

# Estimating Airflow in an Air-Handling Unit Via Electrical Measurements

Christopher Laughman\*, Steven B. Leeb†, Leslie K. Norford‡, Steven R. Shaw§ and Peter R. Armstrong¶

\*Mitsubishi Electric Research Laboratories  
201 Broadway, Cambridge, MA 02139  
email: laughman@merl.com

†Department of Electrical Engineering and Computer Science  
Massachusetts Institute of Technology, Cambridge, MA 02139  
email: sbleep@mit.edu

‡Department of Architecture  
Massachusetts Institute of Technology, Cambridge, MA 02139  
email: lnorford@mit.edu

§Department of Electrical and Computer Engineering  
Montana State University, Bozeman, MT 59717  
email: sshaw@matrix.coe.montana.edu

¶Masdar Institute of Science and Technology, Abu Dhabi, United Arab Emirates  
email: parmstrong@masdar.ae

**Abstract**—This paper describes a method for estimating the amount of airflow travelling through an air-handling unit (AHU) via the use of a model of the system and only electrical measurements..

**Index Terms**—Induction machines, parameter estimation, fault diagnosis.

## I. INTRODUCTION

Ducted ventilation systems are an important component of contemporary buildings, due to the fact that they provide the most direct means to regulate occupant comfort and productivity by regulating the temperature, humidity, and outdoor air entering an occupied space. Precise monitoring of ventilation systems can also be important, such as detecting the presence of natural or artificial contaminants, precisely controlling air conditions for particular processes, e.g., manufacturing processes and food storage, or combinations of these requirements, as are found in hospital systems (NEED CITATION).

Unfortunately, not all ventilation systems always function as intended, resulting in airflow through the system that varies with respect to the design specifications. Such circumstances can arise for a variety of reasons, including blockage or leakage in the ducting system, malfunctioning fans that cannot supply the proper amount of airflow to the occupied spaces, and malfunctioning control systems for the fans and ducting

systems. Moreover, it is often difficult for a building manager to identify these problems, because they are hard to locate physically, and because problems such as severe blockage or leakage are often the cumulative result of a number of smaller problems.

The variety of commonly occurring problems can be seen by reviewing published surveys of airflow faults in buildings. One set of surveys which examined 503 units in 181 buildings from 2001-2004 found that the airflow was out of the specified range in 42% of the units surveyed [1]. Another study of 4,168 commercial air-conditioning units in California reported that the airflow in 44% of the units did not meet the specifications [2]. Duct leakage has been a particular topic of study due to its importance; studies of 29 new homes in Washington State found average duct leakage rates to the exterior ranged from 140 to 687 CFM [3]. Similar results are reported in [4]. Estimates for the total energy lost from such faults that have been formulated by extrapolating from such fault surveys report estimated total energy losses due to duct leakage at \$US 5 billion/year [5]. Field studies have also been conducted on buildings in Belgium and France, from which the authors estimate that substantially reducing duct leakage will result in cumulative energy savings over a period of 10 years of approximately 10 TWh [6].

A variety of different methods have been used to identify airflow through ducts that is out of the specified range. These can be approximately divided into two measurement strategies: methods which directly measure or estimate the airflow, and methods which infer changes in the airflow from other related measurements. Many direct methods exist; these methods use

Essential funding and support for this research was provided by the Grainger Foundation, the National Science Foundation, NASA Ames Research Center, NEMOmetrics, Inc., and the Office of Naval Research under the ESRDC program.

measurements of quantities that can be directly related to the airflow, such as the speed of a propeller placed in the air stream or differences in the pressure at points in the stream. Devices which implement such flow measurement methods include anemometers, pitot tubes, and flow plates [7].

Indirect methods of estimating airflow have also been investigated, due to a desire to avoid installing dedicated flow sensors, since such sensors are themselves often susceptible to mechanical failure and are also often very expensive. Such indirect methods use a set of measurements that are often not directly related to the airflow and construct a fault model, or a model of a quantity that will serve as an effective indicator for a change in the airflow. For example, Li and Braun [8] estimate changes in the airflow through an evaporator by measuring or estimating a set of refrigerant pressures and temperatures in as well as air temperatures, and then infer the presence of a fault from the behavior of a function of these variables. [Must report on how accurate their method is.] Similar methods are discussed in [9]–[13].

While the above methods for detecting deviations in the airflow through an air-handling unit can be useful, they have a few distinct shortcomings. While the direct methods do provide a direct estimate of the quantity of interest, the use of a dedicated sensor to identify this fault is problematic due to the cost of such airflow sensors and the reliability concerns implicit in including an additional sensor in a fault diagnostic system. As mentioned before, indirect methods have been developed largely to address these concerns, but such methods generally have lower sensitivity than the direct methods. Moreover, most extant indirect methods cannot estimate the numerical value of the airflow, but rather only deviations from the baseline. Such methods assume that the system is commissioned properly and that the airflow is within the specifications shortly after its installation. This is often not the case; airflow systems are often built using construction techniques that do not conform to the specifications, so that any reference baseline will not represent the expected behavior of the nonfaulty system. (CITATION NEEDED)

This paper describes a method for obtaining numerical estimates of the airflow through a ducted air-handling system solely through a limited set of measurements and the use of the manufacturer’s specifications for the equipment. Rather than use mechanical sensors to measure the flow at a single point in the duct, this method uses measurements of the current and voltage at the electrical terminals of the fan motor to generate estimates of the average flow through the fan. These electrical sensors, which are usually more robust than mechanical sensors, thus serve multiple purposes, as they can also measure the power consumption of the air-handler. Furthermore, this method generates numerical estimates of the flow through the duct, and thereby can be used as a fault detection method that is not susceptible to commissioning faults. Because of this feature, the method can be used to identify both blockage within the system, e.g., clogged air filters, as well as leaks in the ducts.

Having provided a background to this research in this introduction, the remainder of this paper will proceed as follows: in the following section, the structure of this airflow

estimation method, which can be separated into three distinct tasks, will be thoroughly described. In the succeeding section, the technical background and results relating to these tasks will be discussed in detail. The fourth section will present a set of results relating both to the experimental validation of the efficacy of the method’s components, as well as the performance of the overall airflow estimation method, and these results and possible directions for further work will be reviewed in the final section.

## II. OVERVIEW OF THE AIRFLOW ESTIMATION METHOD

The structure of this airflow estimation method can best be explained by connecting the variables governing the airflow through the AHU to those variables that are measured and estimated. The development of the functional relationship can be seen by manipulating the following constitutive relationship which describes the airflow:

$$W_f = P_f Q_f, \quad (1)$$

where  $W_f$  is the mechanical power supplied to the fan,  $P_f$  is the pressure across the fan, and  $Q_f$  is the volumetric flow rate through the fan. The relationship between these variables, known as a fan curve, is typically measured empirically for most manufactured fans due to its importance in predicting fan performance over a wide range of operating conditions. One representative fan curve expression is

$$P_f = f_1(Q_f, \omega_f), \quad (2)$$

where  $\omega_f$  is the speed of rotation of the fan and  $f_1$  symbolically represents the relationship established by the fan curve. These fan curves can be used to develop a direct relationship between  $Q_f$  and  $W_f$ , i.e.,

$$W_f = f_1(Q_f, \omega_f) Q_f \quad (3)$$

$$= f_2(Q_f, \omega_f). \quad (4)$$

When the fan is in its stable region of operation, a given airflow and speed can only produce one value of mechanical power flowing into the fan, so that the function is single-valued and can be inverted [cite thesis and other books]. This results in the following simplifications:

$$Q_f = f_3(W_f, \omega_f) \quad (5)$$

$$= f_3(\tau_f \omega_f, \omega_f) \quad (6)$$

$$= f_4(\tau_f, \omega_f) \quad (7)$$

This functional relationship suggests that it is possible to obtain a numerical estimate of volumetric airflow through the fan by measuring the torque applied to the fan  $\tau_f$  and the speed of the fan rotor  $\omega_f$ . While it might be otherwise appealing to substitute the electrical power  $W_e$  entering the motor terminals for the mechanical power  $W_f$ , these quantities are generally not equivalent due to the strong dependence of the motor efficiency  $\nu$  on the winding temperature. As this temperature dependence is often not known and is dependent on the particular characteristics of the installation, this airflow estimation method directly estimates the torque developed by the motor *in situ* at a particular speed to compensate for these temperature-dependent effects.

The manipulation of these relationships facilitates the transformation of the method for estimating airflow into a method for estimating  $\tau_f$  and  $\omega_f$  from observations of the electrical terminal variables of the motor, e.g., the voltage  $V_m$  and the current  $I_m$ . The block diagram illustrated in Figure ?? illustrates the structure of this method. The airflow estimation




Fig. 1. Structure of the airflow estimation method.

process is structured as follows: first, the speed  $\omega_f$  and the torque-speed curve  $\tau_f(\omega_r)$  are separately estimated using observations of the electrical variables  $V_m$  and  $I_m$ . The torque at the motor's present operating point is then identified by evaluating the  $\tau_f(\omega_r)$  at the operating speed  $\omega_f$ . With these estimates of  $\omega_f$  and  $\tau_f$ , the operating mechanical power  $W_f$  can be identified, which can be used in turn with the estimate of  $\omega_f$  to identify the point on the fan curve that describes the fan's current state, thereby generating an estimate of the volumetric flow of air  $Q_f$  through the fan.

While the methods for estimating  $\tau_f$  and  $\omega_f$  both use observations of  $V_m$  and  $I_m$ , the physical processes used to generate each estimate are somewhat different. Each of these estimation methods is therefore discussed in its respective section for the sake of clarity, following which the efficacy of the overall method will be presented and discussed.

### III. ROTOR SPEED IDENTIFICATION

The method used to identify the rotor speed analyzes the frequency content of the motor currents and identifies harmonics due to the interaction of the rotor slots with the magnetomotive force (MMF) wave in the air gap of the motor that are related to the rotor speed. Since the variations in the air gap due to the motion of the rotor slots cause corresponding variations in the spatial permeance wave and magnetic flux density in the air gap, corresponding variations can be identified in the motor currents. Hence, the knowledge of the number of rotor slots, the electrical drive frequency, and a few other pieces of information make it possible to identify the rotor speed by tracking the harmonic content of the stator currents. This method is developed in [?] by examining the effect of the rotor slots on the induced stator voltages; an alternate derivation is also developed in [?, ch. 9] that analyzes the effect on the permeance waves as well as the components of the MMF wave. Static and dynamic eccentricities also have an effect on the permeance wave, introducing additional harmonic content in the magnetic flux wave that can also be related to the rotor speed. This interaction can be expressed in the following form,

$$\Phi_{gs}(\phi_m, \theta_e) = \text{MMF}_{gs}(\phi_m, \theta_e) P_{gs}(\phi_m, \theta_e) \quad (8)$$

where  $\Phi_{gs}$  is the flux in the air gap as referred to the stator,  $P_{gs}$  is the permeance of the airgap as referred to the stator,  $\phi_m$

is the mechanical angle of the rotor, and  $\theta_e$  is the electrical angle of the stator.

The MMF wave can be written as having the following components: principal harmonics, stator slot harmonics, and rotor slot harmonics, as discussed in [?], [?], [?]. A simplified expression for the rotor slot harmonics can be derived by neglecting the effect of the rotor slots and considering only the principal harmonics in the MMF wave,

$$\text{MMF}_{gs} = A \cos(pnx \pm \omega t), \quad n = 6k \pm 1, \quad k = 1, 2, 3 \dots \quad (9)$$

where  $p$  is the number of pole-pairs,  $n$  is the order of the harmonic,  $\omega$  is the radian frequency of the electrical excitation, and  $x$  is the linear distance from  $\theta = 0$  in the developed diagram, or linear model of the machine. Under the assumption of no rotor eccentricity, the flux in the air gap takes the form

$$\Phi_{gs} = AP_0 \cos(pnx \pm \omega t) \quad (10)$$

with respect to the stator, where  $P_0$  is the permeance of the air gap. This expression can be modified by transforming the rotor position [?] into a reference frame that rotates synchronously with the rotor via  $x = \omega_r t + x'$ , where  $x$  is in the stator frame of reference,  $x'$  is in the rotor frame of reference, and  $\omega_r$  is the rotational frequency of the rotor in radians, so that

$$\Phi_{gr} = AP_0 \cos[pn(x' + \omega_r t) \pm \omega t] \quad (11)$$

with respect to the rotor.

The rotor slots introduce additional harmonic content at the frequency  $kR$ , where  $k$  is the number of the harmonic and  $R$  is the number of rotor slots, by introducing an additional permeance wave  $\alpha P_0 \cos(Rx')$ , which will multiply the magnetic flux  $\Phi_{gr}$ , generating additional harmonic content in this waveform. By including this additional effect, the magnetic flux in the air gap can be written as

$$\Phi_{grh} = A_1 P_0^2 \cos[Rx' \pm pn(x' + \omega_r t) \pm \omega t - \phi_1], \quad (12)$$

where  $A_1$  represents the set of multiplicative constants in the front of the expression and  $\phi_1$  represents the offset in the spatial phase of the rotor slots. Equation 12 expresses the magnetic flux in the rotor frame of reference; this flux can also be formulated in the stator frame of reference,

$$\Phi_{gsh} = A_1 P_0^2 \cos[(R \pm pn)(x - \omega_r t) \pm pn\omega_r t \pm \omega t - \phi_1]. \quad (13)$$

By making the substitution

$$\omega_r = \frac{(1-s)}{p} \omega, \quad (14)$$

where  $s$  is the fractional slip of the machine, this can be rewritten as

$$\Phi_{gsh} = A_1 P_0^2 \cos \left[ (R \pm pn)x - \left( R \frac{(1-s)}{p} \pm 1 \right) \omega t - \phi_1 \right]. \quad (15)$$

The principal slot harmonics (PSH) can thus be found by analyzing the stator current and looking at the frequency corresponding to

$$\left( R \frac{(1-s)}{p} \pm 1 \right) \omega t. \quad (16)$$

Though this derivation assumes that the air gap is uniform and straight, loading can introduce both static and dynamic eccentricity to the system [14]. Static eccentricity is defined as an eccentricity in which the location of the minimum air-gap distance is fixed in space, while dynamic eccentricity is generally defined as an eccentricity in which the minimum air gap location rotates with the rotor, and can be caused by such phenomena as a bent or misaligned rotor shaft, worn bearings, and so forth. The effect of both of these eccentricities can be seen in the corresponding term for the rotor slot harmonics. The location of the slot harmonics resulting from the rotor slots and the two types of eccentricity is given by

$$f_{sh} = f_e \left[ (kR + n_w) \frac{(1-s)}{p} + n_d \right] \quad (17)$$

where  $f_{sh}$  is the frequency of the slot harmonics,  $f_e$  is the electrical supply frequency,  $n_d = 0, 1, 2, 3, \dots$  is the order of the eccentricity (0 for static eccentricity, and  $1, 2, 3, \dots$  for dynamic eccentricity),  $k$  is any integer, and  $n_w = \pm 1, \pm 3, \dots$  is the order of the stator MMF time harmonic [?]. By using this equation, it is possible to locate harmonics in the current waveform that are related to the rotor speed at a given point in time.

It is worth noting that the complete spectrum of these rotor slot harmonics will not be visible in the current waveform, due to the overlap with other MMF harmonics, the harmonic content of the voltage waveforms driving the motor, and other related phenomena [?]. As long as some of these harmonics can be observed in the current waveform, however, they can be used to identify the rotor speed. Previous research has developed methods of tracking these harmonics in real-time from observations of the current for a single motor [?], [?], but as the objective of this research is limited to the estimation of the rotor speed for the purpose of estimating airflow, this paper is limited to that scope.

#### IV. TORQUE-SPEED CURVE IDENTIFICATION

We need to identify the torque produced by the machine at a given rotational speed for the fan. This could be done by measuring or estimating the motor's torque speed characteristic directly by using a dynamometer or other sophisticated instrumentation apparatus. While this could be conceivably done before using the system, the motor's torque speed characteristic is strongly temperature dependent, due to the system's resistive heating of the motor under load, making the data from such experiments less than reliable. It was therefore important to develop a method by which the parameters could be easily identified for the machine *in situ*, so that the parameters could represent the actual measured behavior of the machine. It was also important that this be performed using low-cost methods, which is another reason that this is performed solely by using voltage and current transducers.

This method utilizes the fact that the effects of the machine's parameters can be readily observed in the current and voltage waveforms during the startup transient. By estimating the parameters of a model of the machine that characterizes its transient behavior, we can characterize the behavior of a particular machine well enough to reproduce the torque speed

curve. So what we will do is collect current and voltage data from the machine, and then use that data to identify the parameters of a model that can accurately describe the machine's torque speed curve.

One model that can be used to describe the motor's torque-speed curve that can also be theoretically identified from electrical data is the fifth-order model of the induction machine. The Park transformation [15] is used to simplify the structure of the induction machine model by transforming the observed variables into a frame which rotates synchronously with the stator voltage. This has the effect of eliminating the time-varying components of the inductances. While this transformation makes the equations describing the behavior of the induction machine much more tractable than they would otherwise be, it is also useful because it can be applied directly to a set of three-phase electrical observations, reducing the three sets of voltages into 2 orthogonal components (assuming a balanced driving voltage).

After applying the Park transformation to the equations describing the dynamics of the induction machine, the machine can be described by the following equations:

$$\frac{d\lambda_{qs}}{dt} = v_{qs} - R_s i_{qs} - \omega_e \lambda_{ds} \quad (18)$$

$$\frac{d\lambda_{ds}}{dt} = v_{ds} - R_s i_{ds} + \omega_e \lambda_{qs} \quad (19)$$

$$\frac{d\lambda_{qr}}{dt} = v_{qr} - R_r i_{qr} - (\omega_e - p\omega_r) \lambda_{dr} \quad (20)$$

$$\frac{d\lambda_{dr}}{dt} = v_{dr} - R_r i_{dr} + (\omega_e - p\omega_r) \lambda_{qr}. \quad (21)$$

where  $\lambda$  denotes the flux linkages with the rotor variables and parameters reflected to the stator,  $\omega_e$  is the frequency of the stator excitation (i.e. the frequency of the drive voltage) in rad/s, and  $\omega_r$  is the rotor speed in rad/s, and  $p$  is the number of pole pairs. The voltages  $v_{ds}$  and  $v_{qs}$  represent the driving voltages in most experimental applications, as  $v_{dr}$  and  $v_{qr}$  are set to zero due to the fact that the rotor bars are shorted together on a squirrel-cage machine. The flux linkages and the currents are related by the following equations:

$$\lambda_{qs} = L_{ls} i_{qs} + L_m (i_{qs} + i_{qr}) \quad (22)$$

$$\lambda_{ds} = L_{ls} i_{ds} + L_m (i_{ds} + i_{dr}) \quad (23)$$

$$\lambda_{qr} = L_{lr} i_{qr} + L_m (i_{qs} + i_{qr}) \quad (24)$$

$$\lambda_{dr} = L_{lr} i_{dr} + L_m (i_{ds} + i_{dr}). \quad (25)$$

The torque of electrical origin produced by the motor is given by

$$\tau_e = \frac{3}{2} p (\lambda_{qr} i_{dr} - \lambda_{dr} i_{qr}). \quad (26)$$

This torque  $\tau_e$  is related to the mechanical load of the fan by the usual force balance equation,

$$\frac{d\omega_r}{dt} = \frac{1}{J} (\tau_e - \beta \omega_r^2). \quad (27)$$

The  $1/J$  term will also be referred to as  $K$ . After (18)-(21) and (27) are integrated over the time interval of interest, the d- and q-axis currents are transformed back into the lab frame by applying the inverse Park transformation.






Fig. 2. Equivalent circuit model for one phase of an induction machine.

The torque-speed curve can be identified from the parameters of this model as well by examining the steady-state behavior of the system. The following equation describes the relation between the mechanical torque produced by the machine and the rotational speed of the machine:

The steady-state behavior of the induction motor can also be described by an equivalent circuit model, as shown in Figure 2. This circuit model does not describe the dynamics of the system in  $dq$ -space, but rather describes the behavior of the induction motor in steady-state operation in the lab frame. This model is used because it can provide intuitively appealing expressions for the steady-state behavior of the induction machine.

As described in [?, p. 317], the mechanical torque developed by the motor is given by

$$T_{mech} = \frac{n_{ph} I_r^2 (R_r/s)}{\omega_s}, \quad (28)$$

where  $n_{ph}$  is the number of phases of the motor,  $I_r$  is the current flowing through the rotor,  $R_r$  is the rotor resistance,  $\omega_s$  is the synchronous speed of the motor, and  $s$  is the value of the fractional slip of the motor. The fractional slip is calculated by

$$s = \frac{n_s - n}{n_s}. \quad (29)$$

where  $n_s$  is the synchronous speed of the motor, and  $n$  is the present operating speed. It is important to again note that the parameters of this steady-state model are the same as the parameters of the transient model described earlier, and that the parameters that describe the transient electrical behavior of the electrical machine will also describe its torque-speed characteristics.

While this appears in general to be the most straightforward and reliable method to describe the machine, there are challenges with using electrical observations of the machine behavior to identify the model parameters. A variety of methods of doing this have previously been developed (cite previous research), but the shortcomings of these methods suggest that another approach to parameter estimation might be more effective.

Overall plan - the difficulty of identifying the parameters of the motor. The definition of the loss function. Using a gradient-based method is fine as long as the initial guess is close; if it is not close, it is a problem, since it has a lot of local minima. An approach was developed for reducing the sensitivity of this problem to the initial guess by effectively taking an approach to smoothing the loss function.

In general, this is a challenging problem. There are a number of parameter estimation methods that can be used; one of the particular issues is that most methods require a high quality




Fig. 3. The basic structure of the simulation process.




Fig. 4. The basic structure of the estimation process.

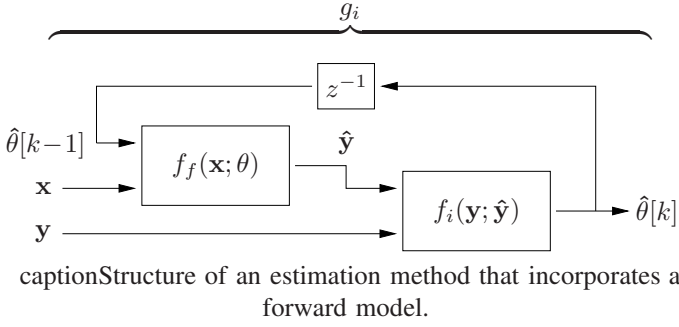
initial guess. If this initial guess is not adequate, the method might identify parameters that are in a local minimum. Bad. In general, speed is often traded off against susceptibility to landing in one of these local minima - sample the space with many initial guesses in order to ensure that you will not end up in a local minimum. Slow.

The method taken in this paper uses a pre-estimation method, or a continuation method. We estimate a preliminary set of parameters which will not be optimal, but which are expected to be close to the correct set of parameters, or in other words a good initial guess for a more detailed method. Then we use this set of pre-estimates to identify the final set of parameters which best represent the data (hopefully).

The parameter estimation method developed in this paper uses the processes of both simulation and estimation, as illustrated in Figures 3 and 4. The process of simulation (denoted in Figure 3 as  $g_f$  to represent the solution of the “forward problem”) generally uses a known functional mapping  $f_1()$  to generate a set of predicted outputs  $\hat{y}$  from a set of measured inputs  $x$ , and the parameters  $\theta$  which govern the behavior of this mapping are known in advance. In comparison, the process of parameter estimation, sometimes referred to as an “inverse problem” or  $g_i$  in Figure 4, uses a set of observations of both inputs  $x$  and outputs  $y$  to construct the functional mapping between the two sets of observations. The output of the inverse problem is thus the set of parameters  $\hat{\theta}$  which best constructs this mapping.

Because the formulation of forward models is usually much more straightforward than inverse models, a variety of different approaches have been studied for the development of these inverse models. An obvious approach is the construction of the exact inverse mapping  $g_f^{-1}$ ; while the simplicity of this method is appealing, major problems in implementing such inverse mappings often arise due to nonlinearities, poor numerical conditioning, or other pathological behavior. Instead of directly inverting this model as in [16], an alternative approach for parameter identification used here embeds the forward model directly in the parameter estimation method, as illustrated in Figure IV.

This diagram illustrates the means by which the simulation is incorporated into the overall parameter estimation method  $g_i$ ; the overbrace  $g_i$  in Figure IV signifies that the function  $g_i$  in this figure represents the corresponding function in Figure 4. The simulation routine is initialized with an initial guess for parameters  $\theta_0$  as well as the inputs to the system  $x$ , the simulation is run, and then the resulting output of this simulation  $\hat{y}(\theta, x)$  is fed into the estimation routine  $f_i$ . This estimation



routine compares the observed data vector  $\mathbf{y}$  and the output of the simulation  $\hat{\mathbf{y}}(\theta, \mathbf{x})$ , and computes a new parameter estimate  $\theta[k]$ , which is then fed back into the simulation routine. After the first iteration, the parameters  $\theta$  of the simulation  $f_f$  are updated with the most recent parameter estimate  $\hat{\theta}[k-1]$ . This cycle is repeated until an exit condition is reached; ideally the difference  $\mathbf{r}(\hat{\theta}) = \mathbf{y} - \hat{\mathbf{y}}(\theta, \mathbf{x})$  is reduced below an established threshold, but other exit conditions typically include exceeding a set number of iterations, no change in the parameters, and no change in the evaluated residual.

The block denoted  $f_i$  uses the Levenberg-Marquardt (LM) algorithm [17] to generate and iteratively refine the parameter estimates based upon the residual  $\mathbf{r}(\hat{\theta})$ . This algorithm is useful because it uses information from previous steps in updating the parameter estimates, and can generate parameter estimates quickly if the problem is formulated correctly. The unmodified LM algorithm does not typically generate good parameter estimates in a power system application, however, because the residual being minimized, i.e.  $r(\theta) = \|\mathbf{y} - \hat{\mathbf{y}}(\theta, \mathbf{x})\|_2^2$ , has a large number of local minima due to the sinusoidal excitation of the system. Because the LM algorithm uses the gradient of the residual to update the estimates of the parameters, the parameter estimates produced often represent the optimum parameter estimate in the region about the initial guess, rather than the globally optimum parameter estimate.

Pre-estimation of the initial guess can be used to improve the convergence of nonlinear least squares [18]. This technique seeks to exploit the fact that a high-quality initial guess will accelerate and improve the performance of nonlinear least squares for many parameter identification problems. The structure of this approach to parameter estimation can best be seen in Figure 5. The poor initial guesses are provided to the system as  $\hat{\theta}_0$ , and are refined through pre-estimation algorithm  $g_{i,pre}$  to be pre-estimates  $\hat{\theta}_{pre}$ . This high-quality initial guess is then iteratively refined by the final parameter algorithm  $g_{i,f}$  to obtain the final parameter estimates  $\hat{\theta}_{final}$ . In effect, the use of pre-estimates trades computational time and complexity for the quality of the initial guess. Each block in this diagram has the form illustrated in Figure IV; that is, each of these blocks has both an embedded forward model  $f_f$  that generates a set of outputs given the current guess for the parameters, and a Levenberg-Marquardt block  $f_i$ , which iteratively refines the parameters. For clarity, the pre-estimation and final estimation steps will be illustrated separately.




Fig. 5. Block diagram illustrating the structure of the pre-estimation method.




Fig. 6. Envelope extraction block.

The motor model described in the previous section (represented as  $f_f$  in Figure IV) used in the pre-estimation block  $g_{i,pre}$  and in the estimation block  $g_{i,f}$  is the fifth-order model of the induction machine [19].

The pre-estimation step was designed to compensate for the fact that the local minima in the residual are largely caused by the sinusoidal component of the current transient. The rejection of this sinusoidal or “carrier frequency” variation helps to eliminate the local minima that degrades the performance of the LM algorithm. Since the envelope of the motor current is sensitive to changes in many of the motor parameters [20], the envelope was extracted from the observed current data by processing the measured currents with a low-pass filter using a standard demodulation technique and Butterworth filters [21]. This technique has the effect of preserving the slow-moving envelope while eliminating the 60 Hz component of the signal. The cutoff frequency of the third-order Butterworth filters was set to 15 Hz.

Figure 7 illustrates the effect of this filter, as the envelope of the startup transient has been obtained from the input current and the sinusoidal component of the current signal has been eliminated. The sensitivity of this envelope to the motor parameters makes this an ideal preprocessing method from which parameter pre-estimates may be obtained.

Since all of the currents in the stator windings are nominally identical for a balanced machine except for phase shifts, the minimization was only performed against one phase of the stator current, referred to as phase A. All three of the voltages  $V_{AB}$ ,  $V_{BC}$ , and  $V_{CA}$  were measured, however, to perform the d-q transformation and obtain the voltages  $V_D$  and  $V_Q$  in the synchronous reference frame as required by Equations 18 and 19.

By incorporating this information from the pre-estimation filter, the first of the two estimation steps may be described. This first step, referred to in Figure 5 as  $g_{i,pre}$ , is illustrated in Figure 8. The user only needs to provide the set of observed




Fig. 7. Plot of the observed phase A motor current along with the preprocessed envelope of the current.




Fig. 8. Motor parameter estimation method: Step 1.

voltages and currents, as well as the approximate set of initial guesses. Most of the initial guesses only need to be within an order of magnitude; the only parameter which must be known to any additional accuracy is the stator winding resistance. This parameter can be determined with a simple measurement when the induction machine is cold by an ohmmeter, and is used to constrain Levenberg-Marquardt's adjustment of the remaining parameters.

The forward model  $f_f$  is first initialized with the initial guess of parameters  $\hat{\theta}_{init}$  and the set of observed voltages  $\mathbf{v}_{obs}$  are input to the parameter estimation method. The forward model is then used to generate a set of predictions of the stator currents  $\hat{\mathbf{i}}$ . These currents are filtered to obtain the current envelope  $\hat{\mathbf{i}}_{filt}$  and this filtered prediction is then compared to the filtered set of observed currents  $\mathbf{i}_{obs, filt}$ . The difference between these two signals is then input to the LM algorithm, which generates a new estimate of the parameters  $\hat{\theta}[k]$ . This estimate  $\hat{\theta}[k]$  of all of these parameters, except for the stator resistance, is then iteratively refined until the signal  $\hat{\mathbf{i}}_{filt}$  closely approximates the signal  $\mathbf{i}_{obs, filt}$  and the residual  $r(\theta) = \|\mathbf{y}_{filt} - \hat{\mathbf{y}}_{filt}(\hat{\theta}, \mathbf{x})\|_2^2$  is minimized. The resulting set of parameters  $\hat{\theta}$  that results after the iterations have stopped is equal to the set of parameter pre-estimates  $\hat{\theta}_{pre}$ .

It is important to note that the model parameters that are pre-estimated are not adequate to describe the unfiltered set of observations of the system; these pre-estimates are only intended to serve as improved initial guesses for the parameter estimation process for the unfiltered set of observations. In order to find the set of parameters that most accurately represents the unfiltered set of observations, a second step




Fig. 9. Motor parameter estimation method: Step 2.

further refines the parameter estimates. In this second step, the simulation  $f_f$  is initialized with the improved pre-estimates  $\hat{\theta}_{pre}$ , rather than the set of rough initial guesses  $\hat{\theta}_{init}$ . All of the parameters, including the stator resistance, are iteratively adjusted using LM in this step to minimize the residual  $r(\theta) = \|\mathbf{y} - \hat{\mathbf{y}}(\theta, \mathbf{x})\|_2^2$  to find the parameters which best describe the observed currents.

Figure 9 illustrates this second step. It is clear from the figure that the structure of the second step is very similar to the first step, with the key difference that neither the output of the simulation nor the observed currents are filtered. Because the parameter estimation method operates on this unfiltered data, the parameters that are iteratively refined will best characterize the performance of the system. Furthermore, because the initial guesses generated by the pre-estimation step are high quality, the parameter estimates generated by Levenberg-Marquardt are more likely to represent the global minimum, rather than a local minimum. Once the method converges, the set of parameters  $\hat{\theta}$  are equal to the machine parameters  $\hat{\theta}_{final}$ .

In both parameter estimation steps, the performance of the standard Levenberg-Marquardt algorithm also depends on the treatment of a few specific numerical considerations. The first of these considerations pertains to the scaling of the parameters: widely spaced parameters, whose size varies over several orders of magnitude, can be very sensitive to the parameter corrections applied by the nonlinear least squares algorithm. For example, if the gradient of the loss function in the direction of one parameter  $K$  is quite steep, the corresponding adjustment to the parameters could cause a large change in another very small parameter, such as  $\beta$ . These large adjustments in  $\beta$  could consequently cause for the convergence of the method as other parameters are adjusted according to the new position on the loss function.

This problem can be mitigated by scaling the parameters in the motor simulation according to their expected magnitudes, which are relatively easy for a user to estimate. This will allow the minimization algorithm to change all of the parameters by the same order of magnitude; for example, if the true parameters of the system are  $\beta = 5 \times 10^{-6}$  and  $K = 140$ , scale factors of  $10^{-6}$  and 100 would multiply the parameters  $\beta$  and  $K$ , respectively, inside of the simulation function, resulting in final parameter estimates of 5 and 1.4. This approach effectively stretches the residual space so that it makes all of the scales of the parameter gradients comparable.

Other important constraints were implemented to improve the performance of the minimization algorithm. One such constraint incorporates the fact that the parameters of the model cannot be negative, due to physical considerations. This constraint was implemented in this research by applying a

sigmoid, or logistic, function to the parameters. The behavior of this function is given by

$$f(x) = \frac{1}{1 + e^{-\alpha x}}. \quad (30)$$

This constraint was implemented in much the same way as the scaling; the parameter estimates generated by nonlinear least squares were transformed using this sigmoid function, and the resulting transformed parameters were used in the motor simulation. This had the effect of transforming the constrained optimization problem, which ensures physically realistic parameter values, into an unconstrained parameter identification problem. Unlike the parameter scaling discussed previously, this constraint was only imposed during the pre-estimation step, as the parameter pre-estimates were sufficiently close to the final parameter estimates that this constraint was not needed during the final step.

The final numerical constraint imposed on the minimization method placed limitations on the size of the adjustment  $\delta^{(i+1)} = \theta^{(i+1)} - \theta^{(i)}$  applied to the parameters after every iteration. Even after applying the appropriate scaling to the parameters, large gradients in the loss function can cause LM algorithm to take large parameter steps. If the loss function is globally convex, as is the case in linear problems, such an adjustment in the stepsize is useful since the gradient will be bigger as the distance  $\theta$  from the minimum increases. This is not the case in nonlinear problems, however, and the resulting movement far away from the previous initial guess will not take advantage of the additional information that is often implicit in the initial guess for the parameters selected by the user.

This information can be incorporated into the minimization algorithm by imposing a set of linear constraints on the parameters, effectively placing a “rubber band” on the parameters so that they do not move far in consecutive iterations. This can be accomplished by adding an additional set of equations to the system of linearized equations that is solved at each step. If the system of equations that represents the linearized behavior of the nonlinear system at the current speed is written by  $\nabla \mathbf{G}(\theta^{(i)})\delta^{(i+1)} = \mathbf{G}(\theta^{(i)})$ , then the additional constraint can be implemented by solving the simultaneous set of equations

$$\begin{bmatrix} \nabla \mathbf{G}(\theta^{(i)}) \\ \mathbf{I} \end{bmatrix} \delta^{(i+1)} = \begin{bmatrix} \mathbf{G}(\theta^{(i)}) \\ \gamma \frac{\delta^{(i)}}{\theta_0} \end{bmatrix} \quad (31)$$

This technique is commonly known as regularization [22], and can be controlled by the size of the constant  $\gamma$  that multiplies the parameters. This parameter is often set by trial and error; values of  $\gamma$  which are too large prevent nonlinear least squares from adjusting the parameters at all, while values that are too small allow the parameters to change more than is desired. The implementation of regularization in the nonlinear least squares algorithm proved to be essential to obtaining useful parameters from the method.

While the parameter estimation method was primarily developed to identify the torque-speed curve of the motor solely on the basis of electrical measurements, the method was also modified so that torque-speed data measured on a dynamometer could be also incorporated into the parameter



Fig. 10. Picture of the overall air handler apparatus and attached duct system.

estimation method. This data was included using an approach that was similar to the addition of the regularization constraints as discussed above; the torque-speed data was concatenated to the set of constraints that describe the system as given by

$$\begin{bmatrix} \nabla \mathbf{G}_1(\theta^{(i)}) \\ \nabla \mathbf{G}_2(\theta^{(i)}) \\ \mathbf{I} \end{bmatrix} \delta^{(i+1)} = \begin{bmatrix} \mathbf{G}_1(\theta^{(i)}) \\ \mathbf{G}_2(\theta^{(i)}) \\ \gamma \frac{\delta^{(i)}}{\theta_0} \end{bmatrix} \quad (32)$$

where  $\mathbf{G}_1(\theta^{(i)})$  represents the gradient of the current residual with respect to each of the parameters  $\theta_i$ , and  $\mathbf{G}_2(\theta^{(i)})$  represents the gradient of the torque-speed residual with respect to each of the parameters  $\theta_i$ . This approach, in which the two different sets of outputs of the same system are simultaneously minimized, generally follows the approach used in [?, p. 112]. Both approaches to the parameter estimation problem were tested, and will be discussed in the results section of this chapter.

## V. RESULTS

### A. Experimental Platform

The architecture of the overall experimental apparatus used, including the air handler, data acquisition system, and length of duct, is illustrated in Figure 10. The air handling unit was purchased from ABCO Refrigeration Supply Corp. in Somerville, MA, and was manufactured by International Comfort Products, LLC. Pictures of this unit were included previously in Figure ?? and Figure ?. This unit is designed to be part of a split air-conditioning system, with an evaporator and filter holder located at one end, and was configured for the duration of the experiments in an “upflow” configuration, meaning that the air intake is located below the air handler, and the exhaust is located at the top of the air handler. It is equipped with a forward-curved double-duct centrifugal fan manufactured by Lau Industries, part no. DD 11-10AT, which is 10 inches in diameter and 11 inches wide, and has 54 blades. This fan is a direct drive unit, with the blower wheel attached directly to the motor shaft via a set screw.

This air handler is sold with a single-phase capacitor-run induction motor driving the fan, but owing to the experimental nature of this work and the fact that effects such as saturation and hysteresis of the magnetic material tend to have a large






Fig. 11. Schematic diagram of the control and data acquisition system.

effect on commercially manufactured single-phase motors, a three-phase induction machine was installed in its place so that the effectiveness of the airflow estimation method could be tested without being obscured by the variety of issues inherent in estimating the parameters of a single-phase machine. This motor was manufactured by Emerson Corp., part no. 1326, and is a delta-connected 3-pole pair induction machine with 48 rotor slots.

A control and data acquisition system was also installed to acquire the variety of signals necessary to implement the fault diagnostic method. This instrumentation necessary to acquire these signals is shown in Figure 11. One can see from this figure that each phase of the electric utility is connected to the motor through a solid-state relay (SSR), and the voltage across each pair of stator windings, as well as the line currents, are input to the data acquisition system. These solid-state relays were used to precisely control when each phase was connected to the electric utility; in traditional mechanical contactors, there can be a substantial delay (perhaps as much as a line cycle) between the time that the first phase of the motor is connected to the utility and the time that the last phase is connected. These arbitrary turn-on relays (manufactured by Crydom, part no. D2450-10) were used to ensure that all three phases of the electric utility were connected to the motor simultaneously to avoid these unusual electrical transients. The fault diagnostic method used would also work if the mechanical contactor were used by monitoring the voltages and currents directly at the motor terminals, but solid-state relays were used to keep the parameter estimation method as simple as possible.

The line-to-line phase voltages  $V_{AB}$ ,  $V_{BC}$ , and  $V_{CA}$  were monitored via three voltage transducers (LEM LV-25P), which are effectively high-accuracy transformers that reduce the phase voltages on the stator windings from  $208 V_{pp}$  to approximately  $3 V_{pp}$ . Similarly, three Hall effect current transducers (LEM LA-55P) were used to monitor and properly condition the phase currents  $I_A$ ,  $I_B$ , and  $I_C$  so that they could be input



Fig. 12. View of ducting apparatus.



Fig. 13. Additional view of ducting apparatus.

to the data acquisition system. The resulting three voltage signals ( $v_{ab}$ ,  $v_{bc}$ ,  $v_{ca}$ ), as well as the three current signals ( $i_a$ ,  $i_b$ , and  $i_c$ ) were input to an Advantech PCI-1710 card, which interfaced directly with a computer running Debian Linux. The seventh signal input to the data acquisition card is the tachometer signal  $v_{tach}$ , which is the 0 – 5 V square-wave output of a US Digital E6S 64 count optical encoder, which was mounted directly to the shaft of the motor. The resulting signal from the encoder had the unfortunate effect of causing noise spikes on the nearby channels, due to its high frequency content; this had no appreciable effect on the torque-speed estimation method, but it dramatically affected the spectrum of the motor currents. Whenever it was necessary to acquire current data for the purposes of analyzing the harmonic content, it was therefore necessary to disable the optical speed encoder.

A length of duct was purchased and installed in the laboratory so that accurate airflow measurements could be taken, and so that the behavior of the air handler would more closely approximate behavior in field-installed equipment. Two views of this ducting system can be seen in Figures The bulk of the ducting system is 30 feet of 16 inch round sheet metal duct, obtained from McMaster-Carr in 3 foot long sections that are taped together with metallized duct tape and suspended from pipes in the ceiling by steel strapping. The transition between the exhaust of the air handler and the 16 inch round duct was accomplished by installing a custom fabricated 90 degree elbow and a 78 inch long converging section that also transitioned from the square cross-section of the elbow to the round cross-section of the remainder of the duct. This particular duct fitting was fabricated by the Medford-Wellington Service Company at a cost of approximately \$250.

The first reason that this length of duct was constructed, i.e. increasing the accuracy of the airflow measurements, stems from the fact that the forward curved fan blades cause the air velocities in the duct to be much higher on the side of the duct that is in the direction of rotation. This fact, in concert with the particular configuration of the ducting system due to space constraints in the laboratory, caused the air velocities at the top of the duct to be much greater than those at the bottom of the duct. A set of adjustable double thickness turning vanes was therefore installed in the elbow so that the velocity profile



Fig. 14. Turning vanes.

could be adjusted. These turning vanes are visible through the access door in Figure 14. This access door was also used to test the leaky duct; when the leaky duct test was conducted, the fan was run with the access door open. The airflow was measured via a set of 4 traverses, each of which included 6 sample points, according to the guidelines specified by ASHRAE [7, p. 14.17]. These traverses were made at a location 77 inches from the end of the duct, or about 5 duct diameters from the end of the duct, to minimize the effects of the diverging flow at the end of the duct. Each of the 4 holes that had to be drilled in the duct was covered with 2 layers of electrical tape and then slit open, in order that the 3 unused holes would mostly be covered when a given traverse was being performed. One of these pieces of duct tape can be seen in Figure 12 at the right end of the duct.

### B. Speed Estimation

The estimation module, as discussed in §??, was the first of the modules to be experimentally evaluated. This data was collected with the fan running in four states of interest: unblocked flow, 30% blocked flow, 50% blocked flow, and leaky flow. The speed estimates were collected and verified by collecting sequential data sets; due to the distortion caused by the high-frequency content of the tachometer signals, tachometer data was collected for 30 seconds while the motor was running, the data collection was stopped and the tachometer channel was disconnected quickly, and then measurements of the motor currents were collected for 30 seconds while the motor continued running. Speed estimates were generated from the tachometer data by using the simple strategy of counting 128 edges in the tachometer waveform, which correspond to the 64 notches in the tachometer dial, and then using the fact that the time in which the 128 edges pass by the detector corresponds to one revolution of the wheel. This algorithm was implemented in C, the source of which is listed in Appendix ??. These speed estimates obtained from the tachometer data have a time-varying component, since the number of samples between the edges is not constant. If there happen to be seven samples before an edge of the tachometer waveform, rather than six, the resulting speed estimate will accordingly change by a few rpm. Accordingly, the speed estimates from the tachometer have a standard deviation around 2 rpm.




Fig. 15. Spectrum of the motor current for the unblocked fan illustrating the locations of the theoretical and experimentally observed slot harmonics.

These estimates from the tachometer were used to validate the estimates of the speed obtained from the observations of the phase A current. A window of  $2^{16}$  datapoints (about 4.5 sec of data at the sample rate of 14.29 kHz) were selected from the datafile, and then a Hanning window was applied to this segment of data to minimize spectral leakage. This data was then filtered twice, first by a third-order elliptical notch filter to remove the 60 Hz component of the signal, and then by a sixth-order elliptical bandpass filter that selected the frequencies between 600 Hz and 1600 Hz where the expected slot harmonics were expected to be. While an algorithm could be designed which locates the peaks in the peaks in the frequency spectrum by using methods such as those suggested in [?], this speed estimation step was carried out by hand in this research to expedite the collection of the speed estimates. The resulting spectrum of the phase A motor current collected when the fan was running in an unblocked state can be seen in Figure 15.

The results of processing the tachometer data that corresponds to this current data yield an average speed over the observed window of 988 rpm with a standard deviation of 2.4 rpm, while the speed predicted from the location of the slot harmonics seen in this figure is 984 rpm. This agreement is reasonably good and indicates the usefulness of the speed prediction algorithm. The discrepancy in the two estimates of the speed can also partially be attributed to the fact that the estimates were not obtained concurrently due to the limitations of the experimental apparatus, so that the fan speed might have changed by a small amount between collecting the two datasets.

The analysis of this current spectrum must also consider the fact that the fan speed in an experimental apparatus typically has short-time variations around a slowly time-varying average speed. The peaks in the Fourier transform of the current spectrum will therefore be “smeared” or wider than they would otherwise be if the motor was only run at a single speed. One can see in Figure 15 that the peaks corresponding to slot harmonics are actually about 5 Hz wide, which corresponds to the small variations present in the fan speed over the measurement window. This variation could be minimized by analyzing a smaller window of data, but this would also reduce the frequency resolution of the spectrum. In developing

an implementation of this system for field use, these two considerations would have to be traded off against one another.

Recalling Equation (17), which identified the expected location of the series of slot harmonics, it is interesting to note that while some of the slot harmonics are clearly present, others are much harder to locate. In particular, the harmonics that correspond to the  $(k, n_w, n_d)$  triplets  $(1, 1, -3)$ ,  $(1, 1, 1)$ , and  $(1, 1, 3)$ , are clearly visible at the frequencies 627 Hz, 866 Hz, and 987 Hz. On the other hand, the harmonics such as those predicted by the triplets  $(1, 0, -1)$ ,  $(1, 1, -1)$ , and  $(1, 0, 1)$ , which would be located at 730 Hz, 747 Hz, and 850 Hz, are not at all discernible from the background noise present in the spectrum of the signal. One can also easily identify the presence of the main MMF harmonics, notably those at 660 Hz, 780 Hz, 900 Hz, and 1020 Hz. Nevertheless, this method is evidently able to identify the motor speed reasonably well.

### C. Motor Parameter Estimation

In order to test the performance of the parameter estimation method on data that most closely resembles the data of interest, the fan was started and run for approximately two hours so that it could reach thermal equilibrium under load. After two hours, a startup transient that represented the torque speed behavior for the fan in continuous operation was obtained by briefly turning the fan off, allowing it to spin down to a standstill, and then recording all three voltages and currents during the following startup transient.

As discussed in §??, the parameter estimation of the motor proceeded in two steps. In the first step, the parameters are pre-estimated by constraining the stator resistance to be a measured value, constraining the remaining motor parameters to be strictly positive, and finding the parameters of the motor which best fit the low-passed envelope of the motor current. The initial guess of the stator current was obtained by measuring the resistance of the A-B stator windings with a Fluke model 87 III multimeter when the motor was cold. While the resistance of the stator windings do change with temperature, as will be seen, this measurement of the stator resistance was used to evaluate the robustness of the parameter estimation method. Since the measured terminal resistance measured was  $8.2\ \Omega$ , the wye-connected model of the induction machine implied that the equivalent resistance  $R_s$  of each stator winding is  $4.1\ \Omega$ . Results using this stator resistance on the pre-estimation fit, as well as estimated scales for the other parameters, are illustrated in Figure 16. It is apparent from this plot that the pre-estimation routine worked fairly well, as the estimated dynamics are captured reasonably well in the estimated envelope. This can also be seen in the generally reasonable parameters that result from this pre-estimation step, as listed in Table I.

Once the pre-estimates of the motor parameters were generated, they were used as initial guesses for the Levenberg-Marquardt minimization of the set of observations of data. This minimization was performed according to the method described in §??, with all of the parameters being minimized in the second step of the algorithm. As with the pre-estimation, the full minimization was tested and validated minimizing only




Fig. 16. Plot of the preprocessed observed phase A motor current envelope along with the estimated current envelope. The observed current envelope is illustrated in blue, while the estimated envelope is shown in red.

Parameter	$R_s$ [ $\Omega$ ]	$R_r$ [ $\Omega$ ]	$X_m$ [ $\Omega$ ]	$X_{ls}$ [ $\Omega$ ]	$X_{lr}$ [ $\Omega$ ]	$\beta$ [ $N \cdot m$ ]
Value	4.10	4.19	90.72	3.78	8.26	$5.69e-04$

TABLE I  
PRE-ESTIMATED MOTOR PARAMETERS.

against the observations of the terminal current  $i_a$ . The results of this minimization are illustrated in Figures 17- 20.

The plots in Figures 17 and 18 are zoomed into the first second of operation of the motor so that the quality of fit can be visually evaluated. The simulation of the induction motor with the final parameter estimates is within 10% of the observed current at all points, and within 5% for most. The size of the residual to the amplitude of the transient is larger in the last 0.2 seconds of the data, which may be caused by unmodeled behavior of the motor, but this amount of misfit remains below 10%. Moreover, the quality of fit improves over time, as one can see from the plots in Figures 19 and 20, which show the same data as in the previous two figures, but on a longer time scale.

These figures show that the quality of fit improves after the initial transient, suggesting that while the model is able to represent the behavior of the system in the transient state fairly well, its description of the system is even better when the system is in steady-state. The final values of the parameters, after this second stage of parameter estimation, are listed in Table II.




Fig. 17. Observed and fitted motor currents for the hot motor; observed data are in blue, and fitted data are in red. This waveform is zoomed in so the agreement between the observations and the predictions can be seen.




Fig. 18. Residual from  $i_{a,fit} - i_{a,obs}$ .






Fig. 19. Observed and fitted motor currents for the hot motor; observed data are in blue, and fitted data are in red. This waveform is zoomed out so that the quality of fit for a larger segment of time can be seen.




Fig. 20. Residual from  $y_{fit} - y_{obs}$ .

$R_{s,init}$	$R_{s,final}$	$R_r$	$X_m$	$X_{ls}$	$X_{lr}$	$K$	$\beta$
0.8	6.24	3.95	54.2	3.72	6.97	4.59e-04	31.0
1.1	6.25	4.12	55.3	2.55	8.48	4.59e-04	31.0
2.1	6.25	4.03	54.7	3.19	7.65	4.59e-04	31.0
4.1	6.25	4.03	54.7	3.14	7.71	4.59e-04	31.0
6.1	6.25	4.03	54.7	3.15	7.70	4.59e-04	31.0
8.1	6.25	4.48	57.7	0.215	11.6	4.59e-04	31.0
10.1	6.25	4.51	57.9	0.0168	11.9	4.59e-04	31.0
20.1	6.25	4.44	57.4	0.458	11.3	4.59e-04	31.0

TABLE IV  
VARIOUS PARAMETERS RESULTING FROM STARTING THE SIMULATION WITH DIFFERENT INITIAL GUESSES OF  $R_s$ .

Parameter	$R_s$ [ $\Omega$ ]	$R_r$ [ $\Omega$ ]	$X_m$ [ $\Omega$ ]	$X_{ls}$ [ $\Omega$ ]	$X_{lr}$ [ $\Omega$ ]	$\beta$ [ $\text{N}\cdot\text{m}/(\text{rad}/\text{s})^2$ ]	$K$ [ $1/\text{kg}\cdot\text{m}^2$ ]
Value	6.25	4.03	57.75	3.14	7.71	4.59e-04	31.00

TABLE II  
FINAL HOT MOTOR PARAMETERS MINIMIZING SOLELY AGAINST THE CURRENT.

The validity of these parameters was evaluated by measuring the impedance of one set of motor terminals before and after running the motor for two hours and comparing the measured impedances with the estimates of the impedances. Such a test provides an independent evaluation of the accuracy of the parameters estimated by the preceding routines. These parameters are reproduced in Table III. It is important to note that these are the measured terminal variables, rather than the model parameters as stated in Table II. These parameters could be converted to the electrical parameters, e.g.,  $R_{meas} = 2R_s$ .

One important consideration regards the dependence of the parameter estimation method on the quality of the initial guess for the motor parameters. As the only parameter explicitly set before the pre-estimation step is  $R_s$ , a number of tests were run in which the value of  $R_s$  was changed over a relatively wide range, from  $0.8\ \Omega$  to  $20.1\ \Omega$ , to ascertain the corresponding impact on the final motor parameters. The set of parameters for each of these initial guesses are collected in Table IV for comparison.

The first fact to note is that all of the parameters that are obtained by setting  $R_{s,init}$  between  $2.1\ \Omega$  and  $6.1\ \Omega$  are nearly identical, which is encouraging in that it suggests that a relatively wide range of initial guesses for  $R_s$ , be they obtained when the motor is either hot or cold, will suffice to

state	$f$ [Hz]	$R_{meas}$ [ $\Omega$ ]	$X_{meas}$ [ $\Omega$ ]
hot	5	11.375	0.00
hot	60	16.173	0.16
hot	100	17.724	0.22
cold	0	8.2	-
cold	60	12.582	0.12

TABLE III  
MEASURED MOTOR PARAMETERS, AS MEASURED ON 11/30/2007.

identify nearly the same set of parameters. Moreover, it is also apparent that some of the parameters are identical in all of the cases; for example, the mechanical parameters  $\beta$  and  $K$  are the same regardless of the initial guess. This is encouraging, as it indicates that the estimation method is able to characterize the mechanical phenomena that govern much of the behavior of the system. The same logic extends to the stator resistance  $R_s$ , and to  $R_r$  and  $X_m$  to a lesser extent.  $X_{ls}$  and  $X_{lr}$  appear to have a more substantial amount of variation upon first glance, but a closer look shows that the sum of the two parameters is nearly the same for all of the initial guesses. This makes some intuitive sense, as the impedance of the rotor leg of the model in Figure 2 is dominated by the rotor leakage reactance  $X_{lr}$  and this term in parallel with the mutual reactance  $X_m$  will be approximately equal to the  $X_{lr}$  as well.

In considering these parameters, however, it is important to remember that they are not necessarily directly indicative of the state of the motor; these parameters suggest that this model is overdetermined with respect to the observable data, since a number of different sets of parameters effectively yield the same observed currents and torque-speed curves (within measurement noise and numerical noise from the process of estimation). While these parameters were all obtained from the same set of initial guesses, the numerical issues surrounding this problem limit their uniqueness to a certain extent. Some of the parameters, such as  $R_s$ ,  $\beta$ , and  $K$ , appear to be well-determined, while this does not appear to be the case for some of the other parameters.

As is also apparent from Table IV, the residual  $\hat{i}_a - i$  was identical for all of the different values of  $R_s$ , but the torque speed curves were relatively different. Visualizing the deviation of these torque-speed curves is instructive to lend additional context to the residuals  $\|\tau - \hat{\tau}\|_2$  tabulated in the previous table, so all of them were compared to the torque-speed curve obtained from an initial guess of  $R_s = 4.1\ \Omega$  and plotted in Figure 21. It is clear that that these deviations, while growing larger for the values of  $R_s$  further away from  $R_s = 4.1\ \Omega$ , are still relatively small with respect to the system's torque-speed curve, and any of these curves will describe the torque-speed curve with sufficient accuracy for most experimental purposes. Nevertheless, the set of initial guesses close to the measured initial guess for  $R_s$  all yield






Fig. 21. Plot of the estimated torque-speed curves for the motor when the parameters of the motor are estimated with different initial guesses for the stator resistance  $R_s$ . The value of the initial guess for  $R_s$  corresponding to each curve is given in the legend.




Fig. 22. Observed and fitted motor currents for the cold motor; observed data are in blue, and fitted data are in red. The minimization was performed with only observations of the motor current.




Fig. 23. Residual from  $y_{fit} - y_{obs}$ .

nearly identical results; this fact is encouraging in that it indicates that the method has some robustness to comparatively poor initial guesses. These estimated torque-speed curves will be compared with torque-speed curves measured on the dynamometer in the following section, providing another means for evaluating the accuracy with which the estimated parameters describe the behavior of the induction motor.

1) *Cold motor; current-based estimation:* The second dataset studied is that in which the motor current is observed when the motor is at room temperature (“cold”<sup>1</sup>). This test was conducted by observing the startup current transient when the motor was turned on after a period of at least 12 hours during which it was at rest. The motor parameters were estimated using only these observations of the motor current; the observed motor current, as well as the predicted current that results from a simulation using the estimated parameters, are shown in Figure 22, with the corresponding residual illustrated in Figure 23. Qualitatively, these plots are quite similar to those of the hot motor with the current-only estimation. The predicted motor current fits the observed motor current quite well, and the residuals are of comparable size. While the airflow estimation method is largely reliant upon the parameters of the hot motor, the performance of the motor parameter estimation for the cold motor suggests that the method can accurately characterize the performance of the

<sup>1</sup>In comparison to the motor that has been running continuously for over an hour, which is most best described as “hot,” in the absence of any other descriptive designation.

Parameter	$R_s$ [ $\Omega$ ]	$R_r$ [ $\Omega$ ]	$X_m$ [ $\Omega$ ]	$X_{ls}$ [ $\Omega$ ]	$X_{lr}$ [ $\Omega$ ]	$\beta$ [ $N \cdot m$ ]
Value	4.96	3.24	53.67	1.64	9.23	$4.91e-04$

TABLE V  
COLD MOTOR PARAMETERS OBTAINED BY MINIMIZING SOLELY AGAINST THE CURRENT.

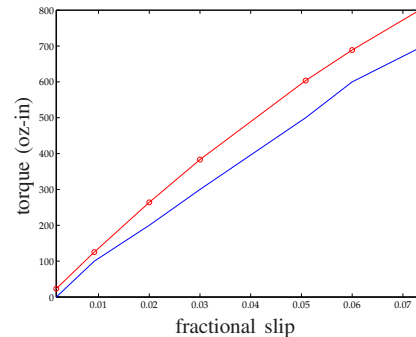


Fig. 24. Plot of the estimated and measured torque-speed curves for motor; the measured curve is in blue, and the estimated curve is in red.

motor regardless of its final temperature.

The parameters of the cold motor identified by the parameter estimation method are listed in Table V. The torque-speed curve generated from these parameter estimates, along with the set of observed torque-speed parameters for the cold motor as obtained from the dynamometer test stand, are plotted in Figure 24. The first observation that might be made by comparing these parameter estimates with the parameter estimates obtained from the hot motor data is that  $R_s$  and  $R_r$  are lower for the cold motor than for the hot motor, as would be expected. One might also note that expectations are also met by the fact that the mechanical parameters  $\beta$  and  $K$  are very similar to those estimated for the hot motor. In addition, the relationship between the predicted and observed torque-speed curves for the cold motor is similar to the same relationship for the hot motor; the predicted torque-speed curve is higher than the observed torque-speed curve.

A set of predicted torque-speed curves for a number of cold motor starts that were collected on different days is illustrated in Figure 25. These plots are also qualitatively similar to the analogous plots for the hot motor, providing additional




Fig. 25. Observed and fitted cold motor torque-speed curves for a number of motor starts on different days; fitted data are in red, and the torque-speed data from the motor is in blue.




Fig. 26. Zoomed torque-speed curves.




Fig. 27. Observed and fitted currents for the cold motor; observed data is in blue, and fitted data is in red. The minimization was performed with observations of both the motor current and torque-speed characteristic.




Fig. 28. Residual from  $y_{fit} - y_{obs}$

Parameter	$R_s$ [ $\Omega$ ]	$R_r$ [ $\Omega$ ]	$X_m$ [ $\Omega$ ]	$X_{ls}$ [ $\Omega$ ]	$X_{lr}$ [ $\Omega$ ]	$\beta$ [N·m]
Value	3.94	4.12	56.32	1.71	8.86	5.39e-04

TABLE VI  
FINAL COLD MOTOR PARAMETERS MINIMIZING AGAINST THE CURRENT AND THE TORQUE-SPEED CURVE.

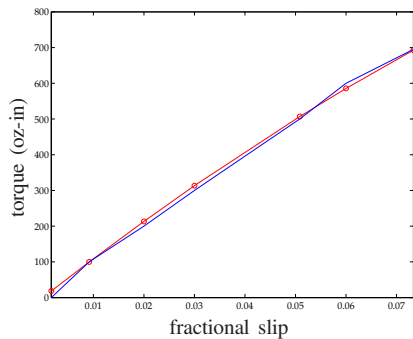


Fig. 29. Plot of the estimated and measured torque-speed curves for cold motor; the measured curve is in blue, and the estimated curve is in red.

evidence for the consistency of the estimated motor parameters for a given experimental configuration.

2) *Cold motor: torque-speed and current-based estimation:* The final set of motor data examined is that in which the parameters of the cold motor are estimated using both the current transient and the set of observed torque-speed datapoints for the cold motor. The prediction of the motor current from the motor parameters and the observed motor current are shown in Figure 27, and the respective residual is also illustrated in Figure 28. As was the case in the estimation of the hot motor parameters from both observations of the current transient and the torque-speed characteristic, these plots show that the prediction of the motor current does not fit the observed motor current as well when performing this type of minimization. The size of the residual in this case is even larger than before; while the fit still appears to be relatively good, the addition of the torque-speed constraints has appeared to have an adverse effect on the fit of the current transient.

The parameters estimated from this data are given in Table VI, and the torque-speed curve generated by these parameters is illustrated in Figure 29. As was the case for the two types of estimates for the hot motor, the sets of parameters given in Tables V and VI are relatively close; their differences can be ascribed to the adjustments required to fit the observed torque-speed datapoints. While the observed torque-speed curve is fit very well by the predicted parameters, it is unclear as to the overall effectiveness of such a parameter estimation method for this case, as the process of obtaining the measured torque-speed datapoints doubtlessly caused some




Fig. 30. Observed and fitted cold motor torque-speed curves for a number of unblocked fan motor starts on different days; fitted data is in red, and the torque-speed data from the motor is in blue. These motor parameters were obtained by fitting to the motor current and the torque-speed curve.




Fig. 31. Zoomed torque-speed curves.

resistive heating of the motor windings, thereby changing the motor's torque-speed characteristic. While the effects of such a process were minimized in the way that the datapoints were obtained, it is nearly impossible to eliminate such effects.

The torque-speed predictions generated by the sets of parameters that are obtained from estimates of three separate and unrelated datasets are shown in Figures 30 and 31, to study the variation existing in the method for a number of unrelated datasets. As would be expected, the resulting torque-speed curves are quite similar, since they are all obtained by minimizing against the same set of torque-speed observations. The predicted torque-speed curves are all very similar to much higher values of slip as well, suggesting that there is some value in extrapolating from the observed torque-speed datapoints, as long as the observed torque-speed datapoints accurately represent the behavior of the system under the desired operating conditions.

#### D. Effects of different operating conditions on parameter estimation

The final set of variation of operating characteristics of the motor that can be evaluated is that in which the air handling unit is operated in a variety of fault conditions. As mentioned previously, each of the faulty conditions is associated with a change in both the load on the fan motor and a change in the airflow cooling the motor; both of these types of changes could conceivably alter the operating torque-speed characteristics of the machine. These predicted torque-speed curves for a number of different air handler fault conditions are shown in Figure 32 in their entirety, and then zoomed in on the particular region in which the fan motor usually operates in Figure 33. Only the torque-speed curves for the current-based estimation are described below, as the parameter estimation




Fig. 32. Fitted hot motor torque-speed curves for a number of motor starts in different conditions; these conditions are listed in the legend. These motor parameters were obtained by fitting to the motor current only.




Fig. 33. Zoomed torque-speed curves.




Fig. 34. Cross-section of the duct illustrating the points where the measurements were made.

State	$R_s$ [ $\Omega$ ]	$R_r$ [ $\Omega$ ]	$X_m$ [ $\Omega$ ]	$X_{ls}$ [ $\Omega$ ]	$X_{lr}$ [ $\Omega$ ]	$\beta$	$K$	$R_r$ [ $\Omega$ ]
Unblocked	6.25	4.03	57.75	3.14	7.71	4.59e-04	31.00	4.26e-04
30% blocked	5.73	4.39	57.39	1.65	9.36	3.84e-04	27.22	5.06e-04
50% blocked	5.23	4.72	60.34	2.51	8.39	5.06e-04	27.41	2.47e-04
Leaky	5.64	4.84	56.53	2.72	8.10	5.06e-04	24.72	2.47e-04

TABLE VII

TABLE OF ALL ESTIMATED MOTOR PARAMETERS FOR DIFFERENT FAULT CONDITIONS.

with the torque-speed constrains exhibits almost no variation in the region of interest.

It is apparent from these different characteristics that the torque-speed curve of the motor is markedly affected by changes in the condition of the air handler. These sources of variation in the torque-speed characteristics can be better understood by considering one representative set of the model parameters for each these different fault conditions, as listed in Table VII. It is important to remember the caveat that some of the parameters appear to be underconstrained, as can also be observed in Table IV. That table showed that the parameters  $R_s$ ,  $\beta$ ,  $K$ , and  $R_r$  (to a lesser extent) are relatively consistent for a given set of data in the face of different initial guesses, while other parameters exhibited less consistency. There are a few plausible explanations for these changes in parameters. For example, the change in the load, due to the different amounts of air being coupled to the fan, could potentially cause the observed changes in the parameters. Another explanation is that the change in the airflow could cause the fractional slip to increase, thereby causing more resistive heating. The unusual curve and datapoint in this respect is that of the leaky duct; one would normally expect that the increased airflow and approximately equivalent slip would equate to a torque-speed characteristic that is comparable to the unblocked duct. Yet another source of variation could simply be changes in the environmental conditions of the laboratory. Nevertheless, while the reasoning behind these explanations is based upon some sense of physical intuition, the precise cause of these changes would ideally be characterized in a more controlled environment to understand the sources of variation in the system.

As this speed and torque modules have been explored and demonstrated fully in previous sections, the only module that remains to be studied is the flow prediction module. It is helpful at this point to recall the overall structure of the airflow estimation method: for a given condition of the air handler, the speed is first estimated from the slot harmonics present in the current, then this speed is used to identify the corresponding torque on the torque-speed curve that describing the motor's present state. With these estimates of the torque and the speed, the mechanical power into the fan can be estimated, and the airflow that corresponds to this measured power and speed can then be determined via a fan curve. Since the torque-speed curve is clearly dependent upon the temperature of the motor, the startup transient was obtained in the same manner as discussed in §??.

The airflow predictions generated by the above method were validated by measuring the airflow through the duct in an unblocked state and with the imposed faulty conditions. These validation tests were performed by measuring the airflow with a hot wire anemometer, as described in §??; for each specified fault condition, four traverses were measured with six samples of the airflow per traverse, so that a total of 24 points were sampled for each fault condition. A general schematic of this measurement process, not drawn to scale to enhance visual clarity, is illustrated in Figure 34. In this figure, the traverses are numbered one through four, and the sample points at which the anemometer measured the airflow are labeled A-F. These measurements were made at a set of points along the given duct diameters that are specified to produce an accurate average of the flow, as enumerated in [?, ch. 19].

As mentioned previously in §??, the speed of the fan varies somewhat substantially over the course of the day, so these airflow measurements were conducted over a period of relative speed stability in the evening. Before conducting these tests, the fan was turned on and allowed to reach thermal equilibrium, so that the airflow would not be affected by changes in the motor impedances. Due to the need for the fan motor to reach a new thermal equilibrium after changing the blockage, the test sequence ran as follows after the motor had been running in its previous state and reached equilibrium in that state: first, the inlet was blocked down by the specified amount (e.g., 30%). The fan was then allowed to run for 50 minutes to an hour to reach thermal equilibrium. After this

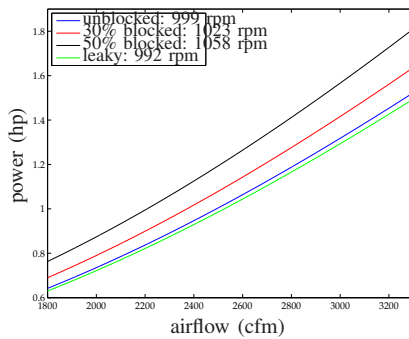


Fig. 35. Fan curves illustrating the power-flow relation at the four speeds of interest.

period of time, the set of traverses were conducted to measure the airflow at this level of blockage. After the traverses were complete, the inlet of the air handler was changed to reflect the next faulty condition under test.

The power-flow fan curves for the different speeds of operation as generated by the fan coefficients are shown in Figure 35 to compare the effects of the different speeds on the corresponding power-flow characteristic. This plot makes it clear that the different blockage conditions will have a significant effect on the flow through the duct.

It was necessary to identify both the mean and the standard deviation of the speed and torque estimates to estimate the flow to an acceptable degree of confidence. These quantities were therefore measured and used to run a series of simulations that could generate distributions of the airflow predictions. These distributions could then be used to quantify the probability that a given airflow measurement represented a given blockage condition. These statistics for the measured speed and torque, as well as the corresponding derived distributions for the power and airflow, are given in Table ??.

This table also includes the numerical values of the predicted airflow. These values were calculated using the following procedure: first, the measured mean and standard deviation for the speed were used to generate a set of 1000 simulated values of the speed that statistically matched the measured speed. Each of these speed values were then used to generate a distribution of torque points according to the statistical distribution of the predicted torque-speed curves. This resulted in a 1000x1000 matrix of torque estimates, in which each row corresponded to a speed estimate. Each element of this matrix was then multiplied by the corresponding speed estimate, which generated a similarly sized matrix of simulated power datapoints. These power datapoints were then used to estimate the corresponding airflow via the fan curve; this set of  $10^6$  simulated airflows could then be used to identify the mean and standard deviation of the airflow predictions, and evaluate them in comparison with the the corresponding velocity and airflow measurements for the unblocked, 30% blocked, and 50% blocked ducts, as given in Table ??. The air velocities in the leaky duct experiment are not included in the table of validation measurements, as they could not be measured upstream of the leak in the experimental apparatus.

As is illustrated in Table ??, both of the motor parameter estimation methods were used to generate airflow predictions




Fig. 36. Illustration of airflow detectability using torque-speed curves that are generated from minimization against only the motor current, as collected for each blockage condition.

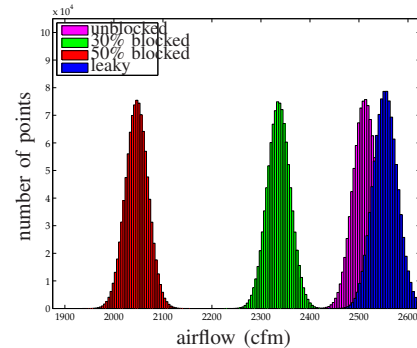


Fig. 37. Illustration of airflow detectability using a torque-speed curve that was generated from minimization solely against the unblocked motor current.

to assess the performance of each method. In the upper half of the table, the torque-speed curves obtained by performing the current-based parameter estimation were used to make the airflow predictions, while the airflow predictions in the bottom half of the table were generated by the torque-speed characteristics from the torque-speed and current-based parameter estimation methods. To facilitate the interpretation of the data in these tables, the histograms of the airflow predictions are also illustrated in three related plots, representing three possible approaches to the use of this FDD method. In the first of these plots, Figure 36, the airflow predictions are generated from four different torque-speed estimates which represent the system in its four faulty states. For each condition of the air handler (i.e. unblocked, 30% blocked, 50% blocked, and leaky), the system was run for two hours, then turned off and immediately turned back on to obtain a startup transient, and then the torque-speed curve used to generate the airflow prediction was obtained from the motor parameters identified by the current-only parameter estimation method. In the second plot of the airflow prediction histograms, Figure 37, all of the airflow predictions are generated from a single torque-speed curve. This torque-speed curve was obtained from a set of parameter estimates identified by the current-only estimation method operating on a current transient collected for the fan running in an unblocked state. In the final plot of the airflow prediction histograms, Figure 38, the airflow predictions were generated in a similar manner to those in Figure 36, with the essential difference that the torque-speed and current-based parameter estimation method was used to identify the



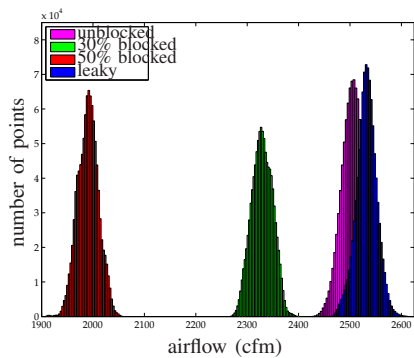


Fig. 38. Illustration of airflow detectability using torque-speed curves that are generated from minimization against the motor current and the torque-speed curve, as collected for each blockage condition.

parameters, rather than the current-only parameter estimation method.

All of these figures illustrate encouraging results. Specifically, it is clear from these figures that the task of differentiating the unblocked flow from either the 30% blocked flow or the 50% blocked flow appears to be relatively simple using either of the motor parameter estimation methods. In particular, the 50% blocked flow is extremely different from any of the other inlet conditions, and should be easily detectable. The 30% flow is also relatively different, and should also be easily detectable; it appears that the tails of the 30% blocked distribution and the unblocked distribution may have some minor overlap.

While the ability to qualitatively identify the presence of blockage is certainly useful, it is even more useful to quantitatively identify the particular airflow in each situation. Examining the unblocked airflow measurement and predictions from Table 36, both the current-based prediction (2511 cfm) and the torque-speed and current-based prediction (2507 cfm) are somewhat higher than the measurement (2305 cfm). The close agreement of the unblocked flows is immediately noticeable, suggesting that little is gained from the minimizing against the torque-speed curves in this scenario. In an absolute sense, these predictions are quite good, as the means of the predictions and measured value differ by about 9%; moreover, the standard deviation  $\sigma_{a,m}$  of the set of airflow measurements made by hand is 33 cfm, while  $\sigma_{a,p}$  for the predictions is approximately 25 cfm. In addition, the datasheet with the fan coefficients for the particular fan used specify that the values of airflow are accurate to within  $\pm 1.25\%$ , while the values of power are accurate to within  $\pm 2.5\%$ . While the prediction and the measurement would ideally be closer, they are sufficiently close to recommend this method for the prediction of airflow, and for the further characterization of the sources of inaccuracy, such as the bias possibly introduced in the process of making the anemometer measurements.

Turning to the values of the 30% blocked airflow, similar results are found: the predicted airflow using the current-based method is 2379 cfm, and the predicted airflow using the torque-speed and current-based method is 2328 cfm, while the measured airflow is 2168 cfm. The two different parameter estimation techniques yield slightly different results for this level of fault, but the two standard deviations of the sets

Fig. 39. Empirically measured fan curve for the blower wheel used in experimental air handler.

of measurements ( $\sigma_{a,p} = 18$  cfm and  $\sigma_{a,p} = 20$  cfm for the current-based and torque-speed and current-based minimization, respectively) suggest that these differences might be attributed in part to measurement noise and bias. One particularly interesting observation is that the difference between the predicted and measured airflows for the no-fault condition for the air handler (206 cfm) is almost the same as the difference between the predicted and measured airflows for the 30% blocked air handler (211 cfm). While this may be partially attributed to a particularly fortuitous set of measurements, it also suggests that the source of bias has a relatively constant value. If this is the case, this bias may be able to be easily calibrated out of the system, thereby increasing the accuracy of the airflow measurements.

The analysis of the 50% blocked fault level is quite similar to that for the 30% blocked fault. It might be particularly noted that both of the airflow predictions are identical (1991 cfm). One change from the 30% condition is that the difference between the predicted and measured flows is only 86 cfm, rather than 210 cfm. One potential cause of this change is the loss of model accuracy due to the fact that the present operating condition of the fan lies slightly outside the region of operation used to generate the fan coefficients, causing a larger residual between the empirical data and the predictions obtained from the fan coefficients in this region. A comparison of the values calculated using the empirical data given in Figure ?? illustrates that the power predicted by the fan coefficients at a given speed and airflow is quite close to the power predicted by the empirically obtained set of curves. Private correspondence with Tim May from Lau Fan Corp. suggests that the prediction from the fan curves at this operating point has a lower accuracy than predictions generated by the fan curves in other regions of operation. Nevertheless, this data indicates that the airflow at this level of blockage can be known with similar accuracy as the other fault and no-fault conditions.

Yet another reason that these differences in the torque-speed curves between the two estimation methods diminish as the fractional slip decreases is that the sensitivity of the model that generates the torque-speed curve decreases as the induction motor runs closer to synchronous speed. This fact is apparent from noticing that the curves tend to converge as the slip gets closer to zero in Figure ?. It is thus reasonable to expect the discrepancy between the predictions from the two methods to get smaller as the fractional slip decreases. It is also striking that the torque-speed curves do not apparently yield significantly better estimates of the airflow; by eliminating the need to obtain *a priori* observations of the torque-speed characteristic, the process of identifying the torque produced by the electric motor is much more straightforward.

Analysis of the leaky duct fault condition for the two different motor parameter estimation methods provides perhaps the most surprising result. If the operating conditions of the motor were identical for the fan in both the leaky fault condition and the unblocked condition, one would expect that the flow through the fan would be greater in the leaky

condition because there would be a smaller pressure drop across the fan due to the open access door. This intuitive expectation is reinforced by the airflow predictions made with the torque-speed and current-based model; since the torque-speed characteristics of the leaky and unblocked motors obtained by this method are nearly identical, the increase in fractional slip results in a higher amount of torque produced, as well as a higher quantity of mechanical power delivered. In comparison, the torque-speed curve for the current-based estimation is sufficiently different from the torque-speed curve for the leaky duct that the airflow through the leaky duct is lower than the airflow through the unblocked duct. This begs the question: what causes the torque-speed curve to change this much? One possible cause is that the torque-speed characteristic is dramatically affected by the change in the load due to the leaks in the ducts, as the increased load on the fan could have thermal effects causing changes in the winding impedances, which in turn could cause the motor to produce less torque at a given speed.

Unfortunately, the particular experimental setup used precludes the empirical measurement of the airflow upstream of the leak and direct determination of the airflow through the fan, as the high degree of turbulence upstream of the access door makes spatial averaging with anemometer measurements in the short upstream length of duct all but impossible. In addition, experimental limitations preventing direct torque measurements being performed in the fan *in situ* make it difficult to explore and identify potential causes for the apparent changes in the torque-speed curve.

Nevertheless, consideration of these results suggests a number of effects upon the implementation of this type of system in a field installation. There are two distinct scenarios in which this airflow estimation method might be used: for validation and commissioning of the air handler when it is installed, and for tracking the airflow delivered to the ducts in normal operation; in the terminology of Chapter ??, the system could be used for detecting either commissioning faults or operational faults. As these results indicate, this airflow prediction method can identify the changes in the flow quite well and can thus identify operational faults. While the assumption that the system is commissioned correctly is not always accurate, many FDD systems are constructed on the basis of such assumptions, and it is useful to consider the capabilities of an airflow prediction system that only needs to monitor ongoing faults. Moreover, such an assumption is often quite useful from the standpoint of implementing a FDD system, as it requires a minimum of *a priori* information for the airflow prediction system, since no correction factor or measurement of a torque-speed characteristic is needed to identify operational faults which result in a change in the volumetric airflow.

The identification of the commissioning faults is more difficult, however, and requires more information than do the operational faults. The structure of possible field implementations of an airflow prediction system which is able to identify commissioning faults depends upon the desired accuracy of the method and the validation of the torque-speed behavior observed in the fan motor from the current-only parameter

estimation. Comparison of the results illustrated in Figures 36 and 38 emphasizes the important function served by the torque-speed curve, and the method is largely dependent upon the accuracy of these torque-speed estimates.

One possible instance of a field implementation of an FDD method is that additional information about the torque-speed curve is either unavailable or too expensive to attain in more detail; however, the results presented previously suggest that the airflow prediction method could still be extremely useful in the absence of any such correction factor. For example, operational faults can be detected by looking for changes in the airflow predictions as in comparison to the system behavior observed in the initial unblocked state, which is established when the air handler is installed and commissioned. While the absence of a correction factor prevents this commissioning data from being completely accurate, the fact that the predictions are within 10% suggests that such estimates still have considerable utility. Different blockage conditions could then be identified and the reduction in airflow determined by the use of the prediction method to identify reduced flow and higher fan speeds, as is clear from the results given in Table ???. One possible application of such a system would be for detecting relatively large commissioning faults. This could be accomplished by comparing the airflow predictions to the results of a traverse taken downstream, but before any branches leave the duct from the air handler.

Leaky ducts could also be detected for a properly commissioned system, though the particular torque-speed curve used changes the way in which the leak is detected. As is clear from Figure 36, if new torque-speed curves are obtained after every period of operation, the presence of a leak will result in a prediction of reduced airflow and a reduced fan speed. Two other possibilities for commissioning and operating the FDD method could also be used, both of which result in the diagnostic indicator for a leak consisting of an increased airflow and a reduced speed: either a single torque-speed curve is obtained from a current-only estimation routine for the unblocked system when it is commissioned (illustrated in Figure 37), or the torque-speed curve is obtained for each operational cycle by using the torque-speed and current-based estimation routine.

While the above discussion suggests that additional accuracy in airflow predictions is not required for the construction of a highly useful FDD method, the application of a correction factor which can compensate for the apparent variation in the torque-speed curves could measurably improve the performance of the prediction method. Such a correction factor would not depend on the particular installation, as it would be largely dependent upon the motor dynamics. One particularly suggestive use of the airflow prediction method could be for the independent verification of commissioning accuracy, as results from the prediction method could be compared to the expected airflow after the installation of the air handler, and any divergence between the two could be identified and repaired.

Such diagnostic tools could have a variety of uses for diagnostics in buildings. For example, the airflow through a ducting system changes as adjustments are made to dampers and so

forth; with prior knowledge of how these changes should be affecting the system, the airflow detection method would be able to use knowledge obtained only by the use of electrical measurements of the system response to determine whether or not those responses are within specifications. Even without the added functionality of detecting commissioning faults, this could be very useful. A variety of different operational faults detection methods and energy monitoring applications could be constructed with such tools, such as methods for identifying time-varying energy consumption of a residential building, and so forth.

## VI. DISCUSSION

PARAGRAPH: Drawbacks/shortcomings of this method, address these problems.

PARAGRAPH: THESIS While this approach to identifying airflow faults has much to recommend it, it is not without its own shortcomings. In particular, it is important to note that this fault detection method is susceptible to diagnostic features indicative of airflow faults that are actually caused by other unrelated changes in the system. Since the diagnostic method evaluates the state of the airflow on the basis of the electrical terminal variables of the motor, it will be difficult to distinguish between electrical behavior related to the airflow fault and electrical behavior which is unrelated to the airflow fault but which has a similar manifestation. Many common motor faults are unlikely to be problematic, however, as such faults (broken rotor bars, shorted windings) affect one phase winding more than the others, while airflow faults will affect all three phases identically. Other motor faults, such as bearing failures, could qualitatively be similar to airflow faults, but even these often have distinguishing frequency characteristics that have been modeled extensively, allowing them to be identified separately [?].

It is also important to note that this method of estimating airflow is largely based upon the correlation of changes in  $P_f$  to changes in  $Q_f$  via the fan curve. This method therefore fundamentally cannot tell whether or not the change in  $P_f$  is caused by an additional pressure drop upstream of the fan, or downstream of the fan. This makes it difficult to identify whether or a change in the fan performance is caused by a change in state of the filter, or by additional duct leakage. In addition, if two simultaneous faults occur, one of which has the effect of reducing  $P_f$  while the other has the effect of increasing  $P_f$ , the diagnostic indicator will remain unchanged, preventing the fault diagnostic method from identifying either fault.

Fortunately, other information may be available in many of the potential circumstances in which these difficulties arise. For example, while simultaneous leakage and blockage are not independently distinguishable, ventilation systems may be more susceptible to leakage during the commissioning phase and more susceptible to blockage that accumulates slowly over time. Information about changes in the system performance over these different timescales could thus be used to differentiate between these types of faults. Such concerns would have to be addressed when testing this fault diagnostic method

in field installations to better characterize its performance in residential or commercial settings.

This diagnostic method was developed using a three-phase induction machine, since this type of machine is generally quite common due to their relatively low manufacturing cost. These types of motors are often used in larger building ventilation situations, where three-phase power is available, while single-phase induction motors are more typically used in residential locations. Due to the wealth of diagnostic techniques and research into models of this type of motor, this research will be focused on developing diagnostic methods for three-phase motors, rather than single-phase motors. While the fact that these motors are not as popular in residential applications may appear to limit the utility of this diagnostic method, three-phase motors were studied because the modeling techniques used are better understood for three-phase than single-phase motors. These motors were also studied because it was desired to the feasibility of the overall diagnostic method in as straightforward a manner as possible, without becoming mired in the technical details of the method's implementation for a single-phase motor. Such an implementation can be developed analogously after the method has been proven to work.

## ACKNOWLEDGMENT

The authors would like to thank Professor James Kirtley at MIT and Professor Robert Cox at UNC-Charlotte for their valuable feedback and input, and Tim May from Lau Fan Co. for being just frickin' awesome.

## REFERENCES

- [1] A. Cowan, "Review of recent commercial rooftop unit field studies in the Pacific Northwest and California," New Buildings Institute, PO Box 653, White Salmon, WA, 98672, Tech. Rep., 8 Oct. 2004.
- [2] R. Mowris, A. Blankenship, and E. Jones, "Field measurements of air conditioners with and without txvs," in *2004 ACEEE Summer Study on Energy Efficiency in Buildings*, 2004, pp. 212–227, panel 1.
- [3] D. Hales, A. Gordon, and M. Lubliner, "Duct leakage in new washington state residences: findings and conclusions," *ASHRAE Transactions*, vol. 109, no. 2, pp. 393–402, 2003.
- [4] J. Cummings, C. Withers Jr., J. McIlvaine, J. Sonne, and M. Lombardi, "Air handler leakage: field testing results in residences," *ASHRAE Transactions*, vol. 109, no. 1, pp. 496–502, 2003.
- [5] K. Srinivasan, "Measurement of air leakage in air-handling units and air conditioning ducts," *Energy and Buildings*, vol. 37, pp. 273–277, 2005.

- [6] F. Carrie, A. Bossaer, J. Andersson, P. Wouters, and M. Liddament, "Duct leakage in european buildings: status and perspectives," *Energy and Buildings*, vol. 32, no. 3, pp. 235–243, Sep. 2000.
- [7] ASHRAE, *ASHRAE Handbook: Fundamentals*. Atlanta, GA: ASHRAE, 2001.
- [8] H. Li, "A Decoupling-Based Unified Fault Detection and Diagnosis Approach for Packaged Air Conditioners," Ph.D. dissertation, Purdue University, Aug. 2004.
- [9] J. House, W. Lee, and D. Shin, "Classification techniques for fault detection and diagnosis of an air-handling unit," *ASHRAE Transactions*, vol. 105, no. 1, pp. 1087–1097, 1999.
- [10] W. Lee, J. House, and D. Shin, "Fault diagnosis and temperature sensor recovery for an air-handling unit," *ASHRAE Transactions*, vol. 103, no. 1, pp. 621–633, 1997.
- [11] W. Lee, C. Park, and G. Kelly, "Fault detection in an air-handling unit using residual and recursive parameter identification methods," *ASHRAE Transactions*, vol. 102, no. 1, pp. 528–539, 1996.
- [12] A. Glass, P. Gruber, M. Roos, and J. Todtli, "Qualitative model-based fault detection in air-handling units," *IEEE Control Systems Magazine*, vol. 15, no. 4, pp. 11–22, Aug. 1995.
- [13] —, "Preliminary evaluation of a qualitative model-based fault detector for a central air-handling unit," in *Proceedings of the Third IEEE Conference on Control Applications*, vol. 3, Glasgow, Scotland, U.K., 24 Aug. 1994, pp. 1873–1882.
- [14] P. Vas, *Parameter Estimation, Condition Monitoring, and Diagnosis of Electrical Machines*. Clarendon Press, 1993.
- [15] R. Park, "Two-reaction theory of synchronous machines: Generalized method of analysis, part i," *Transactions of the AIEE*, vol. 48, pp. 716–727, 1929.
- [16] K. Wang, J. Chiasson, M. Bodson, and L. Tolbert, "A nonlinear least-squares approach for identification of the induction motor parameters," *IEEE Transactions on Automatic Control*, vol. 50, no. 10, pp. 1622–1628, Oct. 2005.
- [17] G. Seber and C. Wild, *Nonlinear Regression*, ser. Probability and Mathematical Statistics. Wiley, 1989.
- [18] S. Shaw, M. Keppler, and S. Leeb, "Pre-estimation for better initial guesses," *IEEE Transactions on Instrumentation and Measurement*, vol. 53, no. 3, pp. 762–769, Jun. 2004.
- [19] P. Krause, O. Wasynczuk, and S. Sudhoff, *Analysis of Electric Machinery*. McGraw-Hill, 1986.
- [20] M. Stocks and A. Medvedev, "Estimation of induction machine parameters at start-up using current envelope," *Conference Record of the Industry Applications Conference*, pp. 1163–1170, 13 Oct. 2002.
- [21] R. Supangat, N. Ertugrul, W. Soong, D. Gray, C. Hansen, and J. Grieger, "Detection of broken rotor bars in induction motor using starting-current analysis and effects of loading," *IEEE Proceedings on Electric Power Applications*, vol. 153, no. 6, pp. 848–855, Nov. 2006.
- [22] H. Engl, M. Hanke, and A. Neubauer, *Regularization of Inverse Problems*. Springer, 2000.