Proceedings of the ASME 2015 International Design Engineering Technical Conferences & Computers and Information in Engineering Conference
IDETC/CIE 2015
August 2-August 5, 2015, Boston, USA

DETC2015-47000

DETECTION OF CHILLER ENERGY EFFICIENCY FAULTS USING EXPECTATION MAXIMIZATION

R Lily Hu*
Mechanical Engineering
University of California, Berkeley
Berkeley, California 94720
Email: lhu@berkeley.edu

Jessica Granderson
Building Technology and Urban Systems
Lawrence Berkeley National Laboratory
Berkeley, California, 94720
Email: jgranderson@lbl.gov

Alice Agogino
Mechanical Engineering
University of California, Berkeley
Berkeley, California 94720
Email: agogio@berkeley.edu

ABSTRACT

To detect degradation in energy efficiency of a chiller in a chiller plant, a multivariate Gaussian mixture model is applied. This classification technique was selected to take advantage of an expected correlation between measurable state variables and equipment and operation specifications and system control targets. The hidden variable is the faultiness of the chiller and can take on one of three possible states. The five observed variables correspond to sensor measurements that are typically available and monitored in commercially available chiller plants. The fault detection algorithm is trained on simulated data for the Molecular Foundry at the Lawrence Berkeley National Laboratory and tested on measured sensor data. The results show that detection of severe faults and no faults are relatively accurate, while detection of moderate faults is sometimes mistaken for severe faults. The computation needs are moderate enough for deployment and continuous energy monitoring. Future research outlines the next steps in regards to sensitivity analyses with alternate probability density functions.

INTRODUCTION

Commercial buildings consume 19% of US primary energy [1]. Of this, an estimated 15% to 30% of energy used in commercial buildings is wasted by poorly maintained, degraded, and improperly controlled equipment [2]. Much of this waste can be prevented with automated fault detection and diagnostics (FDD). Therefore, an enormous opportunity for energy savings is in monitoring of energy use and continuous fault analysis. The median whole building energy savings is 16% and just 13 of the most common faults in U.S. commercial buildings in 2009 were believed to have caused over $3.3 billion of wasted energy [3]. For continuous fault analysis, manual FDD can be tedious. To automate the process of fault detection, a statistical learning technique is applied.

The specific energy application under investigation is the detection of efficiency degradation of chillers in a central chiller plant. These plants are used in the commercial and industrial buildings sectors. A central chiller plant is a large facility that provides centralized chilled water to cool large buildings and multi-building campuses. Within the chiller plant, the chillers are the refrigeration machines that provide the chilled water. Chillers are typically industrial-grade machines and are designed to last more than 20 years. Over such time scales, performance degradation and the resulting costs of energy can build up, hence our focus on this application.

BACKGROUND

Fault detection and diagnostics have been successfully applied in a range of engineering domains including manufacturing, in general, and in the aerospace and automotive industries. For an overview of the broad categories of FDD approaches, readers are referred to a review of FDD in general [4] [5] [6] and a review of FDD as applied to building systems [2] [7].

*Address all correspondence to this author.

These works first classify methods and then compare, and contrast the three FDD categories: qualitative model-based, quantitative model-based, and process history-based methods. Review papers are also available for surveying specific categories of FDD, such as supervisory methods, model-based techniques, and trends in applications of model-based techniques [8].

In addition to reviews of FDD methods, the literature also includes proposed FDD algorithms applied to building systems [9] [10] [11]. The FDD approaches proposed in the literature includes rule-based methods, expert systems, physical models, black box models, etc. What makes chillers an interesting problem is that they operate over a range of efficiencies depending on the required load. Therefore, a good FDD method needs to consider the influences of variables that could impact operating conditions. This paper investigates a statistical learning approach to account for these variations.

Much of the research on chiller applications focuses on internal chiller components, for example, oil levels and refrigerant leaks within the chiller. This paper, on the other hand, analyzes the chiller in the context of sensor data available at the chiller plant level. The benefit of this approach is that newer chillers already have FDD for internal components built in by the manufacturer, but system level FDD is rarely used in current business-as-usual operation and maintenance of commercial buildings.

Furthermore, existing FDD built into equipment, or people themselves, can detect hard faults — failures that occur abruptly and either cause the system to stop functioning or to fail in meeting comfort conditions. On the other hand, soft faults that degrade performance but allow continued operation of the system are more difficult to detect and diagnose, and can continually waste energy [12]. One example of such a fault is degradation in chiller efficiency. For these reasons, this paper focuses on a method of detecting soft faults.

METHODOLOGY

The process used in this project to detect a soft fault is as follows:

1. Select a fault of interest.
2. Hypothesize which sensor data would be useful to detect this fault.
3. Simulate data under faulty and non-faulty conditions. Preferably, data from a range of faulty and non-faulty operations would be acquired from measured sensor points, possibly from maintenance data or by introducing faults to a physical system. However, it is important to recognize that a wide range of operational fault data may be hard to acquire with any statistical accuracy and simulated fault data may be needed.
4. Investigate which sensor data is most useful in detecting this fault and select those sensor points for inclusion in the fault detection algorithm.
5. Chose or develop a mathematical model for fault detection.
6. Define the most critical variables in the mathematical model and select a method to calculate the unknown parameters, e.g., train for the parameters.
7. Train the model on a simulated data set.
8. Test the model on simulated data, different from the training set.
9. Test the model on measured operational data.
10. Analyze the results.
11. Iterate.

These steps are applied to a chiller efficiency application and described sequentially in the next chapters and sections.

Fault Selection

To pick a chiller plant fault to focus on, a literature review of scientific and engineering practice and consultations with chiller plant staff and building researchers resulted in the compilation of a list of approximately 50 candidate faults. These faults were then ranked according to the impact on energy, ease of detection, and frequency of occurrence using input from researchers at LBNL and feedback from operators of a similar plant at the U.S. Washington Navy Yard. Any faults that could already be detected by the chillers own control system were trimmed from the list. Additional considerations in ranking the faults were the availability of data and the ability to generate faulty data through manipulating variables in a physics-based model of the chiller plant.

As a result of these investigations, the fault selected for focus was degradation of chiller efficiency across the range of cooling loads delivered. This fault occurs when the efficiency of the chiller is lower than expected, given the current operating conditions of the chiller. Chillers have efficiency curves that vary with cooling load even under ideal conditions. These efficiency curves are commonly drawn on a plot of kW/ton vs. ton. The chiller efficiency fault is particularly interesting because the efficiency of an ideal chiller naturally depends on the operating conditions, which include entering and leaving chilled water temperatures, condenser water temperatures, and the loading on the chiller. This property can be difficult to capture in simpler fault detection methods such as thresholds on a metric such as power or kW/ton. Thus, this fault is a good candidate for a statistical learning fault detection method.

VARIABLE SELECTION

An initial candidate list of 30 variables corresponding to sensor points was selected. Then, this list was reduced to the most significant five variables through exploratory data analysis.

By plotting the candidate variables, some of them were found to be almost constant or unchanging even when the efficiency of the chiller decreases. Some of these variables have the same pattern and distribution regardless of the efficiency of the chiller. Thus, some of the variables were independent of the efficiency of the chiller and were poor candidates as variables to detect inefficiency.

The choice of the five variables that were selected is consistent with physical intuition of the efficiency degradation. Four of the variables are associated with Chiller 1: power consumption, instantaneous coefficient of performance (COP), temperature of the entering condenser water, and the temperature of the exiting condenser water. All of these values are physically related to the efficiency of the chiller.

The last variable corresponds to the exiting condenser water temperature of the other chillers. Ideally working chillers should have an exiting condenser water temperature that is close to the design specification, and both Chiller 2 and Chiller 3 are the same model. The exiting condenser water temperature would be expected to be higher in the case of a faulty chiller, since more electrical energy is needed to supply the necessary cooling, and hence more heat is generated that results in a larger increase in the condenser water temperature. The exiting condenser water temperature of the faulty chiller can be compared against the corresponding temperatures at the other chillers, which are not faulty. To introduce dependence on historical data, the final variable list included the five sensor points at time t, and at time $t-1$.

TESTING FACILITY

The specific chiller plant used as a case study is the chiller plant at the Molecular Foundry at the Lawrence Berkeley National Laboratory. The Molecular Foundry is a 6-floor nanoscience research facility that is served by a chiller plant that has three chillers operational for 24 hours a day, seven days a week. A simplified schematic of the chiller plant at the Molecular Foundry is shown in Figure 1. Chiller 1 was selected as the location of our investigations of a possible chiller efficiency fault due to the greater availability of sensor data associated with it with which to test the algorithm. Data for many of the sensor points associated with the other chillers was unavailable.

TRAINING DATA

To train the fault detection model, data for normal operations and for operations under faulty conditions were generated using a physics-based model of the plant written in Modelica. It was undesirable to introduce faults in the real plant to collect measured data because introducing faults would interfere with normal operations of the facility and damage the expensive equipment. The physics-based model was created and calibrated by a research partner using the Modelica modeling language within

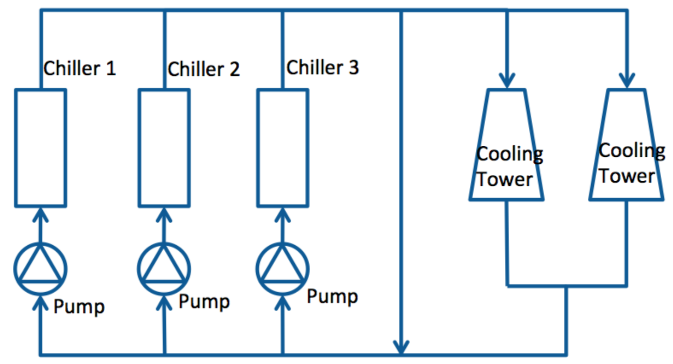

FIGURE 1. SCHEMATIC IMAGE OF THE CONDENSER WATER LOOP OF THE CASE STUDY FACILITY

the Dymola modeling and simulation environment.

The training dataset includes 24 hours of 10-minute data of the chiller plant under nominal conditions, under conditions with a moderate energy efficiency fault in Chiller 1, and under conditions with a severe energy efficiency fault in Chiller 1. Because the fault detection focuses on the performance of Chiller 1, data was only considered for when Chiller 1 was operating. This is because when Chiller 1 was not operating — i.e., not turned on — one cannot deduce information about its operation nor its operational efficiency. The sampling frequency of the training dataset was selected to be 10 minutes with a range from March 11 to March 21 to match the data from sensor measurements at the plant. The model inputs include the setpoints of the plant and program default standard yearly weather profile for San Francisco. The physics-based program simulates the plant and solves for all thermodynamic states for the whole-plant.

The chiller efficiency faults were introduced by modifying Chiller 1's reference coefficient of performance (COP). The COP for cooling systems is an efficiency metric for chillers defined as the ratio between the cooling output of the chiller and the electrical input to the chiller, as shown in Equation 1.

$$COP_{cooling} = \frac{Q_C}{W} \tag{1}$$

where Q_C is the rate of heat removal and W is the electrical input.

The nominal reference COP for Chiller 1 is 5.42. Reference COPs for a moderate energy efficiency fault was chosen as 3.0 and the COP under a severe efficiency fault was 1.0. The COPs chosen for the fault conditions are below the minimum efficiency COP dictated by ASHRAE Standard 90.1-2004.

MEASURED SENSOR DATA

Data from sensors at the Molecular Foundry was collected from July 23, 2012 to April 1, 2013 when continuous data was available during a season when the chillers were operating. Fault conditions were not recorded and thus it is not known when faults occurred during this time period. The only data point not available from the list of sensors of interest was Chiller 1 COP and thus it was estimated from the following data points available: Chiller 1 entering chilled water temperature, Chiller 1 exiting chilled water temperature, Chiller 1 mass flow rate. The resulting dataset covered 750 timestamps.

To coincide with the instantaneous COP in the training dataset, the instantaneous COP was also calculated from the sensor data using Equation 2:

$$COP_{CH1,t} = \frac{\dot{m}_{CH1,CHW,t}(CWT_{CH1,entering,t} - CWT_{CH1,leaving,t})}{w_{CH1}} \tag{2}$$

where all variables are in reference to the chiller: $\dot{m}_{CH1,CHW,t}$ is the mass flow rate of the chilled water, $CWT_{CH1,entering,t}$ is the temperature of the entering chilled water, $CWT_{CH1,leaving,t}$ is the temperature of the exiting chilled water, and w_{CH1} is the power consumed by the chiller.

Fault Detection Algorithm

The fault detection algorithm used is based on a mixture model, which can be used to cluster data into one of the mixture component clusters. In this case, mixture components represent the faultiness of the chiller.

A mixture model is chosen over estimation of a single model parameter because mixture models have not been used in the literature for this specific fault detection problem. It would be interesting to frame fault detection as a mixture model, by viewing fault states as different mixture components. One strength of this framework is that each fault state can have a unique probability distribution of possible observations from sensors. The parameters of the mixture model can then be learned using standard techniques. This model also gives probabilistic assignments of faultiness and can be interpreted as probabilities. In contrast, single model parameter estimation may not be sufficient due to the range of conditions seen even under normal operations. These conditions include start-up and shut-down of chillers, and transients when responding to changes in cooling demand, or due to control sequences, or behaviour of other equipment that is connected in the chiller plant. A mixture model of different probability distributions is more flexible than a single distribution.

A mixture model is a parametric probability density function represented as a weighted sum of component (or subpopulation) probability densities. In this case, the mixture components are modeled as multivariate Gaussian emission probabilities because the observed variables are expected to be centered on a mean; this mean could be, for example, the chiller manufacturers equipment specification, the as-designed chiller plant specification, or the operation control setpoints. Furthermore, these variables are likely correlated because of their physical meaning. That is, underlying physical laws must be obeyed. However, because the system is complex and non-ideal, first principles cannot be used directly. However, covariances can capture some of the dependencies between variables.

To explore how well the data is modelled by a normal distribution, the following figures graph histograms and normal probability plots of the training data. The histograms compare the data to a normal probability curve generated from the mean and standard deviation of the data, as seen in Figures 2, 3, and 4. A normal probability plot is a type of Q-Q plot (quartile-quartile plot) that plots the quartiles of the data against the normal distribution, as seen in Figures 5, 6, and 7. The straight line represents the normal distribution and departures of the data from the line indicate departures from normality. From these graphs, the training data can be seen to somewhat approximate a normal, with larger deviations from normality at the extremes of the data.

A Gaussian mixture model assumes that the data is generated from a mixture of a finite number of Gaussian distributions with unknown parameters. This can be expressed as:

$$p(x \mid \mu, \Sigma) = \sum_i \pi_i N(x \mid \mu_i, \Sigma_i) \tag{3}$$

$$= \sum_i p(Z_i = 1 \mid \pi_i) p(x \mid Z_i = 1, \mu_i, \Sigma_i) \tag{4}$$

where the π_i are the mixing weights and are constrained to sum to one, x is the data, and Z is the component label. N is an abbreviation for the Gaussian function.

HIDDEN VARIABLE

The mixture model components are the energy efficiency fault states of the chiller: *no fault*, *moderate fault*, and *severe fault*. The observed variables are the sensor data. Using this model, sensor data can be classified into one of the fault states of the chiller to detect the faultiness of the chiller. In this application, the fault state Z represents the chiller efficiency. The efficiency of a chiller can take on any value on a continuous range from 0 to a theoretical maximum based on thermodynamics. For simplicity, the efficiency of the chiller has been discretized. Hence, Z becomes a multinomial random variable and the components Z_i are displayed in Table 1.

OBSERVED VARIABLES

The observed variables are the five sensor measurements described in the variable selection process. To create a dependency

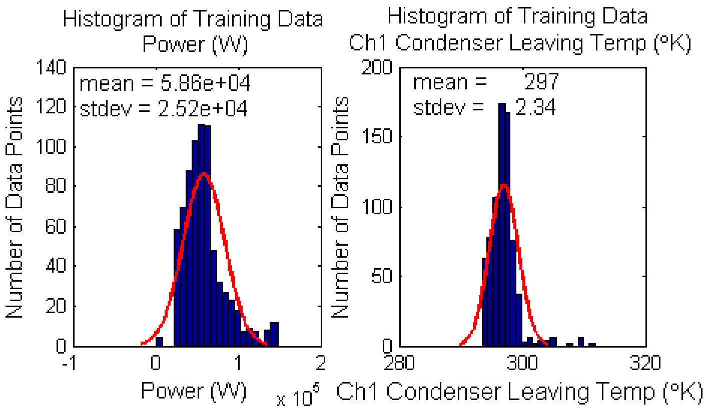

FIGURE 2. HISTOGRAMS OF THE NO-FAULT DATA COMPARED TO GAUSSIAN DISTRIBUTIONS

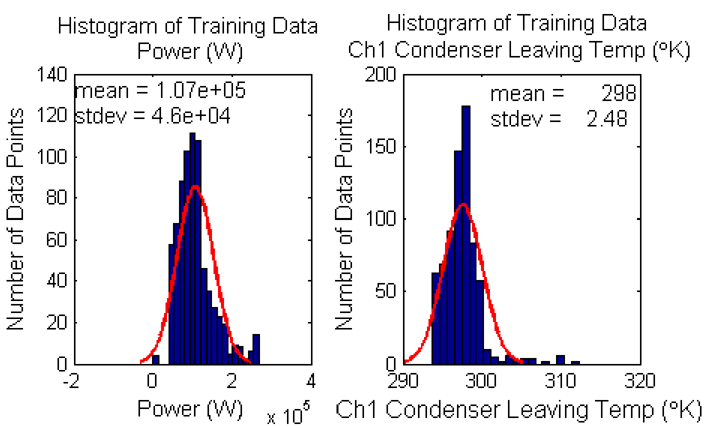

FIGURE 3. HISTOGRAMS OF THE MODERATE FAULT DATA COMPARED TO NORMAL DISTRIBUTIONS

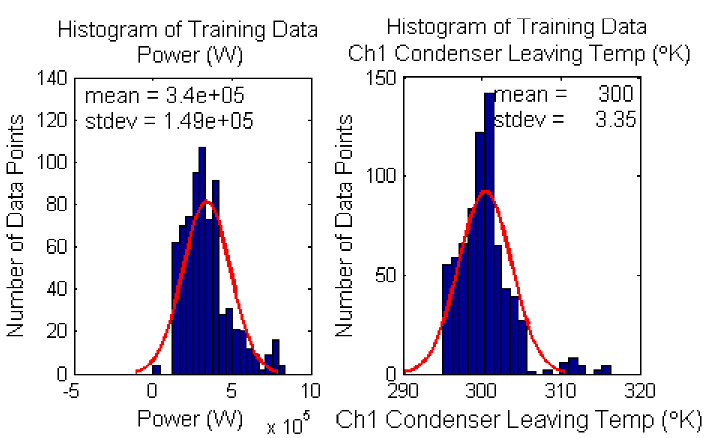

FIGURE 4. HISTOGRAMS OF THE SEVERE FAULT DATA COMPARED TO NORMAL DISTRIBUTIONS

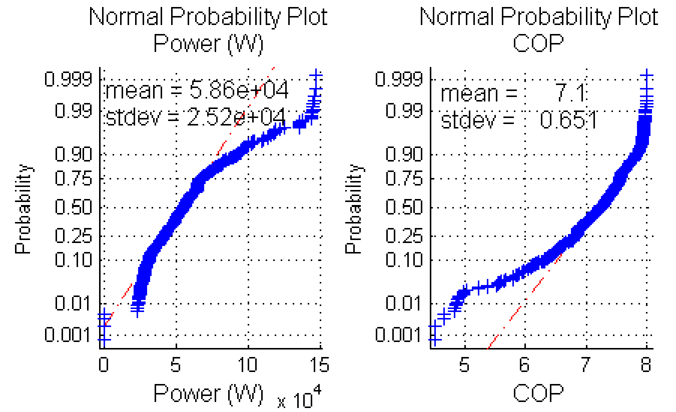

FIGURE 5. NORMAL PROBABILITY PLOTS OF THE NO-FAULT DATA

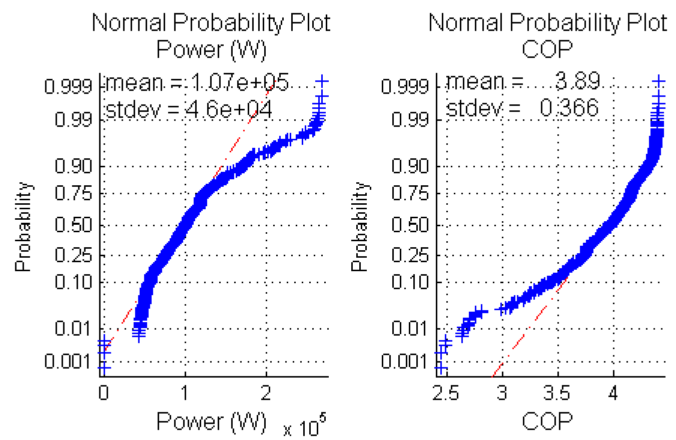

FIGURE 6. NORMAL PROBABILITY PLOTS OF THE MODERATE FAULT DATA

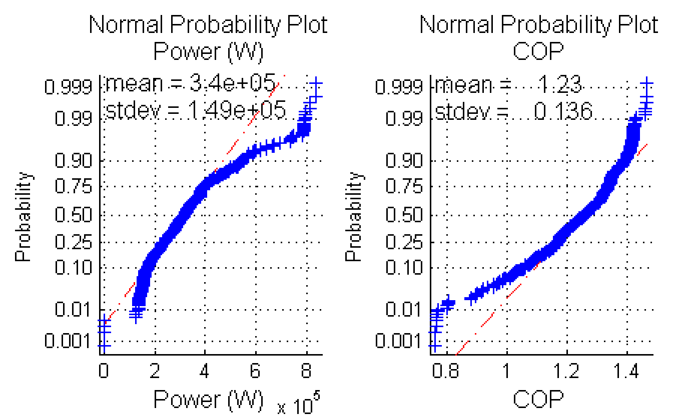

FIGURE 7. NORMAL PROBABILITY PLOTS OF THE SEVERE FAULT DATA

TABLE 1. Components of the hidden variable Z_i and their physical meaning

Colloquial label	Physical interpretation
No energy efficiency fault	COP = 5.42
Moderate energy efficiency fault	COP = 3
Severe energy efficiency fault	COP = 1

on previous sensor measurements in time, the observed variables included measurements at both time t and time $t-1$. Thus, there are a total of ten variables: five sensor measurements at time t and five sensor measurements at time $t-1$. Figure 8 shows the graphical model of the Gaussian mixture model and the interpretation to this application.

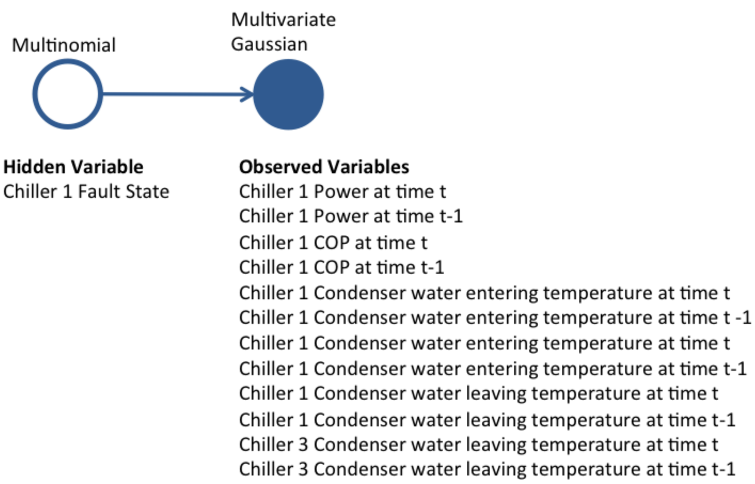

FIGURE 8. GRAPHICAL MODEL AS APPLIED TO THIS APPLICATION

Solving for the Model Parameters Using the Expectation Maximization Algorithm

An algorithm is used to estimate the unknown parameters of the mixture model from data of the chiller plant when operating under faulty conditions and under normal conditions. The estimated parameters are the mixture component weights π_i and the means μ_i and variances Σ_i of the multi-variable normal distributions as shown in Equation 4. The estimates for these parameters are then used in a conditional estimate of the fault state. The determination of the fault state is analogous to detection of the fault.

To solve for the unknown parameters, the Expectation Maximization (EM) algorithm [13] was chosen. This method is an iterative algorithm used to find maximum likelihood estimates of parameters in probabilistic models. It is desirable to maximize the likelihood because the likelihood is a measure for how well the data fits a given model. The likelihood is a function of the parameters of the statistical model, and therefore we want to find parameters that maximize the likelihood.

The likelihood for the Guassian mixture model is:

$$l(\theta \mid D) = \sum_n log(x_n \mid \theta) = \sum_n log \sum_i \pi_i N(x_n \mid \mu_i, \Sigma_i) \tag{5}$$

The goal of the maximization step of the EM algorithm is to find the model parameters that maximize the log likelihood. By setting the partial derivatives of the log likelihood to zero, we obtain the following formulas for the model parameters. This is the maximization step.

$$\mu_i^{(t+1)} = \frac{\sum_n \tau_n^{i(t)} x_n}{\sum_n \tau_n^{i(t)}} \tag{6}$$

$$\Sigma_i^{(t+1)} = \frac{\sum_n \tau_n^{i(t)} (x_n - \mu_i^{(t+1)})(x_n - \mu_i^{(t+1)})^T}{\sum_n \tau_n^{i(t)}} \tag{7}$$

$$\pi_i^{(t+1)} = \frac{1}{N} \sum_n \tau_n^{i(t)} \tag{8}$$

Where τ_i is defined as the conditional probability that the ith component of Z is equal to one.

$$\tau_i \equiv p(Z_i = 1 \mid x, \theta) = \frac{p(x \mid Z_i = 1, \theta_i) p(Z_i \mid \pi_i)}{p(x \mid \theta)} \tag{9}$$

Then, for each iteration of the algorithm, we have

$$\tau_n^{i(t)} = \frac{\pi_i^{(t)} N(x_n \mid \mu_i^{(t)}, \Sigma_i^{(t)})}{\sum_j \pi_j^{(t)} N(x_n \mid \mu_j^{(t)}, \Sigma_j^{(t)})} \tag{10}$$

The expectation step of the EM algorithm calculates the expected complete log likelihood, which in this case is the calculation of the posterior probability τ_n^i. The fault detection algorithm was implemented in Matlab.

RESULTS AND DISCUSSION
MODEL TRAINING

To train the model parameters, the Expectation Maximization algorithm was run for 5 trials of 100 iterations each. From among the trials, the parameters that resulted in the highest likelihood were selected. The initial guess for the mixing proportions was $\pi_i = 1/3$ so that the initial guess for each of the three states is equal. This is because the mixing proportions π_i weigh the mixture components i so equal π_i weigh each component equally. The initial guess for the mean of the variables was randomly generated. A good initial guess was required for the algorithm to run without the covariance matrix becoming non-singular or the log likelihood becoming infinite. Each run of the algorithm took a few minutes to complete on a 1.6 GHz Intel Core i5 processer with 4 GB 1333 MHz DDR3 memory.

Figure 9 plots the training data with the fault detection parameters as calculated by the EM algorithm. The blue diamonds are the means of the Gaussian mixture components — the means of the fault state clusters — as found by the EM algorithm. The black ovals trace the first standard deviation of the covariances of the Gaussian mixture components — the covariances of the fault state clusters. In each plot, two variables are plotted because a two-dimensional space R^2 is easier to visualize than ten dimensional space R^{10}. The mixture model is actually R^{10} Gaussian because the fault detection model consists of ten variables. Thus, the means are actually points in R^{10}, and the one standard deviation covariances are manifolds in R^{10}. Ideally, the mean should fall within the middle of a cluster of data that share the same state; likewise, the one standard deviation ovals should follow the shape of the cluster of data that share the same state.

In Figure 9, data with the same fault state is indicated by the same color. Because the training data was simulated, we know the value of the underlying fault state. In practice, the fault state of the training data is often unknown and it is desirable for a fault detection algorithm to automatically identify data that corresponds to each fault state.

The detection model appears to be better at detecting the severe and no-fault states than the moderate fault state. This is expected because extremes are easier to detect than shades of grey. From Figure 9, this better detection can be seen from the location of the mean closer to the cluster center and the smaller covariance oval. The covariance is largest for the moderate fault state and the oval representing one standard deviation sometimes encloses data points from the severe fault state. When data points with different states overlap, which data point belongs to which state is not as distinguishable.

Figure 9 also shows that for some variables, the EM algorithm found the clusters of data associated with the different fault states better than for other combinations of variables. When the variables are plotted against one another in Figure 9, one can see that some variables are more strongly correlated than others and form distinct clusters for each of the fault states. The variables

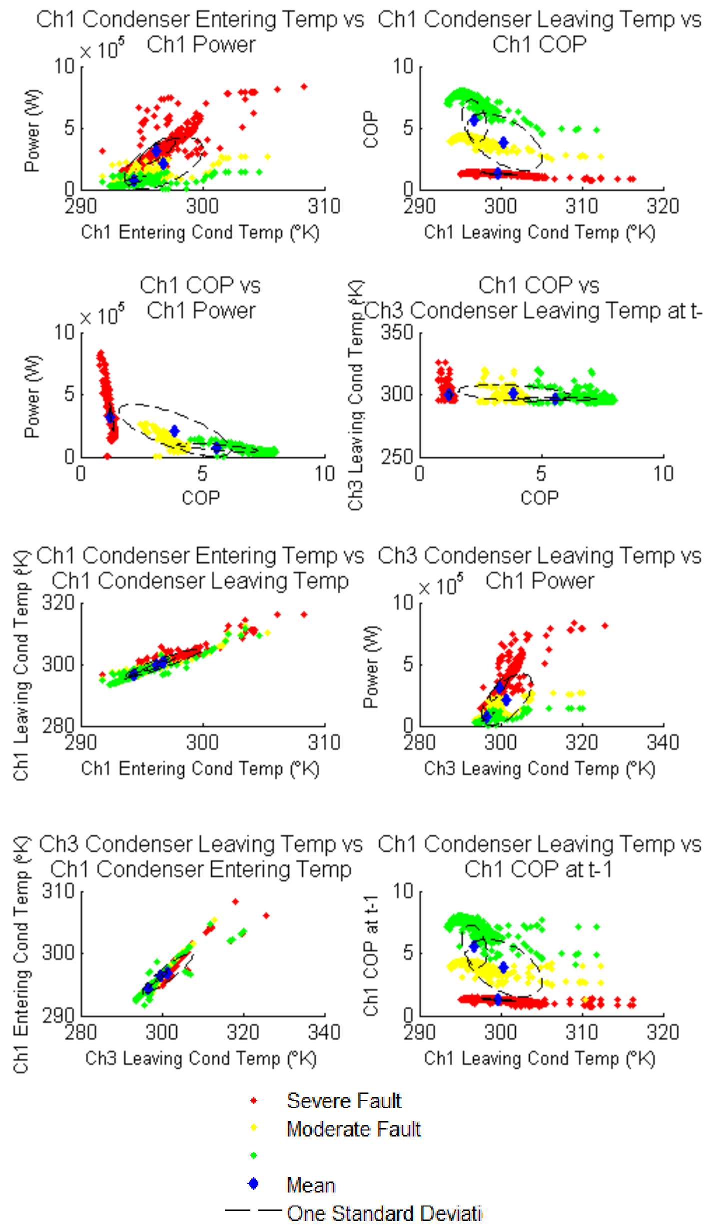

FIGURE 9. MEAN AND CONVARIANCES FOR THE GAUSSIAN MIXTURE MODEL PLOTTED AGAINST THE TRAINING DATA

that have distinct clusters for each of the fault states are more valuable for fault detection. This fault detection model takes this into account with a larger covariance between the variables in the covariance matrix of the Gaussian mixture component. Similarly, variables that are less correlated and less clustered have a smaller value in the covariance matrix.

Some of the poor clustering may be because some data from different fault states are close to one another and may even overlap. For example, the distribution of data points for the condenser leaving temperature is similar between the fault states, which results in a large overlap between the resulting mixture components. Secondly, the large data range between the fault states also affects the detection model for some variables. For example, the Chiller 1 power for the no-fault state is almost constant while the power for the severe fault state ranges across an order of magnitude. Furthermore, some variables are highly linear, as opposed to forming distinct clusters. The temperature variables in particular are highly linearly dependent on the other variables.

Because there are equal amounts of data points for each of the fault states, one would expect that the weighting coefficients π_i would each be equal to 1/3. However, the calculated $\pi = \{0.53, 0.20, 0.27\}$ for {no fault, moderate fault, severe fault} indicates that the algorithm favors classification of data points into the no-fault state more than classification into the other fault states. This may be desirable because it favors decreasing the number of false positives. When data belongs to a fault state, the algorithm expects that the Gaussian mixture component will be evaluated to a high probability; that is, the $N(n \mid \mu_i, \Sigma_i)$ term in Equation 4 will be close to one. Again, this may be desirable because a fault is raised only when the data fits well into the faulty cluster. On the other hand, this could be an indication of over-fitting the model.

MODEL TESTING ON SIMULATED DATA

For each of the data points in the generated data set, the probability of the data point belonging in each fault state was calculated. Then, the data point was classified into the fault state with the highest probability. The runtime on the same laptop was a few seconds. Thus, the computation needs of this approach are practical and relatively inexpensive.

Figure 10 shows how the algorithm classified each of the data points within the training set. Because the data is simulated, we know the underlying fault state. The green bar represents correct classifications while the red bars represent incorrect classifications.

The proposed approach is best at identifying the no-fault state and the severe fault state. Alas, a large percentage of moderate fault data is misclassified as severe faults. The accuracy of the fault state classification can be improved through improved interpretation of the calculated fault state probabilities. This can be achieved by tuning the assignment rules and associated thresholds for assignment of fault states. For example, instead of classifying a data point to the fault state with the highest probability, a data point may be classified into a moderate or severe fault state only if the probability of it being in one of the two fault states is greater than a tunable threshold, such as $> 50\%$. These choices allow for tuning of the fault detection algorithm to different levels of certainty as desired, which greatly improves the flexibility of this approach.

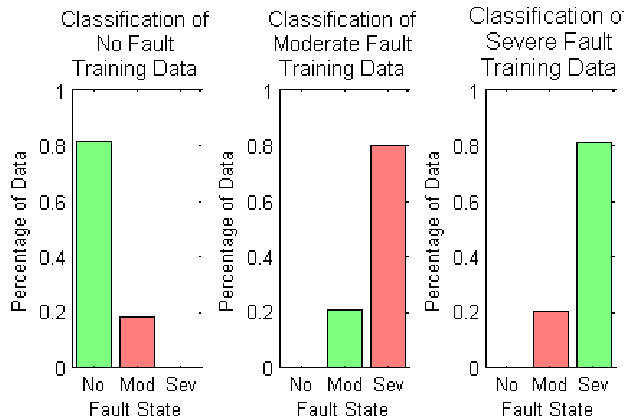

FIGURE 10. CLASSIFICATION OF THE TRAINING DATA AS CLASSIFIED BY THE ALGORITHM

MODEL TESTING ON MEASURED DATA

Measurements collected from the Molecular Foundry are plotted as histograms for each sensor point along with the training data for each of the fault states. The results are show in Figure 11. Training data from each of the fault states is plotted individually in a different color. This is to visually compare the distribution of the measured data with the distribution of each of the fault states. Then, the measured data is plotted as scatter plots with the mean and covariance of the mixture components indicated. The results, shown in Figure 12, have the same variables on each axis as Figure 9.

From Figure 11, it can be observed that the measured data — the grey histogram — is clustered more closely to the data from the no-fault training data — the green histogram — than the data from the other fault states. Thus, one hypothesizes that the Molecular Foundry chiller experiences mostly fault-free operation, and that the fault detection algorithm should classify most data measurements as belonging to the no-fault state. Intuitively, this eight year old chiller plant is relatively new and well instrumented and monitored and one would not expect chiller degradation faults.

From Figure 12, data for some variables seem to better fit into the mixture components than other variables, similar to what was observed from the simulated data. The measured data that corresponds well to the distribution of no-fault training data include: the Chiller 2 condenser leaving temperature, Chiller 1 entering condenser temperature, and Chiller 1 power. There seems to be an offset between the sensor data and the model mixtures, perhaps due to an offset in the physics based model used to generate the training data. Because most of the sensor data seem to lie within and closest to the no-fault mixture component, it seems likely that there are no faults occurring at the Molecular Foundry during this time.

One factor that influences the fit of the real sensor data to

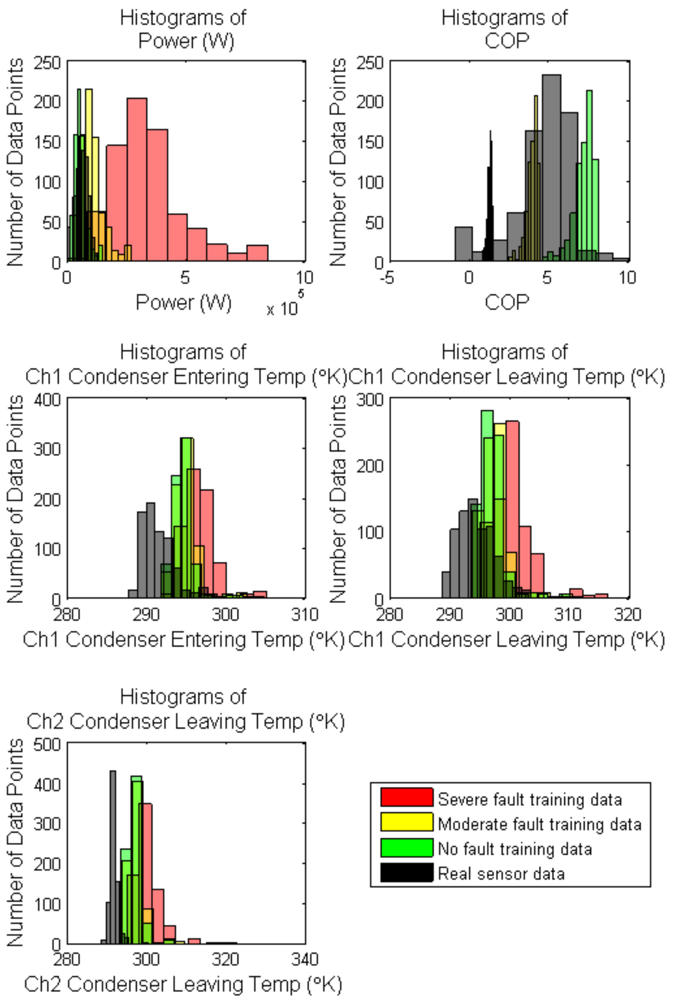

FIGURE 11. HISTOGRAMS OF THE MEASURED DATA COMPARED TO SIMULATED TRAINING DATA

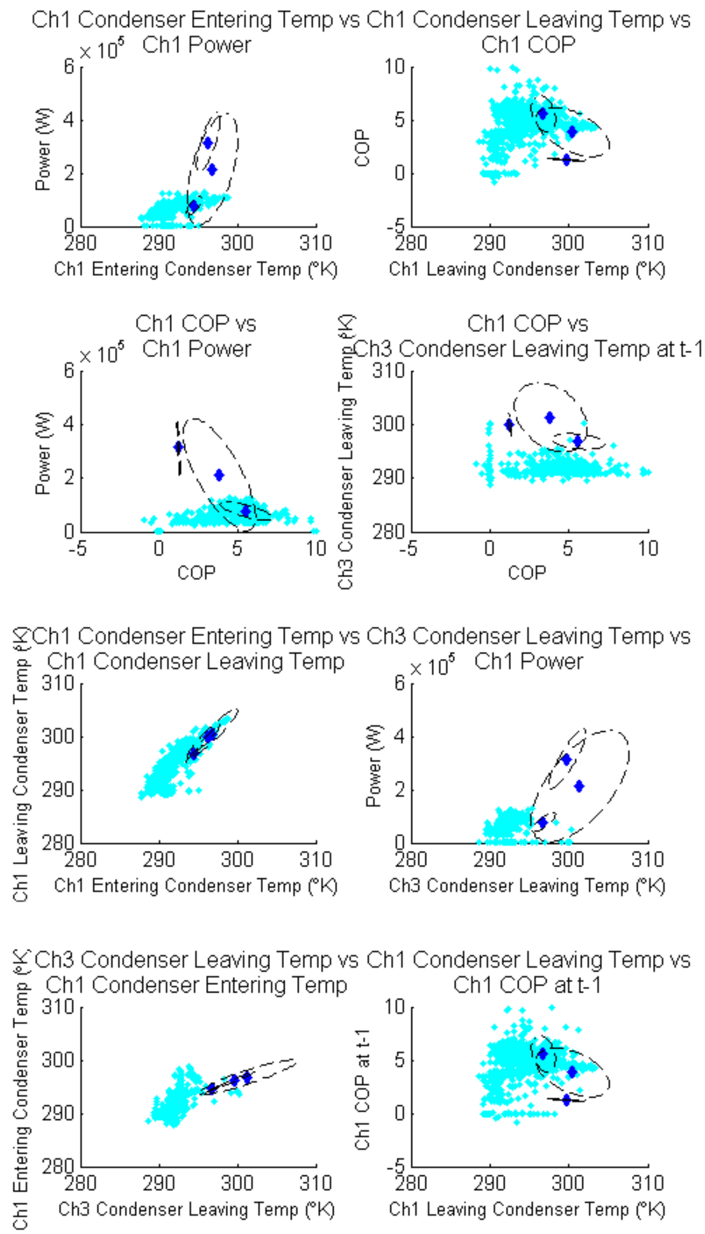

FIGURE 12. REAL DADTA PLOTTED AGAINST THE MIXTURE MODEL COMPONENTS

the mixture components found from the training data is model mismatch between the real chiller plant and the physical model of the chiller plant. One reason for this mismatch is that data used to calibrate the physics-based model does not include data during faulty operations, and thus the training data is an estimate of how the plant will behave based on physical laws. Thus, there is uncertainty in the training data for faulty cases.

DEPLOYMENT

The variables required by this fault detection approach are commonly sensed in chillers and tracked by chiller plant operators. Some systems integration may be required to extract this data from the chiller plant equipment and integration is not always straightforward. Exportation of trend logs from building automation systems is viable, but establishing continuous data acquisition adds complexity. This is not unique to fault detection as unfortunately, getting data to external applications is still costly and time consuming in most buildings projects.

Furthermore, the computation needs of this approach are not very large. This fault detection approach can be implemented in real time to classify incoming data into fault states, or run periodically or on demand to process historical data. Once faults are detected, the information can be relayed to the chiller plant operators through a graphical user interface. When faults are raised, the operator can close the loop by investigating and remedying the fault.

CONCLUSIONS

The multivariate Gaussian mixture model most easily detects the no fault and the severe energy efficiency fault in the chiller, which is understandable given that extremes are the easiest to detect. The moderate fault state is sometimes detected as a severe fault. This is likely due to the large spread of data from the moderate fault state, which overlaps with data points in the severe fault state. This raises the question of what is useful from a decision-analytic prognostics standpoint. What degree of moderate faults would be useful to detect for early warning of a problem that would benefit by maintenance and intervention? For example, instead of three fault states, there could be smaller discretizations on fault levels defined for prognostic value. Then, the algorithm could calculate more detailed probability distributions of fault conditions for each given measurement to understand the uncertainty in the fault assessment.

This fault detection algorithm could be trained using a few days of data that cover a range of fault conditions. Also, the training process is able to distinguish between the different fault states, which means that one does not need to know which training data points correspond to fault conditions, as long as the training data covers a range of faulty and non-faulty operations. The run time of the algorithm to calculate the fault state of measured data is on the order of a few seconds on a laptop, so that the computation needs of this approach are within practical means. Lastly, the fault detection mathematical model and training procedure is flexible so that it can be applied to detect other types of faults with other variables as input.

The current proposed fault detection approach can be improved to more accurately detect faults with less false positives and false negatives. For example, further exploratory data analysis may find parametric probability distributions that describe the mixture components better than the Gaussian distribution. Some preliminary analysis suggests that for the chiller energy efficiency fault, the exponential distribution may be a better choice than the Gaussian distribution. For other faults, other probability distributions may be more appropriate.

To improve the fault detection model, different probability distributions can be explored which may fit the data better than the normal distribution. The normal probability plots and historgrams show that the data deviates from the normal distribution, particularly at the tails. As a result, more accurate fault detection classification may be achieved with a better fitting probability distribution.

To further improve the model, different choices of observed variables, and transformation of variables can be considered. Probabilities could also be conditioned on previous data or on one another. The scope of the detection problem could also be expanded to consider combinations of chillers becoming faulty, other faults, and diagnostics after fault detection.

To train the algorithm, instead of the Expectation Maximization algorithm, other approaches may be used, such as a Gibbs sampler. This may help address the issue of overfitting of the fault detection model to the training data. Overfitting occurs when the fault detection model too closely describes the training data. The issue of overfitting is an area for further exploration in fault detection, particularly because one of the goals of fault detection is to identify deviations from normal operation.

Because the result of probabilistic fault detection is a probability distribution, further work can be done to investigate the interpretation of these probability distributions. These investigations could include, for example, when should a fault be raised, how to assign severity to the fault raised, and how to balance the certainty of a fault occurring for the last question could be through the use of expectation values of the cost of the fault and the certainty of the fault, which would also require the choice of a utility function to value the cost of the fault.

ACKNOWLEDGMENT

The authors would like to thank Prof. Michael Jordan for model advice, Liping Wang for creating the Modelica Model and for the metered data, Rongxin Yin for background information, and Xiupeng Feng, Marco Bonvini, Thierry Nouidui for help with Modelica and acquiring the data, and Professor Auslander for his advice on chiller diagnostics and control.

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Office, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

The authors would also like to acknowledge the Environmental Security Technology Certification Program program and the U.S. Department of Defense.

REFERENCES

[1] Administration, U. E. I., 2013. *Annual Energy Outlook 2013*. U.S. Energy Information Administration.

[2] Katipamula, S., and Brambley, M. R., 2005. "Review article: Methods for fault detection, diagnostics, and prognostics for building systemsa review, part i". *HVACI&R Research,* ***11***(1), pp. 3–25.

[3] Mills, E., 2011. "Building commissioning: a golden opportunity for reducing energy costs and greenhouse gas emissions in the united states". *Energy Efficiency,* ***4***(2), pp. 145–173.

[4] Venkatasubramanian, V., Rengaswamy, R., Yin, K., and Kavuri, S. N., 2003. "A review of process fault detection

and diagnosis: Part i: Quantitative model-based methods". *Computers I& Chemical Engineering,* ***27***(3), pp. 293–311.

[5] Venkatasubramanian, V., Rengaswamy, R., and Kavuri, S. N., 2003. "A review of process fault detection and diagnosis: Part ii: Qualitative models and search strategies". *Computers I& Chemical Engineering,* ***27***(3), pp. 313–326.

[6] Venkatasubramanian, V., Rengaswamy, R., Kavuri, S. N., and Yin, K., 2003. "A review of process fault detection and diagnosis: Part iii: Process history based methods". *Computers I& Chemical Engineering,* ***27***(3), pp. 327–346.

[7] Katipamula, S., and Brambley, M. R., 2005. "Review article: Methods for fault detection, diagnostics, and prognostics for building systemsa review, part ii". *HVACI&R Research,* ***11***(2), pp. 169–187.

[8] Isermann, R., 2005. "Model-based fault-detection and diagnosis status and applications". *Annual Reviews in Control,* ***29***(1), pp. 71 – 85.

[9] Schein, J., Bushby, S. T., Castro, N. S., and House, J. M., 2006. "A rule-based fault detection method for air handling units". *Energy and Buildings,* ***38***(12), pp. 1485 – 1492.

[10] Cui, J., and Wang, S., 2005. "A model-based online fault detection and diagnosis strategy for centrifugal chiller systems". *International Journal of Thermal Sciences,* ***44***(10), pp. 986 – 999.

[11] Liang, J., and Du, R., 2007. "Model-based fault detection and diagnosis of {HVAC} systems using support vector machine method". *International Journal of Refrigeration,* ***30***(6), pp. 1104 – 1114.

[12] Comstock, M. C., Braun, J. E., and Groll, E. A., 2001. "The sensitivity of chiller performance to common faults". *HVACI&R Research,* ***7***(3), pp. 263–279.

[13] Bishop, C., 2006. *Pattern Recognition and Machine Learning*. Springer.