smooth finish required for efficient operation. A final polish with 4000-grit silicon carbide paper provided a mirror-like finish. Lapping preceded bonding to allow thorough degreasing of the stator substrates in a trichloroethylene bath. This step was necessary to ensure the removal of any and all silicon carbide particles. Trichloroethylene is a strong solvent that could be detrimental to the integrity of an epoxy bond if it were to be used post-bond.

While taking care not to scratch the finely lapped teeth, the bonding surfaces of the stator substrate and piezo wafers were then wet sanded with 400-grit paper and again degreased with trichloroethylene to remove any stray fingerprints. The surfaces were carefully cleaned with repeated acid and base scrubs followed by a final wipe with methanol. Similar to the procedure for bonding a strain gage, strict adherence to these steps was necessary to ensure consistently good bonds.

Epotek 301-2 two-part epoxy was chosen for this application because of its low viscosity, good shear strength, and long pot life. In fact, exceptionally thin bond layers on the order of 0.1–0.2mil were easily achieved with the 301-2, while its long pot life eliminated the chances of premature curing before final wafer alignment. Preliminary bonding tests of piezo samples on glass slides were carried out to observe the flow of the epoxy as a function of clamping pressure and to assess the likelihood of trapped air. Fortunately, these tests indicated that any initially trapped air was readily expelled during cure at clamping pressures greater than about 10psi without the need for vacuum. As a result of these tests, all

motor components were bonded under 25psi of pressure following a cure cycle of 80°C for two hours.

A thin, circular bead of epoxy was applied in the middle of each piezo ring so that the majority of the air would be pushed outward as each ring was lowered. The teeth slits were used as visual cues to align the electrode arrays, then flash tape was applied to secure the rings in position. In the symmetric stator design, centering of the piezo ring was crucial to keep a safe distance between the exposed live electrode and the grounded teeth. Therefore, thin strips of teflon were temporarily placed in the 20-mil gap and removed after cure. Pressure during cure was applied by stacked weights on a rubber shim.

All piezoceramic rings have been bonded with their segmented electrodes facing outward. The Shinsei USR60 stator, on the other hand, actually has the piezo bonded face down such that the fully segmented electrode is grounded to the substrate. This requires additional etch lines on the exposed surface to separate the two arrays from each other and the sensor electrode. It's an elegant solution, however, in that the segmented electrode arrays are reconnected during the bonding step. Though this configuration is preferable, and any further prototyping would have followed Shinsei's example, each wafer has, nevertheless, been bonded with the segmented electrode facing out to provide full access to individual sectors for modal testing. Consequently, an additional step was necessary to reconnect the segmented electrode arrays.

In an attempt to minimize resistance and correctly match the Shinsei's electrical characteristics, the electrode sectors on the mock Shinsei stator were reconnected by depositing an overlapping electrode across each of the arrays. The stator was first carefully masked with flash tape then exposed to thermal evaporation of an aluminum target via physical vapor deposition (PVD). The result of this laborious and costly procedure is shown in Figure 3.17(a).

The electrode sectors on the prototype stators, on the other hand, were reconnected with dabs of conductive silver paint. See Figure 3.17(b). This extraordinarily simple measure

(a)   (b)
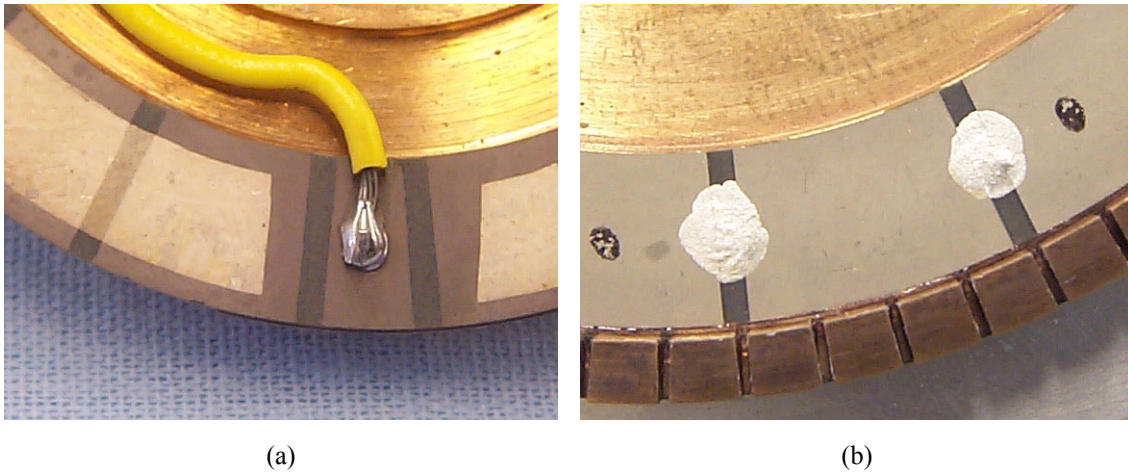
Figure 3.17 Methods for reconnecting electrode sectors: (a) vapor deposited aluminum on the mock Shinsei stator, and (b) conductive silver paint dabs on the prototype stators.

proved to be sufficiently effective and robust during testing as long as the surface had been properly cleaned. Furthermore, in the event that further modal testing was necessary, the conductive paint could be removed with solvents then quickly reapplied. However, electrical resistance is higher with a conductive paint than a deposited metal, potentially leading to increased stator damping. Nevertheless, this disadvantage was deemed acceptable in favor of the convenience of using the paint during iterative prototype testing.

Electrical connections to the mock Shinsei and one-sided prototype stators were straightforward. Signal wires were soldered to each of the two arrays and the sensor electrode. In the mock Shinsei motor, the ground wire was connected directly to the stator hub via a small screw, just like its original counterpart, whereas in the one-sided motor the ground wire was connected to one of the lower housing screws.

Final preparation of the two-sided prototype stator, however, involved the additional steps of press fitting the shaft and applying and connecting brush contacts. Taking advantage of thermal expansion, the oversized shaft was inserted by the heat-shrink technique. Specifically, the shaft diameter was reduced by cooling it in liquid nitrogen to $-196°C$ while the stator hole was expanded by heating the stator to $100°C$. Without hesitation, the shaft was then thrust effortlessly into position, and a tight interference fit was produced as the tem-

peratures equilibrated. Later, copper-coated Kapton bands were bonded to the base of the shaft on either side of the stator with Micro Measurements Group M-Bond 200 cyanoacrylate. The bands serve as brush contacts for the two high-voltage inputs. Soldering short wires between the piezoelectric arrays and their respective contacts finished the assembly.

Rotor preparation

Preparing the rotors involved bonding and lapping the polymer interface material. Fortunately, the sample roll of Ekonol/PTFE was supplied with one side already pretreated and ready for bonding. The material is 80% Teflon and, thus, is highly resistant to bonding in its natural state. By starting with a pretreated sheet, polymer prep was limited to cutting rings and scrubbing the treated bonding surfaces with methanol. The bonding surfaces of the aluminum rotors were roughened and cleaned following the procedures specified for the stator prep above. The same Epotek 301-2 epoxy was again chosen for affixing the polymer rings to the aluminum rotors after preliminary tests demonstrated that the bond would not peel as it was stronger than the polymer itself. The epoxy was applied in a thin bead around the lip of each rotor, and the polymer rings were secured in place by flash tape. Following the procedures developed for the stators, the sandwiched rotors were cured at 80°C for two hours under weights providing 25psi of clamping pressure.

A number of factors were taken into account when tailoring the methods for lapping the rotors. For one, the use of loose-particle compound was out of the question as the abrasive particles would undoubtedly lodge themselves into the soft polymer. Any abrasive particle caught in the contact interface would be catastrophic to motor operation. Instead, the rotors were lapped on fixed-particle silicon carbide paper starting with 400-grit for bulk removal and working down to 1000-grit. Flatness was maintained by firmly securing the paper to a glass slab, and the paper was kept wet to assure consistent removal of material. Finer grits were unnecessary as the soft polymer would naturally achieve its final finish during the initial "wear-in" period of motor operation.

The most important aspect of rotor lapping involved the applied load. When a motor is fully assembled, the rotor and stator bow under the pressure of the axial preload, resulting in a different radial contact profile than the unloaded state. Hence, to achieve ideal mating when fully loaded, the effect of bowing was counteracted by subjecting the rotors during lapping to loading conditions similar to those experienced during operation. Whereas only a minimal pressure was applied to the stators during the lapping process, a significant amount of weight was accordingly stacked on the hub of each rotor. Rather than producing a flat or planar surface, the polymer rings were effectively given a very slight conical finish.

In the two-sided prototype motor, only the rotors bow in response to the preload; the stator does not because it is symmetrically loaded. One can easily deduce that the appropriate amount of weight for loading its rotors is exactly that of the preload force. Consequently, these rotors were loaded with 58lbs while being lapped. In the one-sided motors, however, the stators also bow, though not by as much as the rotors. Since the stators were lapped flat, the extra slack had to be taken up by loading these rotors with weights in excess of the preloads. The appropriate loads were determined through trial and error by monitoring the evenness of radial contact after short periods of operation. As a result, the one-sided prototype rotors were lapped under 77lbs of load, and the Shinsei rotors were lapped under 48lbs, both approximately 33% higher than their respective preloads.

Once again to minimize edge effects, i.e. rounding, the rotors were generally rotated about their axes while being lapped. Proving to be a challenging task on coarse sandpaper, the massive stacks had to be moved by hand rather than by motor. Ultimately, the polymer layers were lapped from the stock thickness of 12mil down to 7mil to match the thickness of the polymer layer in the nominal Shinsei USR60 rotor. The rotors were then finished off by trimming the excess polymer to within a few mil of the inner and outer lip diameter.

3.2 Stator Qualification

The performance of any given motor hinges heavily upon the quality of its stator construction. In this context, quality refers to the degree of stator axisymmetry and the resulting

ability to excite a pure traveling wave. In the ideal case with perfect axisymmetry, the stator modes are degenerate, i.e. they have identical resonant frequencies and the rotational freedom to align with the actuator arrays. But in practice, no stator is perfect. Every real-life stator exhibits some manifestation of asymmetry that causes the degenerate modes to split in frequency and align with preferential axes defined by the asymmetry. This impedes the ability to excite a pure traveling wave and yields less than optimal motor performance.

Fortunately, it is possible to construct a stator that comes close to having perfect axisymmetry or at least one with near-perfect modal characteristics. Except under fine scrutiny, this is the case with the commercially fabricated Shinsei USR60. The near-perfect degeneracy of its modes can be witnessed by the fact that its $B_{0,9}$ modes are coincident in frequency and align as they should with their respective electrode arrays. The in-house stators, disappointingly, though very similar in construction, all demonstrate very noticeable mode splitting despite repeated efforts to improve upon manufacturing techniques. This debatable shortcoming of the experimental research was the most significant contributing factor to the generalization of the motor model to the non-ideal case. In a conceited effort to move forward with the ultimate goal of model-experiment correlation, it was decided to accept the less-than-ideal stator characteristics in the laboratory and instead account for them in the model. The degree to which mode separation was degrading performance, if at all, would have to be determined through simulations.

In this stator qualification section, the modal characteristics of each of the stators are identified and compared. Specifically, the resonant frequencies and quality factors are measured using an electromechanical impedance technique, and the mode shapes themselves are imaged via laser interferometry. Furthermore, visual observations are discussed regarding manufacturing quality control and how they may relate to induced asymmetry. The quantifiable results from these qualification tests will be used appropriately to model the imperfections of each of the stators in the model-experiment correlation studies in Chapter 5.

3.2.1 Electromechanical Admittance

Taking advantage of electromechanical coupling, the modal behavior of an active structure can be characterized by measuring its electrical characteristics as a function of frequency. This technique provides a quick and simple means for identifying mechanical resonant frequencies and for calculating mechanical quality factors, granted that the right equipment is available. Fortunately, this laboratory was in possession of an HP 4194A Impedance/Gain-Phase Analyzer with sweep capability, graphical display, and GPIB output for data storage. Using this device, the $B_{0,9}$ modes of each stator have been characterized by gathering electrical admittance data as a function of frequency for each of its two electrode arrays. Specifically, the data was measured and stored by the HP 4194A in the form of the magnitude and phase angle of the complex electrical admittance, or

$$Y = |Y|e^{i\angle Y}. \tag{3.1}$$

The technique is valid only when the response is linear, so the excitation voltage was limited to $0.1V_{rms}$. Also, note that while testing one array, the other array and sensor electrode, if applicable, were grounded to maintain consistent electrical boundary conditions with those during operation.

Wide frequency sweep examples for the two basic stator configurations are plotted in Figure 3.18. Electrically, piezoelectric actuators are generally capacitive. This is supported by the linear backbones of the admittance traces as a function of frequency, the slopes of which provide a measurement of the dielectric capacitance, C. However, superimposed upon the backbones are pole-zero pairs that arise from electromechanical excitation and sensing of the mechanical modes. These are the features of interest.

Before zooming in, notice the difference in apparent modal content between the two configurations. Beside the $B_{0,9}$ mode in the Shinsei USR60 trace, the neighboring $B_{0,8}$ and $B_{0,10}$ modes are just as conspicuous. This is not the case with the symmetric stator. The reason lies with the degree of orthogonality of the mode shapes to the different electrode patterns. In the symmetric stator, each electrode array constitutes a full circle of eighteen
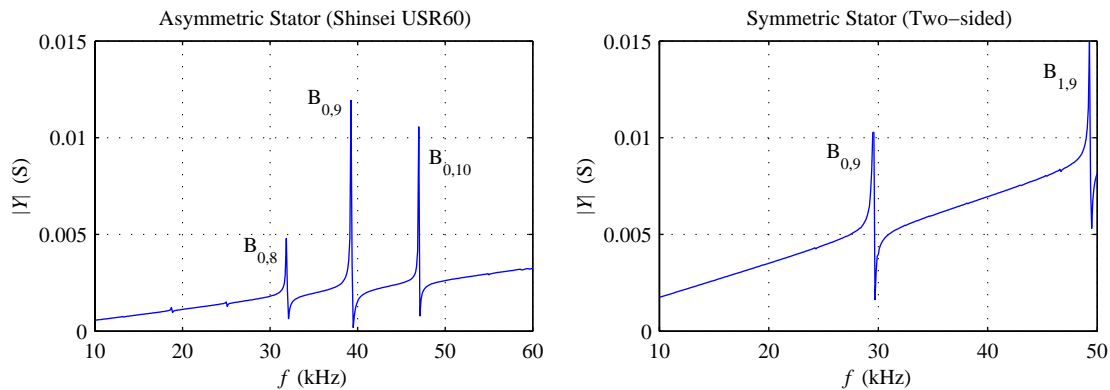
Figure 3.18 Comparison of the electromechanical admittance traces for the asymmetric (Shinsei USR60) and symmetric (two-sided) stator configurations illustrating which modes are controllable and observable by the specific electrode patterns.

equally spaced sectors. Ideally, only modes with nine-fold symmetry should be controllable and observable, all others being spatially orthogonal to the electrode pattern. Consequently, the $B_{0,9}$ mode stands virtually alone in the admittance sweep of the symmetric stator until the $B_{1,9}$ mode is reached. The electrode arrays of the asymmetric Shinsei USR60 stator, on the other hand, do not complete a full circle. Instead, each array constitutes only eight adjacent sectors or just under half of the full ring. The result is that although the modes with nine-fold symmetry are still the most controllable and observable, the finite electrode pattern is not ideally orthogonal to other modes, especially those with one more or less circumferential wavelength. Note that orthogonality is not necessarily related to motor performance, only that it provides some insight into the difference in the observed traces.

The admittance analysis becomes interesting when zooming in on the $B_{0,9}$ resonance. Consider first the room temperature data of the Shinsei USR60 stator shown in Figure 3.19 where the traces from the two electrode arrays have been arbitrarily identified by the colors of the lead wires, namely *red* and *white*. Clearly, the degenerate quality of the two $B_{0,9}$ modes is evidenced by how unmistakably coincident the traces are. In fact, the admittance peaks differ only by 8Hz, or about 0.02% of the 39.3kHz resonant frequency. The stator happened to be bolted to the housing in these tests, but note that the traces are no different when the hub is free. Consequently, the USR60 stator responds as it should, i.e. with degenerate modes that are independent of the hub boundary condition.
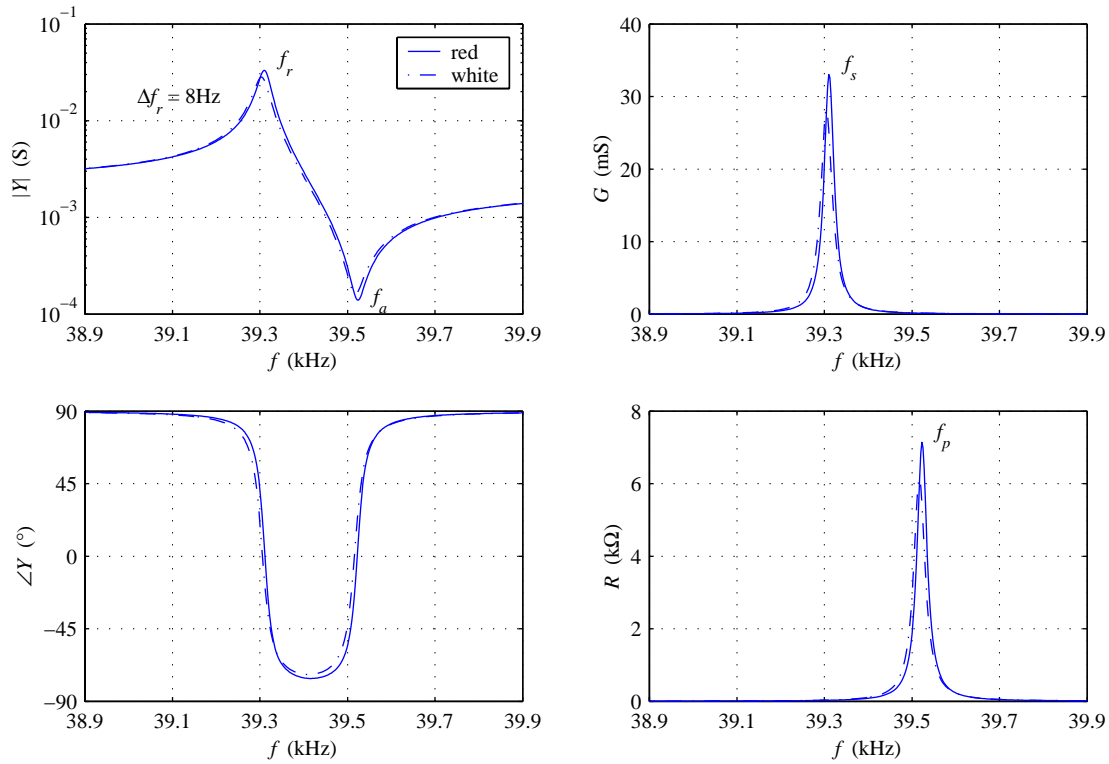
Figure 3.19 Electromechanical admittance of the Shinsei USR60 stator's $B_{0.9}$ modes at 20°C. The conductance, G, and resistance, R, have been calculated from the admittance data, Y.

As mentioned earlier, the electromechanical characteristics of each stator were measured and stored expressly in the form of the absolute electrical admittance, $|Y|$, and its phase angle, $\angle Y$. These are the values plotted in the left two plots of Figure 3.19. The two plots on the right are the associated values of conductance, $G = |Y|\cos\angle Y$, and resistance, $R = \cos\angle Y / |Y|$, as calculated from the admittance data. From the collection of traces, a number of characteristic frequencies can be identified:

- f_r — mechanical resonant frequency (maximum of admittance, $|Y|$)
- f_a — mechanical anti-resonant frequency (minimum of admittance, $|Y|$)
- f_s — series resonant frequency (maximum of conductance, G)
- f_p — parallel resonant frequency (maximum of resistance, R)

Then, as defined by Morgan Matroc, Inc. [Morgan, 1997], the mechanical quality factor, Q_m, of an electromechanical actuator can be calculated from these frequencies by

$$Q_m = \frac{f_a^2}{2\pi f_r Z_r C (f_a^2 - f_r^2)}, \qquad (3.2)$$

where $Z_r = \min(1/|Y|)$ is the minimum impedance value occurring at f_r. This nondimensional parameter provides an overall measure of material damping in the stator. Another useful quantity defined by Morgan is the effective coupling coefficient, k_{eff}, which reflects the frequency separation of the resonance (poles) and anti-resonance (zeros):

$$k_{eff}^2 = \frac{f_p^2 - f_s^2}{f_p^2} \approx \frac{f_a^2 - f_r^2}{f_a^2} \ . \qquad (3.3)$$

Unlike the more common piezoelectric coupling coefficients, k_{eff} is not only a measure of the material constants but also of the effectiveness of the actuator configuration in exciting and sensing the given mode. Basically, k_{eff} can be considered a quantifiable measure of stator design, piezoceramic placement, and material selection, and though not necessary for the model, it provides another figure of merit for comparison. For high mechanical quality factors, $f_r \sim f_s$ and $f_a \sim f_p$, so the former are often used for an approximate calculation of k_{eff}.

In summation, a meaningful description of the stator modes can be determined from the various characteristic frequencies identified in Figure 3.19. In the case of the Shinsei USR60, a negligible amount of mode splitting is identified, i.e. 8Hz, attesting to its symmetry. Then from (3.2), the mechanical quality factors of the two modes are calculated to be as high as 1400. This is indicative of exceptionally low internal damping. Lastly, from (3.3), respectable coupling between the actuator forcing and the modes is denoted by effective coupling coefficients of 10.35%. This last property will have more substantial meaning upon comparison with the other stators. Based solely on the results of electromechanical admittance observations, the preliminary conclusion regarding the quality of the USR60 stator is that it is clearly the product of excellent material selection, careful design, and, most notably, quality manufacturing.
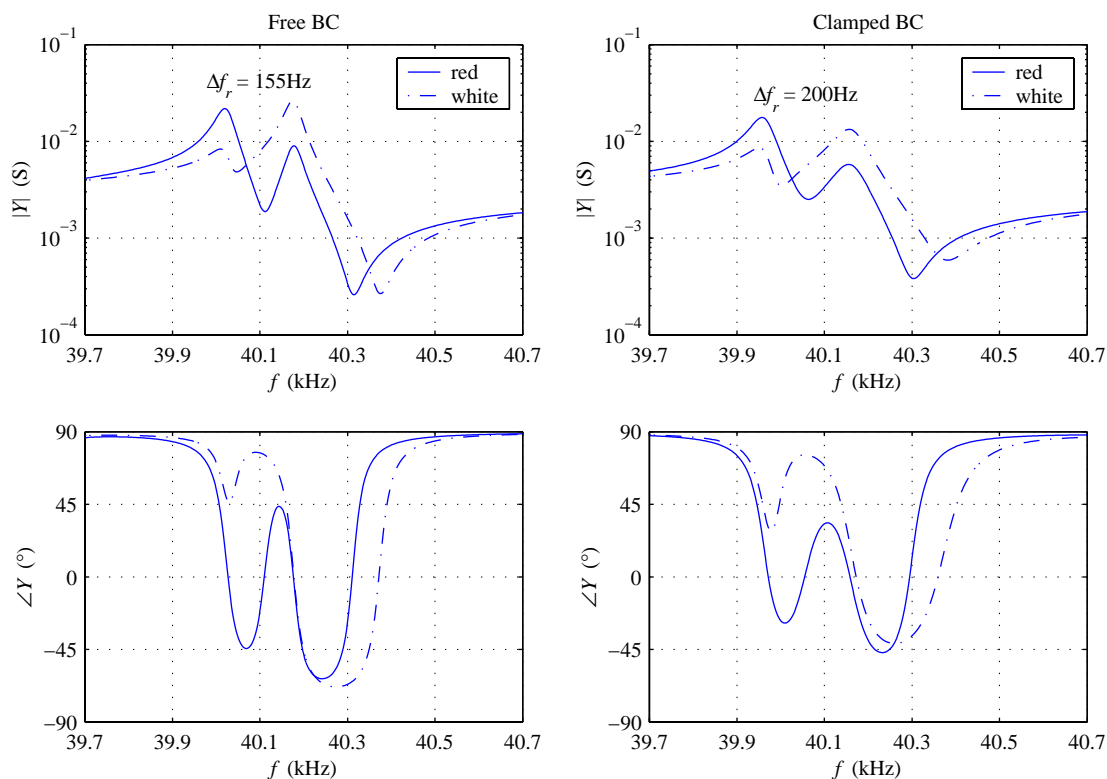
Figure 3.20   Electromechanical admittance of the mock Shinsei stator's $B_{0,9}$ modes at 20°C.

Now compare the modal characteristics of the Shinsei USR60 with those of the mock Shinsei stator shown in Figure 3.20. On the left is admittance data taken with the stator hub resting freely on a block of foam. On the right, the stator has been bolted tightly to the housing. Right away, a number of issues become discouragingly evident. For one, the fact that the boundary condition has an effect on the $B_{0,9}$ modes raises a flag as to the validity of the hub as a natural node. Obviously the modes of the mock Shinsei are corrupt in some way because the dependency on boundary condition indicates that appreciable vibration extends into the hub under free boundary conditions. Modal imaging in Section 3.2.2 verifies this claim.

The mock Shinsei stator also exhibits noticeable asymmetry as primarily evidenced by the separation of the resonant frequencies: 155Hz under free boundary conditions and 200Hz when clamped. The other result of the asymmetry, though maybe less obvious, is that the

$B_{0,9}$ modes have aligned with preferential axes that don't align with the electrode arrays. Consequently, orthogonality has been lost, and each mode is excitable by each array. The misalignment is manifested in the admittance data by the occurrence of both resonant peaks in each trace.

Unfortunately, the flawed mock Shinsei stator does not stand alone as a fluke of manufacturing. In Figure 3.21, a similar degree of mode splitting is observed in the admittance data of both of the symmetric stator prototypes. Here labeling of the electrode arrays has been associated with arbitrary *front* and *back* sides of each stator. The frequency separations are 180Hz for the one-sided stator and 172Hz for the two-sided. Interestingly, not only are the frequency separations about equal between the two, but also the mode rotations appear comparable. This can be concluded from the similar proportionality of the admittance peak amplitudes. Though these non-degenerate modes are anything but desirable, direct perfor-
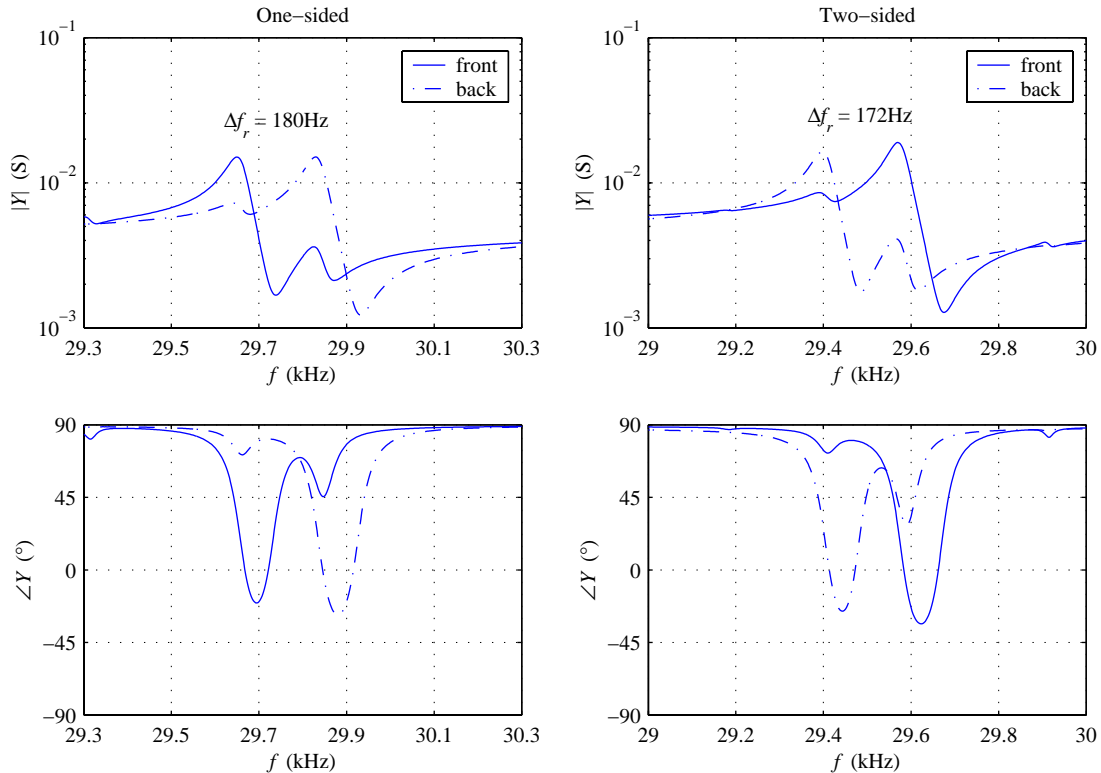
Figure 3.21 Electromechanical admittance of the one-sided and two-sided stators' $B_{0,9}$ modes at 20°C.

mance comparison of the one-sided and two-sided motors can still be considered valid due to the fact that their stators' modal characteristics are nearly identical, especially with regard to mode splitting. It is important to note, however, that like the Shinsei USR60 stator, neither of the symmetric stators exhibited any changes in their admittance traces in response to changes in the hub boundary condition. Luckily, that issue was isolated to just the mock Shinsei stator.

Comprehensive results of the room-temperature stator admittance tests are summarized in Table 3.4. Note that evaluation of (3.2) and (3.3) is valid only when the anti-resonance of the mode in question is not obscured by the resonance of a second mode. Consequently, Q_m and k_{eff} could only be calculated for the higher frequency mode in cases plagued with mode splitting. Nevertheless, a number of conclusions can be drawn from the available numbers. To start, the differences in capacitance, resonant frequency, and quality factor between the Shinsei USR60 and mock Shinsei are strong indications that the materials used are noticeably different. The effective coupling coefficients are very close, however, as the two share identical geometry. The mock Shinsei stator has been listed twice to reflect the change in its modal characteristics due to boundary condition. Notice that in addition to greater mode splitting, clamping of the hub introduces considerable damping.

TABLE 3.4   Stator Characteristics at 20°C from Admittance Measurements

Stator	Array	C (nF)	f_r (kHz)	Q_m	k_{eff} (%)
Shinsei USR60	red	8.88	39.310	1408	10.35
	white	8.81	39.303	1205	10.35
Mock Shinsei (free)	red	11.29	40.019	n/a	n/a
	white	11.37	40.174	920	9.87
Mock Shinsei (clamped)	red	11.29	39.957	n/a	n/a
	white	11.37	40.156	418	10.14
One-sided	front	24.02	29.649	n/a	n/a
	back	23.85	29.830	490	7.67
Two-sided	front	27.75	29.570	514	7.89
	back	25.03	29.397	n/a	n/a

All of the in-house stators, however, have been fabricated from the same materials. Therefore, the differences in modal characteristics between these stators can be attributed to differences in configuration and geometric scale. As expected, the full-ring electrode arrays of the symmetric stators have significantly higher capacitance than those of the mock Shinsei, but notice that increasing the volume of piezoceramic in relation to phosphor bronze has also increased stator damping (compared to the free mock Shinsei stator). This is because the mechanical quality factor of PZT is considerably lower than that of phosphor bronze. Furthermore, as a result of moving the piezoceramic rings inward of the teeth, away from the radius of maximum deflection, the effective coupling of the piezo arrays is noticeably less in the symmetric configuration than in the asymmetric Shinsei design. Lastly, when compared to each other, the two symmetric stators are by and large equivalent except for the difference in capacitance. Remember that the one-sided motor incorporates a sensor electrode, whereas the two-sided motor has all eighteen electrode sectors connected in each array.

The electromechanical admittance measurements have also been taken as a function of stator temperature to account for motor heating during operation. The results of these tests for the Shinsei USR60, mock Shinsei, and one-sided motor stators are plotted in Figure 3.22. As expected, resonant frequencies drop as temperatures rise as a result of material destiffening. The greater sensitivity of the symmetric stator indicates that the stiffness of the piezoceramic is more temperature dependent than that of the phosphor bronze. Temperature studies of the individual materials in Section 3.3 support this deduction.

3.2.2 Laser Interferometry

The stators have been further characterized by imaging their displacement modes shapes with a Polytec scanning laser Doppler vibrometer (SLDV), model PSV-300H. While sweeping the frequency of excitation, this impressive piece of equipment scans the surface of a vibrating object with a heterodyne laser interferometer to produce two-dimensional spatial imaging of the vibration profile. Therefore, in addition to identifying resonant fre-
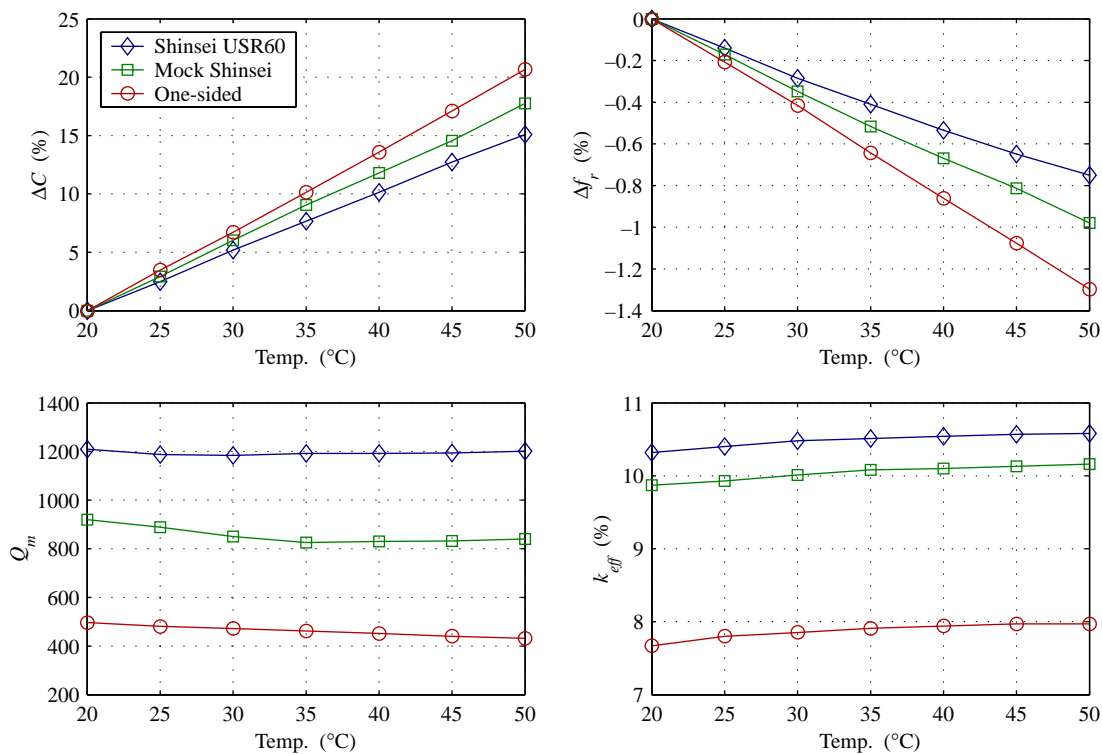
Figure 3.22   Stator characteristics as a function of temperature from admittance measurements.

quencies, the SLDV provided the means to observe the spatial quality and angular orientation of the stator modes.

Figures 3.23–3.26 depict various representations of the vibration data for the Shinsei USR60, mock Shinsei, one-sided, and two-sided stators, respectively. Each column of images is identified by the active array, i.e. red or white or front or back, and the excitation frequency, i.e. the resonant frequency of the prominent modal response. In the middle images, the raw two-dimensional vibration data is depicted by instantaneous displacement contours. Notice that a convenient feature of the PSV is the ability to superimpose the data onto a picture of the target taken by a built-in CCD camera coincident with the scanning laser head. Then in the upper images, the contour values have been interpreted as height to produce three-dimensional renderings of the mode shapes. The overlaid meshes in these images not only assist visualization of the shapes but also reveal the laser scan points. Lastly, in the lower images, the node lines have been isolated by displaying a single isoline

Red Phase @ 39273Hz   White Phase @ 39278Hz
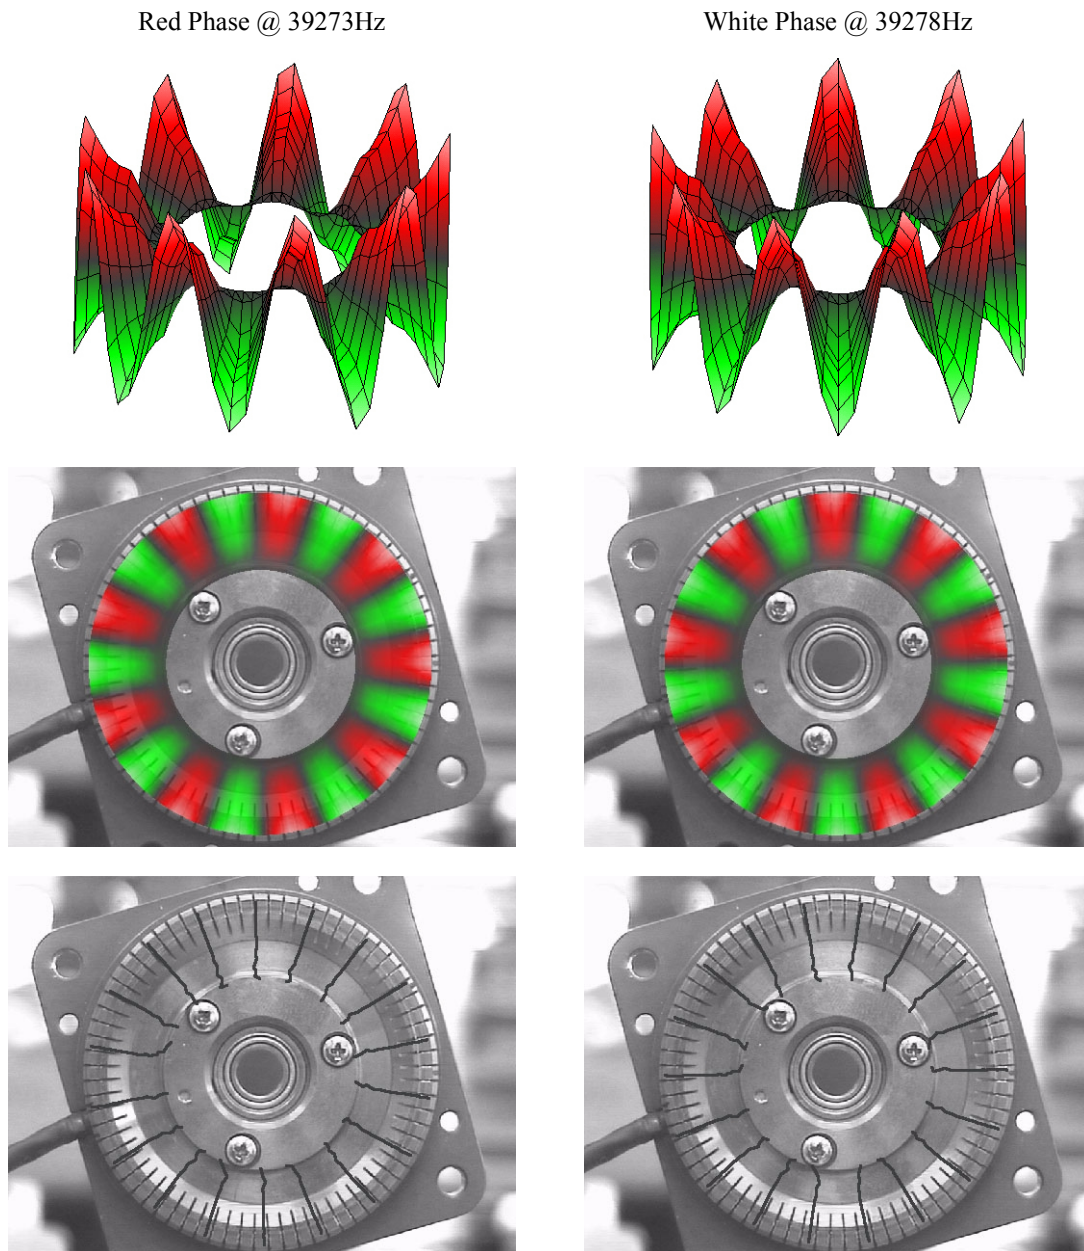
Figure 3.23   Laser vibrometer scans of the Shinsei USR60 stator modes.

where the interpolated value of the vibration amplitude is zero. Though these images are the least impressive, the sharp isolines make it easier to quantify mode rotation. All of these scans were carried out at room temperature while holding the input voltages constant at $0.5V_p$.

As one would expect given the results of the admittance studies, the $B_{0,9}$ modes of the Shinsei USR60 stator appear nearly ideal in Figure 3.23. The amplitudes are consistent over all nine circumferential wavelengths and appropriately null out toward the hub. However, though the two modes do look orthogonal to one another, they do not align perfectly with the electrode etch patterns. Barely visible on the edges of the teeth, the etch locations have been transferred to the front of the stator with small ink marks. By comparing the node lines with the marks in the isoline images, both modes appear to be rotated by about a quarter of a tooth, or $\lambda/32$, with respect to the electrode array axes. The rotation is quite small, though, and could very well be indicative of a misalignment of the USR60's inner and outer electrode etch patterns. Unfortunately, there is no way to verify or disprove that possibility without debonding the wafer. In any case, other than this one inconsistent observation, the laser scans of the USR60 stator further defend the likelihood of it supporting an ideal traveling wave.

In contrast, the mock Shinsei stator modes shown in Figure 3.24 are not so perfect. One can clearly see that the vibration amplitude of the hub is not negligible. In fact, it appears that each $B_{0,9}$ mode may be weakly coupled with a higher-order bending mode, possibly with 5θ periodicity. There is no question that this would cause the modal characteristics of the mock Shinsei stator to be dependent on the hub boundary conditions as was witnessed in the admittance tests. Away from the hub, however, the modes look distinctly familiar, and one could argue that the peak amplitudes appear relatively consistent around the periphery. Nevertheless, the quality of the mock Shinsei mode shapes is highly questionable and will mostly certainly have a negative effect on motor performance. Furthermore, rotational misalignment of the node lines is quite evident at nearly three-quarters of a tooth, or $3\lambda/32$. Remember that the significant amplitudes of the lesser admittance peaks in Figure 3.20 were strong indications that this would be the case.

The PSV images of the two symmetric stators are shown in Figures 3.25 and 3.26. As expected, though the mode shapes aren't perfect, they are better than the mock Shinsei modes with respect to hub vibration. This can best be seen in the three-dimensional images

Red Phase @ 40031Hz   White Phase @ 40210Hz
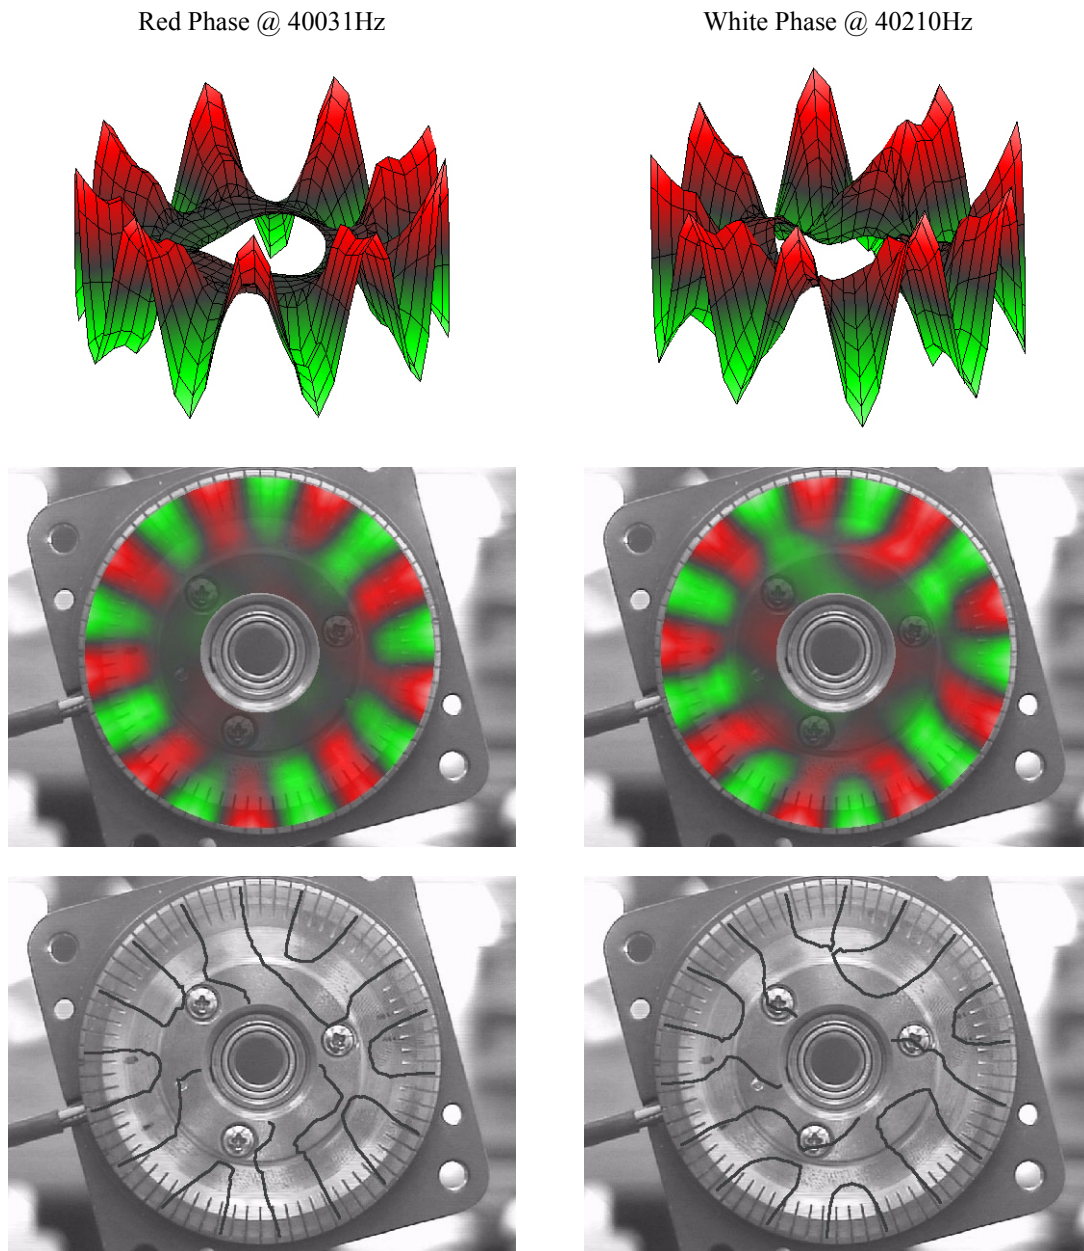
Figure 3.24   Laser vibrometer scans of the mock Shinsei stator modes.

of the two-sided stator modes where the vibration amplitude approaches zero at the inner-most radius of the laser scans. Unfortunately, it was an oversight that the scans weren't extended as far inward for the one-sided stator, but if they had been, the shapes would have looked nearly identical to those of the two-sided stator. There is faint distortion of the modes of both stators inward of their piezoceramic rings, but the likelihood of undesirable

Front Phase @ 29651Hz   Back Phase @ 29828Hz
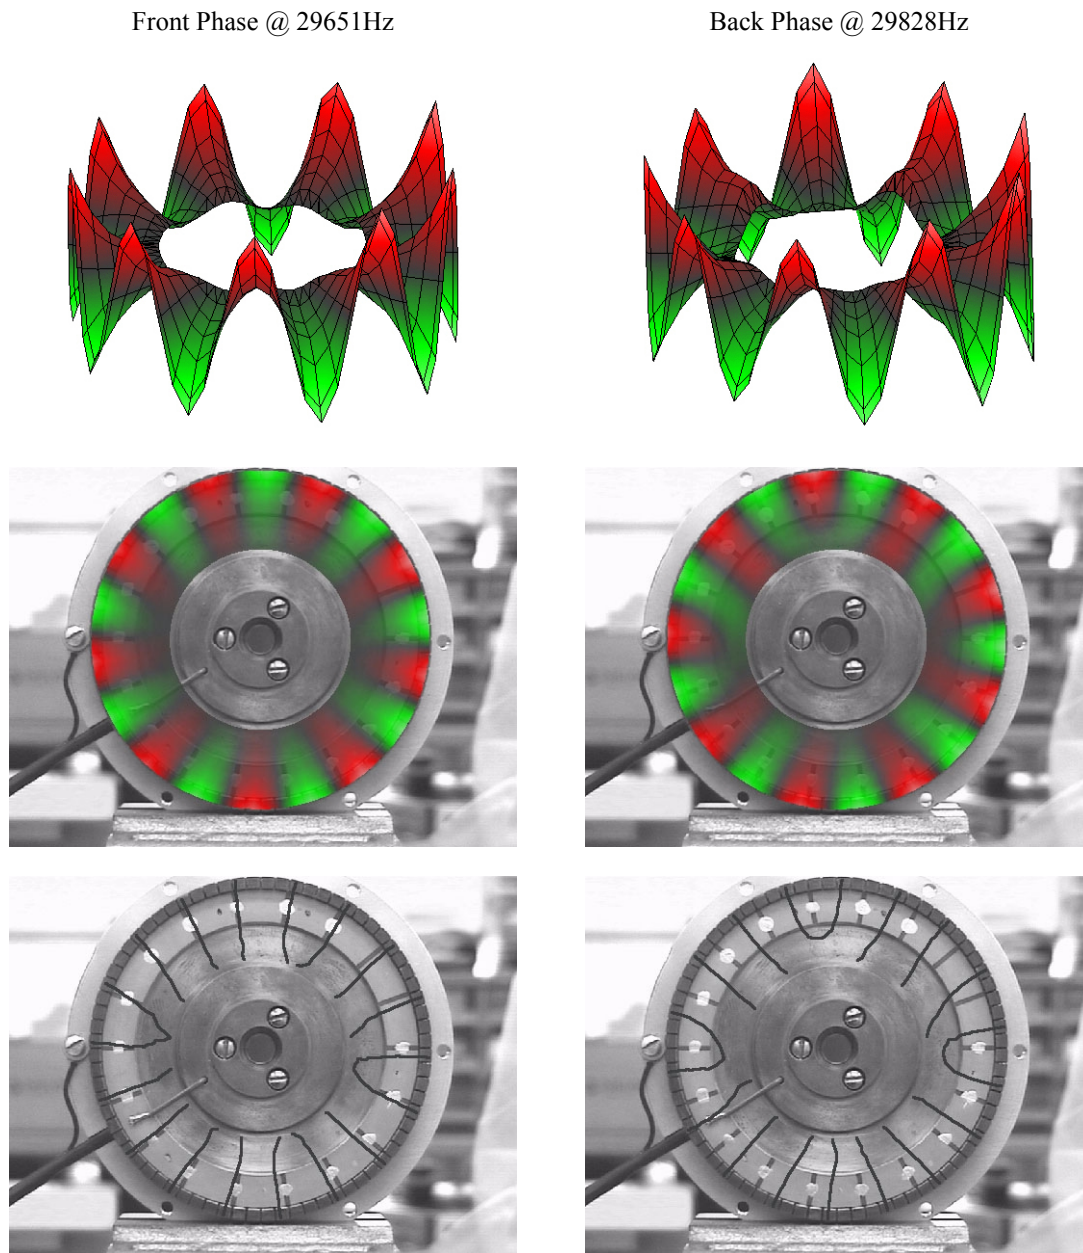
Figure 3.25   Laser vibrometer scans of the one-sided stator modes.

mode coupling is nowhere near as apparent as it is for the mock Shinsei stator. Both stators do, however, exhibit preferential mode orientations. This is much easier to see with the symmetric stators because of the visible electrode patterns in the isoline images. In either case, the modes are oriented about a half of a tooth, or $\lambda/16$, out of alignment with the electrode arrays.

Front Phase @ 29569Hz   Back Phase @ 29398Hz
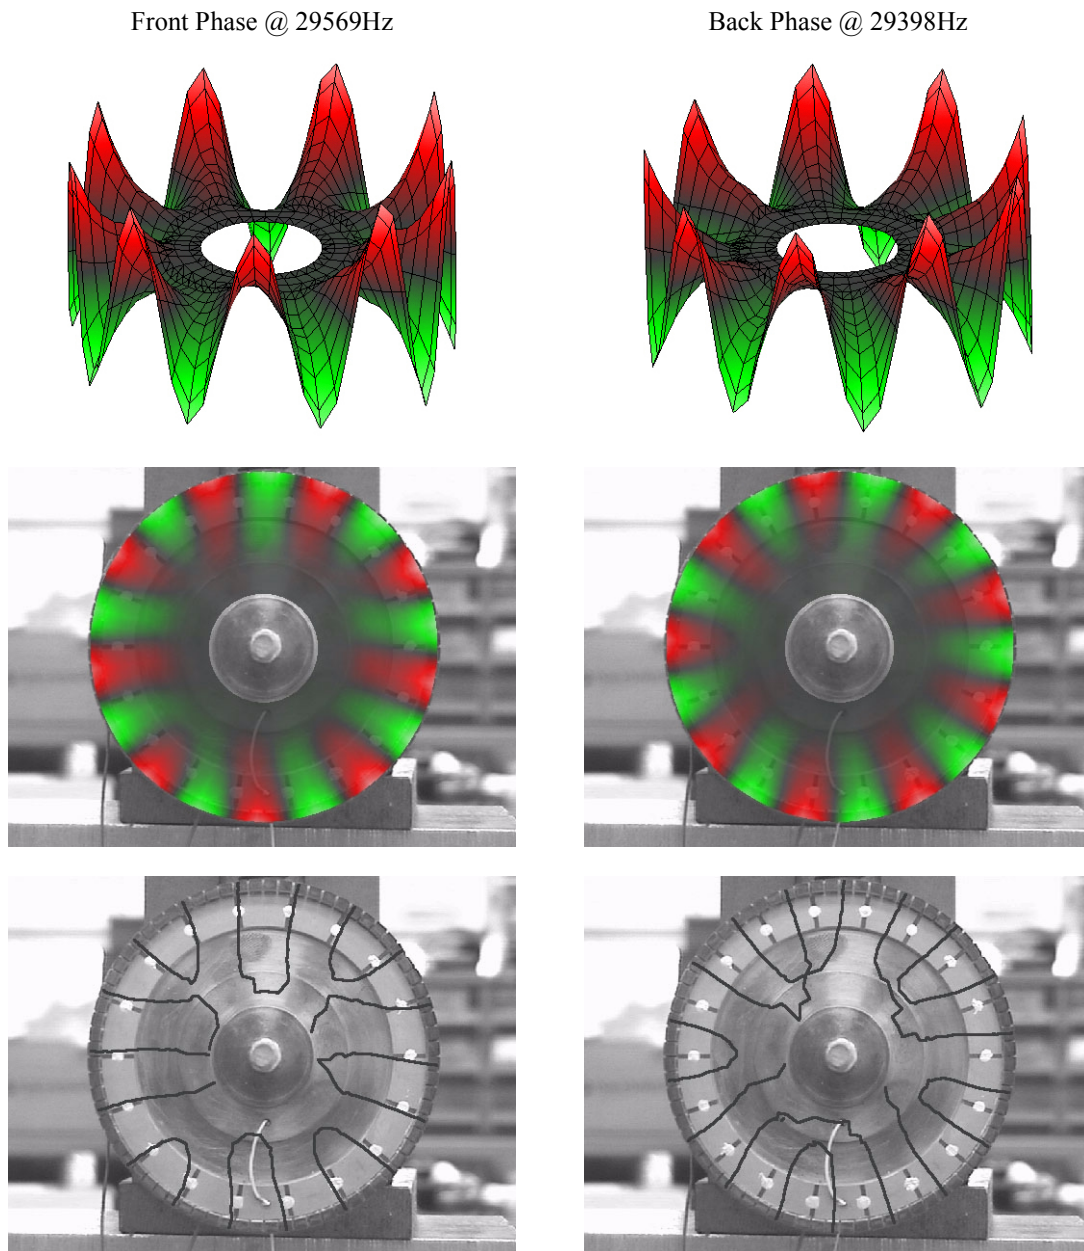
Figure 3.26 Laser vibrometer scans of the two-sided stator modes.

As already mentioned, the PSV is also configured to sweep the excitation frequency while it scans its target. Therefore, in addition to capturing spatial mode shapes, the PSV was used to sample the frequency behavior of each of the four stators. The results are shown in Figure 3.27. For each input, the frequency response is plotted for the vibration amplitude measured at an approximate peripheral antinode and node of the dominant $B_{0,9}$ mode.
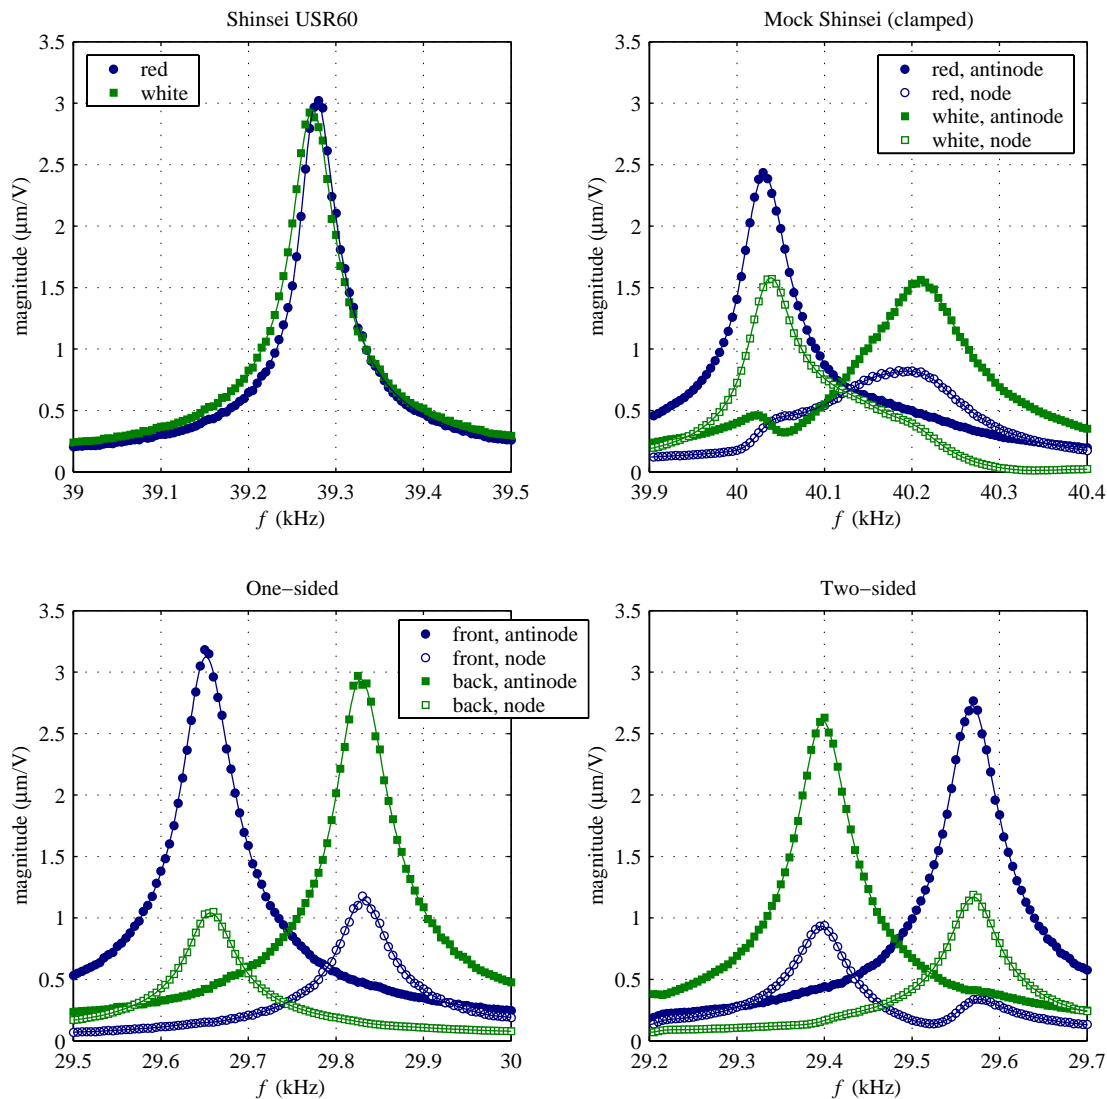
Figure 3.27 Resonant frequency behavior of the various stators as measured by the laser vibrometer.

Except for the USR60 where the incidental nodal responses have been excluded, the result-
ing four curves illustrate the coupling of each mode to each array. Note that both modal
responses are sometimes visible in a single trace because picking of the antinode and node
points was limited by the spacial discretization of the laser scans. The results generally fall
in line with the electrical admittance frequency response data, but surprisingly apparent is
how damped and spoiled the higher mock Shinsei stator mode is due to hub clamping but
not so for the lower mode. In retrospect, the admittance plots in Figure 3.20 do suggest that

the higher mode is more affected by the change in boundary condition. Though similarly plagued by mode splitting, the frequency responses of one-sided and two-sided stators are passably symmetric. They are also very much comparable to one another.

The mechanical quality factors of the stator modes can again be calculated from the vibration data but this time by the half-power bandwidth rule,

$$Q_m \approx \frac{f_r}{f_2 - f_1},$$

(3.4)

where f_1 and f_2 are the frequencies where the amplitude of vibration is 3dB lower than the peak amplitude at f_r. Since the half-power bandwidth approximation does not rely on locating the anti-resonant frequency, quality factor measurement is available for all resonant peaks. The calculations are summarized in Table 3.5 along with the resonant frequencies and mode rotations. Compared to the admittance study results, the quality factors of the symmetric stators are about 10% higher whereas those of the Shinsei USR60 are noticeably lower. The additional information on the mock Shinsei stator further illustrates the unbalanced response of its two modes.

3.2.3 Visual Inspection

Ironically, when visually compared to the in-house manufactured stators, the Shinsei USR60 stator does not appear to have been assembled with strict concern for geometric axisymmetry. Most notably, the single piezoceramic ring rests 5mil off-kilter from the substrate's concentric axis. There is also significant cured epoxy runout unevenly strewn about its outer periphery. In contrast, meticulous care was taken to insure that the piezo rings on each of the in-house stators were correctly positioned within about a mil of the concentric axis. The use of flash tape and teflon shims to immobilize the rings during cure saw to that. Furthermore, the epoxy was precisely apportioned to minimize the excess yet still guarantee that it flowed out in all directions; where possible, any epoxy runout was carefully scraped off post cure. The visual quality of the mock Shinsei stator is compared to that of the Shinsei stator in Figure 3.28. In-house assembly procedures would appear to have been more strin-

TABLE 3.5 Stator Modal Characteristics at 20°C from Vibration Measurements

Stator	Array	low mode		high mode		Mode Rotation
		f_r (kHz)	Q_m	f_r (kHz)	Q_m	
Shinsei USR60	red	n/a	n/a	**39.278**	**1099**	$\lambda/32$
	white	**39.273**	**912**	n/a	n/a	
Mock Shinsei (clamped)	red	**40.031**	**816**	40.194	280	$3\lambda/32$
	white	40.039	809	**40.210**	**417**	
One-sided	front	**29.651**	**559**	29.831	578	$\lambda/16$
	back	29.657	576	**29.828**	**561**	
Two-sided	front	29.397	555	**29.569**	**542**	$\lambda/16$
	back	**29.398**	**541**	29.570	557	

gently controlled than those of Shinsei. The result is that the three in-house stators are geo-metrically more axisymmetric than the Shinsei USR60 stator. Dimensionally, none of the stators exhibit noticeable variations in thickness or diameter beyond expected machining tolerances.

The last observable source of geometric asymmetry concerns the wire connections. After comparison with a second Shinsei USR60 stator that was on hand at the time, neither size nor placement of the solder joints appear to follow any specific pattern in the tested Shinsei USR60. True, the three connections are roughly spaced evenly about the circumference, but the angular spacing is slightly different between the two stators. In addition, the solder blobs are fairly consistent in size but seem larger than necessary. With no reason to do otherwise, the in-house stators were wired very much the same as the Shinsei USR60. Gage sizes are similar and connections have been distributed about the circumference.

For the sake of argument, there is no indication that the Shinsei stator has been altered in any way to correct for its obvious asymmetric imperfections in order to achieve degeneracy of the $B_{0,9}$ modes. This would suggest that modal degeneracy is not unduly sensitive to piezo concentricity or minor localized mass. One must conclude, therefore, that the source of asymmetry in the in-house stators is not geometric in origin. For future consideration, therefore, other causes of the observed mode splitting must be explored. Possible remaining
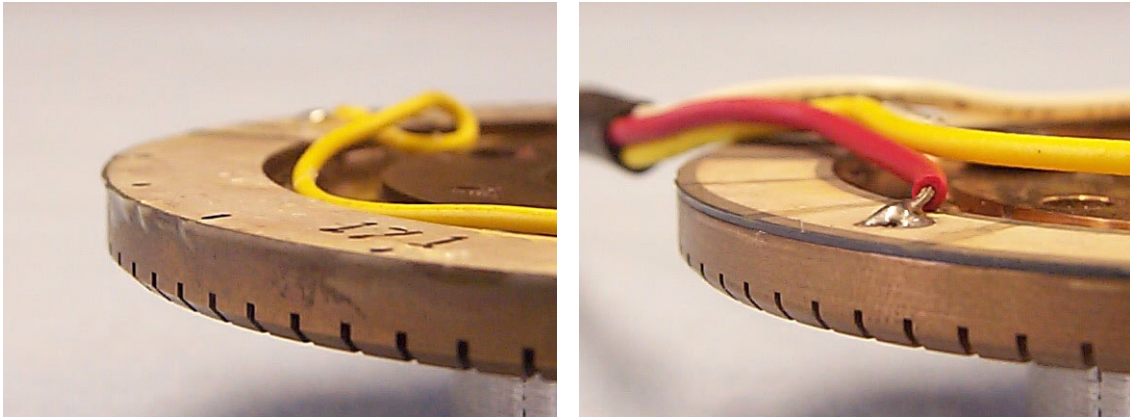
Figure 3.28 Visual comparison of piezo concentricity and epoxy runout between the Shinsei stator on the left and the mock Shinsei stator on the right.

culprits include anisotropic material properties and trapped air in the bond layers. For example, while measuring the resonance characteristics of brass and 6061 aluminum specimens in torsion, Lee *et al*. [Lee, 2000] observed mode splitting on the order of 1.6–2.0%. They were certain that geometric tolerances could only account for 0.1–0.2% splitting and concluded that anisotropy induced during the material forming process must be at fault. Since each in-house stator was machined from the same slug of phosphor bronze, this theory would explain why the degree of mode splitting has been consistent throughout the prototyping process. If so, it is possible that Shinsei subjects the substrate to an annealing process to relieve stress and thus eliminate material anisotropy. In contrast, it seems unlikely that bonding defects could result in such consistent mode splitting, though understandably, just a handful of samples doesn't rule out the possibility of the bond layer being at fault. Whatever the case may be, eliminating mode splitting is left for future study.

3.3 Materials Characterization

Supplying geometric dimensions to the model is straightforward, but many of the required material properties are either not readily accessible in textbooks or product literature or, especially in the case of the piezoceramic, are generally measured and published for moderate operating conditions. More often than not, the published data that is available does not account for the lossy behavior of the materials, piezoceramics being the exception. The pro-

posed model, however, is founded on the use of complex material properties to account for internal losses. Furthermore, the accuracy of the implementation depends on the validity of the supplied properties at the conditions of motor operation. Therefore, the various materials have been characterized experimentally at the appropriate conditions in order generate the necessary complex properties. The utilized measurement techniques are presented here along with the results.

3.3.1 Piezoceramics

Supplier specifications for the NEPEC N-61 and PZT-4 compositions are given in Table 3.6. Notice that although it is common for piezoceramic suppliers to publish the lossy elastic and dielectric constants, Q_m and $\tan\delta$, the real piezoelectric constant, d_{31}, is never accompanied by a lossy counterpart. In line with standard practices, these published values are the results of low-field, room-temperature measurements taken 24 hours after initial poling. Accurate modeling of a high-power actuator, however, necessitates that the piezoceramic properties be measured at the appropriate field level and operating temperature. Not only are piezoceramic properties highly dependent on field and temperature but also on the time history of extreme events involving temperature, stress, and electric field. Therefore, after events such as high-temperature poling, it is important to let test samples age for an appropriate amount of time before taking measurements at less severe operating conditions. Without getting into details here, let it be known that piezoceramic properties continue to stabilize after the accepted 24-hour aging period, and, as such, even the published values can be inappropriate for low-field, room-temperature applications. Accordingly, each of the samples in the following tests have been aged for at least a full week before testing, or retesting, to eliminate any effects of time.

As already mentioned, NEPEC N-61 samples were not obtained for characterization. Comparing the nominal properties of NEPEC N-61 and PZT-4 (Staveley) in Table 3.6, anyone familiar with the broad range of commercial lead-zirconate-titanate ceramics can see that the two compositions are strikingly similar. The only significant difference is in their

TABLE 3.6   Supplier Specifications of NEPEC N-61 and PZT-4 at Room Temperature, Low Field

Property	Symbol	Units	Tokin NEPEC N-61[a]	Staveley (EDO) PZT-4[b]	Morgan Matroc PZT-4[c]
Permittivity	$\varepsilon_{33}^{T}{}'/\varepsilon_0$		1400	1300	1300
Dielectric loss tangent	$\tan\delta$	%	0.3	0.4	0.4
Piezoelectric strain constant	d_{31}'	10^{-12} m/V	-132	-127	-122
Piezoelectric quality factor	Q_p		n/a	n/a	n/a
Mechanical compliance	$s_{11}^{E}{}'$	10^{-12} m^2/N	13.1	12.8	12.3
Mechanical quality factor	Q_m		1800	400	500
Poisson ratio	ν		0.31	0.31	0.31
Density	ρ	kg/m^3	7790	7500	7600
Curie temperature	T_c	°C	315	320	325

a. Tokin Corp., *Piezoelectric Ceramics, Vol. 1*.

b. EDO Corp., *Piezoelectric Ceramic Products*.

c. Morgan Matroc, Inc., *Guide to Modern Piezoelectric Ceramics*.

mechanical quality factors. This explains why the quality factor of the mock Shinsei stator is lower than that of the Shinsei USR60 stator. Also notice that the subtle differences in compliance and density likely account for the higher resonant frequency of the mock Shinsei stator. Nevertheless, the two materials are rather comparable overall and might be presumed to share similar dependencies on field and temperature. Similar temperature dependence of the two stators was already demonstrated in Table 3.4, particularly with respect to resonant frequency and capacitance (dielectric). Consequently, given the unfortunate situation of not being able to characterize NEPEC N-61, it was decided that scaled PZT-4 measurements would have to suffice.

Piezoceramic characterization has been done in two steps. In the first set of tests, a well-accepted resonance technique is adapted to measure the low-field dielectric, elastic, and piezoelectric constants as a function of temperature. In the second set of tests, direct measurements of current and strain are employed to extend the measurements of the dielectric and piezoelectric constants to high field. The elastic constant is assumed independent of

field. On a related note, the characteristic constants of lead-zirconate-titanate piezoceramics have been shown to be insensitive to frequency below several hundred kilohertz [Gonzalez, 1996]. Therefore, only when feasible have the tests been run at frequencies on the order of motor operation. Each of the PZT-4 samples were prepared from wafers purchased from Staveley Sensors, Inc., the same supplier from where the stator rings were obtained.

Low-field

The necessary set of complex constants has been measured at low field by curve-fitting the measured impedance of a sample about its fundamental resonance with the appropriate linear impedance equation derived in the IEEE Standard on Piezoelectricity [IEEE Std 176-1987]. The technique is similar to those developed by Alemany [Alemany, 1993] and Smits [Smits, 1985], except that the calculations are based off hundreds of measurements rather than the measurements at just a few characteristic frequencies. As derived in the IEEE standard from the linear piezoelectric equations, the complex electrical admittance of a bar poled and excited in the direction of its thickness is given by

$$ Y \ = \ G + iB \ = \ i\frac{2\pi flw}{t}\left(\varepsilon_{33}^{T} - \frac{d_{31}^{2}}{s_{11}^{E}}\right) + i\frac{2wd_{31}^{2}}{ts_{11}^{E}\sqrt{\rho s_{11}^{E}}}\tan\left(\pi fl\sqrt{\rho s_{11}^{E}}\right). \tag{3.5} $$

The density, ρ, length, l, width, w, thickness, t, and drive frequency, f, are specified real quantities. The unknown parameters include the dielectric permittivity, ε_{33}^{T}, the piezoelectric coefficient, d_{31}, and the elastic compliance, s_{11}^{E}, each defined in complex form by (2.7). This particular configuration has been chosen because of the convenience of measuring all three necessary piezoelectric constants from a single set of admittance measurements.

Conforming to the aspect ratio requirements of the IEEE standard, a bar of PZT-4 was cut with dimensions $l = 1.130$", $w = 0.139$", and $t = 0.0115$" for excitation in the chosen length-thickness mode. Starting at room temperature, the complex admittance was measured about the sample's fundamental resonance using the HP 4194A Impedance Analyzer at an excitation level of $0.010V_{rms}$. The measurement was repeated at successive temperatures up to 100°C, taking care to stabilize the sample at each temperature for at least five minutes.
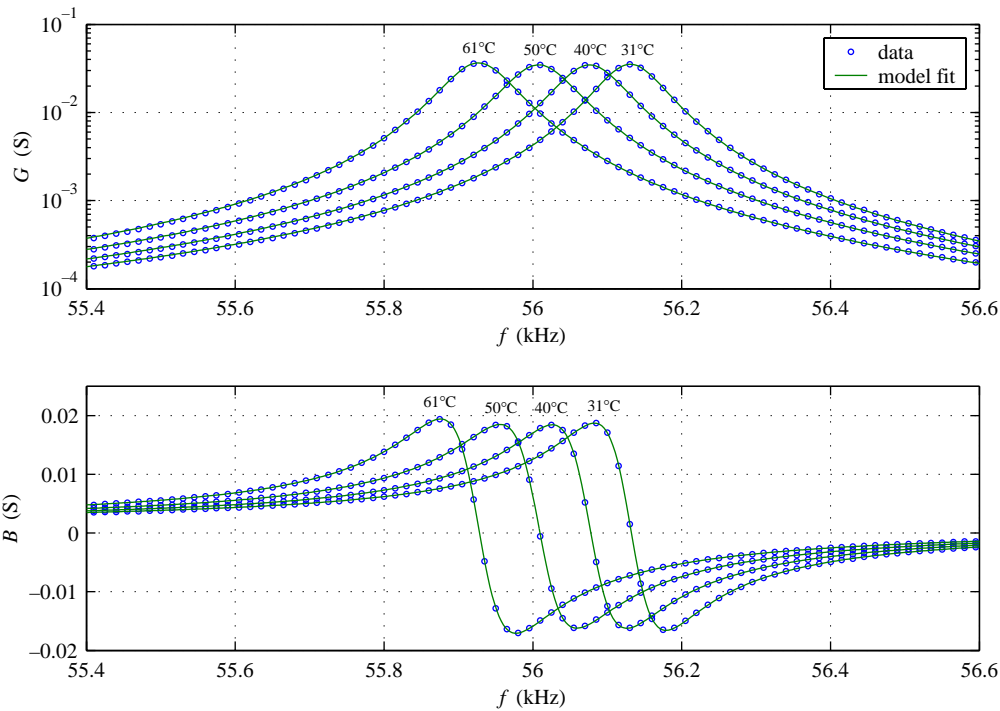
Figure 3.29   Admittance at the fundamental resonance of a PZT-4 sample in the length-thickness mode. Model fits are compared to experimental measurements at several temperatures.

Applying a nonlinear least-squares method to a log-scale cost function, (3.5) was fit to the experimental data at each temperature. The published density of $7500\mathrm{kg/m^3}$ was assumed in the calculations. As seen in Figure 3.29 for selected temperatures, the curve fits are exceptionally good, thus validating that (3.5) has the necessary degrees of freedom to model accurately the admittance of such a sample. The resulting low-field dielectric, piezoelectric, and elastic constants are plotted in Figure 3.30 as a function of temperature. The dependence of each constant on temperature can readily be seen.

Observe that although the compliance is typical, the dielectric and piezoelectric constants are lower than those published in Table 3.6. Some of the difference is due to expected variations in composition, but the shorter aging period of the published values also plays a role. As mentioned earlier, the electromechanical characteristics of piezoceramics, particularly the dielectric and piezoelectric constants, continue to stabilize well after the initial 24-hour aging period.
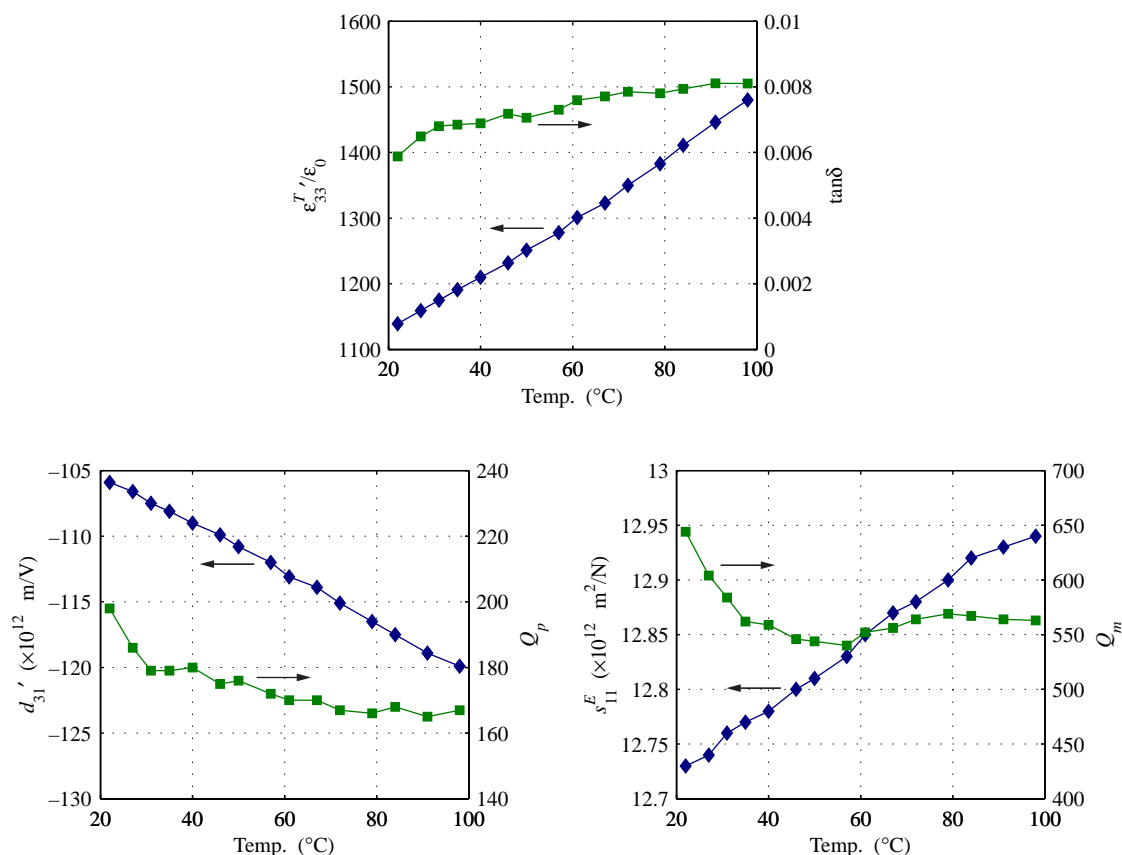
Figure 3.30 Low-field PZT-4 constants vs. temperature as measured by curve-fitting the complex admittance of a sample in length-thickness mode.

High-field

Unfortunately, the resonance technique breaks down at greater field levels as material and displacement amplitude nonlinearities become more apparent [Zelenka, 1986]. Therefore, the high-field complex dielectric and piezoelectric properties have been obtained off resonance, i.e. at quasistatic frequencies, using more direct methods. For these tests, the HP 4194A was used in gain-phase mode to enable the use of an external amplifier and the comparative measurement of two arbitrary signals. Specifically, in gain-phase mode, the HP 4194A outputs a low-voltage excitation signal and records the complex relationship between a reference input, V_{ref}, and a test input, V_{test}. The output can be connected to an external amplifier, and any two measurements can be fed to the inputs as long as the signals have been properly attenuated. Therefore, similar to impedance mode but at higher voltage,

the dielectric constant has been measured by feeding voltage and current sensors to the reference and test inputs, respectively. To measure the piezoelectric constant, the current sensor has been replaced by a strain gage and conditioner. Interpretation of the V_{test}/V_{ref} data, consequently, requires calibration of these external sensors. The two setups are described as follows.

The same PZT-4 sample tested at low field was again used for the high-field dielectric test to provide some continuity between the results. A Yorkville AudioPro 3400 stereo amplifier was used to amplify the low-voltage output of the HP 4194A and drive the 11.5-mil thick sample at voltages up to $130V_p$. The voltage then had to be scaled down by a resistor-bridge voltage divider before being monitored by the reference channel. A small resistor placed in series with the sample provided a current monitor, the differential voltage of which was fed directly to the test channel. Given that the sample's first dynamic resonance was measured around 56kHz, a test frequency of 10kHz was chosen as a trade-off between maximizing the current monitor's signal-to-noise ratio and avoiding dynamic excitation. Gain-phase measurements of V_{test}/V_{ref} were taken as a function of voltage and temperature by sweeping the HP 4194A's oscillation amplitude and gradually stepping the frequency up from room temperature to 100°C.

Given this setup, the applied electric field was related to the reference voltage by

$$E_{33} = \frac{Z_1 + Z_2}{Z_1} \cdot \frac{V_{ref}}{t}, \qquad (3.6)$$

where Z_1 and Z_2 are the complex impedance values of the resistors in the voltage divider. Similarly, the charge density was related to the measured test voltage by

$$D_{33} = \frac{1}{i\omega A} \cdot \frac{V_{test}}{Z_c}, \qquad (3.7)$$

where $A = lw$ is the electrode's surface area, $\omega = 2\pi f$ is the test frequency, and Z_c is the complex impedance of the current monitor resistor. Each of the resistors' impedances were
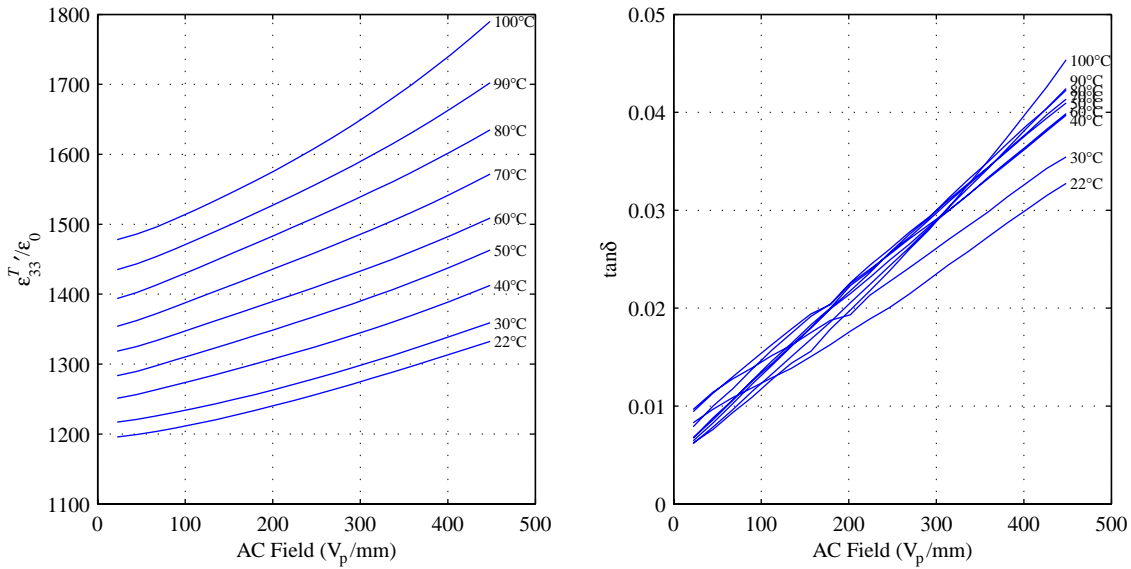
Figure 3.31 High-field dielectric permittivity and loss tangent of aged PZT-4 at 10kHz and elevated temperatures.

measured at the test frequency. Dividing (3.7) by (3.6), the high-field, complex dielectric permittivity was calculated from the gain-phase data by

$$\varepsilon_{33}^{T} = \frac{D_{33}}{E_{33}} = \frac{1}{i\omega A} \cdot \frac{tZ_1}{Z_1 + Z_2} \cdot \frac{V_{test}}{V_{ref}}. \tag{3.8}$$

The results of this test are shown in Figure 3.31. As hoped, the high-field measurements asymptote at low-field to the values plotted in Figure 3.30, thus supporting the validity of the high-field measurements. Though the dielectric loss tangent has been shown to be relatively insensitive to changes in temperature, it is seen here that dielectric loss does, however, dramatically depend on field level.

The high-field piezoelectric constant was measured using a similar procedure but with a few modifications. First, a sample with a larger surface area was necessary to accommodate the bonding of a small strain gage away from edges. For this purpose, a larger PZT-4 sample roughly one-inch square was prepared from a wafer with thickness $t = 0.0205$". The gage was bonded to the grounded electrode with M-Bond cyanoacrylate. A voltage proportional to the strain was fed from a Vishay strain gage conditioner to the test channel on the HP

4194A. An identical dummy sample was also prepared so that the strain gage conditioner could be set up in temperature-stabilized half-bridge mode. Unfortunately, the internal dynamics of the conditioner limited the frequency of this test to 100Hz. To accommodate the lower frequency and thicker sample, the AudioPro was replaced by a Kepco amplifier capable of higher voltages down to DC. Nevertheless, the voltage was again divided by the same resistor bridge and monitored by the reference channel. In order to preserve the integrity of the strain gage bond, the maximum temperature in this test was limited to 60°C.

Per the choice of gage and calibration settings of the strain gage conditioner, the strain was related to the measured test voltage by

$$S_{11} = 1000\mu\varepsilon \cdot \frac{2.0}{F_g V_c} \cdot V_{test}, \tag{3.9}$$

where F_g is the gage factor and V_c is the calibration voltage. Dividing (3.9) by (3.6), the high-field, complex piezoelectric constant was calculated from the gain-phase data by

$$d_{31} = \frac{S_{11}}{E_{33}} = \frac{2\times10^{-3}}{F_g V_c} \cdot \frac{t Z_1}{Z_1 + Z_2} \cdot \frac{V_{test}}{V_{ref}}, \tag{3.10}$$

the results of which are shown in Figure 3.32. Again notice that the high-field measurements asymptote at low-field to the values plotted in Figure 3.30. The scatter in the piezoelectric quality factor data is due to the low signal-to-noise ratio of the strain gage conditioner at these low strains.

The results of the high-field piezoceramic tests demonstrate the importance of measuring the material properties at the appropriate conditions when modeling a high-power actuator. To put things in perspective, each of the motors have been tested nominally at $150V_p$, or $300V_p$/mm, and as high as $200V_p$, or $400V_p$/mm. Through efficient air cooling, the internal temperatures during continuous operation were limited to about 30°C, though without it, temperatures can quickly exceed 50–60°. As one can see, the dielectric and piezoelectric properties of this particular piezoceramic are dramatically different at these conditions than at the moderate conditions usually published.
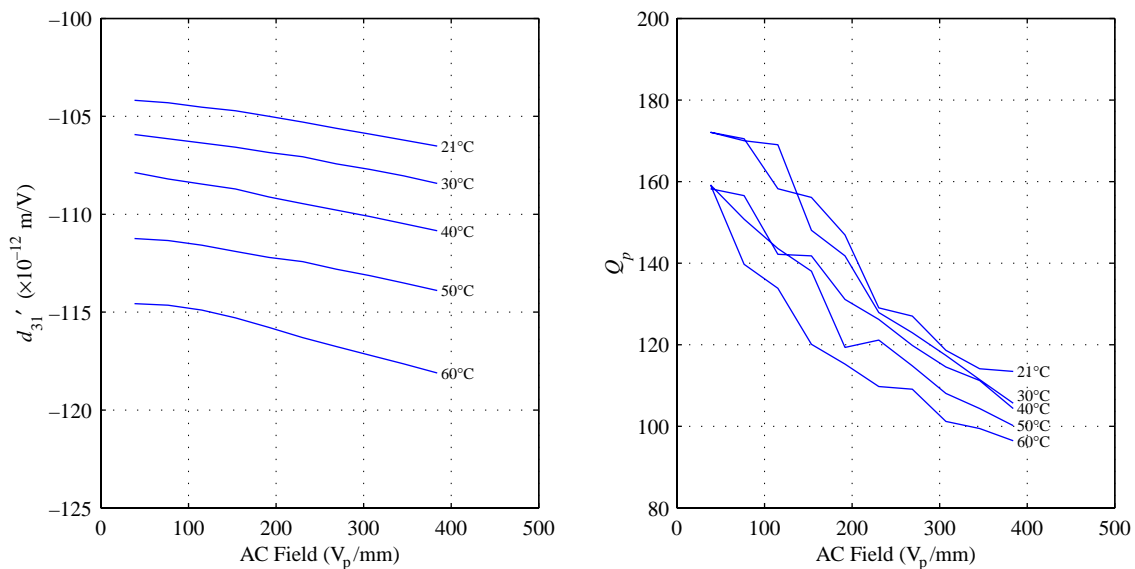
Figure 3.32 High-field piezoelectric constant and quality factor of aged PZT-4 at 100Hz and elevated temperatures.

3.3.2 Metals

Commonly accepted mechanical properties for the metal alloys used in the stators and rotors are listed in Table 3.7 [Clauser, 1967]. As was done with the piezoceramic, measurements could be made to characterize the metals fully as a function of temperature, but since the model-predicted resonant frequencies will require scaling anyway for the sake of direct model-experiment correlation, these absolute values have been used in the model simulations. Unfortunately, typical listings for metals do not include mechanical quality factors, for metals aren't generally regarded as lossy. At the very least, the internal damping of metals is considered negligible compared to other sources of loss in a composite or jointed structure, so most engineers are content with a real modulus. True, the metal substrates in a piezoelectric motor don't contribute as much mechanical damping as, say, the viscoelastic polymer, but still they account for some of the total energy loss and should be characterized for consistency.

The limited references that can be found in the literature regarding metal quality factor measurements mostly involve testing metal resonators near absolute zero for use in gravity-

TABLE 3.7   Published Metal Properties

Property	Symbol	Units	544 Phosphor Bronze	6061 Aluminum
Young's modulus	E	GPa	103.5	69
Density	ρ	kg/m^3	8893	2713
Poisson's ratio	μ		0.33	0.33

wave detection [Aguiar, 1996; de Waard, 2000; Duffy, 1992; Duffy, 1999]. At such cold temperatures, quality factors can far exceed 10^6–10^7, but those numbers have no correlation with the damping of metals near room temperature. Coincidentally, both phosphor bronze and aluminum were among the metals considered in those studies. A mechanical quality factor of 3000 is given for phosphor bronze in an ultrasonic motor book [Sashida, 1993], but without supporting references or statements regarding the conditions under which the number was derived, its accuracy is questioned. Therefore, experimental tests have been carried out to measure the quality factors of both the 544 phosphor bronze and 6061 aluminum alloys.

Based on the simplicity of the stator admittance tests, it was decided to measure the modal quality factors of metal disks via nonintrusive electromechanical coupling. The test pieces were approximately 3" in diameter by 0.2" thick and were supported at their centers by #2-56 threaded rod. As seen in Figure 3.33, a pair of small piezoceramic wafers bonded close to the center of each disk provided weak electromechanical coupling to the vibration modes,
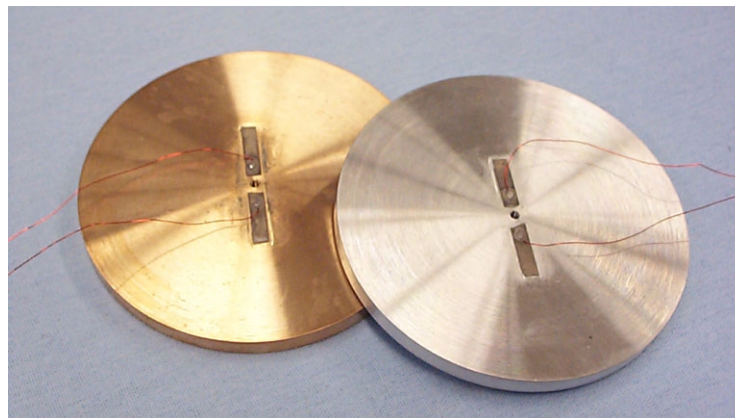
Figure 3.33   Metal test disks.

similar to what was done by Aguiar [Aguiar, 1996]. The central placement of the piezo actuators was necessary to minimize the integration of the piezoceramic properties in the measurements. As evidence of the technique's potential, Aguiar successfully measured quality factors as high as 10^7 for a vanadium disk at cryogenic temperatures. The attachment of the more lossy wafers, therefore, is considered to have negligible influence on the measurement. Prior to bonding, actual densities were calculated from the specimen masses: 8964kg/m^3 for the phosphor bronze and 2720kg/m^3 for the aluminum. Both are in reasonable agreement with the values in Table 3.7.

Admittance traces of various flexural modes were captured using the HP 4194A impedance analyzer, then quality factors were calculated for each mode using (3.2). As the results in Table 3.8 show, there is significant variation in the quality factor measurements from mode to mode. However, the difference in damping between phosphor bronze and aluminum is apparent. For the purpose of supplying the model with usable values, the results roughly translate to $Q_m \sim 3500$ for 544 phosphor bronze and $Q_m \sim 2000$ for 6061 aluminum. Fortunately, Q_m is an inverse measure of damping, so measurement errors for high-Q materials like metals do not greatly affect the prediction of losses in systems with low-Q materials like polymers. Therefore, these approximate values are considered sufficient. Besides, the single book value of 3000 for phosphor bronze provides modest validation to the measurements.

Unexpectedly, these tests offer circumstantial evidence that the perpetual mode splitting of the in-house stators was indeed the unfortunate result of anisotropy. As seen in Table 3.8, a significant degree of mode splitting was observed in several of the phosphor bronze disk modes, yet only very slight mode splitting was observed in the aluminum disk modes. One might want to attribute the asymmetry to the piezo wafers, but as both disks were configured identically, any mode splitting due to the piezos should be equally evident. Obviously, this isn't the case. Furthermore, both disks were similarly toleranced, so the degree of splitting observed in the phosphor bronze disk cannot be ascribed to deviations from perfect geometry. This leaves anisotropy of the phosphor bronze material properties as the only logical

TABLE 3.8 Modal Characteristics of Metal Test Disks at 20°C

Mode	Phosphor Bronze Disk ($\varnothing 2.999\pm.001" \times 0.1995\pm.001"$)		Aluminum Disk ($\varnothing 2.987\pm.001" \times 0.1995\pm.001"$)	
	f_r (kHz)	Q_m	f_r (kHz)	Q_m
$B_{0,2}$	2.869 3.092	3996 4236	4.481 4.482	>1200
$B_{0,4}$	11.551	3644	17.435 17.438	2395 2103
$B_{0,5}$	17.271 17.458	5230 5175	n/a	n/a
$B_{1,2}$	18.416	3292	27.981	1397
$B_{1,3}$	n/a	n/a	40.262 40.271	3048 2553
$B_{1,4}$	36.013 36.166	3110 2242	53.359	1925

conclusion. Conversely, the negligible mode splitting of the aluminum disk attests to the isotropy of its properties and lends support to the assumption that mode splitting is not an issue with the rotors.

Several modes of each metal disk were also characterized at elevated temperatures. The results are plotted in Figure 3.34. Although the modal frequencies declined predictably with increasing temperature, no noticeable trend was observed amid the scatter of quality factor measurements, either from mode to mode or for the same mode. Nevertheless, in addition to providing approximate values of the quality factors, the experiment just at room temperature was invaluable in identifying the source of stator asymmetry.

3.3.3 Polymer

Of the materials used in the motors, the Ekonol/PTFE polymer blend is the least documented, especially with regards to its mechanical properties. Though PTFE by itself is well characterized, Ekonol and its blends have relatively low profiles in the literature. Table 3.9 summarizes the pertinent, yet limited, information that is available from Sumitomo Chemical Co., a thermoplastic processor and supplier of the 20% blend sample used in the experiments, and Norton Performance Plastics Corp., the producer of the Ekonol resin. As usual,
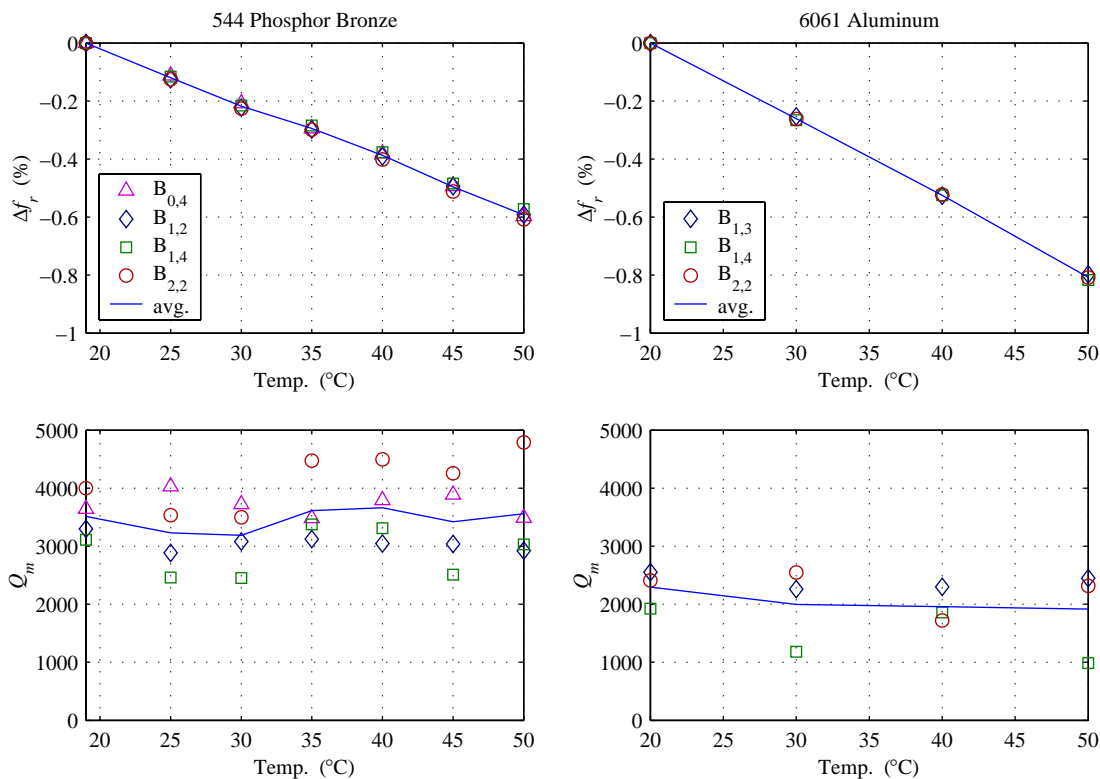
Figure 3.34 Modal characteristics of the metal test disks as a function of temperature.

the sources provide only nominal values for the elastic moduli, giving no indication of their dependence on temperature or frequency. There is also no mention of viscoelastic behavior by way of either a storage modulus or loss tangent. As a significant loss mechanism in an ultrasonic motor is attributable to the viscoelastic deformation of the interface polymer, characterizing the polymer's viscoelastic properties is of utmost importance.

Complicating the issue, the glass transition temperature of Ekonol/PTFE is very close to room temperature, so its viscoelastic properties are especially sensitive to temperature in the range of typical motor operation. And unlike with metals and piezoceramics, the viscoelastic properties of polymers are also highly dependent on frequency. As such, it was deemed necessary to characterize experimentally the viscoelastic behavior of the Ekonol/PTFE polymer as a function of both temperature and frequency. Furthermore, as motor torque and performance are greatly determined by the coefficient of friction, it was important to identify this value specifically for Ekonol/PTFE sliding on phosphor bronze.

TABLE 3.9 Manufacturer Data for Ekonol, PTFE, and Various Blends.

Property	Sym.	Units	PTFE (Teflon)[b]	20% Ekonol/ PTFE[a]	25% Ekonol/ PTFE[b]	30% Ekonol/ PTFE[a]	Ekonol[b]
Elastic modulus	E	GPa	0.56	0.88	0.96	0.98	6.9
Density	ρ	kg/m³	2170	1950	1860	1840	1440
Friction coef.	μ		<0.10	0.18[c]	n/a	0.20[d] 0.21[c]	0.37[e]

a. Sumitomo Chemical Co., *Oil-free Sliding Material: Sumikasuper*, product literature

b. Norton Performance Plastics Corp., website: http://www.ekonol.com

c. counterpart: 5052 aluminum

d. counterpart: 304 stainless steel

e. counterpart: 1040 carbon steel

Coincidentally, the dynamic modulus and friction coefficient of the 20% Ekonol/PTFE blend, along with several other polymers, were measured by Tada *et al.* [Tada, 1991; Ueha, 1993] in a study correlating the performance of an ultrasonic motor with the properties of the interface polymer. In their efforts, they characterized 20% Ekonol/PTFE as having $E' = 0.67\text{GPa}$ and $\tan\delta = 0.089$ ($E'' = 0.06\text{GPa}$) at 50kHz and 80°C and a coefficient of friction of $\mu = 0.24$ when sliding on phosphor bronze. Unfortunately, the measurement temperature is beyond the desired range, and the coefficient of friction appears unquestionably high given the stall torque of the Shinsei USR60 used in the study. Nevertheless, these values provide excellent reference for the measurements to follow.

Viscoelastic Moduli

It is rare that the dynamic properties of a polymer are actually measured at high frequency. Instead, material scientists often rely on an empirically developed technique known as the method of reduced variables [Ferry, 1961] to extrapolate measurements taken over a more practical range of frequencies at several temperatures to a larger effective frequency scale at one "reduced" temperature. The technique has been in practice for several decades and is widely used for obtaining measurements at frequencies outside the operational limits of available measurement equipment. Though originally founded on empirical success, the method has since been supported by numerous theories which reduce time and temperature

to a single independent variable. Thus, the principle upon which the method is based is often referred to as the time-temperature superposition principle. This method is likely what Tada *et al.* used to achieve measurements at 50kHz.

In essence, a series of experimental measurements are taken at several different temperatures over a frequency range usually spanning just a few decades. The measurements at each temperature, T, are first scaled by T/T_0, where T_0 is the temperature of the reference curve to which the other measurements are being reduced. Then, the scaled curves are independently shifted in frequency until they collectively superpose with the reference curve to form one continuous composite curve. It is quite common for the reduced measurement set to span more than ten decades in frequency. As one might expect, however, there are criteria for applying the method of reduced variables. These include: (1) the shapes of adjacent curves must exactly match when superposed, and (2) each viscoelastic function must superpose using the same sequence of frequency shifts. Otherwise, the theory does not support the reduction of measurements in this manner.

As no equipment was readily available for the dynamic testing of polymers at tens of kilohertz, the method of reduced variables has been utilized to extrapolate the viscoelastic moduli of the 20% Ekonol/PTFE blend at the frequencies of interest. Starting with the measurement of the dynamic shear modulus ($G = G' + iG''$), a one-inch diameter sample was tested in a Rheometrics ARES rheometer configured to operate in the parallel-plate oscillatory shear mode. Measurements were performed at five different temperatures (40, 35, 32.5, 30, and 25°C) over a frequency range of 0.1 to 100rad/sec (0.016 to 16Hz). The results plotted in Figure 3.35 show the storage modulus, G', and loss modulus, G'', as functions of shifted frequency, fa_T, reduced to 40°C. Per the second criteria for using the method, simultaneous superposition of each composite curve was achieved with the same sequence of frequency shift factors. These were $\log a_T = 0$, 0.90, 2.40, 3.75, and 6.65 for 40, 35, 32.5, 30, and 25°C, respectively.
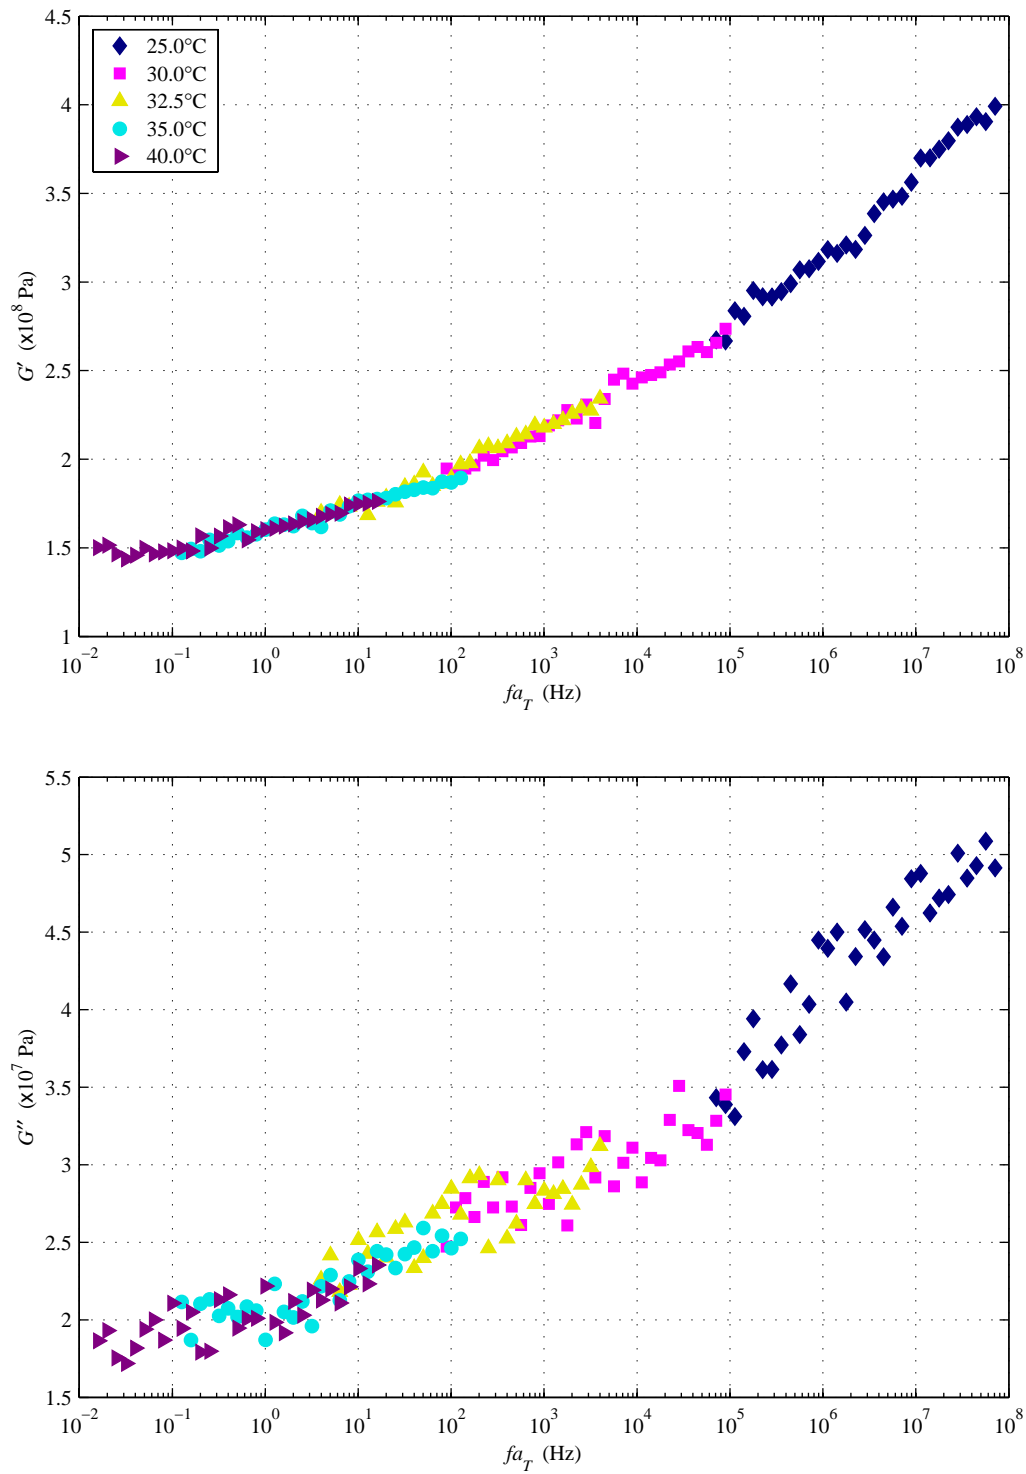
Figure 3.35 Composite curves for the shear storage modulus, G', and the shear loss modulus, G'', of 20% Ekonol/PTFE blend, reduced to 40°C over ten decades in frequency by the method of reduced variables (time-temperature superposition principle).

As is typical for a polymer, the results show that 20% Ekonol/PTFE stiffens with frequency. In fact, the dynamic shear modulus in the range of motor operation is approximately 70% higher than the asymptotic static value. However, the loss tangent in shear, $\tan\delta = G''/G'$, is relatively constant across frequency at 0.13. At 40kHz and 40°C, the dynamic shear modulus is given by $G' = 0.26$GPa and $G'' = 0.033$GPa.

The dynamic elongation modulus ($E = E' + iE''$) was then characterized with a Seiko dynamic mechanical analyzer (DMA) configured to operate in tension. The effective dimensions of the test sample were 22mm x 7mm x 0.3mm. To better identify the temperature dependence of the Ekonol/PTFE polymer, these measurements were instead performed at three frequencies (20, 50, and 100Hz) while sweeping temperature from –40 to 100°C. The results are plotted in Figure 3.36 for the storage modulus, E', loss modulus, E'', and loss tangent, $\tan\delta = E''/E'$, as functions of temperature. Readily apparent in each of the plots is the glass transition temperature around 20°C, at which point the properties change dramatically as the polymer transitions from a stiffer, less lossy, glass-like phase to a softer, more lossy, rubber-like phase. Ekonol/PTFE, therefore, exists in the latter phase during typical motor operation.

Though not a direct application of the method of reduced variables, the elongation modulus data can be extrapolated to other frequencies by the knowledge gained from the shear modulus reduction. As the second criterion states, each viscoelastic function must superpose using the same sequence of frequency shifts. In other words, for the principle of time-temperature superposition to hold true, all moduli of a single material must form logical composite curves using the same frequency-temperature relationship, $a_T(T)$. Hence, the elongation modulus data can be extrapolated to the desired frequency by reverse utilizing the frequency shifts determined from the shear modulus reduction. By interpolating a_T as a function of temperature, an appropriate measurement temperature can be deduced that is equivalent to the desired shift in frequency. Backtracking in this manner indicates that a reduced data point at 40kHz and 40°C could translate to any of 20Hz at 30.8°C, 50Hz at 31.6°C, or 100Hz at 32.1°C, all of which are fairly indistinguishable in Figure 3.36. The
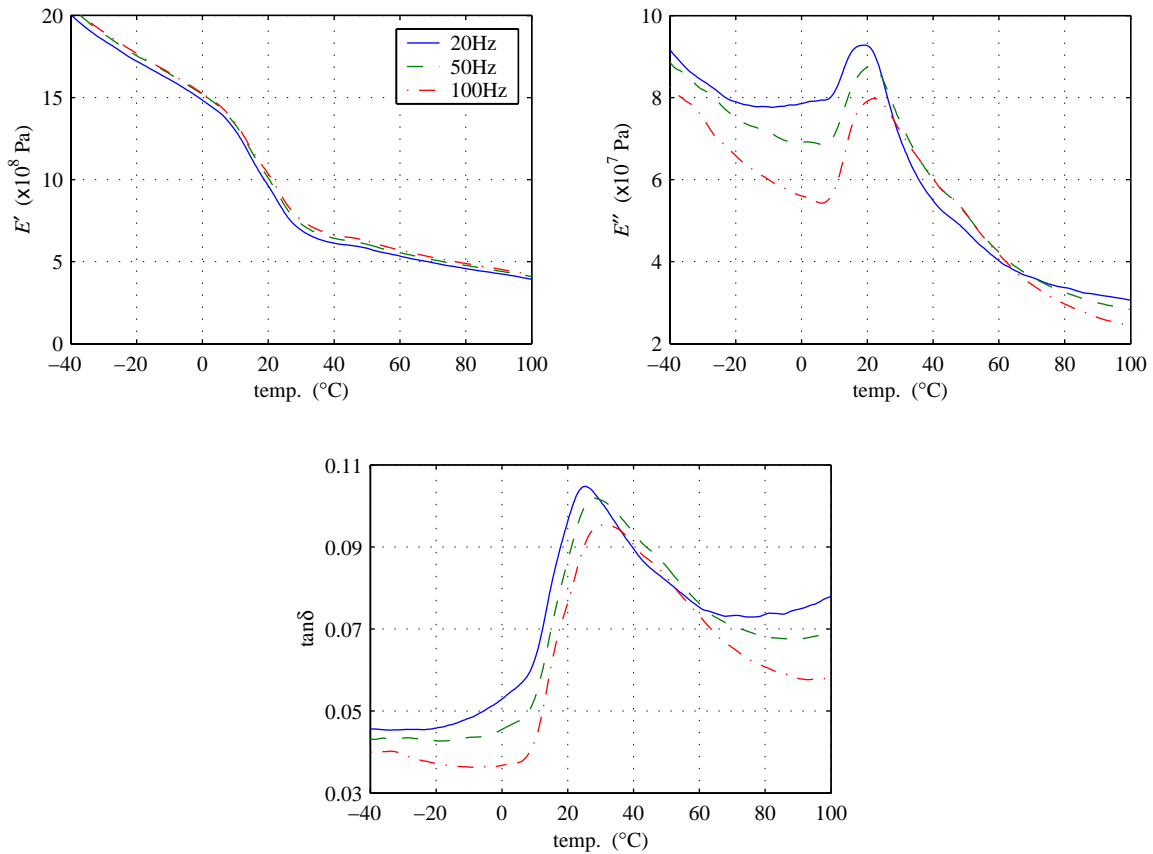
Figure 3.36   Dynamic elongation modulus of 20% Ekonol/PTFE blend as a function of temperature. Separate plots are given for the storage modulus, E', loss modulus, E'', and loss tangent, $\tan\delta$.

dynamic elongation modulus at 40kHz and 40°C is consequently given by $E' = 0.7\text{GPa}$, $E'' = 0.07\text{GPa}$, and $\tan\delta = 0.10$. These values compare rather well with the those measured by Tada *et al*. when the difference in temperature is taken into account.

Friction Coefficient

The friction coefficient, which is sensitive to so many factors including surface preparation and humidity, is rarely published because it depends on the counterface material. In other words, it's not an inherent property of any single material but rather a description of the interaction between two materials. Therefore, published friction coefficients are generally listed as ranges to give approximate indications of what to expect. When an accurate value is needed, the friction coefficient must be measured specifically for a given application tak-

ing care to observe the appropriate surface and environmental conditions. With this in mind, it was decided to measure the coefficient of friction of Ekonol/PTFE on phosphor bronze *in situ* by applying torque to an unpowered motor and measuring the torque at which, once turning, the rotor stops.

The friction tests were carried out using the mock Shinsei stator and both the nominal Shinsei rotor (with unknown Ekonol/PTFE blend) and 0.115" lip rotor (with 20% Ekonol/PTFE blend). In each case, a dozen torque measurements were made in each direction at the full 160N preload, and the coefficient of friction was calculated from

$$\mu = \frac{\tau}{f_{preload} \cdot r},$$ (3.11)

where r is the average contact radius. Within 95% statistical certainty, the friction tests yielded $\mu = 0.152 \pm 0.012$ for the unknown blend, and $\mu = 0.138 \pm 0.011$ for the 20% blend. Without a doubt, the sample that Sumitomo Chemical supplied for this research was not the same blend as used on the Shinsei rotor. Since the difference in the measured friction coefficients is in line with the difference shown in Table 3.9 for the 20% and 30% blends, it can be inferred that Shinsei employed the 30% blend. Though the higher-friction blend would be the preferable option for producing higher stall torque, relative performance is more important to model-experiment correlation than is absolute performance, and either blend performs consistently in this application.

3.4 Motor Testing

Bench-top performance testing of the traveling-wave motors involved three subsystems of supporting equipment: the drive electronics, the performance measurement apparatus, and some basic forced-air cooling paraphernalia. In order to accommodate the testing of several motors over a wide range of drive conditions, the drive system was assembled from off-the-shelf components without concern for compactness or efficiency. In other words, only the metrics of the motors themselves were of interest in this work, not how the power was being

generated. Consequently, the performance measurements were focused on the electrical power into and mechanical power out of each motor, from which the actuator efficiency could be calculated. While combatting overheating, the implementation of forced-air cooling made it possible to run the motors continuously for improved data consistency and reproducibility. It also greatly facilitated the necessary wear-in period of each newly assembled motor. Most importantly, though, the ability to control and maintain temperature was necessary for accurate model-experiment correlation. Each of these support systems are described here.

3.4.1 Drive Electronics

Starting with the former, a schematic of the drive electronics is illustrated in Figure 3.37. Quadrature sinusoids were generated by a pair of Wavetek function generators acting in tandem as master and slave. The master generator provided digital control of the drive frequency of both waveforms, whereas the slave generator was configured in phase-lock mode and triggered by the master to provide analog control of the phase difference between the two waveforms. These two signals were then amplified by a Yorkville AudioPro 3400 high-power stereo amplifier followed a pair of step-up transformers. The AudioPro was specifically chosen for its uncommonly high 80-kHz bandwidth and for its excellent four-quadrant capability of driving reactive loads, e.g. capacitive piezoceramics, near its specified limits and without noticeable distortion. This latter feature is rare among stereo amplifiers as they are usually only intended to drive resistive loads for which current is in phase with the voltage. The transformers, on the other hand, were simply scavenged from surplus switching power supplies. Though they were not being used in their ideal operational frequency range, their outputs were nonetheless sinusoidal. A fan was installed above the transformers to minimize heat saturation of their cores.

The voltage signals running to the respective piezo arrays were probed at v_1 and v_2 in Figure 3.37 and monitored on a digital oscilloscope. Before any given test run, the voltage amplitudes were coarsely set with the gain knobs on the stereo amplifier. Precise setting of
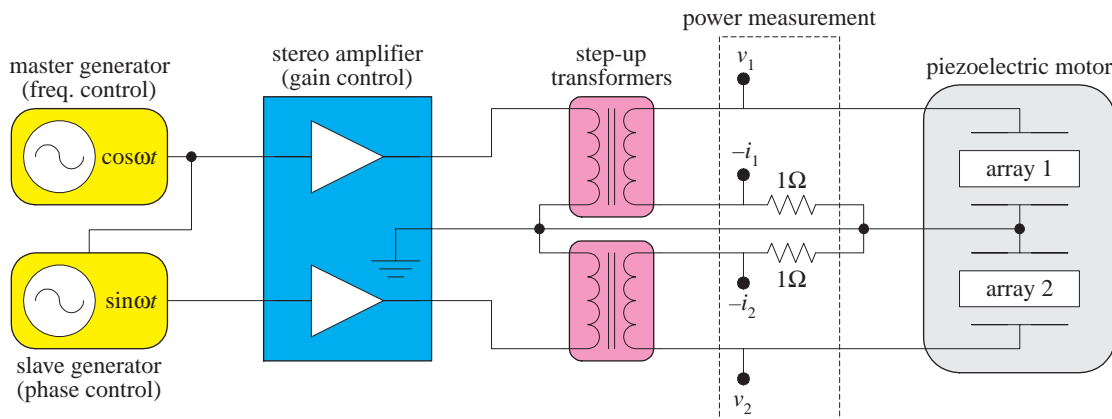
Figure 3.37   Schematic of the drive electronics.

the amplitudes was instead obtained by adjusting the generator outputs. When the motor warms up or its torque load is increased, its electrical impedance changes as seen by the amplification stage. Left unchecked in this open-loop drive scheme, the result could be seen as slight voltage fluctuations during testing. Therefore, fine adjustments were made on the fly to guarantee constant drive amplitudes amid the changing load. The phase between the two high-voltage signals was similarly affected, as witnessed by the tilting of the elliptical X-Y trace of v_1 and v_2. To combat this, the phase offset of the slave generator was continuously adjusted so as to maintain a circular trace ($\pm90°$) at all times.

3.4.2 Performance Measurements

Characterizing motor performance required measurements of rotational speed and electrical input power as functions of the applied torque. Torque also required direct measurement as it was generated open-loop by a braking motor. Mechanical output power was calculated after the fact from speed and torque, as was efficiency from the two powers. Together, these five measurements were used to define the performance of each motor. As the determination of input power was tied closely with the drive electronics, it will lead this discussion.

Power Measurement

Analog multiplier chips were utilized to generate a real-time measurement of instantaneous electrical power from measurements of instantaneous current and voltage. As seen in the

Figure 3.37, the currents, i_1 and i_2, running through the two piezo arrays were read off precision 1Ω resistors in the ground paths between the transformers and the motor. In actuality, these "resistors" were each comprised of fifteen carefully selected 15Ω resistors in parallel, further shunted by appropriate capacitors to cancel the effects of internal inductance. With the assistance of the HP 4194A impedance analyzer, the impedance values of these composite resistors were finely tailored to within 0.1% of 1Ω at 30kHz. Thus, current could accurately be read in amps directly from the voltage readings off the resistors. However, the probed voltages, v_1 and v_2, required scaling below $10V_p$ before being passed into the low-voltage multiplier chips. For this purpose, 20:1 resistor-bridge voltage dividers were put in place. Individual resistors were similarly selected to provide 0.1% accuracy in the scaling of the voltage signals.

Each pair of current and scaled voltage was then passed through its own analog multiplier chip, and the two resulting signals were summed together using the built-in feature of one of the chips. Coincidentally, an additional division by ten was incurred due to internal scaling of the chips' outputs. Collecting the various operations, the instantaneous power monitor signal was related to the currents and voltages by

$$P_{monitor} = \frac{\left(\frac{v_1}{20}\right)(-i_1) + \left(\frac{v_2}{20}\right)(-i_2)}{10} = -0.005(v_1 i_1 + v_2 i_2), \qquad (3.12)$$

such that 5mV measured was equivalent to 1W of power. One should note, however, that instantaneous power is comprised of two components: reactive power at twice the drive frequency and real power at DC (assuming constant motor conditions). Only the latter is representative of the electrical input power used by the motor. To accommodate this, the signal was read off the LED display of a standard bench-top voltmeter, thus requiring no further low-pass filtering to monitor only the DC component. The desired performance metric of real electrical input power, P_{in}, could then be interpreted from $P_{monitor}$ at a calibration of 20W/mV.
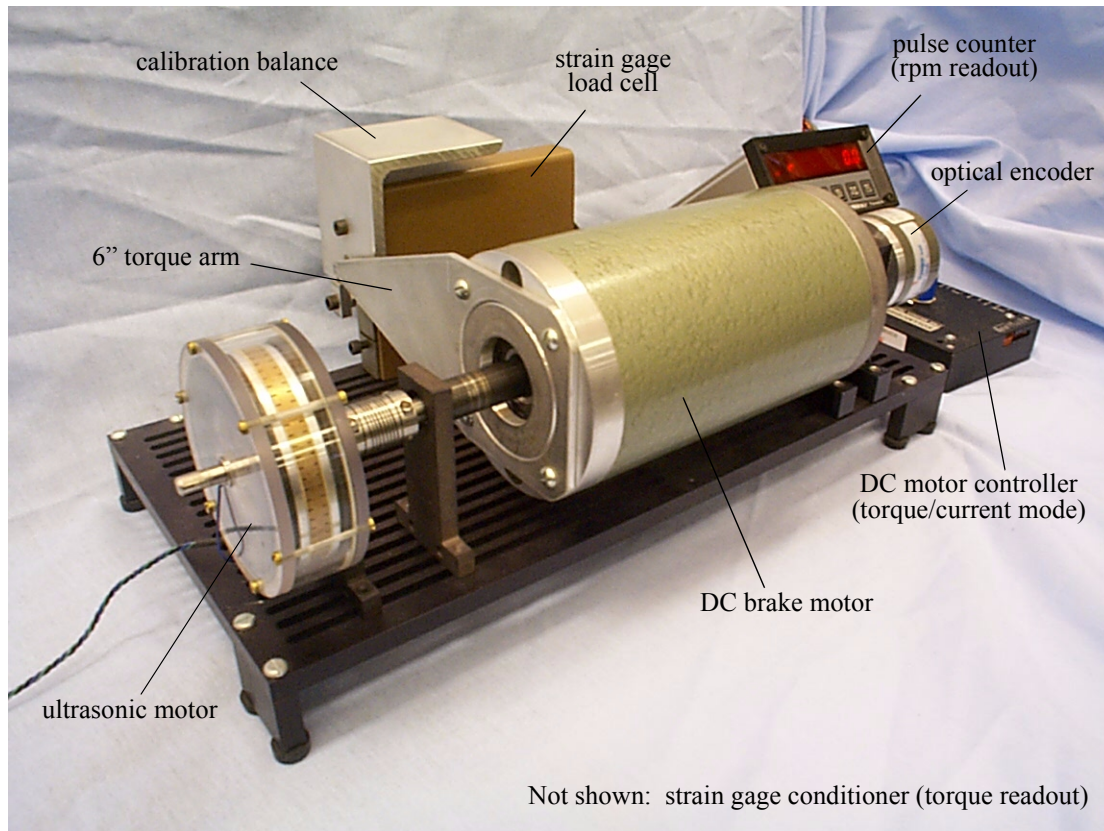
Figure 3.38 Torque-speed measurement apparatus.

Torque and Speed Measurement

A torque-speed measurement apparatus was assembled upon mechanical breadboard components for ease in reconfiguring the components it supported. The various devices of the measurement system are identified in Figure 3.38. Coupled to the shaft of the stationary ultrasonic motor was a relatively large DC motor for applying a resistive torque load. The load was controlled manually by adjusting the current source output of a DC motor controller via a potentiometer. The housing of the brake motor was not fixed to the breadboard supports but, rather, was connected via a 6-inch torque arm and piano wire to a strain gage load cell. Together with a strain gage conditioner, this configuration enabled accurate measurement of the reaction torque acting against the DC motor housing, i.e. the torque being transmitted to the ultrasonic motor shaft. The digital readout of the conditioner was calibrated to yield 10mV/N-cm with the use of a 1-lb weight and the attached load-cell calibration bal-

ance (6lb-in = 67.8N-cm = 678mV). Both gain and zero calibration were adjusted frequently to combat measurement drift. Bearing friction, especially that within the brake motor, was accounted for by specifically zeroing the measurement at speed.

Speed was monitored by an optical encoder mounted on a shaft hanger and coupled to the rear shaft of the brake motor. A pulse counter interpreted the digital signal coming from the encoder and was programmed to display the measurement directly in revolutions per minute (rpm). All of these measurements, including the input power reading, were recorded manually from their respective digital displays.

Forced-air Cooling

A major concern while testing the motors was the inescapable truth of internal heating. Without proper cooling, ultrasonic motors can reach unfavorable temperatures in a matter of minutes, regardless of load. It's not that cooling is necessary, only that ultrasonic motors are typically better suited for intermittent duty than for continuous operation. In an attempt to maintain steady performance as the motor heats up, ultrasonic motor drivers are often configured with a feedback loop to track the temperature-dependent resonance. However, stiffness is not the only property that changes. As the internal temperature rises, material losses increase, and, ultimately, motor performance suffers. Essentially, performance is a strong function of temperature and must be recognized as such. Therefore, to provide validity and consistency to the experimental performance measurements, it was deemed necessary to stabilize temperature by getting internal heating under control.

The solution was to install a forced-air cooling system comprised of an air compressor, a heat exchanger (cooling bath), a regulator valve, and the necessary 1/8" tubing and fittings. As an intermediate stage between the compressor and motor, the heat exchanger consisted of a length of coiled copper tubing submerged in an insulated ice bath. While the flow of air could be regulated with the valve, the temperature of the incoming air (already heated by the compressor) could also be regulated by varying the amount of ice in the bath. However, it was important to maintain a careful balance of flow and ice in order to avoid condensation

from forming inside the heat exchanger stage and blowing through the motor yet still provide the needed cooling.

To accommodate forced-air cooling, each of the motor housings had to be retrofitted with inlet and outlet air paths. In the case of each one-sided motor, as illustrated by the one-sided prototype in Figure 3.39, an inlet fitting was installed in the upper housing cap such that the high-velocity intake air would impact the upper rotor surface then disperse around its periphery. One set of exit holes was then drilled though the lower housing cap outside the radius of the stator hub. This creates a flow of air running around and under the stator. A second set of exit holes was instead drilled inside the radius of and through the stator hub to create a second flow of air through the upper teeth slits and across the upper stator surface. This combination of air flows for the typical one-sided assembly is illustrated in Figure 3.40(a).

In the case of the two-sided motor, all surfaces are cooled by the single path of air depicted in Figure 3.40(b). The inlet fitting was again positioned in one of the housing caps such that the high-speed air first impacts a rotor, but unlike in the other motors, it disperses inward and through the cavity occupied by the brushes. The air escapes via the other brush cavity through exit holes drilled at a similar radius in the opposing housing cap. In effect, the cooling air is forced to flow across the flat stator surface, out through the first ring of teeth,
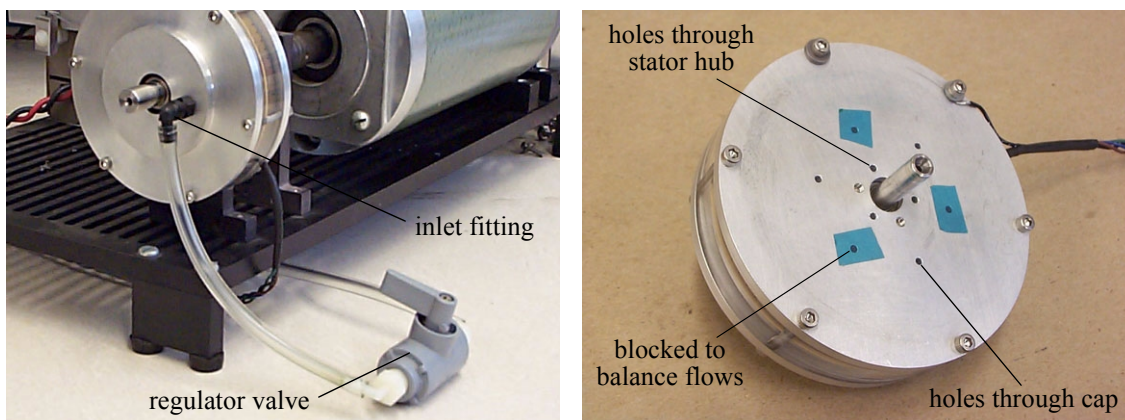
Figure 3.39   Implementation of air cooling for the one-sided prototype. The regulator valve and inlet fitting are shown on the left. Exit hole placement is shown on the right.
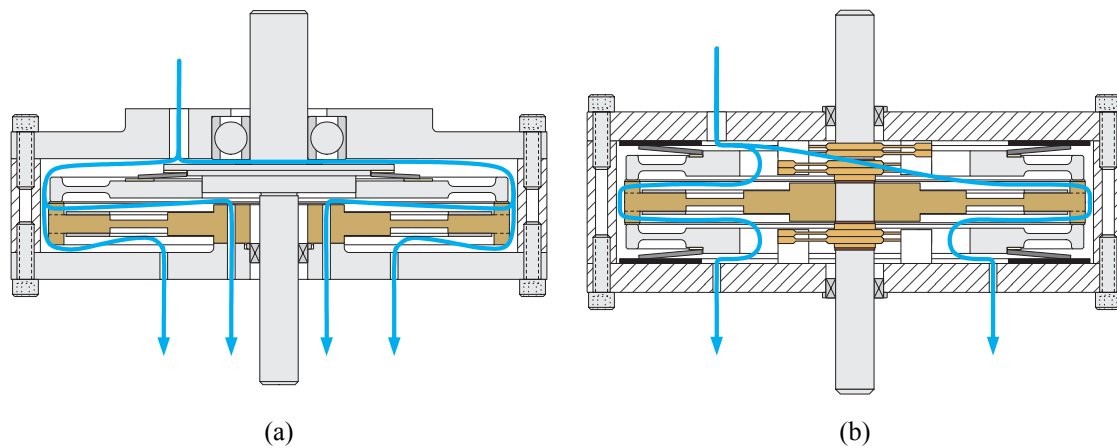
Figure 3.40   Cooling-air flow paths through (a) the one-sided motor and (b) the two-sided motor.

around the stator's periphery, back through the other ring of teeth, and across the other stator surface.

As an example of the benefits of cooling, Figure 3.41 compares time histories of the Shinsei USR60 motor with cooling and without cooling. For the demonstration, the USR60 was assembled with the 0.115" lip rotor and powered by the supplied USR60 motor driver. The most notable difference between the Shinsei driver and the open-loop, rack-mount setup described in Section 3.4.1, aside from compactness, is that the Shinsei driver has an internal feedback loop for compensating for temperature changes. Specifically, it monitors the signal from a sensor electrode on the stator and maintains its amplitude by appropriately tracking the drive frequency with the falling resonance. As was done for all motors in later tests, the temperature was monitored by potting a thermocouple with beeswax into one of the stator hub air holes. This configuration provided the most direct thermal contact possible with the heat-generating stator without interfering with its vibration. The drive frequency was probed using a standard frequency meter.

Looking first at temperature versus time in Figure 3.41, one can see that the uncooled case was showing no sign of stabilizing when the air valve was opened at the five-minute mark. In fact, temperatures as high as 65°C had been measured on the housing's outer surface during earlier runs of extended operation. As expected, the nearly linear temperature rise was
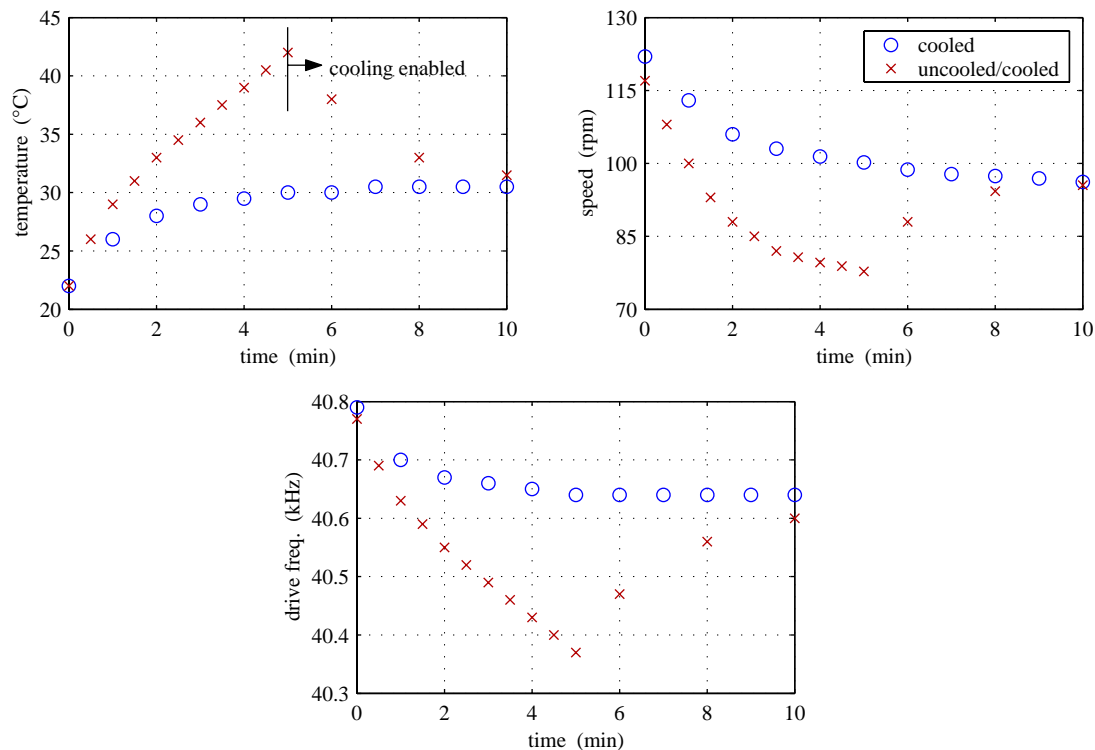
Figure 3.41   Demonstration of the benefits of forced-air cooling. Test case was the Shinsei USR60 motor with 0.115" rotor. Power was provided by the supplied Shinsei motor driver, as evidenced by the tracking of the drive frequency.

matched by a similar fall in drive frequency. Even with frequency tracking, however, motor speed took a significant hit before appearing to taper off. In contrast, the cooled case quickly stabilized within the first four or five minutes at around 30°C. Again, frequency and speed followed the trend. Notice that the when the valve was opened after five minutes during the uncooled test, each metric tended toward the cooled case within approximately the same time constant.

The results of this test demonstrated that the ultrasonic motor could be run continuously at a temperature not far from ambient using a simple system of forced-air cooling. By operating at a relatively cool and constant temperature for long periods of time, motor performance was extremely reproducible, and a significant amount of data could be recorded in a single sitting with greater confidence of consistency. As an obtainable goal, the cooling system was sized such that each and every experimental run could be sustained at 30°C.

Just as the signal amplitudes and phase were continuously adjusted during testing, so was the amount of flow and the temperature of the ice bath. The required balance of cooling system settings depended on the motor being tested and how hard it was being driven. Though maybe not indicative of how a motor would operate in real service, holding temperature constant reflected how the motor was being modeled.

Chapter 4

EXPERIMENTAL PERFORMANCE

Using the setup described in Section 3.4, steady-state motor performance was measured under carefully controlled conditions. As already discussed, a forced-air cooling system was implemented to counter internal heating and thus enable continuous motor operation at near constant operating temperature. In each of the one-sided assemblies, temperature was monitored at the stator hub by passing a thermocouple through one of the exit air holes. The hole was blocked on either end with beeswax to insulate the thermocouple from the cooler air flow. With the stator rotating in the two-sided motor, however, the temperature had to be probed elsewhere. Instead, the thermocouple was inserted through an exit air hole and wedged between a Belleville spring and the outer black rubber shim, again isolating it from the air flow. Neither location was in the direct vicinity of the heat-generating piezoceramics or contact interface, but both were the least intrusive choices given the two motor configurations. Throughout all experimental testing, the temperature and flow of the forced air were continually regulated to maintain the measured temperature at 30°C.

Another rather important control protocol was to subject each newly assembled motor to an initial wear-in period. When a motor is first assembled, performance can change dramatically over time as the newly mated polymer surface temporarily undergoes significant wear in localized areas of high contact pressures. As the mating improves, the contact pressures eventually even out, preferably below the limiting PV (pressure times velocity) value of the polymer. Wear then slows to a negligible rate and motor performance ultimately stabilizes.
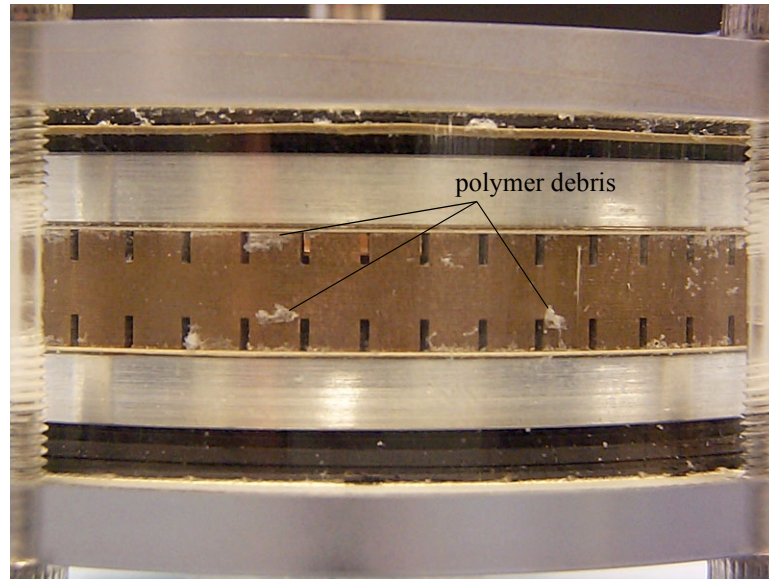
Figure 4.1 Polymer debris generated during the initial wear-in period, shown here for the two-sided motor.

Consequently, before any valid measurements could be documented, each newly assembled motor was driven continuously for an appropriate period of time until its speed and input power were observed to be steady. The procedure typically lasted several hours. Dramatic changes in performance were most noticeable for freshly lapped components, but the wear-in protocol was honored each and every time a motor was reassembled whether or not the components had been pre-worn during an earlier assembly. Figure 4.1 shows an example of the particles generated during the initial wear-in period of a freshly lapped polymer.

In addition to testing various rotor designs, parameter studies included varying drive frequency and voltage. However, because of the painstaking preparation involved in properly lapping a rotor for any given preload (see Section 3.1.6), each motor was tested under only its nominal preload as specified in Section 3.1.3. Specifically, the Shinsei USR60 and mock Shinsei motors were always preloaded with 36lb (160N) whereas the two prototypes were preloaded with 58lb (258N). As already described, the preload was set by varying the thickness of the appropriate shim in each of the assemblies. Weights were used to determine the force necessary to close the housing, leading to an estimated preload measurement error of

about 2lb (9N). The performance measurements of each of the four motors are presented in the following sections, one motor at a time.

4.1 Shinsei USR60

As a point of reference, the Shinsei USR60 was first tested with the supplied rotor at the nominal drive voltage of $150V_p$. Though the factory specifications list a drive voltage of $100V_{rms}$, or $141V_p$, the voltage amplitudes of the Shinsei motor driver were measured at 5–10% higher. The measured performance curves at a number of drive frequencies are plotted in Figure 4.2. As expected, motor speed increases as the drive frequency is brought closer to the nonlinear motor resonance. However, the nominal USR60 motor could only be driven down as far as 40.6kHz while still providing stable performance through to stall. At 40.5kHz, the abrupt onset of contact instability invalidated any performance measurements above a relatively low torque load. Easily recognized by an audible squeal, this phenomenon was common among all motor assemblies and dictated the lower limit to the valid range of drive frequencies.

The first observation to make is that none of the curves quite recreate the speed-torque curve in Figure 3.4. Though a comparable output power of 3.4W was achieved at 40.6kHz, it was done so with a performance profile with higher no-load speed and lower stall torque. In fact, because the measured speed-torque curves are more linear than the published curve, achieving the same output power required driving this particular USR60 considerably faster than expected. At 40.6kHz, the no-load speed was 145rpm and the stall torque was just higher than 50N-cm. These are compared to the factory specs of 100rpm, 62N-cm, and 3.5W.

What Shinsei does not provide are numbers for input power or efficiency. For the drive frequency of 40.6kHz, input power ranged from 9–15W thus yielding a maximum efficiency of 24%. As one can see, efficiency had yet to reach an absolute maximum as a function of the drive frequency. Unfortunately, the contact instability prohibited the motor from being driven any lower in frequency, as was generally the case in all the performance studies.
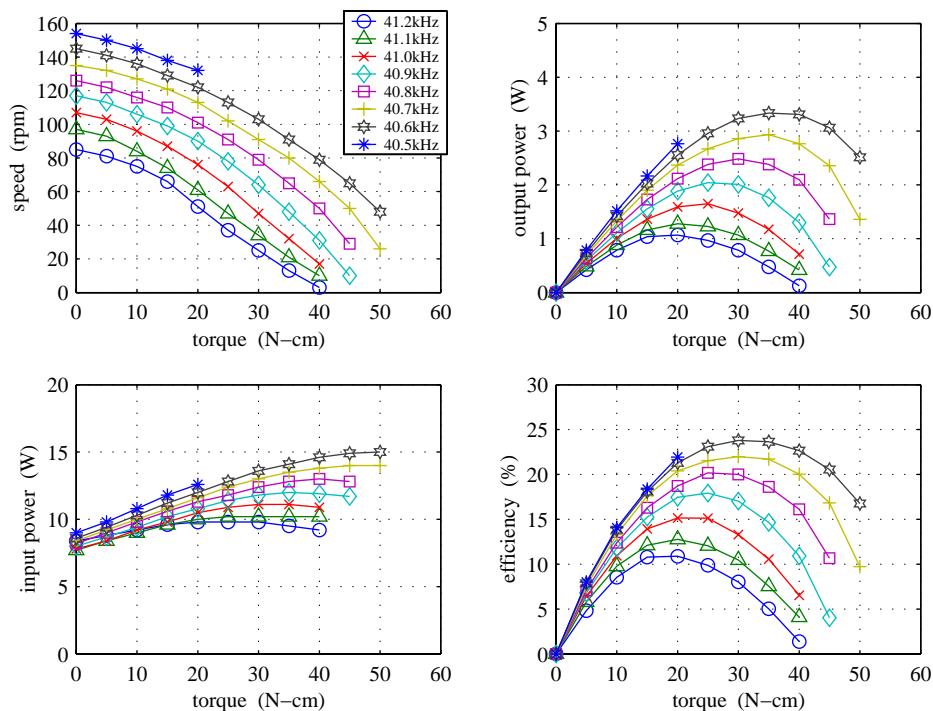
Figure 4.2   USR60 performance vs. frequency with nominal rotor (0.115") at $150V_p$.

Consequently, maximum efficiency was most often observed at the upper threshold of stable motor performance.

Figure 4.3 shows the performance of the nominal USR60 motor at 40.75kHz for varying drive voltage. Notice how motor performance is not a linear function of the voltage amplitude. In fact, the no-load speed is only about 25% higher at $200V_p$ than at $100V_p$. However, increasing the drive voltage not only increases speed but also the maximum achievable torque, so the benefit of increasing voltage is most noticeable by an increase in peak output power. What is more interesting, though, is that unlike with varying frequency, peak efficiency is rather constant over a wide range of drive voltages. Only when the amplitude was lowered to $100V_p$ did efficiency and torque suffer. There is a practical upper limit to the drive voltage, however, that is once again attributable to the contact instability at large vibration amplitudes. This can be seen in the $200V_p$ curve where stable operation fell just short of the anticipated stall torque.
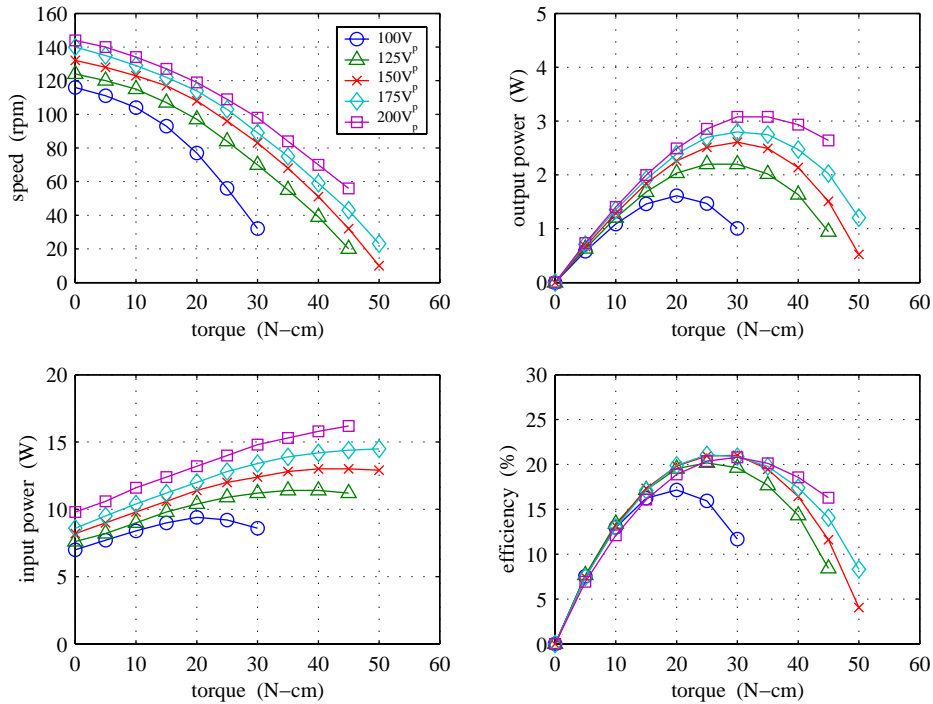
Figure 4.3 USR60 performance vs. voltage with nominal rotor (0.115") at 40.75kHz.

The criteria for choosing an appropriate drive voltage becomes clearer when looking at the additional voltage study in Figure 4.4. For this series of tests, the USR60 was assembled with the 0.125" lip rotor and driven at 40.8kHz. Again, peak efficiency was constant across the board, even to the point where the curves themselves are almost indistinguishable from one another. More noticeable than in Figure 4.3, however, are the commonly experienced upper and lower bounds of suitable drive levels. On the high side is the contact instability, most obvious at $200V_p$ but also slightly limiting at $175V_p$. On the low side is a somewhat less understood phenomenon where operation is not sustainable beyond a torque which falls short of the anticipated stall load. For example, at $100V_p$ the motor repeatedly stopped above 35N-cm without any prior warning of the speed beginning to waver. Notice that there is no sign of the $100V_p$ speed-torque curve diverging from the others like it did in the previous study. One might argue that the cause could be real-world factors such as surface roughness that come into play as the vibration amplitude falls too low. Whatever the case may be, the inability to characterize performance out to stall at any frequency can usually be attributed to too low of a drive voltage.
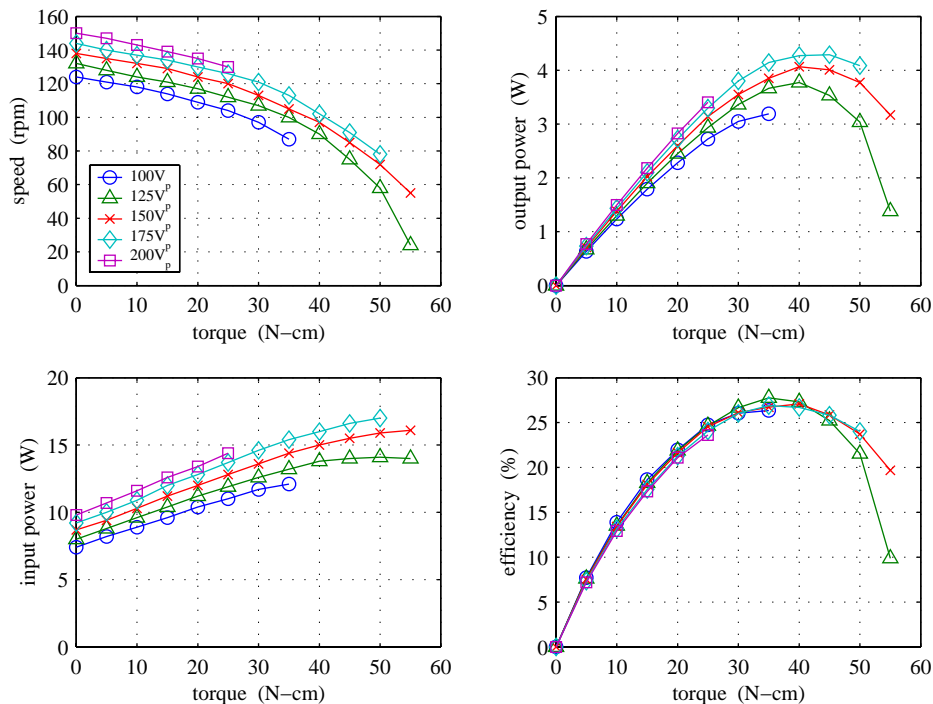
Figure 4.4   USR60 performance vs. voltage with 0.125" rotor at 40.8kHz.

These two voltage studies provide supporting evidence as to why $150V_p$ was a judicious choice for driving the Shinsei USR60. Although speed and output power would be higher at higher voltage, the marginal gains are not necessarily worth running the risk of losing smooth performance characteristics through to stall. Hence, the remaining USR60 rotor studies were all carried out at $150V_p$. After all, the results shown in Figure 4.4 are a preview to the fact that proper rotor design contributes more to performance than does increasing the voltage.

As discussed in Section 3.3.3, the polymer used by Shinsei is a slightly different blend of Ekonol and PTFE than the 20% blend bonded to all the in-house rotors. Most notably, the Shinsei polymer has an unmistakably higher coefficient of friction, so all else being equal, the nominal Shinsei rotor has that advantage over the others. The difference in the polymers becomes quite evident when the performance of the nominal rotor in Figure 4.2 is compared to that of the identically dimensioned 0.115" lip rotor in Figure 4.5. Both cases were operated under identical preloads and drive levels, but the lower coefficient of friction of the lat-
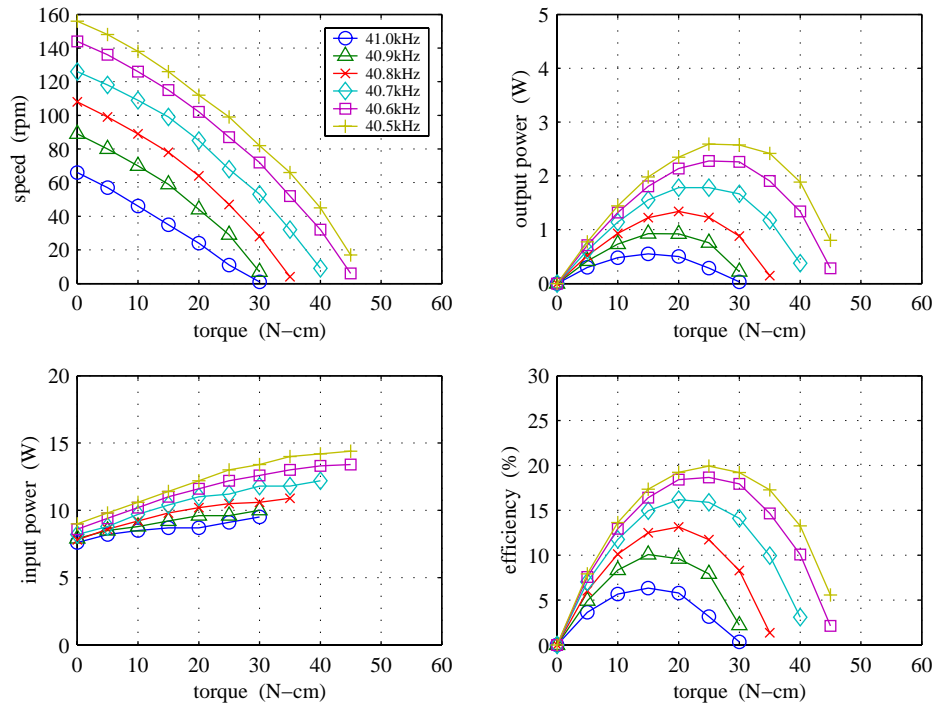
Figure 4.5 USR60 performance vs. frequency with 0.115" rotor at $150V_p$.

ter is clearly evidenced by a significant performance drop. The input powers are reasonably similar, which indicates that the material losses are comparable between the two blends. The achievable torques, however, have been dramatically lessened by the use of the 20% blend, thus taking a hit on output power and efficiency. Comparing the 40.6kHz curves, stall torque and output power dropped by about 10N-cm and 1.1W, respectively, and efficiency fell an absolute 5%. Interestingly, some of the performance loss was regained by the fact that the 0.115" rotor could be driven stably at 40.5kHz.

The performance with the 0.115" rotor sets the bar for the USR60 rotor study. It shares identical geometry with the nominal rotor but accounts for the performance loss due to the different polymer blend. As laid out in Section 3.1.5, four other rotors were manufactured with varying lip thicknesses: two thinner (0.095", 0.105") and two thicker (0.125", 0.135"). The dimensional modifications effectively change the dynamics of each rotor in relation to the stator.
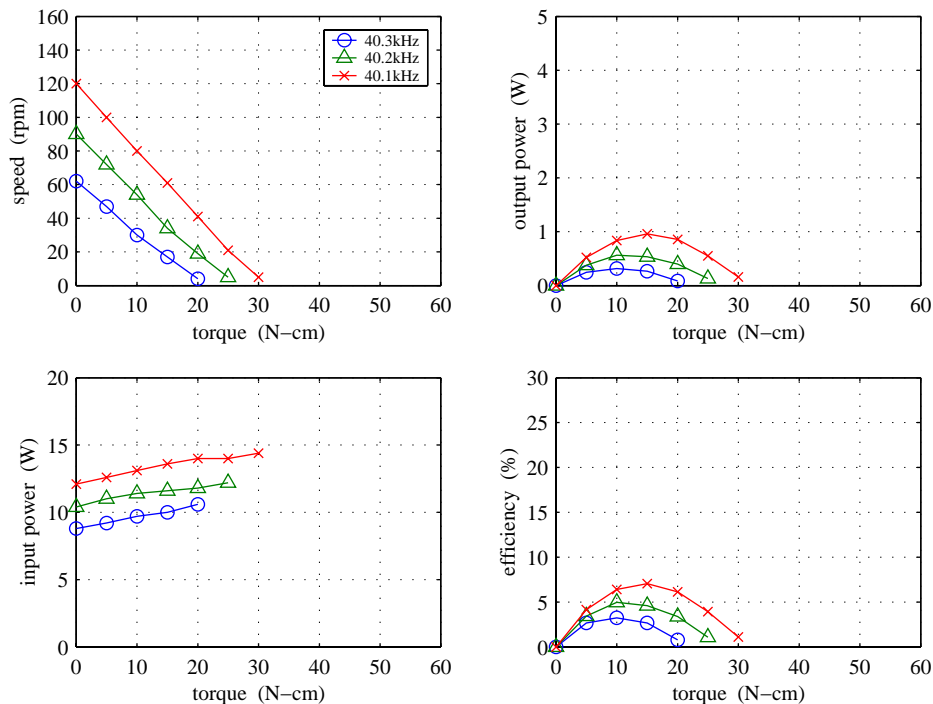
Figure 4.6   USR60 performance vs. frequency with 0.105" rotor at $150V_p$.

Looking first at a thinner design, the performance characteristics of the USR60 with the 0.105" rotor are shown in Figure 4.6. Immediately evident is that rotor design is very important and, more specifically in this case, that decreasing the nominal rotor stiffness was not the right direction in which to go. The motor would not run below 40.1kHz, and even then it quickly stalled at just 30N-cm. Performance was so bad across the board with the 0.105" rotor that only 1W of output power and 7% efficiency were achieved. Contributing to the low efficiency was the fact that the thinner rotor also demanded noticeably more input power. As one might anticipate, the 0.095" rotor was nearly impossible to drive with any measure of consistent results, so it was discarded from the study without hesitation.

On the other end of the spectrum, performance was actually greatly improved by making the rotor thicker, i.e. stiffer. As seen in Figure 4.7, increasing the lip thickness by just 10mil to 0.125" significantly extended the operational torque limit and broadened the useful range of drive frequencies. Speeds were also less inclined to drop linearly toward stall, thus yielding greater output powers. With the maximum practical no-load speed being only slightly
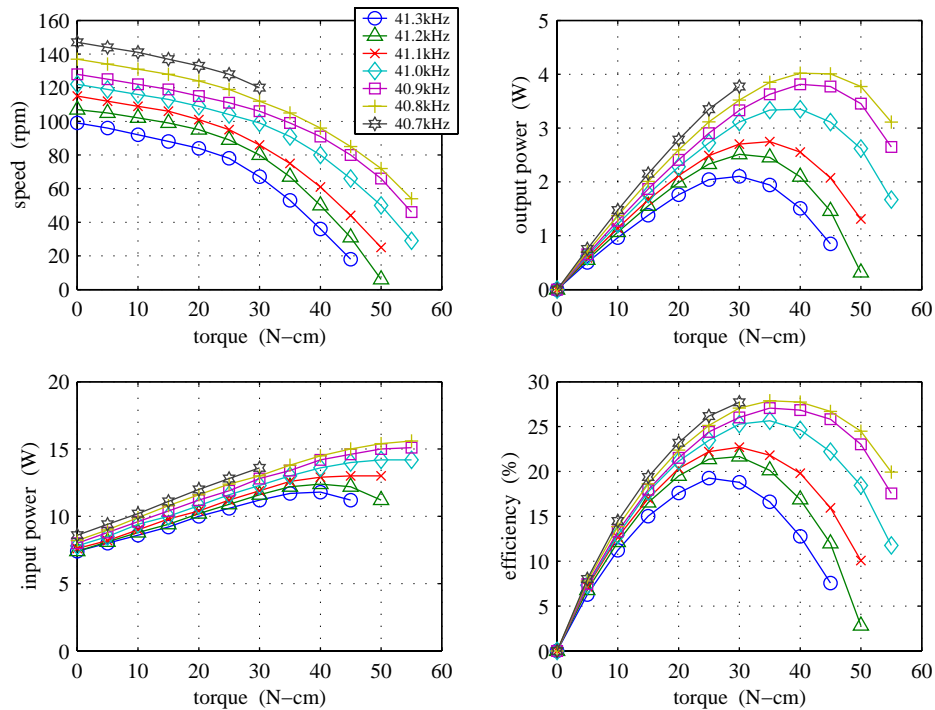
Figure 4.7   USR60 performance vs. frequency with 0.125" rotor at $150V_p$.

lower than that with the 0.115" rotor, peak output power shot up to 4W. Stall torque was increased to better than 55N-cm and peak efficiency rose to 28% at a drive frequency of 40.8kHz. Notice how efficiency was showing signs of saturating at the point where the contact instability was reached.

Performance was not, however, improved any more by increasing the lip thickness to 0.135". As seen in Figure 4.8, peak output power and efficiency were actually lowered somewhat, though they were still noticeably better than those with the nominal 0.115" rotor. The improved stall torque was maintained, but with lower overall speeds, peak output power was limited to 3W. Nevertheless, efficiency still reached a respectable 25% because the USR60 drew similarly less input power with the thicker rotor. Interestingly, the speed-torque characteristics with the 0.135" rotor look much more like the published curve in Figure 3.4 than do either of the 0.115" rotors.
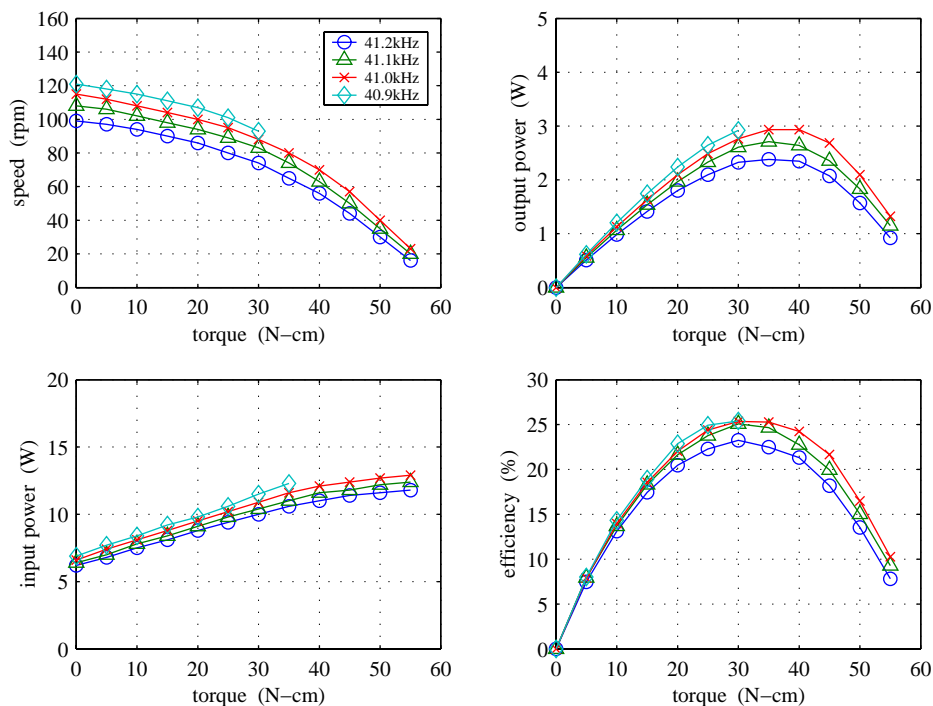
Figure 4.8   USR60 performance vs. frequency with 0.135" rotor at $150V_p$.

The importance of the rotor's dynamics is unmistakable after witnessing the dramatic variation in performance in Figures 4.5–4.8. At the extremes, performance has been shown to suffer by using a rotor that is either too flexible or too stiff. The optimal design, therefore, is somewhere in between. Of the rotors constructed for the USR60 study, the 0.125" rotor certainly stands out as the best overall performer. Interestingly, this study has also shown that the popular Shinsei motor can be dramatically improved upon with a simple dimensional modification to the rotor. In fact, one would expect that performance would be even more impressive than in Figure 4.7 if the 0.125" rotor profile were combined with the better Shinsei polymer blend.

4.2 Mock Shinsei

With the 0.125" rotor already having been identified as the best performer, it was chosen for the performance testing of the mock Shinsei stator. Together they were assembled within the Shinsei housing under the same preload as the USR60 above and similarly driven at
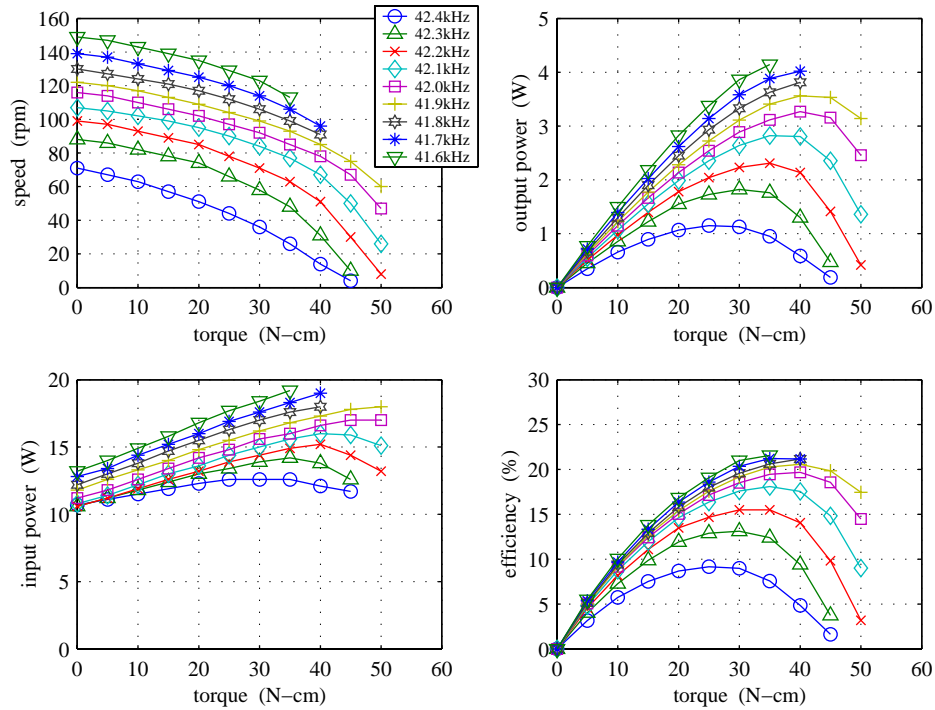
Figure 4.9 Mock Shinsei performance vs. frequency with 0.125" rotor at $150V_p$.

$150V_p$. However, as the results show in Figure 4.9, the higher resonant frequency of the mock Shinsei stator demanded driving this combination slightly higher in frequency than the USR60 in Figure 4.7. Consequently, the dynamic stiffness of the 0.125" rotor in the case of the mock Shinsei stator was effectively "softer" than in the previous case, though not by much. This is one factor to consider when comparing the results.

In light of the significant mode splitting of the mock Shinsei stator, the resulting speed-torque characteristics are surprisingly similar to the original. However, valid operation did not extend quite as far out in torque and was impacted sooner by the contact instability. Thus, a somewhat lower legitimate output power of 3.5W was achieved at 41.9kHz, though powers exceeding 4W were reached at lower frequencies near the threshold of the contact instability. In spite of the comparable speeds and output power, efficiency took a hit as a result of a dramatic increase in input power. Consequently, peak efficiency was only 20% for the 41.9kHz drive frequency.

The higher electrical power draw was undoubtedly due to the higher damping of the mock Shinsei stator and higher capacitance of the PZT-4 arrays (see Tables 3.4 & 3.5). Mode splitting likely contributed as well, though a degraded traveling wave is considered to have a more prominent effect on speed than input power. Adding to the list, the mock Shinsei stator had some very unique issues regarding its modal characteristics that were brought to light in Section 3.2. Because of this, it is unknown whether the same dominant source of loss is attributable to the two similarly manufactured prototype stators or whether excess loss can be attributed to the mock Shinsei stator because of its unique flaw. This should be kept in mind when making comparisons to the following prototype results.

4.3 One-sided Prototype

As was done for the USR60 motor, the one-sided prototype was tested at various voltages to determine the appropriate drive levels. For this study, the one-sided motor was assembled with the 0.170" stator and driven at 31.4kHz, a combination that had been shown to yield good results during preliminary testing. The results are shown in Figure 4.10, again for drive voltages ranging from 100–200V_p. Interestingly, the one-sided motor was capable of being driven full range at 200V_p without triggering the contact instability. Consequently, the highest achieved stall torque and output power went hand-in-hand with the highest drive level. Peak efficiency, on the other hand, had already begun to drop after the mid-range voltages. Therefore, in spite of the lower stall torque, it was decided to perform the parametric rotor study at 150V_p in order to maximize efficiency and also maintain consistency with the USR60 tests. However, the performance benefit of operating at higher voltage was not easily dismissed and indeed prompted further investigation. Those results are presented following the rotor study. On a final note on the drive voltage, it was no surprise that performance was limited at low voltage by the same seemingly premature stopping that was observed with the USR60.

The results of the one-sided motor rotor study are presented in Figures 4.11–4.14 for the 0.150", 0.160", 0.170", and 0.180" rotors, respectively. Though none of these performed as
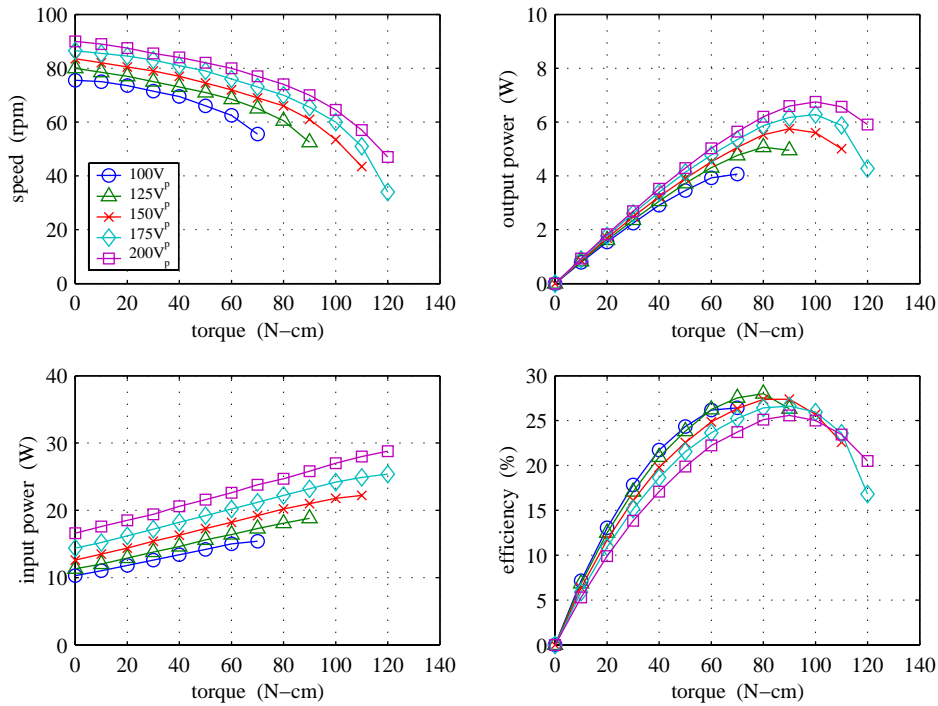
Figure 4.10   One-sided prototype performance vs. voltage with 0.170" rotor at 31.4kHz.

poorly as the thinnest of the USR60 rotors, a performance trend similar to the USR60 rotor study can be seen as a function of rotor thickness. Actually, for this particular array, the trend was more pronounced. As thickness was increased, overall speeds dropped, the torque range was extended, and less electrical power was consumed. Furthermore, performance became less sensitive to frequency. The inverse trends of speed and torque resulted in the greatest output power being attained with the 0.160" rotor. Peak efficiency was shared by the less power hungry 0.170" and 0.180" rotors, albeit the latter was incapable of carrying those drive levels out to full stall.

The 0.170" rotor was decidedly the best balanced performer, having produced the best overall metrics over a relatively wide range of drive frequencies. With this rotor, the one-sided motor reached a peak output power of 6.8W and a very respectable efficiency of 29% at a drive frequency of 31.2kHz. Stall torque was better than 110N-cm, and the no-load speed was 93rpm. Fortunately, the contact instability remained in check long enough that efficiency appears to have reached a maximum at this point. It should be noted, however, that

Figure 4.11   One-sided prototype performance vs. frequency with 0.150" rotor at $150V_p$.

Figure 4.12   One-sided prototype performance vs. frequency with 0.160" rotor at $150V_p$.

Figure 4.13 One-sided prototype performance vs. frequency with 0.170" rotor at 150V$_p$.

Figure 4.14 One-sided prototype performance vs. frequency with 0.180" rotor at 150V$_p$.

although the output power achieved with the 0.180" rotor was noticeably lower than that with the 0.170" because of its lower speeds, it nevertheless extended the operational torque range of the one-sided motor to as high as 130N-cm.

As was learned by the voltage study in Figure 4.10, the torque range of the 0.170" rotor could be extended by increasing the drive voltage beyond $150V_p$. Therefore, it was decided to investigate the potential of the 0.170" rotor at $175V_p$ at a lower frequency. As such, the previous performance at $175V_p$ and 31.4kHz is repeated in Figure 4.15 with the addition of data measured at 31.2kHz. Just as had been hoped given the results at $150V_p$, even better performance was achieved with this combination. Compared to $150V_p$ and 31.2kHz, peak output power was raised another 0.8W to 7.5W while still retaining 29% efficiency. Together with a no-load speed of 97rpm and a stall torque of 120N-cm, these are considered the defining performance specifications of the one-sided motor.
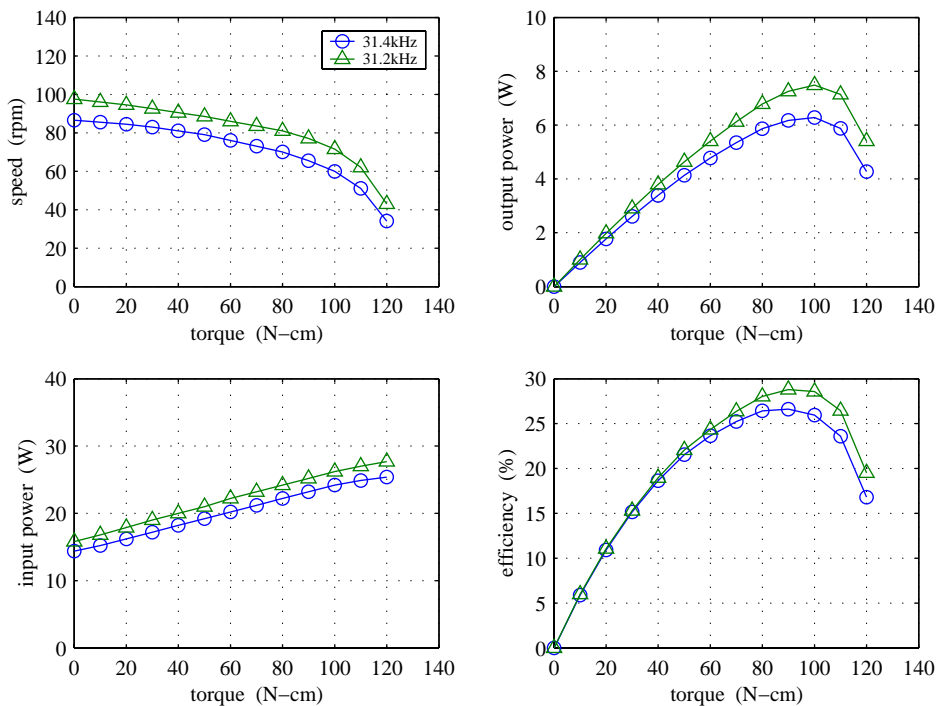
Figure 4.15 One-sided prototype performance vs. frequency with 0.170" rotor at $175V_p$.

4.4 Two-sided Prototype

The final motor to be tested was the two-sided prototype. To remain consistent with previous motor tests, performance measurements were first attempted at $150V_p$. Not too surprisingly, however, it soon became evident that higher drive voltages were going to be necessary to counteract the increased contact forces acting on the stator. To investigate this, the two-sided motor was subjected to a brief voltage study using the 0.180" rotor pair at 33.0kHz. The results are shown in Figure 4.16. Though not optimal, this combination proved in early tests to provide fairly respectable performance. As seen by the curves, the torque capability of the two-sided motor was cut short at the nominal drive level of $150V_p$ by the unknown low-voltage phenomenon discussed earlier. Torque output improved dramatically at $175V_p$ and $200V_p$, though even then it was a far cry from the expected capability of the symmetric design. Accordingly, $150V_p$ was rejected as a valid operating voltage for the two-sided motor, and all subsequent tests were carried out at $175V_p$.
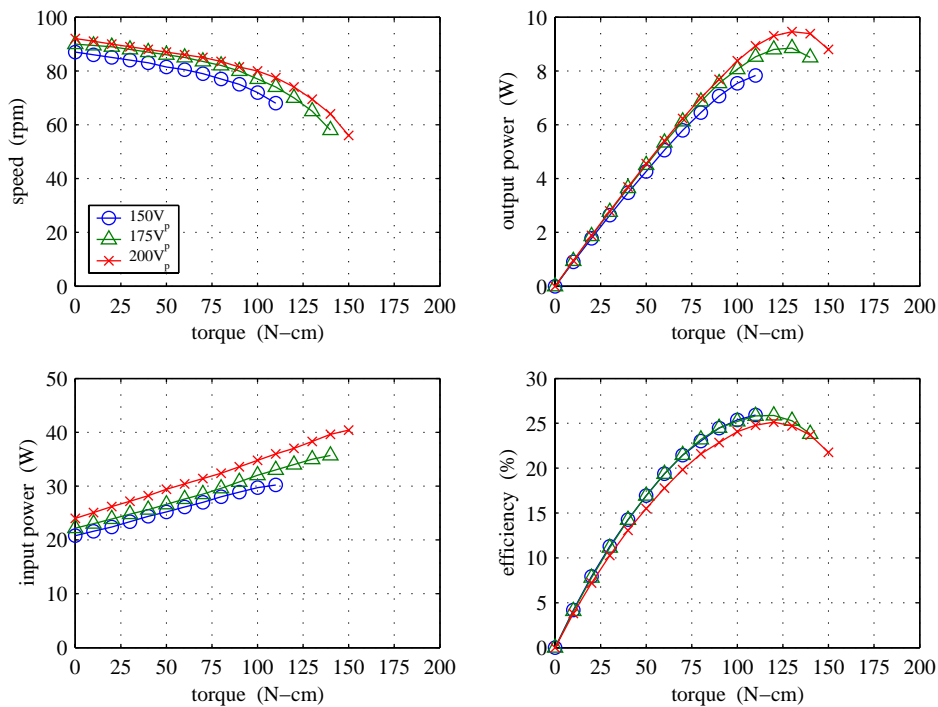
Figure 4.16 Two-sided prototype performance vs. voltage with 0.180" rotors at 33.0kHz.

The two-sided prototype was not tested with the full set of rotors originally defined in Table 3.3. Instead, given the success of the 0.170" and 0.180" rotors in the one-sided motor, only those dimensions were investigated in the two-sided motor tests. With the importance of the rotor already well characterized, simply maximizing the performance of the two-sided motor was of greatest interest. Consequently, the thinner rotors were not considered.

The measured performance curves with the 0.170" and 0.180" rotor pairs are presented in Figures 4.17 & 4.18, respectively. Quite unexpectedly, the speeds achieved with the two-sided motor were noticeably higher than those of its one-sided counterpart. Consequently, even though the achieved torque output was significantly lower than twice that of the one-sided motor, the power output of the two-sided motor, though not quite double, was considerably higher. In fact, when driven optimally at 32.0kHz, the two-sided motor put out an astounding 12.2W of output power with the 0.170" rotors compared to the one-sided motor's 7.5W. Its greater mechanical output was also naturally accompanied by greater demands for electrical input power, reaching 50W at one extreme. At those power levels, it was a challenge keeping the internal temperature down at 30°C with the existing air-cooling system. Favorably, input power scaled in proportion with output power with the implementation of two-sided operation such that efficiency remained high, peaking out at 27% with the 0.170" rotor. The no-load speed at that optimal drive frequency of 32.0kHz was 116rpm with a maximum operable torque of 160N-cm.

As before, the 0.180" rotor pair was slightly better suited for high torque at the cost of lower output power due to lower speeds. Efficiency also dropped a little with the use of the thicker rotors. Nevertheless, the two-sided motor was conceived with the desire to increase torque, and 160N-cm was well below expectations. Therefore, an even thicker pair of rotors was machined with a lip thickness of 0.190" to ascertain whether the upper torque limit could be extended any farther at this voltage. The disappointing results are shown in Figure 4.19. Performance was generally as expected, but maximum operable torque actually went down. It is an interesting result, nonetheless, as it confirms that even torque output reaches a peak as rotor stiffness is increased, something only vaguely concluded in the earlier rotor studies.
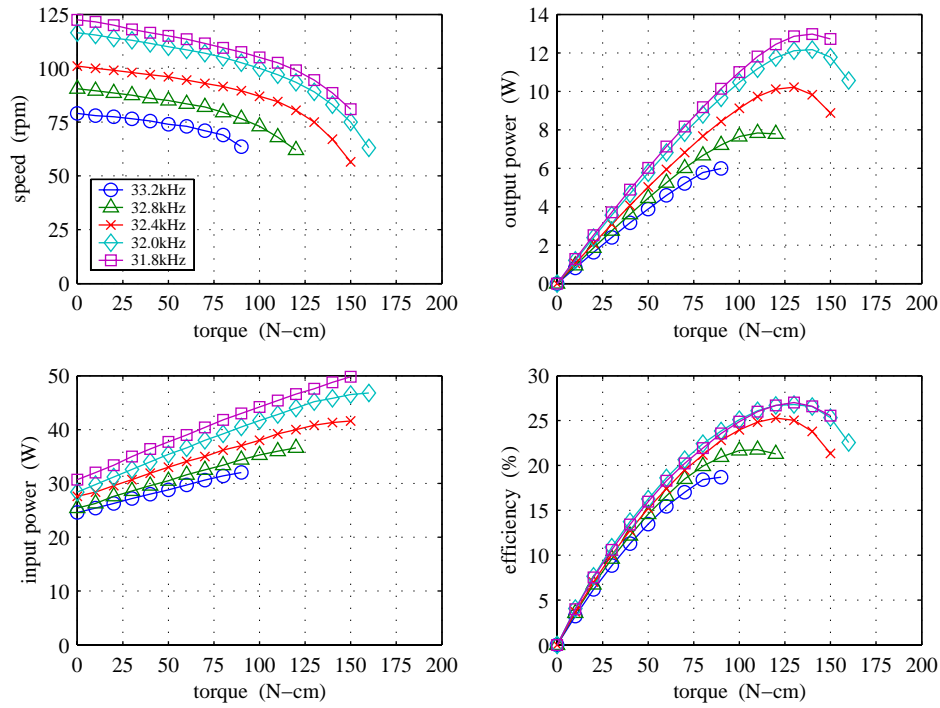
Figure 4.17 Two-sided prototype performance vs. frequency with 0.170" rotors at 175V$_p$.
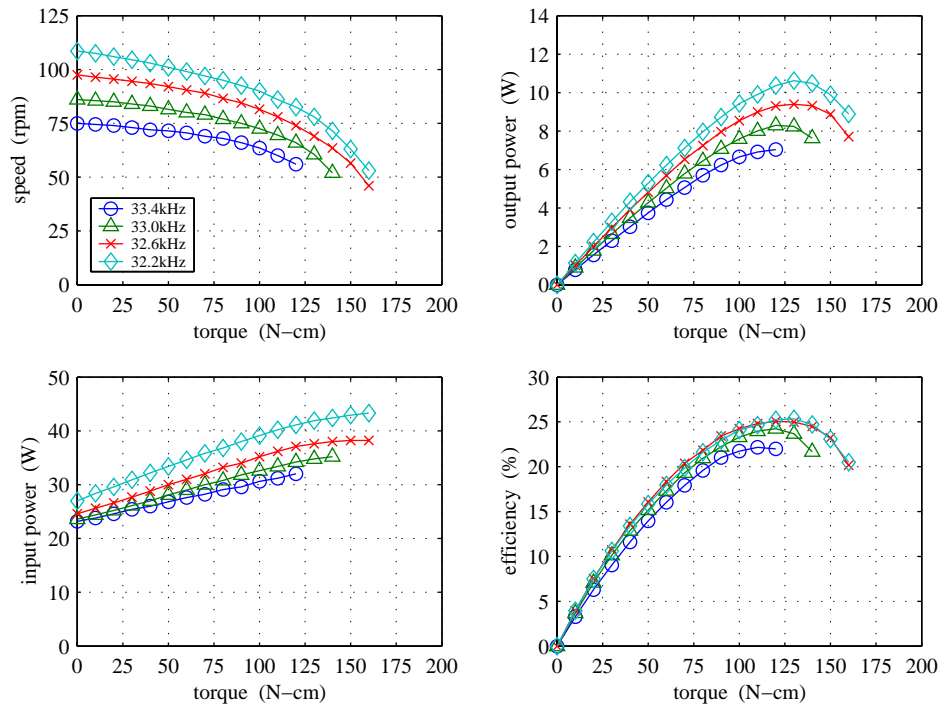
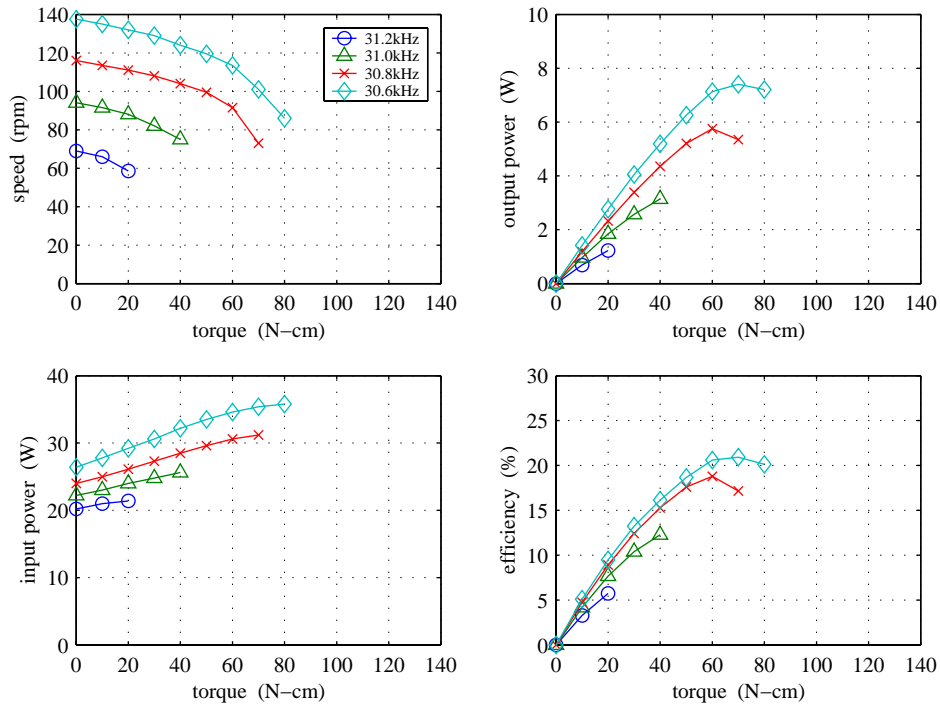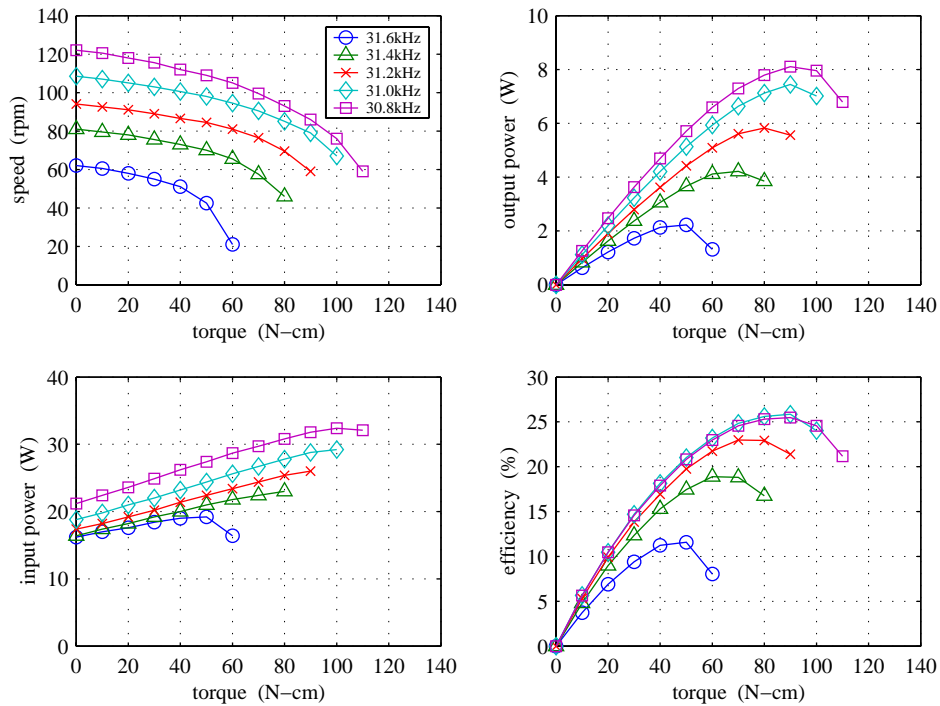**Figure 4.18** Two-sided prototype performance vs. frequency with 0.180" rotors at 175V$_p$.
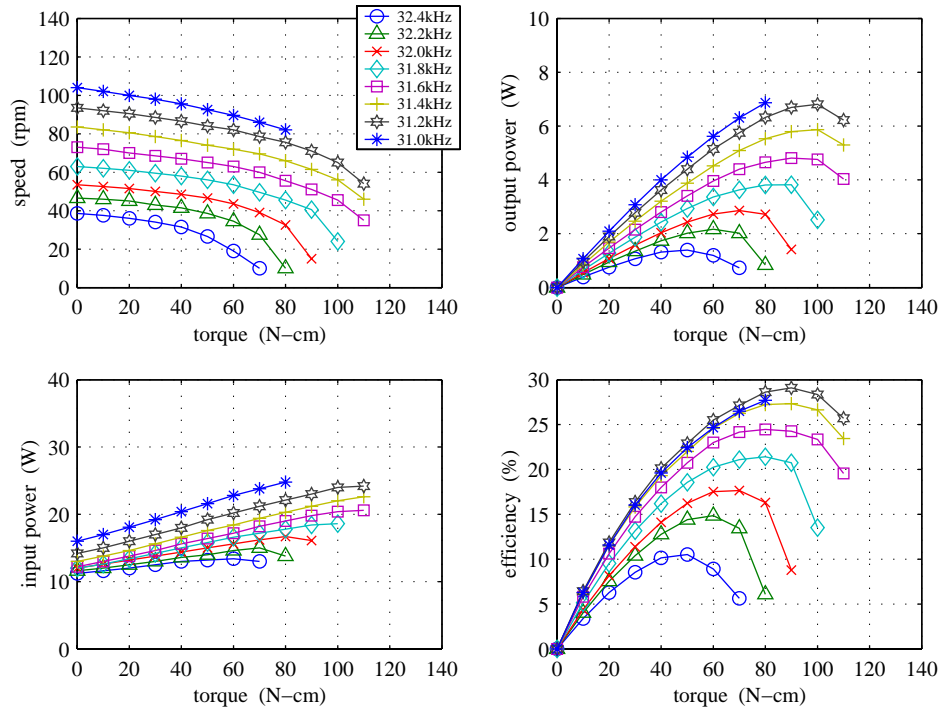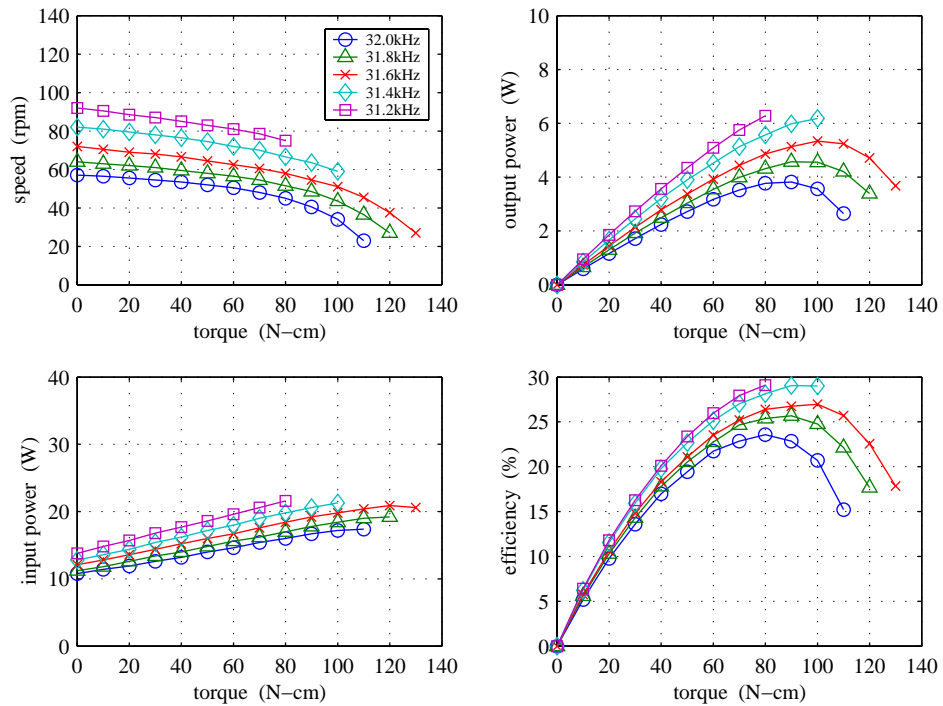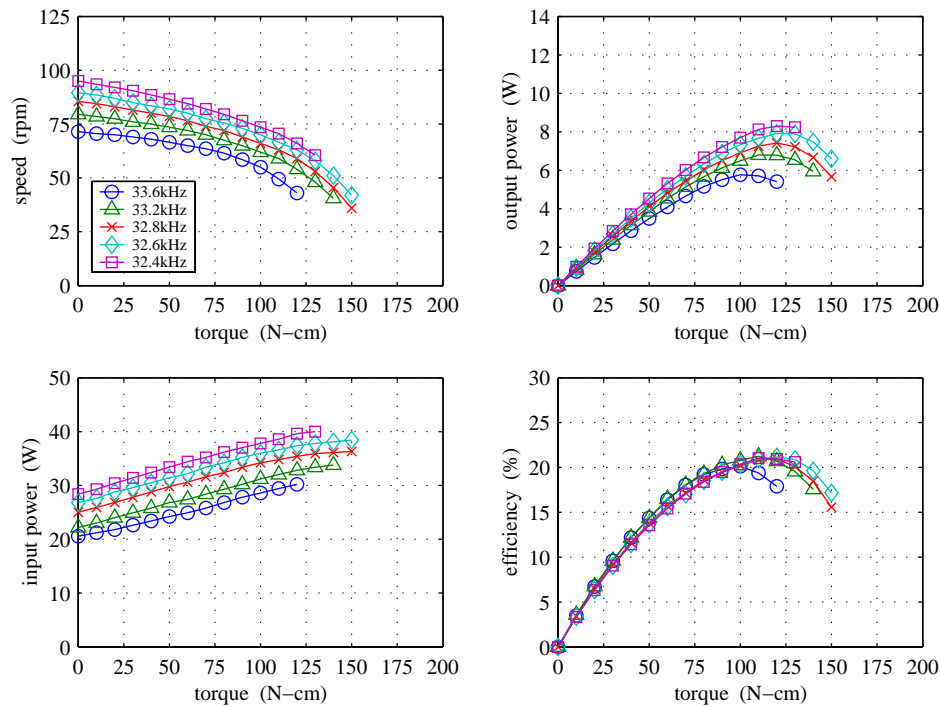
Figure 4.19   Two-sided prototype performance vs. frequency with 0.190" rotors at 175V$_p$.

4.5 Summary

A number of noteworthy observations were made during the experimental performance studies. For one, as seen in Figures 4.3, 4.4, 4.10, and 4.16, speed was not as sensitive to the drive voltage as one might have supposed. The maximum operable torque, on the other hand, was greatly influenced by the drive voltage, though understandably only up to the maximum stall torque defined by the friction coefficient. At too low a voltage, each motor was shown to stall prematurely, well below when expected. Operation simply could not be sustained at higher loads without increasing the voltage. At the other extreme, increasing the drive voltage too far resulted in lower efficiency as the dielectric losses began to dominate. Thus, a reasonably optimal drive voltage was chosen in each case that yielded the highest torque and output power without noticeably degrading peak efficiency. Ultimately, the maximum achievable output of each motor was limited by the onset of contact instability at high drive levels. Coincidentally, it should be noted that peak performance was generally achieved at a drive level verging on this instability.

Although selecting an appropriate drive level was important for optimizing performance, it was clearly demonstrated that no combination of voltage and frequency could make up for poor rotor design. This was most notable in the USR60 rotor study where the thinnest rotor could not even be tested reliably, let alone with any measure of success. In each of the rotor studies, the measured speed-torque characteristics of a given motor were strongly dependent on the dynamic stiffness of the rotor, a property which was controlled by varying rotor thickness. Speeds were generally higher with more flexible rotors, whereas higher torques were achieved with stiffer rotors. As predicted, efficiency dropped off in the direction of both extremes, thus demonstrating that efficiency can be optimized with a suitably "tuned" rotor. The dynamics of the rotor play a major role in the performance of a traveling-wave motor, and their inclusion in the proposed model is clearly warranted and, more appropriately, vital.

The results of the experimental performance studies are summarized in Table 4.1, where the best performing combination of rotor, voltage, and frequency is highlighted in bold for each motor. In the case of the Shinsei USR60, stall torque, output power, and efficiency all peaked simultaneously with the 0.125" rotor, making it the unequivocal choice. This wasn't the case with the one-sided prototype, however, where torque was greatest with the 0.180" rotor, output power with the 0.160" rotor, and efficiency with the 0.170" rotor. The latter was therefore selected as it yielded the best balanced performance. Not too surprisingly, that same rotor design produced the best results in the two-sided motor as well.

The most unexpected outcome of all these tests was how the two-sided motor performed in comparison to the one-sided motor. The logical hypothesis was that the implementation of two-sided operation would double the stall torque and output power while not affecting speed. Though the two-sided motor clearly outperformed its one-sided counterpart, albeit at a slightly reduced efficiency, it did not do so in line with the theory. Torque fell well short of expectations, and speed actually increased. While the increased speed can possibly be attributed to delayed contact instability, the resulting torque output of the two-sided motor cannot be explained here. Being the first of its kind, it's understandable that the hypothesis

TABLE 4.1 Performance Summary

Motor	Mass	Rotor Dim.	Drive Volt.	Drive Freq.	Stall Torque	No-load Speed	Max. Power	Peak Eff.	Torque Density	Power Density
	(g)	(in)	(V_p)	(kHz)	(N-cm)	(rpm)	(W)	(%)	(N-cm/kg)	(W/kg)
Shinsei USR60	240	nom.	150	40.6	50	145	3.3	24	208	13.9
		0.105	150	40.1	30	120	1.0	7	125	4.0
		0.115	150	40.5	45	156	2.6	20	188	10.8
		0.125	**150**	**40.8**	**55**	**137**	**4.0**	**28**	**229**	**16.8**
		0.135	150	41.0	55	115	2.9	25	229	12.2
		0.125	175	40.8	50	144	4.3	27	208	17.9
Mock Shinsei	240	**0.125**	**150**	**41.9**	**50**	**122**	**3.6**	**21**	**208**	**14.8**
One-sided prototype	420	0.150	150	30.6	80	137	7.4	21	190	17.6
		0.160	150	30.8	110	122	8.1	26	262	19.3
		0.170	150	31.2	110	93	6.8	29	262	16.2
		0.180	150	31.6	130	72	5.3	27	310	12.7
		0.170	**175**	**31.2**	**120**	**97**	**7.5**	**29**	**286**	**17.8**
Two-sided prototype	350	**0.170**	**175**	**32.0**	**160**	**116**	**12.2**	**27**	**457**	**34.8**
		0.180	175	32.2	160	108	10.6	25	457	30.3
		0.190	175	32.6	150	90	7.9	21	429	22.6

did not hold true to this one isolated example. Perhaps future work on other two-sided motors will provide an explanation, or better yet, produce a two-sided motor that holds up to its potential of doubling the stall torque. Nevertheless, this particular two-sided prototype did produce higher torque, speed, and output power than the similarly designed one-sided prototype via an assembly that was only incrementally larger and that was actually lighter. Its normalized performance metrics of torque and power density stand out among the others, thus exemplifying the potential of two-sided operation.

Chapter 5

SIMULATED PERFORMANCE

To assess the validity of the presented traveling-wave motor model, numerical simulations have been performed for direct comparison with several of the experimental performance case studies. As a prelude to examining the results of the simulations, however, a thorough discussion of the model implementation procedures is first detailed along with a cumulative review of the pertinent input parameters. These preparatory segments are then followed by a preliminary evaluation of the model's fundamental capabilities. Specifically analyzed are the Rayleigh-Ritz calculated modal characteristics and a sample of the simulated contact dynamics and forces. Finally, the results of the motor performance simulations are presented, and the strengths and limitations of the model approach are identified. In closing, the input power calculation is revisited for the purpose of identifying specific loss contributions in the simulated motor system and relating how capturing each promotes accurate performance prediction.

5.1 Model Implementation

Aside from the hysteretic friction force, the motor model in Chapter 2 was formulated analytically without regard for specific implementation procedures. This section, therefore, discusses how the model was numerically implemented in MATLAB[1] and addresses the building of the system matrices and the details of the solution process. A number of

---

1. The MathWorks, Inc., 3 Apple Hill Drive, Natick, MA 01760-2098, http://www.mathworks.com.

assumptions are made to simplify the model implementation including how to handle non-idealities and two-sided operation. Nevertheless, keep in mind that the analytical model is more general than the resulting implementation presented here, and that other implementations may be more suitable for other studies.

5.1.1 System Matrices

The first matter of business is to calculate the system matrices in the equations of motion. In Section 2.2.1, a Rayleigh-Ritz approach is employed involving the use of assumed shape functions for the midplane displacements and electric potential. The system matrices are calculated by integrating products of these functions and the material properties over the volumetric geometries of both the stator and rotor. Appropriate shape functions and numerical integration procedures must be chosen to provide acceptable convergence of the matrices and the eigenvalue problem.

Another factor to be discussed in the building of the system matrices is the handling of disk asymmetry which leads to mode splitting and non-ideal motor behavior. Rather than fully characterize the allusive cause of asymmetry, it is argued that mild asymmetry can be accounted for by integrating the effects into the system matrices. This will be explained in detail following the discussions on shape functions and numerical integration.

Assumed Shape Functions

As an integral step in the Rayleigh-Ritz formulation of the equations of motion, the spatial variances of the midplane displacements are replaced by linear combinations of assumed shape functions. The validity and accuracy of the Rayleigh-Ritz approach depends on the choice and number of these functions, or degrees of freedom, included in the formulation. Of primary concern, proper selection is governed by a number of factors including the desired types of displacement fields, the geometric boundary conditions, and the forces.

The desired displacement fields can be classified into three types: rigid-body, modal, and static. In Section 2.2.2, the external forces and rigid-body modes were handled separately

from the strain-energy modes, so assumed shape functions are not necessary to capture the rigid-body displacements. Similarly, static displacements have been neglected in the formulation of the modal equations of motion, though a more general approach could have been taken. An external mechanical work term could have been added beside δW_m in (2.13) to represent the work done by external forces acting on the rotor, i.e. the static axial preload and torque load. Doing so would allow static warping of the rotor to be captured if shape functions without circumferential periodicity are assumed. No external forces act on the stator, on the other hand, so the contact forces alone are sufficient to model static warping of the stator. Warping is a real response to the preload and affects the radial contact force profile. However, the effect of warping is counteracted as the contact surfaces wear during operation. Furthermore, the polymer surfaces of the experimental rotors have been lapped conically such that the stator and rotor contact surfaces mate properly under preload upon initial assembly. Given these conditions, assumed shape functions are not included for the purpose of capturing static displacements.

Therefore, only modal displacements are considered when selecting the appropriate shape functions. The desired modal excitation consists of degenerate mode shapes with $N\theta$ circumferential periodicity where N is the number of wavelengths. As the piezo arrays are configured to excite only these shapes and have little or no authority to excite modes of different periodicity off resonance, only shape functions with $N\theta$ circumferential periodicity need to be included. Furthermore, higher spatial harmonics are neglected, so the assumed shape functions are chosen to vary circumferentially only by the fundamental sinusoids $\cos N\theta$ and $\sin N\theta$. The validity of ideal $N\theta$ circumferential periodicity depends on the purity of system axisymmetry. Even the worst-case mode splitting observed in the experimental stators resulted in relatively pure mode shapes, so the assumption is considered valid for mild asymmetry.

Utilizing the separation of variables, the assumed out-of-plane shape functions, ϕ_w, are then given by

$$\phi_w(r, \theta) = \left[\left[\psi_w^1(r) \cdots \psi_w^n(r) \right] \cos N\theta \; \left[\psi_w^1(r) \cdots \psi_w^n(r) \right] \sin N\theta \right]. \qquad (5.1)$$

In the radial direction, a large number of shape functions ($n = 30$) are used to accommodate the sharply varying radial profiles. Clamped-free beam functions [Pilkey, 1994] have been chosen to match the geometric boundary conditions of the center-clamped stator. Defining the normalized radial position

$$\xi = \frac{r - a}{b - a}, \qquad (5.2)$$

where a and b are the inner and outer radii, the shape functions are given by

$$\psi_w^i(\xi) = \psi_{beam_{c\text{-}f}}^i(\xi) = \cosh\lambda_i\xi + \cos\lambda_i\xi - \beta(\sinh\lambda_i\xi + \sin\lambda_i\xi), \qquad (5.3)$$

where

$$
\begin{aligned}
\lambda_i &= 1.875104,\ 4.694091,\ 7.854757,\ 10.995511,\ 14.137168 & i \le 5 \\
&= (2i - 1)\pi/2 & i > 5 \\
\beta &= \frac{\cosh\lambda_i - \cos\lambda_i}{\sinh\lambda_i - \sin\lambda_i}.
\end{aligned}
\qquad (5.4)
$$

The rotor is unconstrained, assuming ideal loading by the preload spring, so the first two assumed out-of-plane radial functions are $\psi_w^1 = 1$ and $\psi_w^2 = 1 - 2\xi$. Then free-free beam functions are used:

$$\psi_w^{i+2}(\xi) = \psi_{beam_{f\text{-}f}}^i(\xi) = \cosh\lambda_i\xi - \cos\lambda_i\xi - \beta(\sinh\lambda_i\xi - \sin\lambda_i\xi), \qquad (5.5)$$

where

$$
\begin{aligned}
\lambda_i &= 4.730041,\ 7.853205,\ 10.995608,\ 14.137166,\ 17.278760 & i \le 5 \\
&= (2i + 1)\pi/2 & i > 5 \\
\beta &= \frac{\sinh\lambda_i - \sin\lambda_i}{\cosh\lambda_i + \cos\lambda_i}.
\end{aligned}
\qquad (5.6)
$$

Appropriate in-plane shape functions, ϕ_u and ϕ_v, can then be related to the specified out-of-plane shape functions. Since the purpose of including extensional midplane displacements is to compensate for the unknown location of the neutral plane in the case of an asymmetric disk with coupled extensional-bending modes, the most efficient shape functions are formed by the respective derivatives in (2.16):

$$\phi_u^i(r, \theta) \propto \frac{\partial \phi_w^i}{\partial r} \quad \text{and} \quad \phi_v^i(r, \theta) \propto \frac{1}{r}\frac{\partial \phi_w^i}{\partial \theta}. \tag{5.7}$$

So, in terms of the assumed out-of-plane shape functions,

$$\phi_u(r, \theta) = \left[\left[\frac{\partial \psi_w^1}{\partial r} \cdots \frac{\partial \psi_w^n}{\partial r} \right] \cos N\theta \quad \left[\frac{\partial \psi_w^1}{\partial r} \cdots \frac{\partial \psi_w^n}{\partial r} \right] \sin N\theta \right]$$

$$\phi_v(r, \theta) = \left[\left[\frac{\psi_w^1}{r} \cdots \frac{\psi_w^n}{r} \right] \sin N\theta \quad \left[\frac{\psi_w^1}{r} \cdots \frac{\psi_w^n}{r} \right] \cos N\theta \right]. \tag{5.8}$$

Actually, fewer in-plane shape functions (used $n = 10$) than out-of-plane functions are needed for convergence of the eigenfrequencies of interest. The first five assumed shape functions for each of the u, v, and w displacements are shown in Figure 5.1. Each function has been unit normalized both in the plot and in the model.

Numerical Integration

Different numerical integration schemes were utilized for the calculation of the system matrices involving assumed displacement or strain shapes. Each was chosen to provide the best trade-off between implementation complexity, evaluation speed, and numerical accuracy. For each algorithm, enough integration points were used to achieve numerical convergence to several significant figures.

Integration over the radial direction was accomplished via the Romberg Integration method [Burden, 1993]. Though the most difficult to implement, the method quickly converges
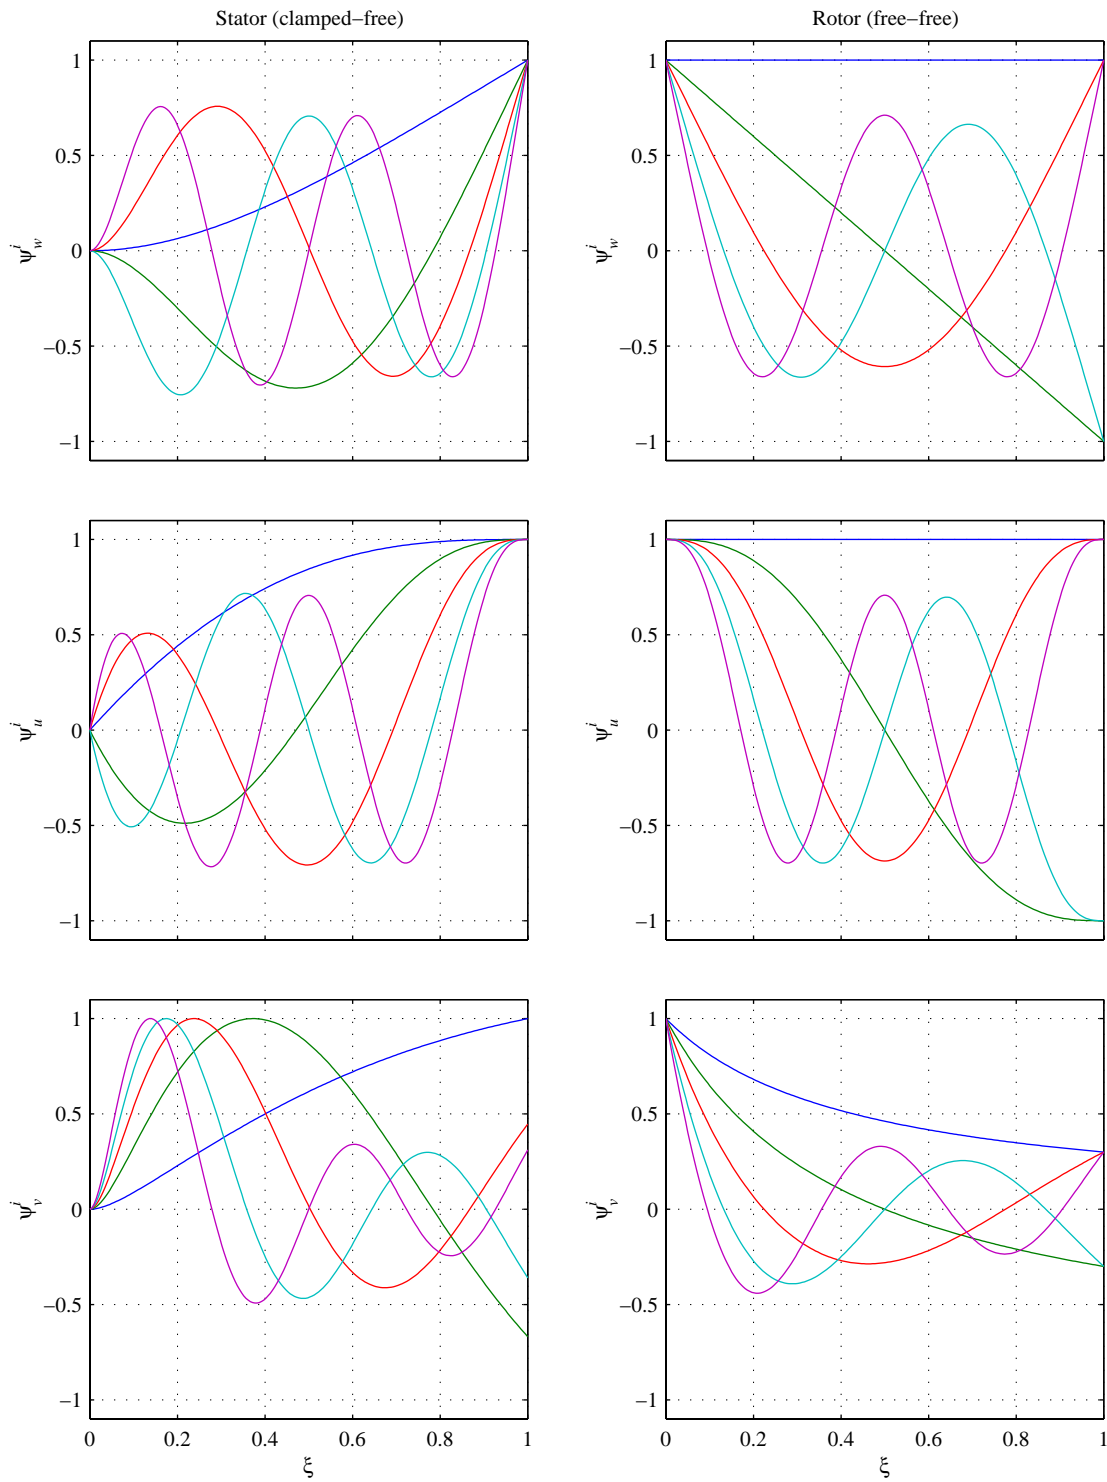
Figure 5.1 First five assumed shape functions for each of the u, v, and w displacements.   Shapes   are shown for both the clamped-free stator and free-free rotor. Amplitudes are unit normalized.

with considerably fewer integration points than simpler integration schemes when a large number of beam functions are used.

The much simpler Simpson's Rule was used instead for integration in the circumferential direction. Periodicity of the sinusoidal functions and geometric axisymmetry allow the integration to be represented by as little as a quarter wavelength. Indeed, the quadratic Simpson's Rule is well suited for integrating sinusoidal functions; only a few integration points are needed for convergence over the fractional wavelength. Where the integrand involves squares of the shape functions, i.e. the **M** and **K** matrices, integration over θ can actually be evaluated exactly at a single point by taking advantage of the fact that the integral of $\cos^2 N\theta$ or $\sin^2 N\theta$ over an integral number of periods is equal to 1/2. Consequently, the Simpson's Rule need only be applied to the calculation of Θ.

Integration through the thickness, however, does not require numerical methods as the shape functions do not depend on z. Instead, the integrations were done analytically upon manipulating the matrix formulations to group each power of z. The resulting formulations were then numerically integrated over r and θ via the above schemes.

Disk Asymmetry

If the *cause* of unwanted asymmetry is known, say a localized solder mass or misalignment during manufacturing, the integrations in (2.33) and (2.34) can be used directly to account for such asymmetries during formulation of the system matrices. However, the factors contributing to preferential alignment of the modes and splitting of the eigenvalues can be very difficult to identify, let alone quantify. This is the case with all the in-house stators: though they appear to be as geometrically symmetric as the Shinsei USR60, if not more so, consistent mode separation has persisted throughout all the manufacturing iterations.

On the other hand, the *effect* that asymmetry has on the modes is quite apparent and, more importantly, quantifiable. Rather than comprehend the cause of the asymmetry, it is relatively more straightforward to account for the effect of the asymmetry. For the following implementation of this approach to be valid, it is assumed that the asymmetry is mild

enough that the mode shapes are not noticeably affected and that orthogonality of like modes is still depicted by 90º spatial phasing. Put another way, the asymmetry is assumed to manifest itself simply by the splitting in frequency of each degenerate eigenvalue pair and a preferential alignment of the mode shapes to axes defined by the asymmetry.

First, consider resonant frequency separation. Upon reduction of the modal equations, the eigenvalue matrix for an ideal component is given by

$$\Lambda = \begin{bmatrix} \lambda & 0 \\ 0 & \lambda \end{bmatrix}, \tag{5.9}$$

where the repeated eigenvalue, λ, is complex and the Rayleigh-Ritz calculated resonant frequency is given by $f^{R\text{-}R} = \Re e(\sqrt{\lambda})$. Loss is represented by the imaginary component of the eigenvalue, and an equivalent modal damping coefficient can be defined by $\zeta = \Im m(\lambda)/\Re e(\lambda)$, though this relationship is valid only at resonance. In practice, the paired resonant frequencies do not coincide but, rather, are separated by some finite difference, $\Delta f^{non\text{-}ideal}$. To capture this effect of frequency separation in the non-ideal simulations, the repeated eigenvalues could be scaled independently to yield unique resonant frequencies of $f^{R\text{-}R} \pm \Delta f^{non\text{-}ideal}/2$.

However, further updating of the eigenvalue matrix is desired at this point to account for errors in predicting the ideal resonant frequencies. Specifically, in order for a direct comparison study of simulated data versus experimental data to be possible, it is argued that the simulation resonant frequencies should match the best known values. Otherwise, not only would the appropriate simulation drive frequencies not coincide with those predetermined during the experimental performance tests, but also errors in the relative spacing between the stator and rotor resonant frequencies would introduce further ambiguity as to which simulated data should be compared to any given experimental data. This decision to manually update the calculated values is arguably justified by the fact that more accurate resonant frequency predictions could be achieved with better displacement assumptions such as those accompanying higher-order laminated plate theories. The mode shapes, however, would remain relatively unchanged.

To these ends, each of the eigenvalue matrices have been updated using the following matrix operation:

$$\Lambda^{updated} = \mathbf{W}^\mathsf{T}\Lambda\mathbf{W} \quad \text{where} \quad \mathbf{W} = \begin{bmatrix} \dfrac{f^{ideal} + \Delta f^{non\text{-}ideal}/2}{f^{R\text{-}R}} & 0 \\ 0 & \dfrac{f^{ideal} - \Delta f^{non\text{-}ideal}/2}{f^{R\text{-}R}} \end{bmatrix} . \quad (5.10)$$

In the case of the stator, f^{ideal} refers to the mean of the experimentally measured values and $\Delta f^{non\text{-}ideal}$ to their difference. Actual measurements, however, were not readily obtainable for the passive rotors. Instead, the most suitable updates for f^{ideal} have been chosen to be the converged finite-element model predictions in Section 3.1.5, and $\Delta f^{non\text{-}ideal}$ has been presumed zero. The frequency measurements of the metal test disks in Section 3.3.2 indicated that splitting of the stator modes was likely due to anisotropy of the phosphor bronze material. As negligible anisotropy was exhibited in the aluminum sample, it is logical to conclude that mode splitting is less of an issue with the rotors.

Second, mode splitting in the offending stator results in preferential orientation of the eigenmodes to principle axes that are not necessarily aligned with the electrode arrays. This is equivalent to imposing a circumferential angular phase difference, $\Delta\theta$, between the vibration shape functions, Φ_s, in (2.22) and the potential shape functions, Φ_φ, in (2.27). As a matter of numerical convenience, it was decided to maintain orientation of the vibration shape functions with the axes of integration and shift the potential shape functions, i.e. the electrode patterns, by $\Delta\theta$. In the ideal case where vibration modes align with the electrode arrays, there is no cross-coupling and $\tilde{\Theta}$ is diagonal. However, as the vibration axes rotate out of alignment with the arrays, orthogonality is lost and $\tilde{\Theta}$ becomes fully populated.

5.1.2 Two-sided Operation

To reduce confusion, the model formulation in Section 2.2 was worded and expressed as if only a single rotor were in contact with the stator. However, the model can easily be imple-

mented in one of two ways to handle the case of two-sided operation: (1) include a second full set of rotor degrees of freedom and evaluate all interface force integrations over both contact surfaces or (2) assume symmetry of the rotor responses and simply double the appropriate interface forces. Naturally, the latter approach is considerably less computationally intensive and is the one used in the two-sided motor simulations.

The two rotors are assumed to be identical and symmetric about the midplane of a symmetric stator. Both rotors are pressed to the stator with the same axial preload, and they are constrained to rotate together. Under these conditions, the interface force profiles are identical at each of the two contact surfaces but with a rotational offset of a half wavelength. The modal forces acting upon the stator are then twice those of a one-sided motor. Hence, the integration in (2.38) can be evaluated over a single contact surface, S_c, then scaled by a factor of two to give

$$\mathbf{f}_{sc}(t) \,=\, 2 \int \int_{S_c} \mathbf{N}_u^{(s)\mathsf{T}}(-\mathbf{f}_c(r, \theta, t)) r dr d\theta \,. \tag{5.11}$$

The modal forces acting upon a single rotor, on the other hand, are evaluated over a single contact surface in either case, so $\mathbf{f}_{rc}$ does not require scaling.

Similarly, the axial rigid-body interface forces react in series with the axial preload, so the calculation of f_{int} in (2.41) is valid as is. However, the rotational rigid-body torques react in parallel to equilibrate the external torque load in (2.40). Either τ_{load} should be designated as half the load applied to the shaft or the equation can be taken to represent the rotation of the coupled rotor system. In the case of the latter,

$$\tau_{int}(t) \,=\, 2 \int \int_{S_c} f_v(r, \theta, t) r^2 dr d\theta \,. \tag{5.12}$$

The two representations are identical, differing only in the placement of the factor of two.

5.1.3 Simulation Details

As with constructing the system matrices, implementing the model into a numerical simulation requires a number of choices. Before describing the solution method, spatial and temporal discretization must be addressed regarding numerical evaluation of the interface force integrations. Accuracy and speed of the solution depend on the spatial resolution of the contact areas and the time step involved in calculating the hysteretic friction force via the time-marching routine. Then, the infinite Fourier series representations must be replaced by reasonably finite series to keep the number of variables within reason yet still capture the relevant dynamics. Studies show that, indeed, very few harmonics are needed to maintain simulation accuracy. Finally, with all the simulation details in place, an iterative routine for solving the system of nonlinear equations can be discussed.

Spatial and Temporal Discretization

Calculation of the interface forces in (2.37) is done in the time domain over a surface in two dimensions. The modal states are first inverse transformed over two periods in time, as explained in Section 2.2.7, and then expanded back to displacements at the contact surfaces. For numerical calculation of the forces, therefore, these steps require mapping the states over a discrete mesh in both space and time.

A time-marching routine was presented in Section 2.2.3 for the calculation of the hysteretic friction forces. The appropriate time step size depends on the stiffness of the interface and frequency content of the interface dynamics. A good balance between computation speed and convergence was achieved by a time discretization of 250 points per period in each of the simulations.

As described in Section 2.1.2, the polymer layer is modeled by a mesh of distributed normal and tangential springs and dampers. That results from the spatial discretization of the contact surfaces. In the ideal case, each point along the circumference follows the same trajectory but with differing phase, so in accordance with the Nyquist Criterion, only a few integration points per wavelength are required for exact calculation of the modal force.

However, each point follows a unique trajectory in the non-ideal case, and a larger number of circumferential integration points are required for convergence. In practice, fifty points per wavelength provided a good balance between convergence and computation time.

Because the radial dimension of the contact surface is so small compared to the circumference, the forces are often evaluated in a line at a single radius. However, the mode shapes are generally very steep at these radii, and contact pressures can vary dramatically over the contact width. Therefore, it is important to calculate the interface forces over the full contact area in both directions. Radial integration was done over seven discrete points.

Circumferential periodicity of the displacements and forces allows the spatial interface forces to be calculated over a single wavelength. They are then converted to modal coordinates by (2.38) before transformation back to the frequency domain. Numerical integration about the circumferential direction is done simply by the Trapezoidal Rule whereas by Simpson's Rule over the radial direction.

Fourier Series Truncation

Obviously the Fourier series expansions of the complex variables cannot be infinite in the numerical implementation of the motor model, nor do they have to be to attain excellent accuracy in the simulation. Therefore, valid truncation of the infinite series must be addressed in relation to the dynamics and frequency content of the solution. Aggressive truncation is preferable to minimize computation, but significant terms must be retained to maintain simulation accuracy.

In the case of a perfect traveling wave excited about the stator via sinusoidal voltage signals, i.e. $\mathbf{v}(t) = \mathbf{V}e^{i\omega t}$, the spatial profiles of the contact forces remain constant in the rotating frame of the wave. Mathematically, this means that the forces can be expressed as functions of the traveling-wave variable $e^{i(\omega t - k\theta)}$. With this in mind, it is easy to show that the contact forces in modal form are harmonic solely at the drive frequency, i.e. void of static terms and higher harmonics. For example, considering the simplest case of a traveling wave in one dimension, the arbitrary contact force, f, can be written

$$f(x, t) = f(e^{i(\omega t - kx)}).$$ (5.13)

As a power series,

$$f(x, t) = \sum_{n=0}^{\infty} f_n \cdot (e^{i(\omega t - kx)})^n = \sum_{n=0}^{\infty} f_n e^{in(\omega t - kx)}.$$ (5.14)

Then, transforming the spatial force to modal form for the mode $\phi(x) = \cos kx$,

$$\tilde{f}(t) = \int \phi(x) f(x, t) dx$$

$$= \int \cos(kx) \left\{ \sum_{n=0}^{\infty} f_n \cdot (\cos(nkx) - i\sin(nkx)) e^{in\omega t} \right\} dx$$ (5.15)

$$= \sum_{n=0}^{\infty} f_n \cdot \left\{ \int (\cos(kx)\cos(nkx) - i\cos(kx)\sin(nkx)) dx \right\} e^{in\omega t}$$

Utilizing the orthogonality of non-identical sines and cosines, the integration reduces to

$$\tilde{f}(t) = f_1 e^{i\omega t}.$$ (5.16)

By this argument, the modal contact forces for the ideal traveling-wave motor are void of higher harmonics, so any Fourier series expansions can be truncated to the fundamental term without loss of information. Thus, dropping the superscript notation for convenience ($n = 1$ is assumed), (2.65) and (2.75) reduce to, in the ideal case,

$$\mathbf{Z}_s = (\Lambda_s - \omega^2 \mathbf{I})^{-1} (\tilde{\Theta} \mathbf{V} + \tilde{\mathbf{F}}_{sc})$$
$$\mathbf{Z}_r = (\Lambda_r - \omega^2 \mathbf{I})^{-1} \tilde{\mathbf{F}}_{rc}$$ (5.17)

and

$$\langle P_{in} \rangle = \frac{1}{2}\Re e(i\omega \mathbf{V}^H \tilde{\Theta}^T \mathbf{Z}_s) + \frac{1}{2}\Re e(i\omega \mathbf{V}^H \mathbf{C}_p \mathbf{V}).$$ (5.18)

In general, however, the higher harmonic terms are non-zero, becoming increasingly significant for greater standing-wave content. To demonstrate this, simulations of the Shinsei USR60 under typical operating conditions were carried out for varying degrees of stator mode separation and rotation. For this study, the first nine harmonics ($n = 1 \rightarrow 9$) were included in the transformation of the modal forces to the frequency domain. The resulting Fourier coefficients of the two stator modes are shown in Figure 5.2, normalized by the magnitude of the fundamental coefficients in the ideal case ($\Delta f = 0$, $\Delta\theta = 0$). As expected, the solution consists only of odd harmonics ($n = 7$ & 9 have been omitted from the plots for clarity).

In Figure 5.2a, the difference in resonant frequency, Δf, between the two stator modes is varied from the ideal degenerate case to a relatively high degree of eigenvalue separation, i.e. $0 \rightarrow 500$Hz in this study. Spatial alignment with the electrode arrays is maintained. The most significant result of increasing separation is divergence of the two modal amplitudes, as seen by the dominant fundamental coefficients. On top of that, higher harmonics come into play due to the interaction of the non-ideal wave with the nonlinear contact dynamics. Nevertheless, the amplitudes of the higher harmonics remain small relative to the fundamental for reasonable separation. Experimentally observed mode separation for the in-house manufactured stators was fairly constant around a couple hundred hertz, where in Figure 5.2a the 5th harmonic is barely distinguishable. Therefore, it was decided that truncation to the first few odd harmonics provides a sensible balance between simulation accuracy and computation efficiency.

Then, in Figure 5.2b, mode separation is instead held constant at 200Hz while the angular orientation, $\Delta\theta$, of the modes with respect to the electrode arrays is swept through a 1/4 wavelength. This amounts to mode rotation over the full width of a single electrode sector. At $N\Delta\theta = \pi/2$, the modes realign with the opposing arrays, so this is identically an ideal orientation. At $N\Delta\theta = \pi/4$, the modes are equally out of phase with either array; resulting in what was expected to be the worst-case scenario.
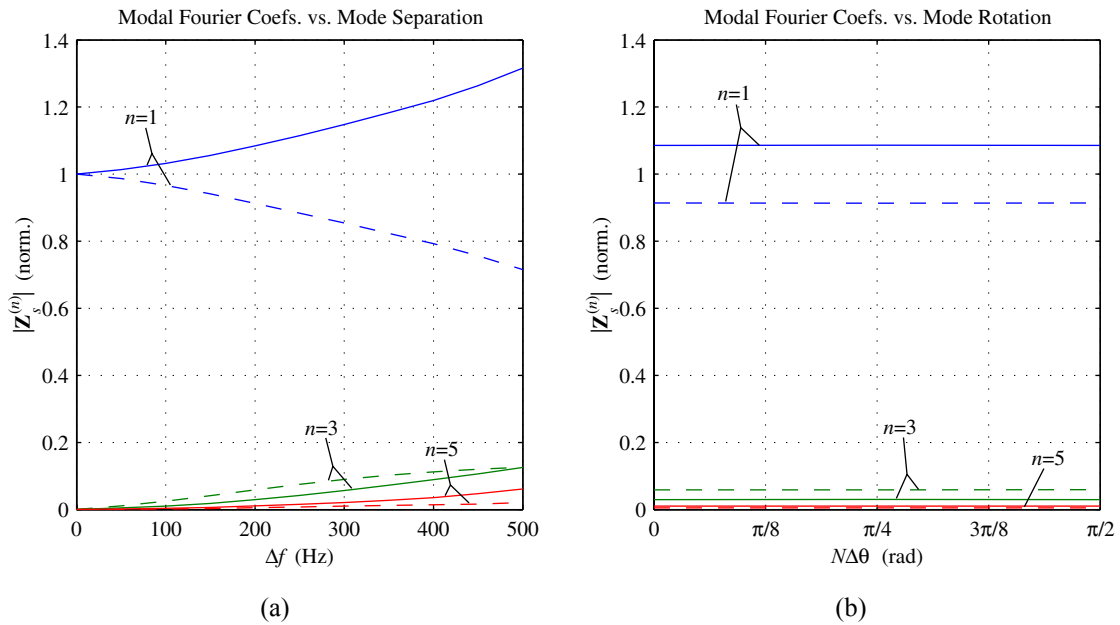
Figure 5.2 Amplitudes of the modal Fourier coefficients as a function of (a) mode separation ($\Delta\theta = 0$) and (b) mode rotation ($\Delta f = 200$Hz), normalized by the ideal case: — high mode, – – low mode. Case study is for the USR60 with 0.125" rotor at 40.8kHz, 150V_p and no load.

Surprisingly, however, varying the angular orientation has no effect on the magnitudes of the modal coefficients, even when the modal eigenvalues are split. Rather, preferential misalignment of the modes due to stator asymmetry merely changes the phase of the progressive wave in relation to the voltage inputs. Analytically, the effect is to rotate the Fourier coefficients in the complex plane while maintaining identical magnitudes. This is trivial to show for ideally degenerate modes, but a rather complex analysis is necessary to prove this for the general case of separate eigenvalues. The results of this simulation study will suffice as indication that preferential misalignment need not be a factor in the model nor a point of worry in manufacturing.

Nonlinear Equation Solver

To solve a system of nonlinear equations such as (2.76), a number of iterative search methods are available. In this MATLAB implementation, the root of (2.76) is found with the function `fsolve`, an iterative algorithm based on a least squares method for simultaneously solving multiple equations of multiple unknowns. Convergence has best been

achieved by utilizing the Gauss-Newton method option with 'quadcubic' line search, a safe-guarded mixed quadratic and cubic polynomial interpolation and extrapolation method.

A limitation to the `fsolve` algorithm, however, is that it only handles real variables. Accordingly, the complex variables and equations in (2.76) must be split into real and imaginary parts. In the ideal case where only the fundamental Fourier terms are non-zero, the system is contained to four complex and two real equations of four complex and two real unknowns. Upon conditioning for the solver, then, the ideal system grows to ten equations and unknowns.

The system of equations is generally much larger in the non-ideal case where each complex variable is expressed as a Fourier series. Fortunately, as discussed above, each series can be truncated aggressively to a very manageable number of harmonics. For series truncated to N terms, (2.75) represents $4N$ complex and two real equations of just as many unknowns. Again, separating the real and imaginary parts, the problem to be solved by `fsolve` consists of $8N+2$ equations and unknowns. As only the first three odd harmonics were kept during the simulations, this meant solving for twenty-six real variables. A flowchart of the solution process was depicted in Figure 2.2.

5.2 Input Parameters

A significant fraction of the experimental preparation efforts in Chapter 3 was dedicated to compiling the necessary set of material properties, many of them complex, for input into the motor model. Though a handful of properties have been retained from the literature, most were characterized in-house using various measurement techniques to provide the most uniquely relevant values possible given the anticipated conditions of the materials during motor operation. The fruits of these efforts are summarized in Tables 5.1–5.3 for the piezo-ceramic, the metals, and the polymer, respectively. These are the material parameters used throughout the simulation case studies. For each of the tabulated properties, the source of the data in Chapter 3 is identified for quick reference.

TABLE 5.1 Input Piezoceramic Properties

Property	Symbol	Units	Value for PZT-4 @ 30°C and:				Source
			$100V_p$	$150V_p$	$175V_p$	$200V_p$	
Permittivity	$\varepsilon_{33}^{T}{}'/\varepsilon_0$		1260	1300	1320	1340	Figure 3.31
Dielectric loss tangent	$\tan\delta$	%	2.1	2.6	2.9	3.3	(high-field)
Piezoelectric strain constant	$d_{31}{}'$	10^{-12} m/V	−107	−108	−108	−108	Figure 3.32
Piezoelectric quality factor	Q_p		135	120	115	110	(high-field)
Mechanical compliance	$s_{11}^{E}{}'$	10^{-12} m²/N	12.75				Figure 3.30
Mechanical quality factor	Q_m		580				(low-field)
Poisson ratio	ν		0.31				Table 3.6
Density	ρ	kg/m³	7500				(Staveley)

TABLE 5.2 Input Metal Properties

Property	Symbol	Units	Nominal Value for:		Source
			554 Phosphor Bronze	6061 Aluminum	
Young's modulus	E' $(1/s_{11}^{E}{}')$	GPa	103.5	69	Table 3.7
Mechanical quality factor	Q_m		3500	2000	Section 3.3.2
Density	ρ	kg/m³	8964	2720	Section 3.3.2
Poisson's ratio	ν		0.33	0.33	Table 3.7

TABLE 5.3 Input Polymer Properties

Property	Symbol	Units	Value @ 40°C, 40kHz for:	Source
			20% Ekonol/PTFE Blend	
Elongation storage modulus	E'	GPa	0.7	Figure 3.36
Elongation loss modulus	E''	GPa	0.07	(extrapolated)
Shear storage modulus	G'	GPa	0.26	Figure 3.35
Shear loss modulus	G''	GPa	0.033	
Density	ρ	kg/m³	1950	Table 3.9
Coefficient of friction	μ		0.14	Section 3.3.3

Note that the polymer properties have been reduced to 40°C instead of the controlled oper-
ating temperature of 30°C. Though somewhat arbitrary in its selection, it was hypothesized
that the significant friction and viscoelastic loss mechanisms at the contact interface would
logically result in a higher operating temperature of the polymer layer than that measured
remotely. This would be difficult to validate experimentally and merits future investigation.

In comparison to the material properties, the remaining model parameters were understand-
ably trivial to set. These included accurate representations of the three-dimensional stator
and rotor geometries and straightforward input of the drive voltage amplitudes and static
axial preloads to mimic the conditions of the experimental test cases. Appreciation of the
specified drive frequencies, however, requires further knowledge of the calculated modal
characteristics investigated in the following section. Specifically, only after appropriate
scaling of the eigenvalue matrices was it deemed appropriate to run the simulations at drive
frequencies identical to those during the experimental testing.

5.3 Calculated and Updated Modal Characteristics

With all the material properties and geometric parameters specified, the first course of
action is to process the Rayleigh-Ritz formulation of the system matrices and to transform
the equations of motion to uncoupled modal coordinates. These procedures were developed
analytically in Sections 2.2.1 & 2.2.4, and a detailed discussion on their implementation
was given in Section 5.1.1. In the end, the calculated modal characteristics of the stator and
rotor are defined by the resulting eigenvalue and eigenvector matrices. The resonant fre-
quencies of the system are related to the eigenvalues, and the associated mode shapes can
be constructed by multiplying the assumed shape function matrix, $\Phi(r,\theta)$, by their respec-
tive eigenvectors.

5.3.1 Mode Shapes

The calculated mode shapes of the mock Shinsei and two-sided motor components are
depicted in Figures 5.3 & 5.4, respectively. In the upper plots, each of the three centerline

(a)

(b)
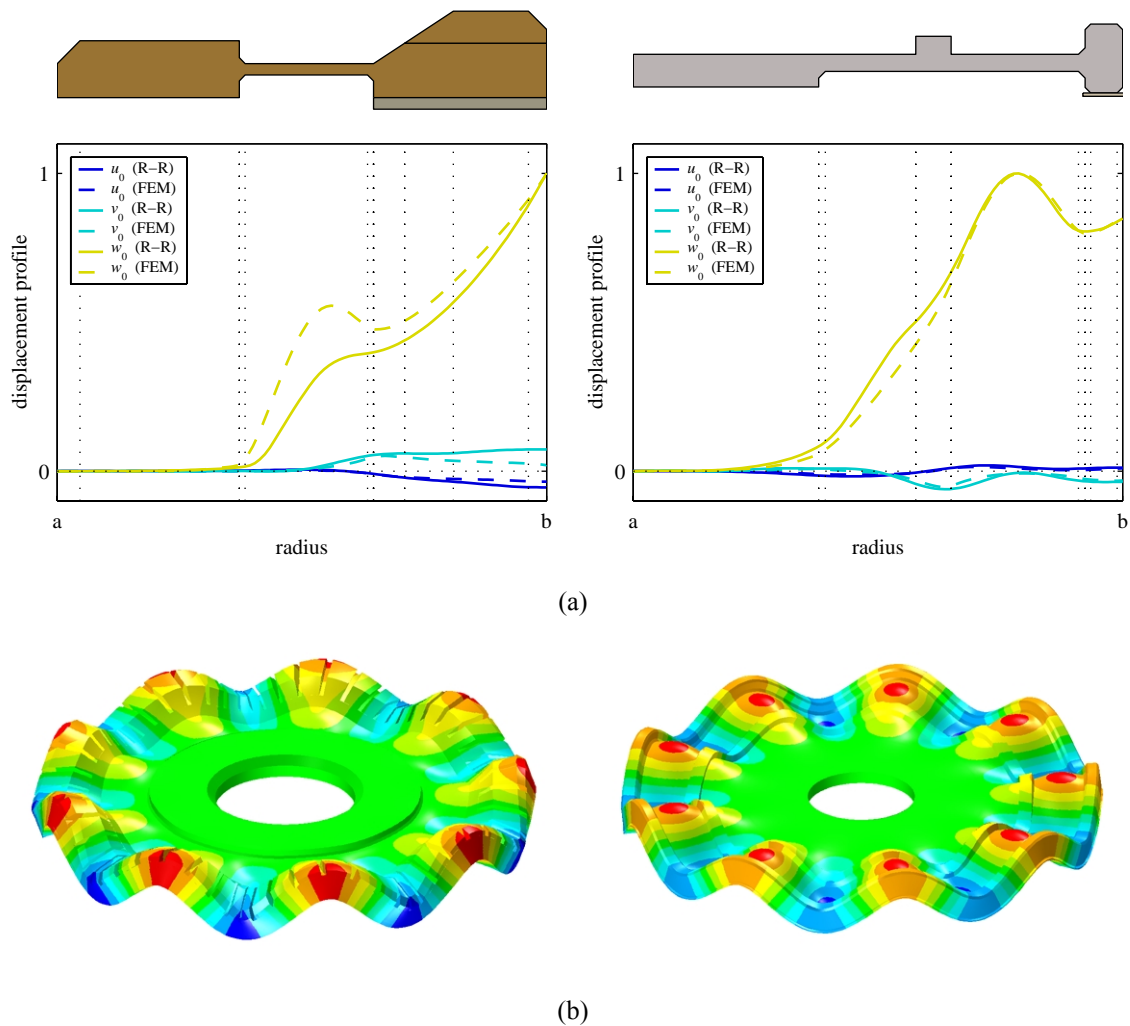
Figure 5.3 Rayleigh-Ritz calculated mode shapes of the mock Shinsei stator (left) and 0.125" rotor (right): (a) radial dependence of the centerline displacement profiles compared to the those obtained from converged finite element models, and (b) 3-D renderings of the full displacement fields superposed upon the component geometries (contours reflect w-displacement).

displacement shapes are plotted as a function of radius. The amplitudes have been unit normalized relative to the maximum out-of-plane displacement, and radial cross sections of the respective components are displayed above the plots for geometric reference. Below the radial shape plots, the three-dimensional displacement fields have been generated from the centerline displacement shapes according to (2.16) and then superposed upon the component geometries for realistic, though exaggerated, representations of the stator and rotor

(a)

(b)
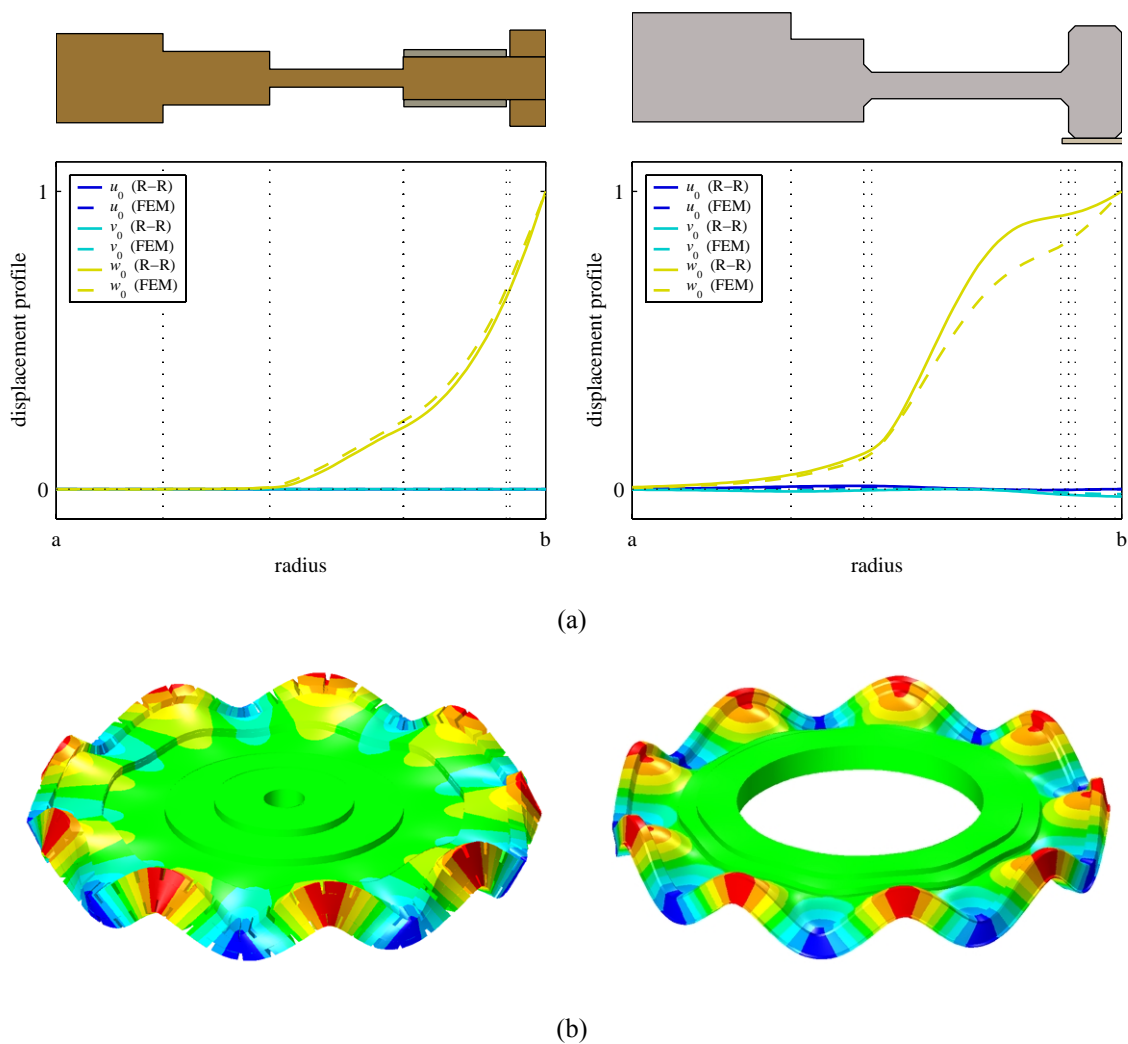
Figure 5.4 Rayleigh-Ritz calculated mode shapes of the two-sided stator (left) and 0.170" rotor (right): (a) radial dependence of the centerline displacement profiles compared to the those obtained from converged finite element models, and (b) 3-D renderings of the full displacement fields superposed upon the component geometries (contours reflect w-displacement).

vibration modes. Note that the overlaid color contours in the 3-D renderings are indicative of the out-of-plane displacement, i.e. $w(r,\theta,z)$.

To qualify the accuracy of the Rayleigh-Ritz formulation and the assumed Kirkhhoff displacement relations (zeroth-order laminated plate theory), the results of converged finite-element modal analyses have been added to the radial shape plots for comparison. Excellent agreement is demonstrated for the mock Shinsei's rotor and the two-sided motor's sym-

metric stator, whereas minor discrepancies can be seen for their counterparts. The greatest deviation is exhibited by the mock Shinsei stator as a result of neglecting tooth stiffness in the stator stiffness matrix calculations. Though the inertia of the teeth was fully accounted for in the mass matrices, the frequency of the slits and protruding nature of the teeth led to the decision to assume rigid body motion of the teeth, i.e. zero strain. In light of the agreement in Figure 5.4(a) for the symmetric stator prototype, the assumption seems logically grounded. However, in contrast to the blocky teeth of the symmetric stator, the teeth of the mock Shinsei stator span a significantly greater fraction of the radius, and the assumption of zero radial strain begins to lose its validity. Consequently, a mode shape is predicted which sees less deflection of the inner flange and a slightly overestimated slope of the tooth surface.

The limitations of the assumed Kirkhhoff displacement relations become more noticeable for the thick lipped rotors like the one shown in Figure 5.4(b). Because shear is not modeled by the relations, the complex strain field at the sharp transition from the thin flange to the thick lip cannot be accurately captured. A recommendation for future implementations of the Rayleigh-Ritz formulation would be to utilize the improved displacement assumptions of higher-order laminated plate theories that include shear terms.

5.3.2 Resonant Frequencies

The Rayleigh-Ritz calculated resonant frequencies are listed in Table 5.4 for comparison with those obtained from the finite-element analyses and, in the case of the stators, to actual measurements (from Table 3.5). As a result of not including tooth stiffness, the calculated stator frequencies are actually a bit low, understandably most so for the mock Shinsei stator. The rotor frequencies, on the other hand, are noticeably high. This is presumably due to the neglect of shear in the assumed displacement relations and the associated ramification of the upright geometry of the lip cross section as compared to the wide geometry of the stator ring cross section.

TABLE 5.4   Calculated and Updated Resonant Frequencies

Motor	Component	Rayleigh-Ritz $f^{R\text{-}R}$ (kHz)	Measured @ 20°C or predicted by FEM $f_1^{actual}/f_2^{actual}$ (kHz)	f^{FEM} (kHz)	Scaled for 30°C and used to update Λ f^{ideal} (kHz)	$\Delta f^{non\text{-}ideal}$ (Hz)
Mock Shinsei	Stator	38.975	40.031/40.210	39.742	39.960	178
	Rotor – 125"	59.402	—	53.049	52.890	0
One-sided	Stator	28.935	29.651/29.828	29.134	29.621	176
	Rotor – 150"	44.396	—	39.789	39.670	0
	Rotor – 160"	46.924	—	41.764	41.639	0
	Rotor – 170"	49.529	—	43.709	43.578	0
	Rotor – 180"	52.176	—	45.588	45.451	0
Two-sided	Stator	28.935	29.398/29.569	29.134	29.366	170
	Rotor – 170"	50.255	—	44.115	43.983	0

Nevertheless, it was argued in Section 5.1.1 that the calculated eigenvalue matrix would be updated via (5.10) so as to reflect the best known values of the system resonant frequencies. The objective in doing so is to synchronize the modeled and actual uncoupled modal characteristics so that the two coupled systems can be compared legitimately under identical operating conditions, particularly with regard to drive frequency. The calculated mode shapes have been shown to be in good agreement with the results of converged finite-element analyses, so scaling the eigenvalues for this purpose is considered a justifiable, and preferable, alternative to increasing the complexity of the displacement assumptions for better frequency prediction. Similar performance predictions should be achievable with or without the scaling, but synchronization of the modal characteristics makes feasible direct comparison with experimental results for a clearer model validation study.

In accordance with the eigenvalue matrix scaling operation in (5.10), the actual stator frequencies and FEM-calculated rotor frequencies have been restated in Table 5.4 for entry as $f^{ideal} \pm \Delta f^{non\text{-}ideal}/2$. However, note that one additional scaling of the frequencies has been applied to account for operation at 30°C rather than 20°C. As both the stator and rotor are made mostly of metal, their selected frequencies have been scaled proportionately to reflect the respective temperature dependencies as measured for the metal test disks in Figure 3.22.

Specifically, a 0.4% drop has been applied to each of the actual stator frequencies and 0.3% to the degenerate FEM rotor frequencies. The eigenvalue matrices were then updated to reflect these resonance characteristics in each of the following simulation studies.

5.4 Simulated Contact

Although model validation will be judged solely on the prediction of motor performance metrics, it's prudent to inspect the states and forces being calculated at the contact interface for verification that the contact dynamics are being captured as intended. For this purpose, a series of simulated contact results are plotted in Figure 5.5 for varying torque load. This particular case study focuses on the one-sided motor with 0.170" rotor driven at 31.2kHz and $150V_p$. So as to streamline the results for coherent, static plotting, the simulations were carried out for an ideal traveling wave, i.e. $\Delta f^{non\text{-}ideal}$ was set to zero. In the presence of an ideal wave, the contact profiles maintain their shapes in the rotating frame, and, consequently, the x-axes in Figure 5.5 can be construed as either a single period in time or a single wavelength about the circumference (for the leftward traveling wave). Needless to say, this would not hold true in the non-ideal case and would make static portrayal of the profiles more difficult to comprehend. Animations have been generated for the non-ideal case, however, that provide some interesting insight into the contribution of the standing-wave component. Nevertheless, several of the key observations from the simulated contact results can be seen in this simplified example.

Radial Variation. The three rows of plots in Figure 5.5 illustrate the following contact profiles: (top) out-of-plane displacements, w_s and $w_r + h_r$, of the stator and rotor, respectively; (middle) normal contact force, f_w, vs. overlap, Δw, of the stator and rotor displacements; and (bottom) tangential contact force, f_w, vs. tangential deformation of the polymer, v_c. The three columns represent different torque loading conditions: (left) no load; (middle) two-thirds of stall where efficiency happens to peak; and (right) at the limit of stall. As in all the simulations, the contact surface was integrated at seven discrete radial locations, but to reduce clutter, only curves for the inner, center, and outer contact radii are shown in each
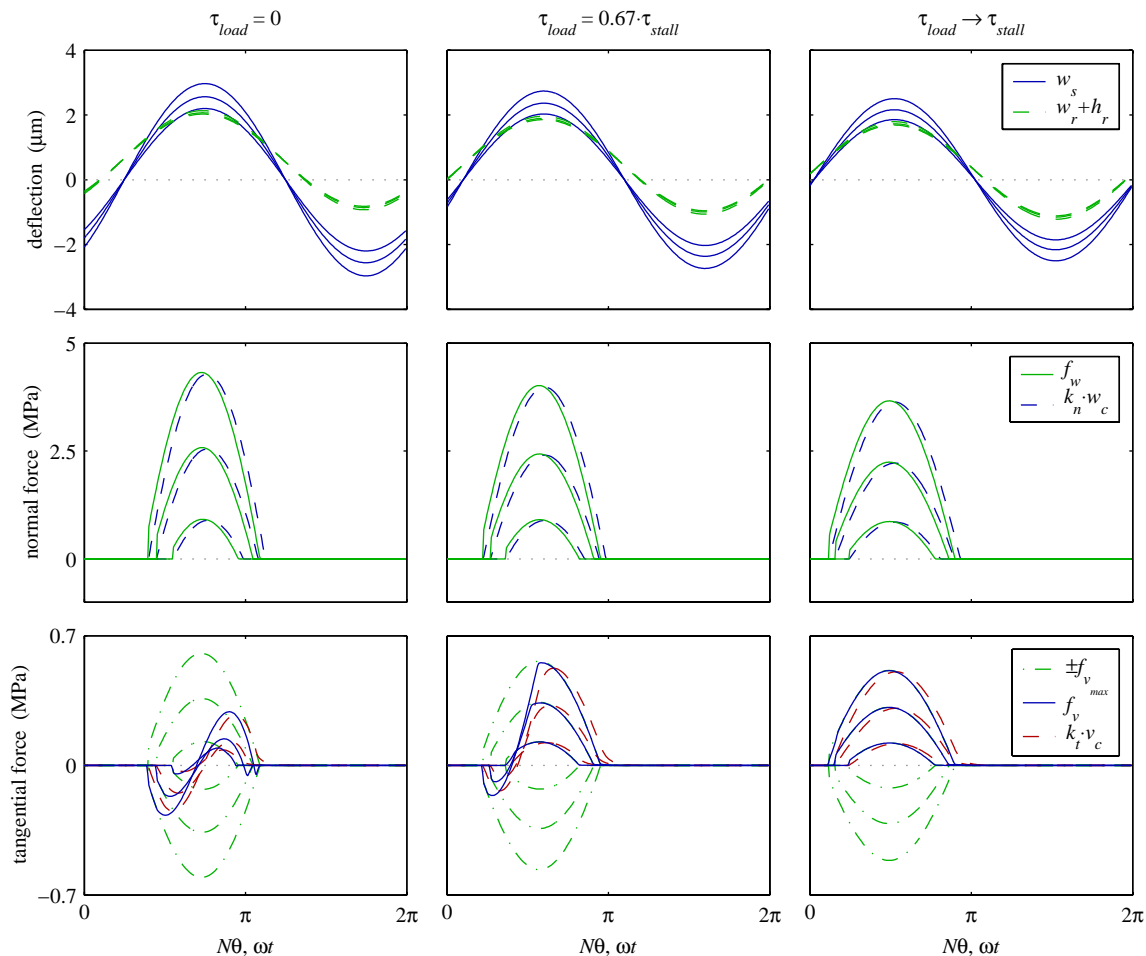
Figure 5.5 Simulated contact profiles for varying torque load. Case study is the one-sided motor with 0.170" rotor at 31.2kHz and $150V_p$ for ideal traveling-wave excitation.

plot. Clearly, the contact state and force profiles vary significantly over the small radius of contact, thus demonstrating the importance of simulating contact in two dimensions.

Rotor Dynamics. Of primary noteworthiness, it is seen that the rotor does indeed respond dynamically to the periodic contact forces and, furthermore, that its vibration amplitude is substantial to say the least. In this particular example, the rotor is shown to vibrate at about half the amplitude of the stator, similar to experimental findings of Kawai *et al.* for the optimal scenario [Kawai, 1995]. As the drive frequency is increased away from stator resonance, the ratio has been shown in the simulations to grow even larger. This can be

explained by the observation that while the stator amplitude decreases for increasing frequency, the rotor amplitude is governed by approximately quasistatic response in the useful frequency range and, thus, stays relatively constant. This may also explain why the relative phasing between the rotor and stator vibrations remained small throughout the simulation studies.

Load Dependency. As torque is increased through to stall, the stator amplitude experiences noticeable attenuation. This is accompanied by widening of the contact area and a slight decrease in the normal contact pressures. Though this trend may not seem particularly remarkable, it's a clear demonstration of the interaction between the applied static forces and the dynamic states and why the amplitudes cannot legitimately be prescribed. Curiously, however, the rotor maintains a more constant amplitude than the stator as a function of torque, again likely due to the implication of quasistatic excitation discussed above.

Viscous Damping. A number of distinguishing features can be observed in the normal and tangential force profiles that are associated with the viscoelastic interface model. For example, the simulated normal and tangential forces each lead their respective polymer deformation states as a result of the viscous loss mechanisms introduced in either direction (see Section 2.1.2). Consequently, each of the forces also experience a sudden jump at the leading edge of contact (left edge). Observed phase between a force and its associated displacement means that not all the stored elastic energy is returned. As captured by the viscoelastic model, some of this energy is dissipated in the polymer. The significance of the viscoelastic loss is quantified in relation to each of the other loss mechanisms later in Section 5.6.

Stick-Slip. Finally, as purposed by the modeling of tangential compliance in the polymer layer, distinct stick-slip regions have been captured in the simulations. To help identify these, the limits of the friction force, $\pm f_{v_{max}}$, have been added to the tangential force plots in Figure 5.5: where $|f_v| = f_{v_{max}}$, the interface is slipping, and where $|f_v| < f_{v_{max}}$, the interface is sticking. At no load, a broad stick region is observed, and the force profile takes on an asymmetric balance. As the torque load is increased, the latter slip region grows to

accommodate a greater integration of positive traction forces so as to react the increasing load. Ultimately, as the torque load approaches the stall limit and all available traction force is utilized, the slip region grows to encompass the entire contact surface.

As one can imagine, there exist many potential avenues for studying the contact problem in greater detail. For example, many of the observed performance trends for varying drive frequency or rotor stiffness could reasonably be explained to some degree by their influence on the contact profiles. A considerable amount of time could also be spent investigating the more complex contact problem of the non-ideal case. However, much of this is left for future studies, and attention is focused here on the main objective of determining the performance prediction capabilities of the model in the following case studies.

5.5 Simulated Performance Case Studies

So as to avoid unknown variables, the simulation case studies have been limited to the three in-house constructed motors, i.e. those known to be made of the materials characterized in Chapter 3. In other words, simulated performance will not be compared to the measured performance of the authentic Shinsei USR60 motor because no samples of its piezoceramic composition were available for testing, and the use of PZT-4 properties would undoubtedly produce inaccurate results. However, this shortcoming was the motivation for constructing the mock Shinsei stator as an acceptable substitute. In fact, its flawed stator asymmetry is the subject of the first case study on performance degradation due to mode splitting. As each of the in-house stators exhibited asymmetry, examining the performance loss due to non-ideal conditions seemed the logical place to start. The following two studies involve simulating the one-sided prototype to investigate the accuracy of varying voltage and rotor stiffness in the model. In the final study, the procedures described in Section 5.1.2 have been implemented to enable simulation of the two-sided prototype.

In each of the studies, the operating conditions have been specified to coincide with those of the experimental measurements for the purpose of a fair and equitable comparison. This was made possible by scaling the eigenvalue matrices to reflect the updated frequencies in

Table 5.4. As a matter of thoroughness, specific details of the numerical simulations, such as discretization or truncation, were discussed earlier in Section 5.1.3. The remaining discussions, therefore, focus on the results themselves.

5.5.1 Degradation Due to Mode Splitting

Modal characterization of the mock Shinsei stator in Section 3.2 identified a notable separation in frequency of the two $B_{0,9}$ modes that was not exhibited by the Shinsei USR60 stator. It has been stated throughout this work that splitting of the modes degrades the quality of the excited traveling wave and undoubtedly reduces the potential performance of the motor. However, the degree to which this characteristic flaw hinders performance would be difficult to examine experimentally. Instead, a specific objective of this thesis has been targeted at the generalization of the model formulation to account for non-ideal wave generation and at the prediction of the loss through simulation. To this end, the flawed mock Shinsei stator with 0.125" rotor has been chosen as the testbed for which to run the first comparative simulation study. Specific attention will be given in this study to the difference in performance characteristics between the ideal and non-ideal wave cases.

While specifying drive conditions identical to those of the experimental measurements presented in Figure 4.9, simulations were performed for both the ideal case ($\Delta f^{non\text{-}ideal} = 0$) and the appropriate non-ideal case ($\Delta f^{non\text{-}ideal} = 178\,\text{Hz}$). The results are compared in Figure 5.6 along with the actual measurements. Immediately observed is how remarkably good the model-experiment correlation is. In particular, not only are both the shapes and spacing of the various calculated curves consistent with experiment, but the absolute speed-torque characteristics are also predicted with excellent accuracy. Still, the simulations yield input powers that are lower than measured, resulting in inflated predictions of the efficiency. The error may not be in the simulations, however, but rather due to the unexplained excess damping of one of the mock Shinsei's modes when clamped at the stator hub (see Section 3.2). As material parameters used in the simulations were based on independent measurements, the source of this excess damping was not captured by the simulation. Con-
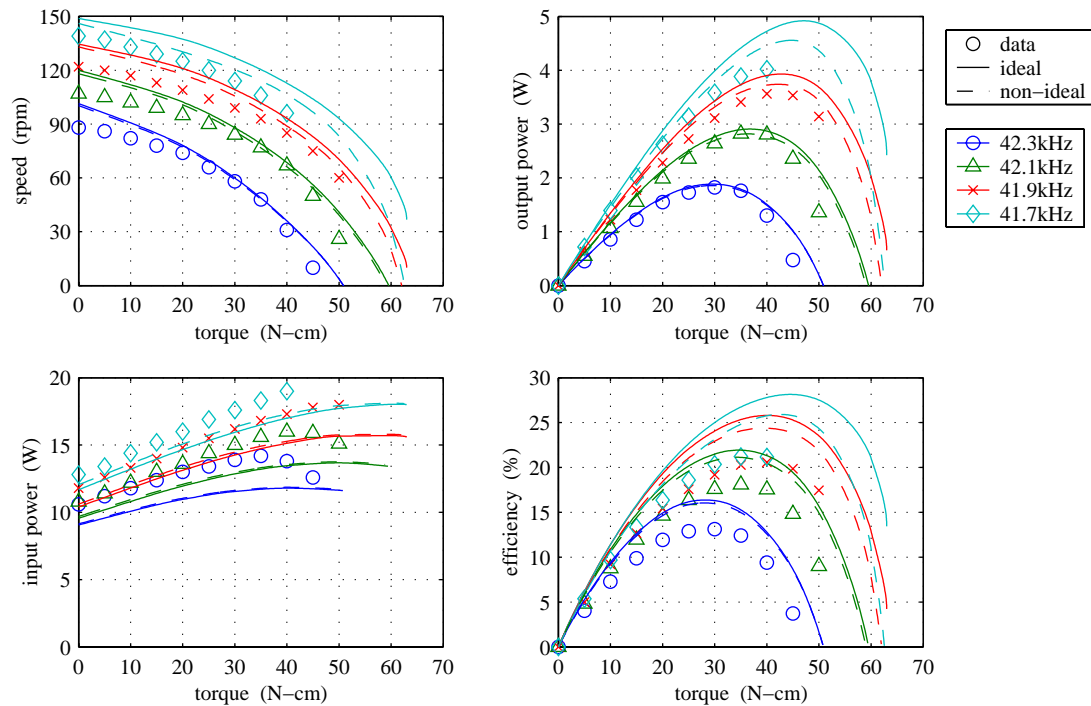
Figure 5.6 Simulated vs. measured performance of the mock Shinsei stator with 0.125" rotor at $150V_p$.

sequently, the realized input power is likely higher than what the model was allowed to predict given the supplied information. The results of the later case studies help to support this conclusion.

Such good correlation with the data promises great potential for this modeling approach as an accurate design tool, but almost as remarkable is how the ideal and non-ideal simulations compare. As it turns out, the extent of mode splitting observed in each of the in-house stators may not have as much negative effect on motor performance as originally thought. In fact, the ideal and non-ideal curves are almost indistinguishable for the higher drive frequencies. This high-frequency convergence is, of course, understandable given that the difference in phase response of the separated modes diminishes as the drive frequency is shifted away from resonance. Nevertheless, even the lowest frequency curves do not differ significantly. At 41.7kHz, the decreased speed and increased power draw of the non-ideal case are realized by only a 2% drop in peak efficiency.

The most noticeable result of split modes appears to be a premature tapering of speed near stall for the highest drive levels, i.e. for drive frequencies approaching resonance. Notice how the non-ideal speed curve for 41.7kHz diverges near stall from the ideal curve and actually asymptotes toward the non-ideal curve of the next higher frequency. It was already demonstrated through experimental testing that this particular motor combination could not be operated at 41.7kHz across the full torque range. Whether or not the contact instability that comes into play at this high drive level is a related phenomenon would be an interesting issue for a future study. However, the fact that the relatively ideal Shinsei USR60 stator was limited similarly for this same rotor combination (see Figure 4.7) indicates a likely coincidence.

In spite of the outlying divergence, the non-ideal and ideal simulation curves aren't remarkably different in the valid ranges of operation, i.e. those populated by measured data. In addition to the impact this realization may have on future manufacturing tolerancing, it's a welcome verdict as it pertains to these computational efforts. Specifically, the non-ideal simulations require significantly more computation time than the ideal simulations and can often be difficult to converge because of the increased number of simultaneous equations and unknowns for which to solve. With such little difference in the two theoretical results, the added cost of computation time discourages the need to simulate mild asymmetry.

To put things in perspective, as stated in Section 5.1.3, the ideal system consists of ten equations of ten unknowns. For this, typical simulation times run about thirty seconds per point on a 1GHz AMD Athlon processor. However, for the non-ideal case, the system of equations grows as higher harmonics must be included. Even with the system truncated to just three odd harmonics, the number of unknowns for which to iteratively solve jumps to twenty-six. With this many degrees of freedom passed to the nonlinear, iterative search algorithm, the solver is dramatically more prone to getting hung on false solutions. Furthermore, if and when satisfactory initial conditions are finally identified that do lead to a converged solution, the simulation times are rather non-ideal themselves at around thirty minutes per point.

It can be argued that relatively little accuracy is lost with the ideal simulation predictions in the presence of mild stator asymmetry. Therefore, as long as potential asymmetry is limited via sufficient manufacturing tolerances, commercial applications of this modeling approach would benefit greatly by avoiding the computational burden of accommodating the non-ideal case. Nevertheless, both ideal and non-ideal wave conditions have been simulated in the each of the following case studies for academic thoroughness. Consistent trends observed throughout the case studies support a likely correlation between the onset of contact instability and the combined frequency-torque threshold where the non-ideal case significantly diverges from the ideal case.

5.5.2 Effect of Varying Drive Voltage

To investigate the model accuracy when varying drive voltage, the one-sided prototype with the 0.170" rotor was simulated at a fixed drive frequency of 31.4kHz for three different drive voltages: $100V_p$, $150V_p$, and $200V_p$. The results are compared in Figure 5.7 to the equivalent experimental performance data from Figure 4.10. As with the previous study, the shapes and spacing of the predicted curves are in excellent agreement with the real measurements. Even in the simulations, speed is only minorly sensitive to the drive voltage.

However, in terms of absolute accuracy, both speed and input power were predicted consistently high. Given the accuracy of the previous study, a possible explanation for this could be a slight mismatch in the actual and modeled resonant frequencies. The prototype motors are only slightly larger than the Shinsei configuration, yet they dissipate notably more absolute input power as heat. As the temperatures of the mock Shinsei and one-sided motors were both monitored at the stator hub, it is reasonable to speculate that the true operating temperature of the outer ring may have been higher in the one-sided motor at this high drive level. A lower stator resonance than considered would have resulted, and any simulations founded on the higher resonant frequency would over-predict both speed and power. Whatever the source of error may be, efficiency has been predicted comparatively well in this study.
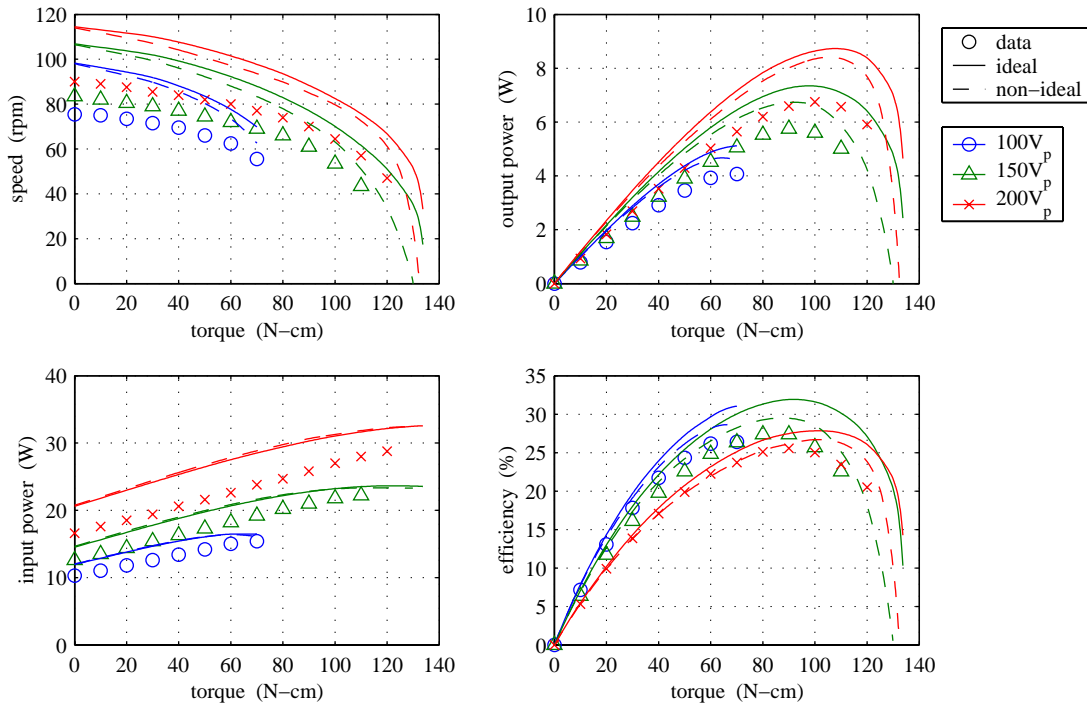
Figure 5.7 Simulated vs. measured performance of the one-sided motor with 0.170" rotor at 31.4kHz.

A rather curious result of the simulations can be seen in the $100V_p$ curves of Figure 5.7: beyond a specified load of 70N-cm, the iterative solver simply would not converge upon a solution for any of the various supplied initial conditions. Coincidentally, the actual motor would not operate at $100V_p$ above this torque either. A thorough examination of the contact profiles at this threshold might provide insight into the premature stall phenomenon at low drive levels. However, determining whether these two events are at all related is left as a recommendation for future investigation.

5.5.3 Effect of Varying Rotor Stiffness

With the most exhaustive set of experimental rotor study data among the in-house assembled motors (see Figures 4.11–4.14), the one-sided motor was again selected for this simulated rotor study. A drive voltage of $150V_p$ was chosen to correlate with the experimental study in Section 4.3, and simulations were performed at various frequencies for each of the 0.150", 0.160", 0.170", and 0.180" rotors. The results are presented in Figures 5.8–5.11,

Figure 5.8   Simulated vs. measured performance of the one-sided motor with 0.150" rotor at $150V_p$.

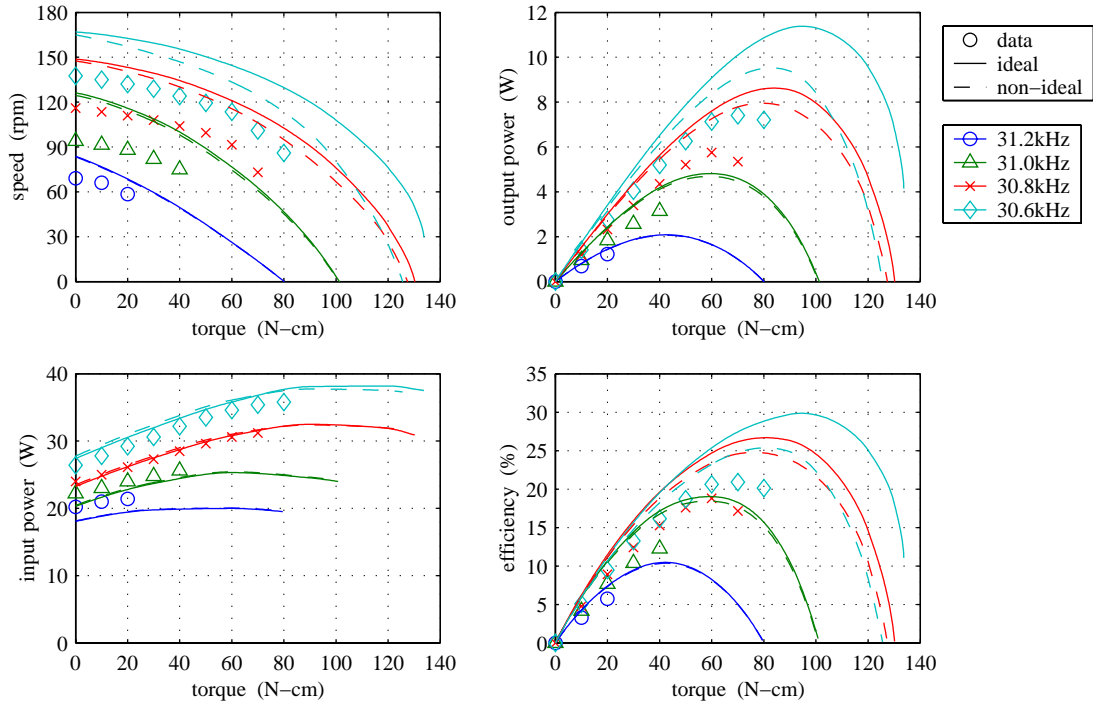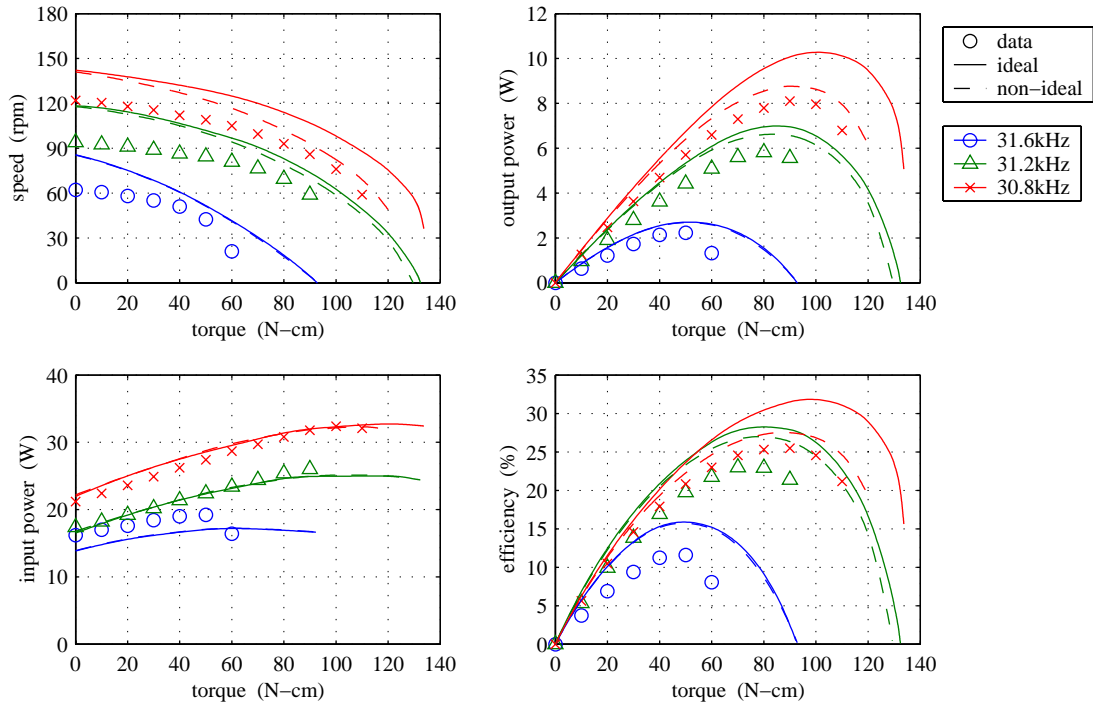**Figure 5.9**   Simulated vs. measured performance of the one-sided motor with 0.160" rotor at $150V_p$.
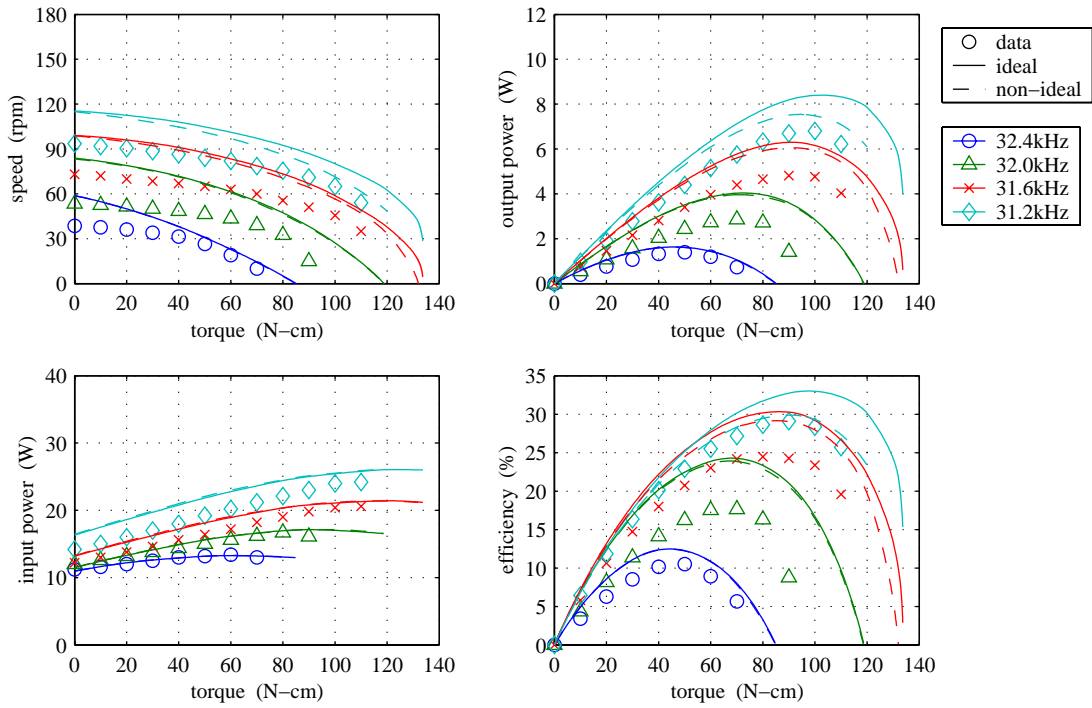
Figure 5.10   Simulated vs. measured performance of the one-sided motor with 0.170" rotor at $150V_p$.
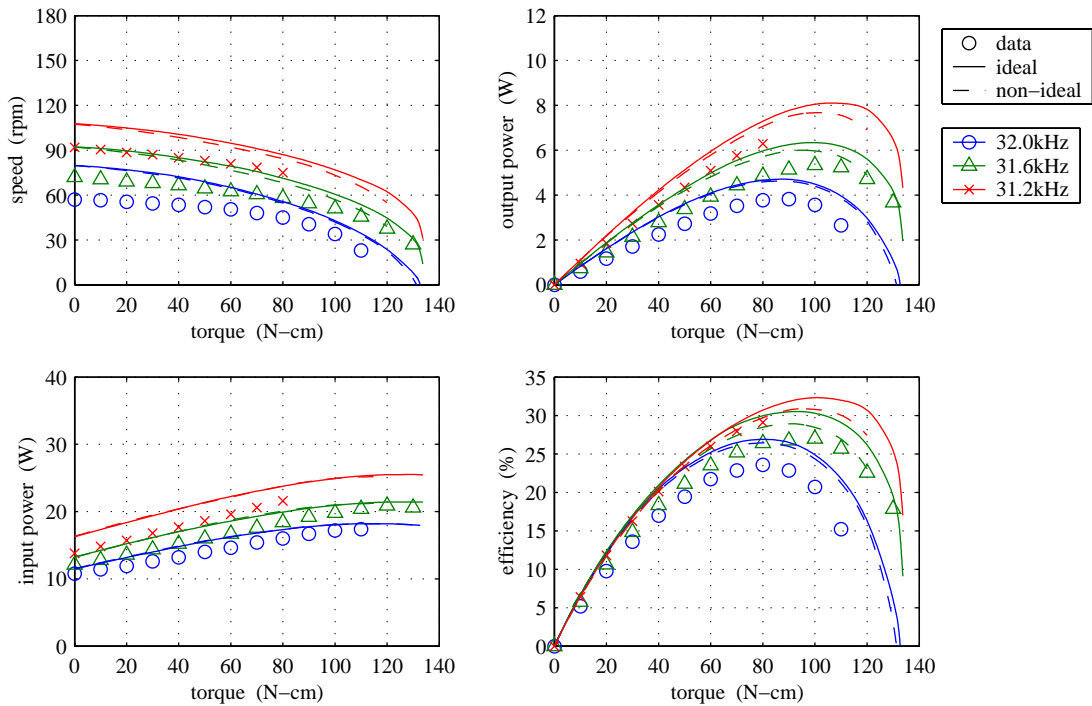

Figure 5.11   Simulated vs. measured performance of the one-sided motor with 0.180" rotor at $150V_p$.

respectively, for comparison with each particular set of measured performance data. Upon review, a clear trend in error is perceived as a function of the rotor thickness.

Looking at the 0.170" and 0.180" rotor data in Figures 5.10 & 5.11, the simulation accuracy for the thicker rotors appears to be consistent with the observed trends of the two previous studies. In particular, excellent correlation of the speed and input power data is achieved for low drive levels while a slight, yet proportionate, overestimation of both is observed for high drive levels. The simulated efficiencies for the these rotors achieve respectable correlation with the measurements. Unfortunately, as seen somewhat for the 0.160" rotor in Figure 5.9 and more so for the 0.150" rotor in Figure 5.8, speed is increasingly overestimated in proportion to the input power as the rotor thickness is decreased. Furthering the error is a broadening of the spacing between the input power curves to the effect of underprediction at low drive levels. As a result, simulated efficiency does not decline as dramatically with decreasing rotor thickness as was observed experimentally.

Though the cause of the increased error in simulated performance for the more flexible rotors is not understood, it is hypothesized that the validity of the contact model may be weakened as the rotor/stator amplitude ratio approaches unity. The absolute stator amplitude under typical operation is greater as well for the thinner rotors, so it is also possible that some developing amplitude nonlinearity is not being captured by the linear dynamics model. Nevertheless, even for the extreme case of the 0.150" rotor with its partially realized torque range, moderate correlation between the simulations and measurements is achieved throughout the range of rotor thicknesses.

5.5.4 Two-Sided Operation

The last of the performance simulation case studies investigates the potential of the two-sided motor configuration. As observed by the experimental test results in Section 4.4, the two-sided prototype did not perform as well as expected. Specifically, although speeds were achieved that slightly exceeded those of the one-sided motor, the operable torque range fell exceedingly short of twice that of the one-sided motor. Therefore, upon imple-

mentation of the appropriate model modifications outlined in Section 5.1.2, simulations of the two-sided prototype with 0.170" rotors at $175V_p$ were carried out to provide a baseline for the believed potential of two-sided operation. Figure 5.12 compares the results of these simulations with the respective experimental measurements from Figure 4.17.

As hypothesized, the simulated performance curves of the two-sided motor exhibit twice the stall torque and output power of the one-sided motor predictions in Figure 5.10. With a near doubling of the input power to match, the predicted efficiency of the two-sided motor remains on par with that of the similarly configured one-sided motor, if not slightly better. Allowing for the possibility of an apparent 0.4kHz error in matching the actual and speci-fied frequencies due to temperature inconsistencies, excellent agreement is observed for the measured and simulated performance data in the lower part of the predicted torque range. However, correlation of the speed-torque characteristics breaks down above 120N-cm due to a premature drop in measured speed. As a result, each of the peak performance metrics saturate below their predictions. The cause of the premature torque barrier is not presently
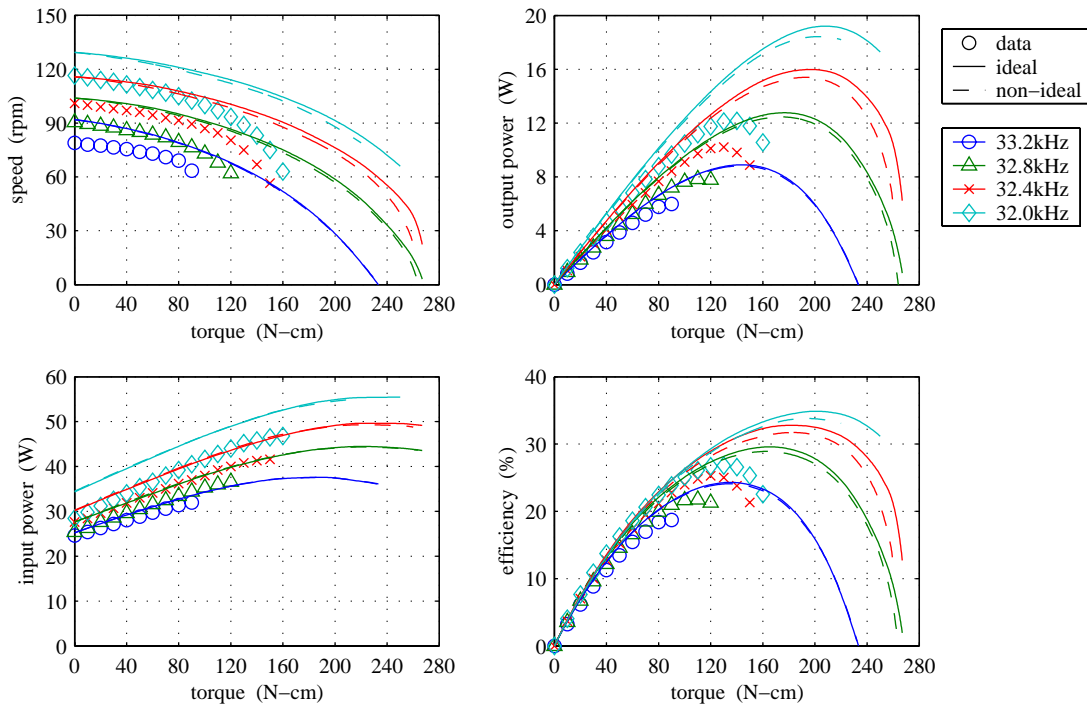

Figure 5.12   Simulated vs. measured performance of the two-sided motor with 0.170" rotors at $175V_p$.

understood, but with such good model-experiment correlation at lower torques, there is evidence that supports the potential for much greater performance from the two-sided motor implementation.

Extension of the experimental performance at 32.0kHz along the 32.4kHz simulated curves provides a baseline for the believed potential of the two-sided prototype motor. Of course, this assumes that the premature torque barrier would be surmountable upon further investigation. Normalizing the foreseen stall torque of 260N-cm and peak output power of 15W by its mass of 350g, the two-sided prototype could conceivably achieve a torque density of 740N-cm/kg and a power density of 43W/kg. These metrics would be outstanding compared to the performance of the other existing traveling-wave motors listed in Table 1.1. Furthermore, the 32.4kHz simulation suggests that the two-sided motor could achieve a peak efficiency greater than 30%, i.e. more than the one-sided motor. This is understandable given that the piezoelectric and stator volumes are not doubled in the two-sided configuration. In other words, losses in the two-sided system should be less than twice those of the one-sided system for equal drive and vibrations amplitudes. As a result, not only should the two-sided motor produce twice the torque and output power, it should do so with better efficiency.

5.6 Investigation of Loss Contributions

In Section 2.2.6, the electrical input power was derived from the converged state solution in such a fashion that only two components of the input power were identified. Restated here for reference, the time-averaged input power to the motor system was expressed

$$\langle P_{in} \rangle = \langle P_{electromech.} \rangle + \langle P_{dielectric} \rangle, \tag{5.19}$$

where

$$\langle P_{electromech.} \rangle = \frac{1}{T} \int_0^T \mathbf{v}^\mathsf{T} \tilde{\Theta}^\mathsf{T} \dot{\mathbf{z}}_s dt = \frac{1}{2} \Re e \left(\sum_{n=0}^{\infty} \mathbf{V}^{(n)\mathsf{H}} \tilde{\Theta}^\mathsf{T} (in\omega \mathbf{Z}_s^{(n)}) \right), \tag{5.20}$$

and

$$\langle P_{dielectric} \rangle = \frac{1}{T}\int_0^T \mathbf{v}^\mathsf{T}\mathbf{C}_p\dot{\mathbf{v}}dt = \frac{1}{2}\Re e\left(\sum_{n=0}^{\infty} \mathbf{V}^{(n)\mathsf{H}}\mathbf{C}_p(in\omega\mathbf{V}^{(n)}) \right). \qquad (5.21)$$

The latter term in (5.19) represents the fraction of electrical power that is dissipated in the electrical regime due to the lossiness of the dielectric. This leaves $\langle P_{electromech.} \rangle$ as the remaining fraction that is converted via the piezoceramics to mechanical power. Though the calculation of input power above is efficient, it does not provide much insight into the distribution of losses in the mechanical regime. Therefore, an expanded power analysis is presented here for identification and quantification of each of the mechanical loss contributions.

Utilizing first the stator modal equation of motion, (2.58) is rearranged to isolate the electromechanical forcing term:

$$\tilde{\Theta}\mathbf{v} = \ddot{\mathbf{z}}_s + \Lambda_s\mathbf{z}_s - \tilde{\mathbf{f}}_{sc}. \qquad (5.22)$$

Then, premultiplying by $\dot{\mathbf{z}}_s$ and applying a time average integral over one period yields the following power relationship within the stator body:

$$\frac{1}{T}\int_0^T \dot{\mathbf{z}}_s^\mathsf{T}\tilde{\Theta}\mathbf{v}dt = \frac{1}{T}\int_0^T \dot{\mathbf{z}}_s^\mathsf{T}\ddot{\mathbf{z}}_s dt + \frac{1}{T}\int_0^T \dot{\mathbf{z}}_s^\mathsf{T}\Lambda_s\mathbf{z}_s dt - \frac{1}{T}\int_0^T \dot{\mathbf{z}}_s^\mathsf{T}\tilde{\mathbf{f}}_{sc}dt. \qquad (5.23)$$

The term on the left has already been defined by $\langle P_{electromech.} \rangle$ in (5.20) as the fraction of electrical input power actually available for driving the mechanical system. The first term on the right represents an inertial term and integrates to zero. The remaining two terms are identified and expanded as

$$\langle P_{stator\ vib.} \rangle = \frac{1}{T}\int_0^T \dot{\mathbf{z}}_s^\mathsf{T}\Lambda_s\mathbf{z}_s dt = \frac{1}{2}\Re e\left(\sum_{n=0}^{\infty} (in\omega\mathbf{Z}_s^{(n)})^\mathsf{H}\Lambda_s\mathbf{Z}_s^{(n)} \right), \qquad (5.24)$$

and

$$\langle P_{stator\ contact} \rangle = -\frac{1}{T}\int_0^T \dot{\mathbf{z}}_s^\mathsf{T} \tilde{\mathbf{f}}_{sc} dt = -\frac{1}{T}\int_0^T \dot{\mathbf{z}}_s^\mathsf{T} (\mathbf{V}_s^\mathsf{T} \mathbf{f}_{sc}) dt = -\frac{1}{T}\int_0^T (\mathbf{V}_s \dot{\mathbf{z}}_s)^\mathsf{T} \mathbf{f}_{sc} dt$$

$$= -\frac{1}{T}\int_0^T \dot{\mathbf{p}}_s^\mathsf{T} \mathbf{f}_{sc} dt = -\frac{1}{T}\int_0^T \dot{\mathbf{p}}_s^\mathsf{T} \left(\iint_{S_c} \mathbf{N}_u^{(s)\mathsf{T}}(-\mathbf{f}_c) dS_c \right) dt$$

$$= \frac{1}{T}\int_0^T \left(\iint_{S_c} (\mathbf{N}_u^{(s)} \dot{\mathbf{p}}_s)^\mathsf{T} \mathbf{f}_c dS_c \right) dt = \frac{1}{T}\int_0^T \left(\iint_{S_c} \dot{\mathbf{u}}_s^\mathsf{T} \mathbf{f}_c dS_c \right) dt \qquad (5.25)$$

$$= \frac{1}{T}\int_0^T \left(\iint_{S_c} \dot{v}_s f_v dS_c \right) dt + \frac{1}{T}\int_0^T \left(\iint_{S_c} \dot{w}_s f_w dS_c \right) dt \ .$$

$\langle P_{stator\ vib.} \rangle$ is the loss due to vibration damping in the stator, and $\langle P_{stator\ contact} \rangle$ is the power flowing out of the stator through the contact interface. In summary, the power relationship within the stator is given by

$$\langle P_{electromech.} \rangle = \langle P_{stator\ vib.} \rangle + \langle P_{stator\ contact} \rangle \ . \qquad (5.26)$$

A similar relationship for power conversion within the rotor body can be derived from the rotor modal equation of motion. Premultiplying (2.59) by $\dot{\mathbf{z}}_r$ and applying a time average integral over one period yields

$$\frac{1}{T}\int_0^T \dot{\mathbf{z}}_r^\mathsf{T} \ddot{\mathbf{z}}_r dt + \frac{1}{T}\int_0^T \dot{\mathbf{z}}_r^\mathsf{T} \Lambda_r \mathbf{z}_r dt = \frac{1}{T}\int_0^T \dot{\mathbf{z}}_r^\mathsf{T} \tilde{\mathbf{f}}_{rc} dt \ . \qquad (5.27)$$

As before, the first term represents an inertial term and integrates to zero, and the latter two are identified and expanded as

$$\langle P_{rotor\ vib.} \rangle = \frac{1}{T}\int_0^T \dot{\mathbf{z}}_r^\mathsf{T} \Lambda_r \mathbf{z}_r dt = \frac{1}{2}\mathfrak{Re}\left(\sum_{n=0}^\infty (in\omega \mathbf{Z}_r^{(n)})^\mathsf{H} \Lambda_r \mathbf{Z}_r^{(n)} \right) , \qquad (5.28)$$

$$\langle P_{rotor\ contact} \rangle = \frac{1}{T}\int_0^T \dot{\mathbf{z}}_r^\mathsf{T} \tilde{\mathbf{f}}_{rc} dt$$

$$= \frac{1}{T}\int_0^T \left(\iint_{S_c} \dot{v}_r f_v dS_c \right) dt + \frac{1}{T}\int_0^T \left(\iint_{S_c} \dot{w}_r f_w dS_c \right) dt . \qquad (5.29)$$

$\langle P_{rotor\ vib.} \rangle$ is the loss due to vibration damping in the rotor, and $\langle P_{rotor\ contact} \rangle$ is the power flowing into the rotor through the contact interface. Therefore, the power relationship within the rotor is given simply by

$$\langle P_{rotor\ vib.} \rangle = \langle P_{rotor\ contact} \rangle . \tag{5.30}$$

The power conversion at the stator-rotor contact interface can be derived by relating $\langle P_{stator\ contact} \rangle$ and $\langle P_{rotor\ contact} \rangle$ via the contact kinematics. Rearranging (2.45) and (2.42) and taking the derivative with respect to time, the contact kinematics are expressed

$$\dot{v}_s = (\dot{v}_r + \omega_r r) + \dot{v}_c + \Delta\dot{v} \quad \text{and} \quad \dot{w}_s = \dot{w}_r + \Delta\dot{w} . \tag{5.31}$$

Substituting (5.31) into (5.25), $\langle P_{stator\ contact} \rangle$ is expanded as

$$\langle P_{stator\ contact} \rangle = \frac{1}{T}\int_0^T \left(\int\int_{S_c} ((\dot{v}_r + \omega_r r) + \dot{v}_c + \Delta\dot{v}) f_v dS_c \right) dt$$
$$+ \frac{1}{T}\int_0^T \left(\int\int_{S_c} (\dot{w}_r + \Delta\dot{w}) f_w dS_c \right) dt, \tag{5.32}$$

so that the following terms can be identified:

$$\langle P_{rotor\ contact} \rangle = \frac{1}{T}\int_0^T \left(\int\int_{S_c} \dot{v}_r f_v dS_c \right) dt + \frac{1}{T}\int_0^T \left(\int\int_{S_c} \dot{w}_r f_w dS_c \right) dt , \tag{5.33}$$

$$\langle P_{normal\ def.} \rangle = \frac{1}{T}\int_0^T \left(\int\int_{S_c} \Delta\dot{w} f_w dS_c \right) dt = \frac{1}{T}\int_0^T \left(\int\int_{S_c} \dot{w}_c f_w dS_c \right) dt , \tag{5.34}$$

$$\langle P_{shear\ def.} \rangle = \frac{1}{T}\int_0^T \left(\int\int_{S_c} \dot{v}_c f_v dS_c \right) dt , \tag{5.35}$$

$$\langle P_{friction} \rangle = \frac{1}{T}\int_0^T \left(\int\int_{S_c} \Delta\dot{v} f_v dS_c \right) dt , \tag{5.36}$$

$$\langle P_{out} \rangle = \frac{1}{T} \int_0^T \left(\int \int_{S_c} \omega_r f_v r \, dS_c \right) dt = \tau_{int} \omega_r = \tau_{load} \omega_r. \tag{5.37}$$

$\langle P_{rotor\ contact} \rangle$ is easily recognized from the previous derivation in (5.29). $\langle P_{normal\ def.} \rangle$ and $\langle P_{shear\ def.} \rangle$ are the viscous losses in the polymer layer associated with the local normal and shear deformations, respectively. In (5.34), the overlap of the stator and rotor normal profiles, Δw, can be substituted by the normal polymer deformation, w_c, because f_w is zero outside the contact region where Δw and w_c differ. The loss due to friction is given by $\langle P_{friction} \rangle$ as this term accounts for slipping at the interface, i.e. Δv. The final term represents the product of the rigid-body interface torque in (2.41) and the constant rotational speed of the rotor. As long as the numerical damping coefficient in (2.66) is negligible, τ_{int} and τ_{load} are assumed equal, and the final term can be expressed as the mechanical output power, $\langle P_{out} \rangle$, of the system. Upon identifying each of the above terms, the power relationship at the contact interface is given by

$$\langle P_{stator\ contact} \rangle = \langle P_{rotor\ contact} \rangle + \langle P_{normal\ def.} \rangle + \langle P_{shear\ def.} \rangle$$
$$+ \langle P_{friction} \rangle + \langle P_{out} \rangle. \tag{5.38}$$

Finally, by combining the four power relationships in (5.19), (5.26), (5.30), and (5.38), the complete breakdown of power utilization reduces to

$$\langle P_{in} \rangle = \langle P_{losses} \rangle + \langle P_{out} \rangle, \tag{5.39}$$

where

$$\boxed{\begin{aligned} \langle P_{losses} \rangle = &\langle P_{dielectric} \rangle + \langle P_{stator\ vib.} \rangle + \langle P_{rotor\ vib.} \rangle \\ &+ \langle P_{normal\ def.} \rangle + \langle P_{shear\ def.} \rangle + \langle P_{friction} \rangle. \end{aligned}} \tag{5.40}$$

The losses are summarized by

$$\langle P_{dielectric} \rangle = \frac{1}{2} \Re e \left(\sum_{n=0}^{\infty} \mathbf{V}^{(n)\mathsf{H}} \mathbf{C}_p (in\omega \mathbf{V}^{(n)}) \right)$$

$$\langle P_{stator\ vib.} \rangle = \frac{1}{2} \Re e \left(\sum_{n=0}^{\infty} (in\omega \mathbf{Z}_r^{(n)})^{\mathsf{H}} \Lambda_s \mathbf{Z}_s^{(n)} \right)$$

$$\langle P_{rotor\ vib.} \rangle = \frac{1}{2} \Re e \left(\sum_{n=0}^{\infty} (in\omega \mathbf{Z}_r^{(n)})^{\mathsf{H}} \Lambda_r \mathbf{Z}_r^{(n)} \right) \qquad (5.41)$$

$$\langle P_{normal\ def.} \rangle = \frac{1}{T} \int_0^T \left(\int \int_{S_c} \dot{w}_c f_w dS_c \right) dt$$

$$\langle P_{shear\ def.} \rangle = \frac{1}{T} \int_0^T \left(\int \int_{S_c} \dot{v}_c f_v dS_c \right) dt$$

$$\langle P_{friction} \rangle = \frac{1}{T} \int_0^T \left(\int \int_{S_c} \Delta \dot{v} f_v dS_c \right) dt$$

Note that whereas the dielectric and vibration losses can be calculated from the frequency-domain states, the polymer deformation and friction losses must be calculated within the time domain of the contact model.

To illustrate the relative contributions of each of the identified losses, a graphical example is presented in Figure 5.13 for varying torque load. The particular case study involves the simulation shown previously in Figure 5.10 for the one-sided motor with the 0.170" rotor at 31.6kHz and $150V_p$ under ideal traveling-wave excitation. On the left, the losses are plotted on an absolute scale along with the output power. On the right, the losses and output power are stacked to show their relative contributions to the total electrical input power.

A number of noteworthy conclusions can be drawn from Figure 5.13. First, the dielectric loss constitutes a significant fraction of the total loss. As the dielectric loss has been captured as a direct result of the use of complex material properties, its unmistakable presence
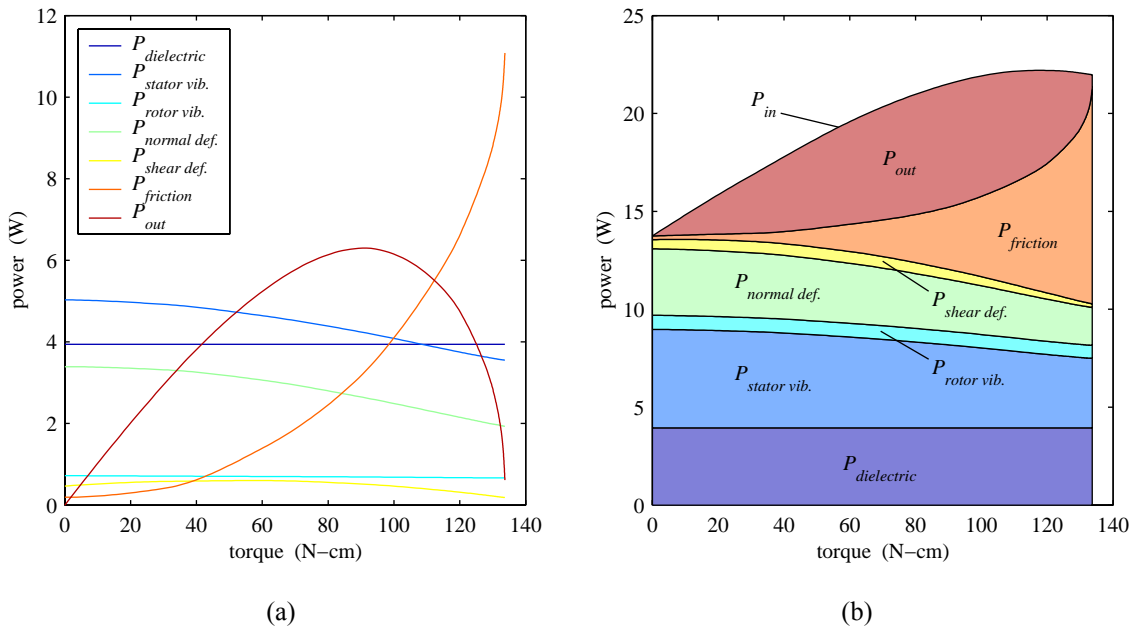
Figure 5.13 Contributions to the input power: (a) absolute, (b) stacked. Case study is the one-sided motor with 0.170" rotor at 31.6kHz and $150V_p$ for ideal traveling-wave excitation.

confirms the inaccuracy of a loss model based solely on the use of lumped modal damping coefficients. The only components of loss that should be captured by a modal damping loss model are the stator and rotor vibration losses. Seen here, these constitute only about a third of the total loss, mostly from the stator due to the higher mechanical damping of the piezo-ceramics. The decrease in stator vibration loss as a function of torque reflects the decrease in stator vibration amplitude seen in Figure 5.5.

The viscoelastic deformation losses of the polymer layer also contribute significantly to the total power dissipation. Most of the deformation loss is incurred in the normal direction and drops with torque as the stator amplitude decreases. The less prominent shear loss tapers toward zero near stall because the onset of uniform slipping leads to unidirectional shear and slower shear transitions in the contact region. Neither polymer deformation loss would have been captured with a purely elastic foundation model.

Finally, the dramatic dependence of the friction loss on torque load is directly related to hav-ing captured the stick-slip behavior of the contact interface. As seen in Figure 5.5, most of

the contact region sticks at no load, so friction losses start out low. As load is increased, the slip region grows, and friction losses become more significant. By stall, friction is shown to consume half the total input power in this particular example. Though a pure-slip contact model would have yielded the same friction loss at stall, it would have severely overpredicted the loss at lower loads. Therefore, the accuracy of a pure-slip contact model cannot begin to compare to the accuracy of a stick-slip contact model when it comes to simulating friction losses.

In summary, the investigation of the loss contributions provides quantifiable evidence of the importance of several elements of this modeling approach. By introducing a complex dielectric constant, a significant contribution of power dissipation in the motor system has been identified that had not been considered before. Local polymer deformation losses have been shown to be comparable to the component vibration losses, thus demonstrating the importance of capturing the viscoelasticity of the polymer layer. A small but valid fraction of overall loss has been attributed to the rotor vibration which otherwise would have been neglected by the assumption of a rigid rotor. Lastly, the magnitude of the friction loss has been shown to vary dramatically with torque load as the contact condition transitions from stick to slip. A pure-slip contact model would not be capable of predicting this behavior. Each of these elements contribute to the ability of this model approach to predict accurately the input power and efficiency of the motor system.

Chapter 6

SUMMARY AND RECOMMENDATIONS

6.1 Thesis Summary

The objectives of the present work were twofold: to develop an advanced modeling approach for the accurate prediction of steady-state performance of piezoelectric traveling-wave motors, and to improve upon the typical piezoelectric traveling-wave motor configuration by investigating the novel implementation of two-sided operation. In response to the growing consideration of traveling-wave motors as suitable replacements for electromagnetic motors, the modeling objective addressed the need for an efficient design tool to complement or even supplant the costly process of prototype iteration. Similarly, to expand the viable commercial application of the traveling-wave motor as a direct-drive actuator, the two-sided configuration has been suggested as a means of doubling torque output without significantly increasing size or mass.

Model Development

Previous modeling endeavors have all relied on the arbitrary adjustment of one or more model parameters to achieve adequate fit of simulated performance with limited experimental measurements. Though useful for qualitative studies or even extrapolation of performance relative to an existing design, the approaches lacked the completeness to provide accurate prediction of absolute motor performance solely from parameters physically attributable to the component materials and geometry. Only a handful of these model formula-

tions actually addressed the fully coupled motor system, while many others focused specifically on the contact dynamics at the stator/rotor interface. Those that did account for the coupled interaction between the states, contact forces, and electromechanical forces unfortunately relied on oversimplified contact models to accommodate the specific formulations. Conversely, those that focused their efforts on developing advanced contact models simply prescribed the global vibration states and ignored the electromechanical forcing of the piezoceramic. Common among all the previous model approaches were the false assumption of a rigid rotor and, where included at all, a lumped representation of system damping. Consequently, none of the previous modeling approaches had fully captured the nature of the piezoelectric traveling-wave motor system and thus were incapable of implementation as an *a priori* design tool.

Recognizing the individual accomplishments and shortcomings of several of the previous approaches, a unified model formulation was developed in Chapter 2 that integrated seamlessly each of the most advanced modeling elements from those approaches with several new developments. Inspired by the work of Hagood and McFarland [Hagood, 1995], a Rayleigh-Ritz approach was used to assemble the modal equations of motion of the motor system including the electromechanical forcing of the piezoceramic actuator arrays. However, the rotor was allowed the same modal degrees of freedom as the stator rather than assuming it as rigid. Furthermore, for the first time ever, an advanced contact model was implemented in conjunction with fully coupled modal forcing terms, i.e. without prescribing the states. Specifically, a hysteretic stick-slip friction algorithm was developed to function with the viscoelastic polymer model of Schmidt *et al*. [Schmidt, 1996].

The success of the model approach was founded on two key factors: (1) the handling of system damping through the use of complex material properties, and (2) the implementation of a mixed-domain formulation. The inability of previous models to predict performance accurately was largely due to their reliance on lumped modal damping coefficients to account for damping in the motor system. As a property of the assembled system, specifying a modal damping coefficient requires measurements from a existing device, precluding

the objective of *a priori* performance prediction. To remedy this shortcoming, complex material property representation was embraced as a means of attributing modal damping directly to the materials themselves. In other words, damping was ascribed as a measurable property of the materials, not the assembled components. This approach to damping was the key contributing model element to enabling performance prediction.

The use of complex properties predestined the formulation of the modal equations of motion in the frequency domain. However, requiring knowledge of the state history for analysis of the hysteretic stick-slip friction model, a mixed-domain solution procedure was developed for efficient analysis of the modal dynamics in the frequency domain while retaining the ease and accuracy of simulating the nonlinear contact dynamics in the time domain. This mixed-domain formulation was the key to integrating seamlessly each of the desired modeling elements into one cohesive model. Where others had sacrificed completeness for analysis in a single domain, the mixed-domain approach promoted completeness by embracing the benefits of both domains.

Last but not least, by representing the frequency-domain modal states and forces as Fourier series expansions and retaining higher harmonic terms through harmonic balance, the model was generalized to account for non-ideal traveling-wave excitation. Previous models had assumed ideal wave conditions and were incapable of capturing performance degradation due to either intentional or unavoidable, real-world non-idealities. As a specific objective of the modeling approach, generalization to the non-ideal case was directed toward an investigation of the performance losses in the experimental prototypes due to observed stator asymmetries.

Motor Development

Insufficient direct-drive torque output from existing piezoelectric traveling-wave motors was identified as a major limiting factor to their consideration in a broader marketplace, particularly the automobile industry. Unfortunately, the use of vibration modes presents a challenge to scaling the technology any larger than what has already been achieved. Whereas

direct scaling would result unacceptably in audible drive frequencies, increasing the vibration mode number to maintain resonance in the ultrasonic range would reduce the likelihood of supporting the necessary degenerate modes. Instead, a potential source of untapped motive force in the typical configuration was hypothesized, and in Chapter 3 the first two-sided piezoelectric traveling-wave motor was developed. This novel implementation of the traveling-wave motor offered the potential to double torque output without increasing the dimension, thus avoiding the drawbacks of increased scale. Moreover, with anticipated speeds and mass on par with those of the standard configuration, the two-sided motor presented the opportunity to produce not only twice the torque density but also twice the power density.

In parallel with the development of the two-sided prototype, an essentially identical one-sided version of the motor was fabricated for a controlled demonstration of the performance benefits of two-sided operation. Accordingly, both prototypes incorporated the same symmetric stator configuration with full piezoelectric actuators arrays bonded to each side, and both stators were mated to dynamically similar rotors. However, aside from the obvious second rotor in the two-sided configuration, the two implementations did differ greatly in their execution. Whereas the stator in the one-sided motor was mounted to the housing, a shaft-mounted stator setup had to be devised for the two-sided motor to accommodate the second rotor. This required the introduction of brushes to deliver the electrical power to the piezoceramics in the rotating frame. Nevertheless, the two testbeds were equally matched for the purpose of the direct comparison study.

The experimental performance results of the one-sided and two-sided motors were presented in Chapter 4. As hypothesized, the two-sided motor significantly outperformed the one-sided motor with regard to peak torque and power output, though at a slight decrease in efficiency. In fact, compounded by a less massive housing requirement, the implementation of two-sided operation achieved 60% greater torque density and almost 100% greater power density than the typical configuration. Unfortunately, the two-sided prototype exhibited a premature torque barrier well below expected stall that was not witnessed in the one-

sided prototype. Its cause was not identified, and, consequently, the two-sided prototype was inhibited from reaching its full performance potential. Nevertheless, the novel implementation demonstrated great promise as a means for increasing the torque output without increasing scale.

In addition to the development of two-sided operation, particular attention was paid to the dynamic design of the rotor. Though previous modeling endeavors had assumed the rotor to be rigid, recent experimental studies had begun to provide evidence to the significance of the rotor's dynamic response. Therefore, to investigate the effect of rotor design on motor performance, an array of rotors with varying lip thickness was fabricated for each motor. The thickness parameter was chosen as a means of varying the rotor resonant frequency in relation to the stator frequency and, in effect, as a means of varying the vibration response of the rotor in relation to that of the stator. Performance optimization was believed possible with proper tuning of the rotor dynamics.

The results of the parametric rotor studies were also presented in Chapter 4. As anticipated, motor performance was shown to be highly sensitive to the dynamic stiffness of the rotor. Whereas higher no-load speeds were generally achieved with the thinner rotors, both greater power draw and lower torque output accompanied decreasing rotor thickness. Conversely, increasing the rotor thickness extended the torque range to the point of maximum stall and drew less electrical input power at the cost of lower operating speeds. Consequently, an optimal rotor thickness could be identified in each case with respect to maximizing output power and efficiency, though not necessarily with the same rotor selection.

For further demonstration of the importance of good rotor design, an array of rotors was also tested in the highly recognized, commercially available Shinsei USR60 motor. Its nominal rotor geometry was reproduced using similar materials, and both thinner and thicker rotors were fabricated for comparison. The study showed that by decreasing the rotor thickness by just 8.7%, efficiency was reduced from the nominal peak of 20% to just 7%. Increasing the nominal thickness by the same amount actually raised efficiency to a peak of 28%.

Although the Shinsei USR60 has a well engineered stator made from carefully selected materials, this study has demonstrated that there is considerable room for improvement in the USR60's performance through more careful attention to the design of the rotor. More importantly, the study provides unmistakable proof of the importance of including the rotor dynamics in the model of a piezoelectric traveling-wave motor.

Model-Experiment Correlation

Model simulations of several of the experimental case studies were presented in Chapter 5. In addition to providing performance predictions under ideal wave assumptions, the modeled eigenfrequencies were scaled independently in each case to mimic the degree of mode splitting observed in the experimental motors. Simulating the non-ideal wave conditions in this manner provided an updated prediction of the performance metrics given knowledge of the actual modal behavior.

Excellent agreement between the experimental measurements and simulated data was demonstrated for a broad range of drive voltages, drive frequencies, and dynamic rotor stiffnesses for three unique motor configurations. Particularly impressive was the consistent correlation of the shapes and relative spacing of the speed-torque and input power characteristics over the full parametric space. This attested specifically to the validity of the fundamental dynamics included in the model approach. Even more exciting was the competent prediction of absolute performance metrics throughout the case studies. Agreement with the measured efficiencies confirmed that the losses in the motor system were being captured accurately by the use of appropriate complex properties. A detailed breakdown of the simulated loss contributions further clarified the relative significance of each of the modeled elements in determining the overall power draw and efficiency.

Correlation was further improved with the updated predictions, particularly at high drive levels, which took into account the non-ideal wave conditions due to stator asymmetry. Still, the results of the simulations indicated that the observed degree of mode splitting did not dramatically degrade motor performance in the valid range of drive frequencies. How-

ever, an investigation of the simulated contact dynamics in the non-ideal case suggested that mode splitting likely influenced the onset of contact instability. Consequently, the lowest valid drive frequency is considered to be directly related to the degree of non-ideality.

As demonstrated by the ability of the model to accurately predict motor performance over a broad range of parameters and operating conditions, the developed approach is decidedly viable for implementation as an effective *a priori* design tool. For the first time, successful correlation of traveling-wave motor performance has been achieved without the use of model adjustment factors or the need for measurements taken from an existing device.

6.2 Contributions

Analytical

- Developed a piezoelectric traveling-wave motor model founded solely on parameters physically attributable to the component materials and geometry, thus enabling all model parameters to be specified *a priori* for the purpose of performance prediction.

- Directly captured multi-field material losses throughout the motor system by adopting complex material property representation, thus overcoming the need to fit simulated data via the arbitrary selection of modal damping coefficients.

- Dispensed of the previously accepted, yet obviously false, assumption of a rigid rotor by accounting for the modal dynamics of the rotor in the model formulation.

- Developed a mixed-domain model formulation and solution technique for efficient analysis of the modal dynamics in the frequency domain while retaining the ease and accuracy of simulating the nonlinear contact dynamics in the time domain. The formulation enabled the first successful merging of fully coupled state-force interaction with an advanced stick-slip contact model.

- Developed a finite-time simulation routine for analyzing the hysteretic, stick-slip contact dynamics within the alternating domain formulation.

- Generalized the model formulation to account for non-ideal traveling-wave excitation by expanding the modal states and forces as Fourier series and applying harmonic balance to the frequency-domain modal equations of motion.

- Compared ideal and non-ideal simulated performance predictions of real-world motors and quantified the performance loss due to mode splitting.

- Demonstrated excellent model-experiment correlation for several case studies over a wide range of parameters, not just a single curve, without the need for model "adjustment factors".

Experimental

- Developed a novel traveling-wave motor implementation employing two-sided operation for the purpose of improving torque density.

- Demonstrated greatly improved torque and power density of a two-sided prototype through direct comparison with an otherwise identical one-sided prototype under essentially identical operating conditions.

- Provided conclusive experimental evidence of the importance of rotor dynamics and the existence of optimal design criteria.

6.3 Recommendations for Future Work

Analytical

- Adapt and validate the model approach for other motor configurations such as the rod actuator traveling-wave motor.

- Improve the Rayleigh-Ritz frequency predictions by accounting for shear in the displacement assumptions, i.e. through application of a higher-order laminated plate theory.

- Investigate other iterative solution algorithms for more robust convergence and improved speed.

- Implement a heuristic routine such as a genetic algorithm for application of the model as a tool for automated design optimization.

Experimental

- Further investigate the two-sided configuration to determine the cause of the premature torque barrier and hopefully extend the torque and power outputs of the two-sided motor to its full potential

- Validate the simulated stator and rotor vibration profiles with experimental data, especially for the non-ideal case.

- Apply a ceramic coating on the polymer interface lining to decouple the compliance and low loss requirements from the wear and friction requirements.

REFERENCES

[Aguiar, 1996] Aguiar, O. D., "Acoustic quality factor of vanadium," *Cryogenics* **36**, No. 1, pp. 53–55, 1996.

[Alemany, 1993] Alemany, C., Pardo, L, Jimenez, B., Carmona, F., Mendiola, J., and Gonzalez, A. M., "Automatic iterative evaluation of complex material constants in piezoelectric ceramics," *Journal of Physics D: Applied Physics* **27**, pp. 148–155, 1993.

[Bar-Cohen, 1998] Bar-Cohen, Y., Bao, X, and Grandia, W., "Rotary Ultrasonic Motors Actuated by Traveling Flexural Waves," *Smart Structures and Materials 1998: Smart Structures and Integrated Systems*, Marc E. Regelbrugge, Editor, Proceedings of SPIE Vol. 3329, pp. 794–800, 1998.

[Barth, 1973] Barth, H. V., "Ultrasonic Driven Motor," *IBM Technical Disclosure Bulletin* **16**, No. 7, p. 2263, December1973.

[Burden, 1993] Burden, R. L. and Faires, J. D., *Numerical Analysis*, 5th ed., PWS-Kent Publishing Co., Boston, 1993.

[Cameron, 1989] Cameron, T. M., and Griffin, J. H., "An Alternating Frequency/Time Domain Method for Calculating the Steady-State Response of Nonlinear Dynamic Systems," ASME *Journal of Applied Mechanics* **56**, pp. 149–154, 1989.

[Cao, 1995a] Cao, X., and Wallaschek, J., "Estimation of the tangential stresses in the stator/rotor contact of travelling wave ultrasonic motors using visco-elastic foundation models," *Proceedings of the 2nd International Conference on Contact Mechanics*, Ferrara, pp. 53–61, 1995.

[Cao, 1995b] Cao, X., and Wallaschek, J., "Modelling and analysis of dynamic contact problems in travelling wave ultrasonic motors," *Proceedings of the 3rd International Congress on Industrial and Applied Mathematics, ICIAM 95*, Hamburg, 1995.

[Clauser, 1967] Clauser, H. R. (Ed.), Materials Selector Issue 1967, *Materials Engineering*, Reinhold Publishing Corp., New York, 1967.

[de Waard, 2000] de Waard, A., Frossati, G., Zendri, J. P., Coccia, E., and Fafone, V., "New technique to measure the mechanical quality factor of metals using spherical samples," *Physica B* **280**, pp. 535–536, 2000.

[Duffy, 1992] Duffy, Jr., W., "Acoustic quality factor of molybdenum and tungsten at low temperatures," *Journal of Applied Physics* **72**, No. 12, pp. 5628–5634, December 1992.

[Duffy, 1999] Duffy, Jr., W., "Acoustic quality factor of copper alloys from 50mK to 300K," *Journal of Applied Physics* **86**, No. 5, pp. 2483–2489, September 1999.

[Ferry, 1961] Ferry, J. D., *Viscoelastic Properties of Polymers*, 1st ed., John Wiley & Sons, New York, 1961.

[Flynn, 1995] Flynn, A. M., *Piezoelectric Ultrasonic Micromotors*, MIT Ph.D. Thesis in Electrical Engineering and Computer Science, June 1995.

[Gonzalez, 1996] Gonzalez, A. M. and Alemany, C., "Determination of the frequency dependence of characteristic constants in lossy piezoelectric materials," *Journal of Physics D: Applied Physics* **29**, pp. 2476–2482, April 1996.

[Gromakovskii, 1978] Gromakovskii, V. A. *et al.*, "On the possibility of using a piezo-electric motor for direct actuation of the drive shaft of a video tape recorder," *Tekhnika Kino i Televideniya*, No. 5, pp. 33–43, 1978.

[Guyomar, 1996] Guyomar, D., Mazoyer, T., Ching, G., Sauzade, J. D., Lagarrigue, O., "High torque/High power piezomotor — influence of the rotor vibration," *Proceedings of the 5th International Conference on New Actuators, ACTUATOR 96*, Bremen, Germany, pp. 160–164, June 1996.

[Hagedorn, 1998] Hagedorn, P., Sattel, T., Speziari, D., Schmidt, J., and Diana, G., "The importance of rotor flexibility in ultrasonic traveling wave motors," *Smart Materials and Structures* **7**, No. 3, pp. 352–368, June 1998.

[Hagood, 1990] Hagood, N., Chung, W., and von Flotow, A., "Modeling of Piezoelectric Actuator Dynamics for Active Structural Control", *Journal of Intelligent Material Systems and Structures* **1**, No. 3, pp. 327–354, July 1990.

[Hagood, 1995] Hagood, N. W. and McFarland, A. J., "Modeling of a Piezoelectric Rotary Ultrasonic Motor," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **UFFC-42**, No. 2, pp. 210–224, March 1995.

[Holland, 1967] Holland, R., "Representation of the Material Constants of a Piezoelectric Ceramic by Complex Coefficients," *IEEE Transactions on Sonics and Ultrasonics* **SU-14**, pp. 18–20, 1967.

[Hosoe, n/a] Hosoe, K., "An Ultrasonic Motor for Use in Autofocus Lens Assemblies," Canon, Inc., translated by Japanese Translation Team (no year given).

[IEEE Std 176-1987] IEEE Standard on Piezoelectricity, ANSI/IEEE Std 176-1987.

[Inaba, 1987] Inaba, R., Tokushima, A., Kawasaki, O., Ise, Y., and Yoneno, H., "Piezo-electric Ultrasonic Motor," *Proceedings of the 1987 IEEE Ultrasonics Symposium*, pp. 747–756, 1987.

[Inman, 1994] Inman, D. J., *Engineering Vibration*, pp. 99–101, Prentice Hall, Englewood Cliffs, NJ, 1994.

[Kawai, 1995] Kawai, Y., Asai, K., Naito, S., Fukui, T., Adachi, Y., Handa, N., Ikeda, K., and Tsuda, K., "High Power Travelling-Wave Type Ultrasonic Motor", *Japanese Journal of Applied Physics* **34**, pp. 2711–2714, May 1995.

[Kawasaki, 1993] Kawasaki, O., Nishikura, T., Sumihara, M., Ohtsuchi, T., Takeda, K., Nojima, T., and Imada, K., "A Small Size Ultrasonic Motor," *Proceedings of the 1993 IEEE International Conference on Systems, Man, and Cybernetics*, pp. 435–440, October 1993.

[Kumada, 1985] Kumada, A., "A Piezoelectric Ultrasonic Motor", *Japanese Journal of Applied Physics* **24**, Supplement 24-2, pp. 739–741, 1985.

[Lau, 1982] Lau, S. L., Cheung, Y. K., and Wu, S. Y., "A Variable Parameter Incremental Method for Dynamic Instability of Linear and Nonlinear Elastic Systems," ASME *Journal of Applied Mechanics* **49**, No. 4, pp. 849–853, 1982.

[Lee, 2000] Lee, T., Lakes, R. S., and Lal, A., "Resonant ultrasound spectroscopy for measurement of mechanical damping: comparison with broadband viscoelastic spectroscopy", Review of Scientific Instruments **71**, No. 7, pp. 2855-2861, July 2000.

[Maas, 1995a] Maas, J., Ide, P., Fröhleke, N., and Grotstollen, H., "Simulation Model for Ultrasonic Motors Powered by Resonant Converters," *Proceedings of the 1995 IEEE Industrial Applications Conference, IAS 95*, Orlando, pp. 111–120, 1995.

[Maas, 1995b] Maas, J., Ide, P., and Grotstollen, H., "Characteristics of Inverter-Fed Ultrasonic Motors—Optimization of Stator/Rotor-Interface," *Proceedings of the 5th International Conference on New Actuators, ACTUATOR 96*, Bremen, Germany, pp. 241–244, June 1996.

[Maeno, 1990] Maeno, T., Tsukimoto, T., and Miyake, A., "The Contact Mechanism of an Ultrasonic Motor," *Proceedings of the 1990 IEEE 7th International Symposium on Applications of Ferroelectrics*, pp. 535–538, 1993.

[Maeno, 1992] Maeno, T., Tsukimoto, T., and Miyake, A., "Finite-Element Analysis of the Rotor/Stator Contact in a Ring-Type Ultrasonic Motor," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **UFFC-39**, No. 6, pp. 668–674, November 1992.

[Morgan, 1997] *Guide to Modern Piezoelectric Ceramics (Rev. 7/97)*, Morgan Matroc, Inc., Electro Ceramics Division, Bedford, Ohio, July 1997.

[Nayfeh, 1979] Nayfeh, A. H., and Mook, D. T., *Nonlinear Oscillations*, John Wiley and Sons, New York, 1979.

[Panasonic, n/a] "Ultrasonic Motor" Technical Reference, Panasonic Industrial Co., Electric Motor Division of Matsushita Electric Corporation of America (no year given).

[Petit, 1998] Petit, L., Briot, R. Lebrun, L., and Gonnard, P., "A Piezomotor Using Longitudinal Actuators," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **UFFC-45**, No. 2, pp. 277–284, March 1998.

[Petit, 2001] Petit, L., Gonnard, P., and Gréhant, B., "A piezomotor for the small actuator market," *Proceedings of the 2000 12th IEEE International Symposium on Applications of Ferroelectrics* **1**, pp. 33–36, 2001

[Pilkey, 1994] Pilkey, W. D., *Formulas for Stress, Strain, and Structural Matrices*, p. 564, John Wiley and Sons, New York, 1994.

[Sashida, 1980] Sashida, T., *Supersonic Vibration Driven Motor Device*, U.S. Patent No. 4,325,264, filed March 24, 1980.

[Sashida, 1982] Sashida, T., "A prototype ultrasonic motor — principles and experimental investigations," *Oyobutsuri (Applied Physics)* **51**, No. 6, pp. 713–720, 1982.

[Sashida, 1983] Sashida, T., *Mechanical Automation of Japan* **15**, No. 2, pp. 31–35, 1983.

[Sashida, 1993] Sashida, T., and Kenjo, T., *An Introduction to Ultrasonic Motors*, Clarendon Press, Oxford, 1993.

[Schmidt, 1996] Schmidt, J. P., Hagedorn, P., and Bingqi, M., "A Note on the Contact Problem in an Ultrasonic Travelling Wave Motor," *International Journal of Non-Linear Mechanics* **31**, No. 6, pp. 915–924, 1996.

[Schreiner, 2000] Schreiner, M., Huber, W., Schöner, H.-P., and Dörmer, M., "Robotic Wrist Actuator with High Torque Piezoelectric Travelling-Wave-Motor," *Proceedings of the 7th International Conference on New Actuators, ACTUATOR 2000*, Bremen, Germany, P3, June 2000.

[Sherrit, 1998] Sherrit, S. and Mukherjee, B. K., "The Use of Complex Material Constants to Model the Dynamic Response of Piezoelectric Materials," *Proceedings of the 1998 IEEE Ultrasonics Symposium* **1**, pp. 633–640, 1998.

[Siebert, 1986] Siebert, W. M., *Circuits, Signals, and Systems*, p. 380, The MIT Press,

Cambridge, MA, 1986.

[Smits, 1985] Smits, J. G. "Iterative Method for Accurate Determination of the Real and Imaginary Parts of the Materials Coefficients of Piezoelectric Ceramics," *IEEE Transactions on Sonics and Ultrasonics* **SU-23**, pp. 393–402, November 1976.

[Tada, 1991] Tada, Y., Ishikawa, M., and Sagara, N., *Polymer Preprints* **40**, 40th SPSJ International Polymer Conference, Japan, p. 1408, 1991.

[Ueha, 1993] Ueha, S., and Tomikawa, Y., *Ultrasonic Motors: Theory and Applications*, Clarendon Press, Oxford, 1993.

[Vasiliev, 1979] Vasiliev, P. E., Klimavichjus, P.-A. R., Kondratiev, A. V., Matsjukyavich-jus, J. J., Beksha, B.-V. L., and Kaminskas, V. A., *Drive Means and Force Control for Piezoelectric Vibration Motor*, U.S. Patent No. 4,249,100, filed April 12, 1979.

[Vishnevsky, 1975] Vishnevsky, V. S., Kavertsev, V. L., Kartashev, I. A., Lavrinenko, V. V., Nekrasov, M. M., and Prez, A. A., *Piezoelectric Motor Sturctures*, U.S. Patent No. 4,019,073, filed August 12, 1975.

[Wallaschek, 1998] Wallaschek, J., "Contact mechanics of piezoelectric ultrasonic motors," *Smart Materials and Structures* **7**, No. 3, pp. 369–381, June 1998.

[Williams, 1948] Williams, A. L. W., and Brown, W. J., *Piezoelectric Motor*, U.S. Patent No. 2,439,499, 1948.

[Zelenka, 1986] Zelenka, J., *Piezoelectric Resonators and Their Applications*, Elsevier Science Publishing Co., New York, pp. 179–181, 1986.