A Finite State Machine-based Characterization of Building Entities for Monitoring and Control

Gowtham Bellala, Manish Marwah, Amip Shah, Martin Arlitt, Cullen Bash
Hewlett Packard Laboratories
Palo Alto, CA
firstname.lastname@hp.com

Abstract

Cyber physical systems such as buildings contain entities (devices, appliances, etc.) that consume a multitude of resources (power, water, etc.). Efficient operation of these entities is important for reducing operating costs and environmental footprint of buildings. In this paper, we propose an entity characterization framework based on a finite state machine abstraction. Each state in the state machine is characterized in terms of distributions of sustainability or performance metrics of interest. This framework provides a basis for anomaly detection, assessment, prediction and usage pattern discovery. We demonstrate the usefulness of the framework using data from actual building entities. In particular, we apply our methodology to chillers and cooling towers, components of a building HVAC system.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: *Miscellaneous*; J.7 [**Computing in Other Systems**]: *Industrial Control*

General Terms

Algorithms, Management, Performance

Keywords

Energy, Buildings, Prediction, Anomaly detection, HVAC

1 Introduction

Buildings account for about 40% of all energy use in the U.S., almost evenly split between residential and commercial buildings. This translates to 8% of the global carbon dioxide emissions [12]. In light of climate change, dwindling natural resources and rising energy prices, there is an increased focus on making buildings more energy efficient.

Facility managers are increasingly instrumenting building infrastructure to find ways of reducing resource consumption (and thereby decreasing operational costs). One approach to savings is to monitor equipment (or appliances/devices) over time, to identify equipment requiring maintenance, or to detect inefficient operation. Such equipment have been reported to waste 15 to 30% of energy in commercial buildings [4]. However, given the various factors that can impact energy consumption of equipment in a building, it is challenging to detect if its energy consumption, under the current conditions, is excessive. This is especially true if the functionality of the equipment has not deteriorated in any way. For instance, if the coefficient of performance (COP) of a chiller, part of a building HVAC system, decreases, it is unlikely it will be detected, even though the corresponding increase in power consumption may be significant.

In this paper, we present a data driven framework to characterize building appliances or devices or their aggregate (which we refer to as entities) in terms of the consumption of resources such as power or water, and/or a sustainability metric like carbon footprint or toxicity. We use that characterization as a basis for assessment, anomaly detection, prediction, and control. The characterization is based on a finite state machine abstraction of the entity's operation, where each state is associated with the corresponding distributions of one or more sustainability metrics. We demonstrate the feasibility of our approach by applying it to five months of data collected from three chillers that provide chilled water to 300,000 ft^2 of office and data center space.

The remainder of the paper is organized as follows. Section 2 introduces our building entity characterization framework. Section 3 provides a use case involving building chillers, while Section 4 describes other use cases, mainly cooling towers. Section 5 discusses related work. Lastly, Section 6 concludes the paper with a summary of our work and future directions.

2 Entity Characterization Framework

2.1 Overview

The goal of the framework is to characterize the operation and resource consumption (e.g., power, water) of entities such as, HVAC components, lighting, etc., in a building. These characteristics can be used to compare the current operation of an entity with its past operation, in terms of sustainability metrics of interest. Similarly, comparisons can be performed between an entity and its "peer group". A key mechanism for the characterization is a finite state machine abstraction augmented with distributions of sustainability performance metrics.

The architecture of the framework is shown in Figure

Buildsys'12, November 6, 2012, Toronto, ON, Canada.

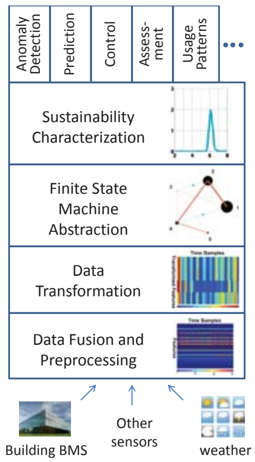

Figure 1. Entity characterization framework.

1. The raw data associated with an entity is obtained from sources such as the building management system (BMS), a weather data service, and any additional sensors related to the entity. Next, the *data fusion* step merges (or "joins") data from multiple sources. Data quality issues are also addressed in this step, which include removing outliers and invalid values, denoising, and imputing missing values. This is followed by *data transformation*, which consists of feature selection and dimensionality reduction; *finite state machine abstraction*, where operational dynamics of an entity are discretized into states and transitions; and, *sustainability characterization*, where each state is associated with probability distributions of sustainability metrics. A number of applications can be built on top of this formulation as described later.

2.2 State Machine Abstraction

The dynamics of operation of most devices or entities in a commercial building are highly complex. Abstracting their operation into a finite state machine allows for their modes of operation to be better understood in terms of discrete states and transitions. The states, in essence, are a partition of the feature sub-space that impacts the entity's operation.

An entity is treated as a black box and all the external factors that influence its operation are used as features to determine the underlying states, or the modes of operation. An initial feature selection is performed to pick significant features, which can be combined with domain expertise, if available, to select or discard features. To construct the states, the selected feature space could be linearly partitioned. However, that is likely to result in a large number of states, and further, choosing the "correct" number of partitions along each feature can be difficult. We address these problems using low dimensional embedding of the feature space, by applying dimensionality reduction techniques such as principal component analysis (PCA), followed by clustering to arrive at the states. One of the goals in the construction of the states is that they are meaningful to a domain expert.

Formally, an entity finite state machine can be defined by the following 4-tuple: (S,A,R,T), where

- S is the set of all states, and $|S| = k$.
- A is a $k \times k$ transition matrix such that a_{ij} specifies the probability of the entity transitioning from state i to state j, i.e., $\forall i, j \in \{1,\cdots,k\}$,

$$a_{ij} = \Pr(s[n+1] = j | s[n] = i),$$

  where $a_{ij} = 0$ implies that a transition from state i to state j has not been seen. Further, $\sum_j a_{ij} = 1$.
- R is a set of m symbols which correspond to events (or causes) that could lead to a state transition.
- Finally, T is a $k \times k \times m$ matrix, where t_{ijl} is the probability that event l causes transition from state i to j to occur, i.e., $\forall\ i, j \in \{1,\cdots,k\}$, and $l \in \{1,\cdots,m\}$,

$$t_{ijl} = \Pr(s[n+1] = j | s[n] = i, r[n] = l)$$

The state machine parameters (S,A,R,T) can be derived from data available for the entity. The transition matrix A is determined by mapping the data onto states and counting. However, determining the set of events R would usually require some domain knowledge. The events related to an entity could be logged by the entity itself, or manually by an operator. These events could also correspond to changes in the value of a feature. Once the events are determined, they could be correlated with the state changes to determine the matrix T. Note that an event need not result in a state transition, or could be responsible for several transitions.

2.3 Sustainability Characterization

Once the state machine is created for an entity, each state can be characterized in terms of one or more sustainability metrics. The sustainability metric could be a measure of resource consumption (e.g., power, water) or environmental impact (e.g., carbon footprint, toxicity), or any other quantifiable measure of interest, including a monetary cost.

To add the sustainability metrics, we augment the state machine to a 5-tuple: (S,A,R,T,D). D corresponds to a $k \times p$ matrix of distributions of p sustainability metrics associated with each of the k states. The matrix D is computed from the data by mapping all the points to the states and computing each of the p sustainability metrics for each of the states.

2.4 Applications

2.4.1 Anomaly detection

The state machine can be used for detecting anomalies that result in degradation of the sustainability metrics. This is done by computing the current sustainability metrics of the entity and comparing it with the metrics distributions associated with the corresponding state in the state machine. Based on where the current point lies on this distribution, an anomaly alert can be issued. Similarly, for operational data collected over a longer period, the distribution of the sustainability metrics can be compared with that of the state, and if the difference is statistically significant, an anomaly can be flagged. Note that the two distributions can be compared based on a number of techniques such as: degree of overlap, Kullback-Leibler divergence, etc.

2.4.2 Assessment

The above is a temporal assessment of an entity, that is, a comparison of its efficiency to the efficiency of the same entity under similar conditions in the past. If data is available for two or more similar entities, then an entity can be compared or assessed with its "peer" group. If the number of such entities is small, then a pair-wise comparison can be performed, where the distributions of the sustainability metrics are compared for the same states for the two entities. If a sufficient number of entities are available, then reference distributions for a particular state for that kind of entity can be established, and distributions of individual entities can be compared against the reference. This enables quantitative answers to questions like, "At what percentile does a particular entity lie in a population of peers in terms of a (specific) sustainability metric?"

2.4.3 Control

The state machine representation allows a user/operator to reason about the entity in terms of a number of discrete states, characterized by the sustainability metrics. Based on the sustainability metrics, each state can be ranked, and this ranking can be used to compare states, and also determine desirable transitions (to a better ranked state) and, similarly, undesirable transitions. The events (set R) causing a transition can be categorized into those within the control of a user (e.g., some parameter setting) and those outside a user's control (e.g., weather conditions). Such a framework can be used to select user settable parameters such that the entity is likely to transition to better ranked states (or at least is prevented from transitioning to worse ranked states).

2.4.4 State Prediction

A state in the state machine captures all of the features used to determine that state. Thus, assuming that the error introduced by state machine abstraction is acceptable, predicting future states is equivalent to multivariate time series prediction for the entity features considered. In fact, the state machine abstraction essentially converts a multivariate time series prediction problem into a univariate time series prediction problem. Since the states are associated with sustainability metric distributions, knowing future states provides information on future performance of the entity in terms of these sustainability metrics. We propose a random forest based model to predict the future states based on statistics of the past state sequence and weather forecast information.

3 Use Case: Chillers

In commercial buildings, heating, ventilation and air-conditioning (HVAC) constitute a significant portion of the power consumption, accounting for nearly 32% of their total energy usage [12]. We selected chillers to demonstrate our framework, since within the HVAC system, chillers are often the largest consumers of power, and show complex behavior.

3.1 Background

A chiller provides a cooled liquid, typically water, that can be circulated through a heat exchanger to provide air conditioning for buildings and IT infrastructure such as data centers. Chillers can be broadly classified into two categories, air-cooled and water-cooled chillers. Figure 2 demonstrates a schematic for a typical water-cooled chiller system. The gray, dashed box in the center of this figure corresponds to a chiller unit, which is composed of four basic components, namely, an evaporator, a compressor, an economizer, and an air-cooled or water-cooled condenser. The heat from the chiller load is dissipated to the atmosphere through

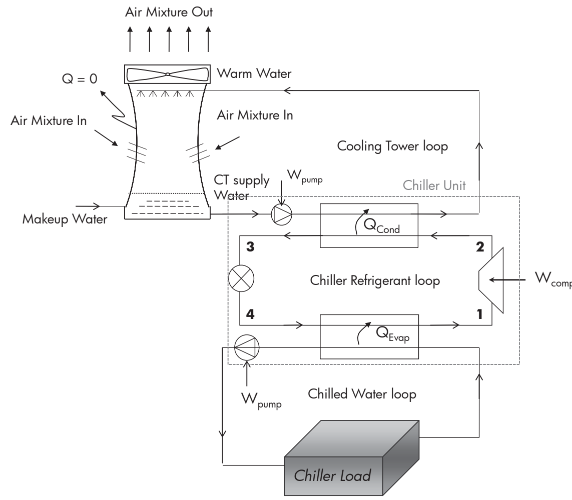

Figure 2. Schematic of a water-cooled chiller system.

	Capacity (Tons)	Mfg Year
Chiller 1	650	2001
Chiller 2	600	2005
Chiller 3	600	2010

Table 1. Details of Chiller Infrastructure on site

a series of steps involving a refrigerant loop, a chilled water loop, and a cooling tower water loop [7].

We define three terms used in the context of chiller units. *Chiller load* corresponds to the amount of heat that is generated (and thus needs to be dissipated) on the site. It is commonly specified in Tonnes (Tons). *Chiller power consumption* reports the power consumed by the chiller unit. It is commonly measured in kilowatts (kW). *Chiller COP* (Coefficient Of Performance) of a chiller unit indicates how efficiently the unit provides cooling, and is defined as $COP = 3.517 \times \frac{L}{P}$, where L corresponds to the chiller load (in tons), and P denotes its power consumption (in kW) [7]. While one chiller unit may be sufficient for a small load, several units working as an ensemble are usually required to meet the cooling demand of a large commercial campus.

3.2 Test Bed

We consider three large, two-storey buildings on a commercial campus as an initial test bed for our analysis. The campus is a mixed use (commercial and industrial) facility consisting of office spaces, laboratories, a data center, a cafeteria, restrooms, and other shared indoor and outdoor spaces. The three buildings on this campus has a total footprint of $300,000$ ft^2, and hosts about 500 occupants.

The site is equipped with three TRANE (model CVHF) CenTraVac water-cooled liquid chillers. Table 1 lists their capacity along with their manufacturing year. The primary load on these chillers correspond to the cooling requirements (air-conditioning) for the three buildings, as well as cooling services for labs, clean rooms, and a data center.

Figure 3 shows the chiller load on this site. The site is equipped with multiple chillers and sufficient spare capacity to ensure business continuity and to be able to meet the site cooling demands, in the event of a unit becoming unavailable as a result of failure or required maintenance.

The chillers on site are designated as primary, secondary

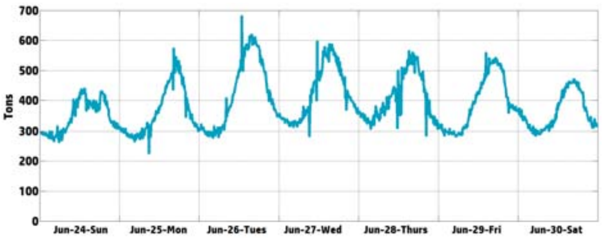

Figure 3. Total Chiller Load on site.

and back-up, where the primary chiller acts as the main chiller with the secondary chiller being switched on whenever the load exceeds the primary chiller's capacity. These designations are usually rotated among the three chillers every few months for even wear and tear.

3.3 Data Collection and Cleanup

The site is equipped with a BMS system that has about 6,000 data points corresponding to various sensors and meters installed throughout the campus, of which only a small subset are related to the site chiller system. However, identifying the set of relevant data points is a challenging task as most points are labeled in an ad-hoc manner, with little or no description. With the help of a domain expert (a building administrator), we identified around 100 data points to be related to the site chiller system. These BMS parameters span flow rate, temperature, pressure, power, and other parameters related to the site air handling equipment such as cooling towers, pumps, etc., as well as ambient weather information such as humidity, ambient temperature, and ambient pressure. Figure 4(a) provides a sample list of these parameters along with their BMS name tags.

We maintain a historical log of these parameters sampled every 5 minutes, and collected over a period of 5 months. However, the resulting time series data has missing values caused by hardware and software failures. Treating these missing values as zeros could lead to erroneous results, and hence need to be filled. We adopt a weighted global average strategy to impute the missing values. This method could be used to impute blocks of missing values, while preserving the local structure. Specifically, if $x[n]$, $n = 1, \cdots, N$ denotes a time-series curve sampled at N different time points, for any time index $1 \leq m \leq N$ with $x[m]$ missing, its value is imputed by $x[m] = \frac{\sum_{k=1}^{N} w[k]x[k]}{\sum_{k=1}^{N} w[k]}$, where the weights $w[k]$ are chosen such that they decrease as a function of their distance from the missing value. For example, the weight function can be chosen to be $w[k] = 1/|m-k|^2$. On average, less than 4% of the values were missing in our data.

3.4 State Machine Abstraction

One of the goals of a state machine abstraction is to be able to efficiently summarize the dynamics of the operation of a device using a few distinct states. To identify the underlying operating states of a device, we propose a cluster based approach on a low-dimensional embedding of the feature space. The proposed algorithm consists of three steps.

1. **Feature Selection:** The first step is to identify the features that affect the operating behavior of a device. We treat the device as a black box and take a control volume approach, where the selected features correspond to the input and output parameters to this black box. Figure 4(b) demonstrates these parameters in the case of a chiller, where the features correspond to chilled water supply temperature (T_{CHWS}), chilled water return temperature (T_{CHWR}), chilled water supply flow rate ($\dot{f}_{CHWS}$), condenser water supply temperature (T_{CWS}), condenser water return temperature (T_{CWR}), and condenser water supply flow rate ($\dot{f}_{CWS}$).
2. **Low-dimensional embedding:** The features selected in the first step can be correlated. In this module, we remove such redundancies by projecting the data onto a low-dimensional space. Furthermore, reducing the dimensionality of the feature space aids in improving the performance of the next step.

   We perform dimension reduction in two stages. In the first stage, we use domain knowledge to reduce the feature dimensions, followed by projection using principal component analysis (PCA). We choose PCA over other dimensionality reduction algorithms such as multidimensional scaling or Laplacian Eigenmaps, as PCA is a simple, linear projection method that is both computationally fast and feasible even on large datasets.

   In the case of chillers, we use domain knowledge to reduce the feature space from the initial six features to the following four features, T_{CHWR}, $(T_{CHWR} - T_{CHWS})\dot{f}_{CHWS}$ (which is proportional to the amount of heat removed from the chilled water loop, i.e., chiller load), T_{CWS}, and $(T_{CWR} - T_{CWS})\dot{f}_{CWS}$ (which is proportional to the amount of heat removed from the condenser water loop). The obtained feature space is further reduced using PCA, where we chose the first two principal dimensions which capture around 95% of the variance in the feature data.
3. **Clustering:** The final step is to partition the projected data into clusters, where each cluster represents an underlying operating state of the device. The clusters are determined using the k-means algorithm based on the Euclidean distance metric.

The output of this algorithm corresponds to a state sequence $s[n], n = 1, \cdots, N$, where $s[n] \in \{1, \cdots, k\}$ with k denoting the number of clusters (or states). Using this state sequence, we can estimate the *a priori* probability of a device operating in state i, as well as the probability of the device transitioning from state i to state j, $\forall i, j \in \{1, \cdots, k\}$:

$$\pi_i = \frac{\sum_{n=1}^{N} \mathbf{1}_{\{s[n]=i\}}}{N}, \text{ and } a_{ij} = \frac{\sum_{n=1}^{N-1} \mathbf{1}_{\{s[n+1]=j,\ s[n]=i\}}}{\sum_{n=1}^{N-1} \mathbf{1}_{\{s[n]=i\}}}, \quad (1)$$

where $\mathbf{1}_{\{E\}}$ is an indicator function that takes the value 1 when the event E is true, and 0 otherwise.

Figure 5(a) shows the state transition diagram for Chiller 3 based on three months of training data. The feature data has been partitioned into five clusters leading to five different states. The nodes in this figure correspond to the operating states of the chiller, where the size of a node determines its frequency of occurrence (i.e., π_i). The edges denote the state

BMS Name Tag	Description from Domain Expert	Units
B1L_CH3.CHWRT	Chilled water return temp	^oF
B1L_CH3.CHWS.FLOW	Chilled water supply flow	gpm
B1L_CH3.CHWS.STPT	Chilled water supply temp set point	^oF
B1L_CH3.CHWST	Chilled water supply temp	^oF
B1L_CH3.COMP.TEMP	Compressor temp	^oF
B1L_CH3.COND.PRESS	Condenser pressure	psi
B1L_CH3.COND.TEMP	Condenser temp	^oF
B1L_CH3.CWRT	Condenser water return temp	^oF
B1L_CH3.CWS.FLOW	Condenser water supply flow	gpm
B1L_CH3.CWST	Condenser water supply temp	^oF
B1L_CH3.EVAP.PRESS	Evaporator pressure	psi
B1L_CH3.EVAP.TEMP	Evaporator temp	^oF

(a)

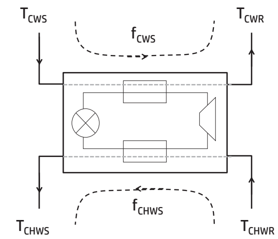

(b)

Figure 4. (a) A sample list of the parameters related to Chiller 3 (b) Features that determine operating states of a chiller.

transitions[1], where bi-directional transitions are represented by red lines, and uni-directional transitions by blue arrows. The thickness of these edges correspond to the frequency of occurrence of the transition.

3.5 Sustainability Characterization

The operating behavior of the chiller in each of these states can be characterized in terms of its power consumption, and its efficiency of operation as measured by COP. Figure 5(b) shows the probability density function (pdf) of the chiller power consumption and COP in each of the 5 states. The density functions are estimated using the kernel density estimate with a Gaussian kernel, as described below.

Let X_i denote a random variable corresponding to the chiller power consumption in state i, and let $\mathcal{X}_i$ denote the corresponding set of power consumption values observed in the training data. Then, the distribution of X_i is estimated by

$$\widehat{f}_{X_i}(x) = \frac{1}{|\mathcal{X}_i|} \sum_{z \in \mathcal{X}_i} \frac{1}{\sqrt{2\pi\sigma^2}} e^{\frac{-(x-z)^2}{2\sigma^2}}, \tag{2}$$

where $|\mathcal{X}_i|$ denotes the size of the sample set, and the kernel bandwidth $\sigma = 1.06\widehat{\sigma}|\mathcal{X}_i|^{-1/5}$, $\widehat{\sigma}$ being the standard deviation of the samples. This choice of σ is known to be optimal for estimating normal distributions [11].

Figure 5(b) shows that the chiller operates at a lower efficiency in states 3 and 5 with a mean COP value of 4.74 and 5.43, as compared to states 1, 2, and 4 whose mean COP values are 6.12, 6.26, and 6.09, respectively. Using these efficiency values, the states can be characterized into "good" (higher efficiency) and "bad" (lower efficiency) states. Ideally, the chiller would operate only in the "good" states. The cause for a transition from a "good" state to a "bad" state can be identified via the transition characteristics.

The state transitions capture the dynamics of the operation of a device. Each transition exhibits a unique characteristic in terms of the input features responsible for the transition. Figure 6 shows the change in feature distribution in transitioning from state 5 (gray dotted curves) to state 3 (orange solid curves). This figure reveals that the chilled water return temperature (T_{CHWR}) and the condenser water supply flow rate ($\dot{f}_{CWS}$) are the features that are most likely responsible for this transition. We quantify this visual observation

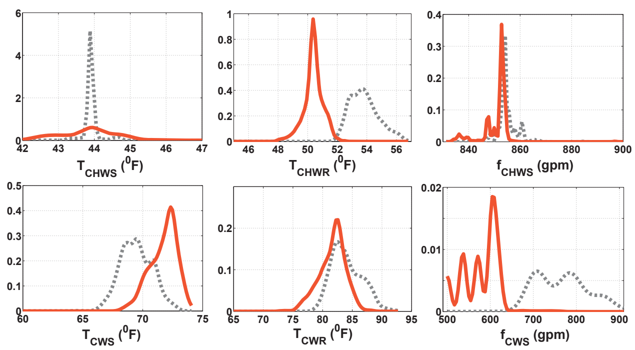

Figure 6. Input features that are responsible for the transition from state 5 to 3.

by means of an overlap measure defined below. For any input feature X, we define the probability of that feature being responsible for a transition from state i to j as

$$\nu_{ij}^X := 1 - \int \min(\widehat{f}_{X_i}(x), \widehat{f}_{X_j}(x))dx, \tag{3}$$

where $\widehat{f}_{X_i}(x)$, $\widehat{f}_{X_j}(x)$ correspond to the kernel density estimate of the feature distribution in states i and j, respectively. The integral value in the expression for ν_{ij}^X corresponds to the area under the overlap between the two distributions, whose value ranges from 0 to 1. Hence, ν_{ij}^X takes a value close to 0 when there is significant overlap between the two distributions, and close to 1 when there is no overlap. The ν_{ij}^X values for the six features in the above transition correspond to 0.5679, 0.9858, 0.5199, 0.5929, 0.3273, and 0.9786. This corroborates our earlier observation that T_{CHWR} and $\dot{f}_{CWS}$ are the features most likely responsible for this transition.

3.6 Anomaly Detection/Assessment

We will now demonstrate the power of state machine models in assessing the performance of a chiller with respect to its past performance, as well as with respect to its peers. As mentioned earlier, one of the main advantages of assessing the performance of a device within each state is that it ensures comparison under identical input/external conditions, thereby allowing for a fair and unbiased assessment.

We partition the chiller data into two sets. We train the state machine model based on three months of training data, and used the remaining two months of chiller data for performance assessment within each state. We further partition the two months of test data into six different test samples, where each sample consists of ten consecutive days of chiller data.

[1] The self transitions (i.e., transition within the same state) are not shown in this figure.

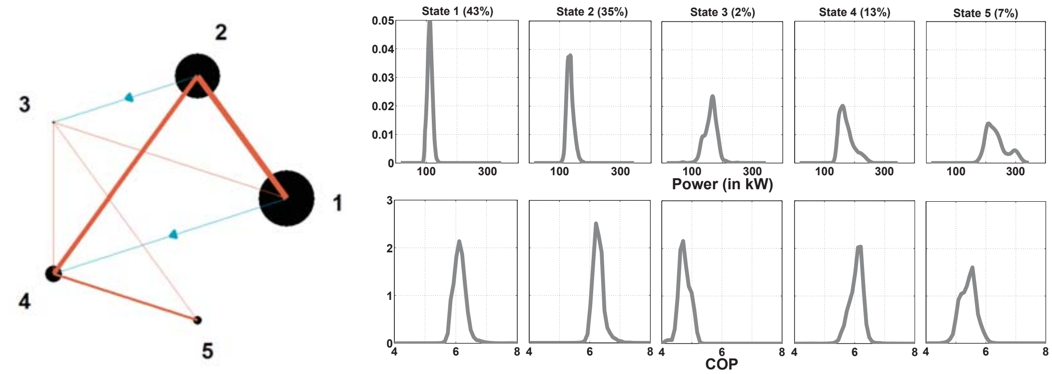

(a) (b)

Figure 5. (a) State Transition Diagram for Chiller 3 (b) Power consumption and COP characteristics of the chiller in each state. Each plot corresponds to a probability density function of the corresponding random variable.

For each sample we project the feature data onto the principal dimensions learned during the training phase, and assign each projected data point to its nearest state (or cluster). We then compare the distribution of the chiller COP in the training data with that of the test data, for each state. We raise an anomaly flag if these two distributions are significantly different, as quantified by the Kullback-Leibler divergence [1], or the overlap measure defined in (3).

Figure 7 demonstrates the performance assessment results for 4 different test samples, where we show the performance assessment results in one state for each case. The (gray) dotted curves correspond to the chiller COP or feature distribution in the training data, and the (orange) solid curves correspond to that of the test data.

Figure 7(a) demonstrates a normal scenario, where the chiller COP behavior in the test phase is similar to that during the training phase. Figure 7(b) demonstrates a scenario where the chiller COP distribution in the test phase is significantly different from that of the training phase. To identify the cause for this anomalous behavior, we examine the distribution of the input features, and look for features that have a significantly different distribution in the test data as compared to the training data. In this case, we identified the chiller load to have a significantly different distribution as shown in Figure 7(b). On further examination, we identified the cause for this change in load distribution to be that of a sensor error, where the sensor monitoring the chiller load temporarily stopped refreshing its readings, resulting in the spike at around 300 Tons. However, the true load during this period could have been different, and hence the time points that have been assigned to state 5 could correspond to other states. This example is an instance of a temporal anomaly, and it can be further categorized into "sensor malfunction" or "hardware issues" anomaly category.

Figure 7(c) demonstrates a second anomalous scenario where the chiller's performance has improved in the test sample as compared to that of the training period. To identify the cause for this anomalous behavior, we once again compare the feature distributions in the training data with that of the test sample. In this case, we identified the chilled water supply temperature T_{CHWS} (which serves as a proxy to the set point temperature) to have been increased over this period, resulting in an improved performance.

These three examples correspond to the scenario where the chiller's performance is assessed with respect to its past performance. We will now demonstrate performance assessment of the chiller with respect to its peers, under similar conditions. Out of the three chillers on site, chillers 2 and 3 are identical (same brand, model and capacity). Hence, we compare the performance of these two chillers in each state, i.e., under identical input conditions. Figure 7(d) demonstrates the COP behavior of chiller 3 (gray dotted curve) and chiller 2 (orange solid curve) in state 2. This figure reveals that chiller 2 has a significantly higher COP than that of chiller 3. We observe a similar difference in the COP behavior of chillers 2 and 3 in the remaining four states. This anomalous behavior could have been caused due to reasons such as different internal settings within the chillers, or due to the continuous operation of chiller 3 over a long period resulting in a degradation of its performance.

Identifying anomalies that correspond to the chiller performance degradation is extremely critical, as timely detection of such anomalies could result in huge savings in their power consumption. For example, by identifying the cause for the anomaly in Figure 7(d), and improving the COP of chiller 3 to that of chiller 2 could potentially result in 25.8% savings in the power consumption of chiller 3. The above savings has been estimated as follows,

$$\% \text{ savings} = \frac{\eta_{\text{chiller 3}} - \eta_{\text{chiller 2}}}{\eta_{\text{chiller 3}}} \times 100,$$

where $\eta = 3.517/\text{COP}$, which corresponds to the amount of power consumed (in kW) to provide 1Ton of cooling load.

Finally, note that this is an unsupervised problem with no labeled data to validate the trained models. Hence, temporal assessment can only detect deviations relative to past behavior. Identifying the true versus anomalous behavior can be done through peer assessment (if data from a set of similar peers is available), or by a domain expert.

3.7 State Prediction

Each state of the chiller is a representation of the distribution of the input features, as well as sustainability met-

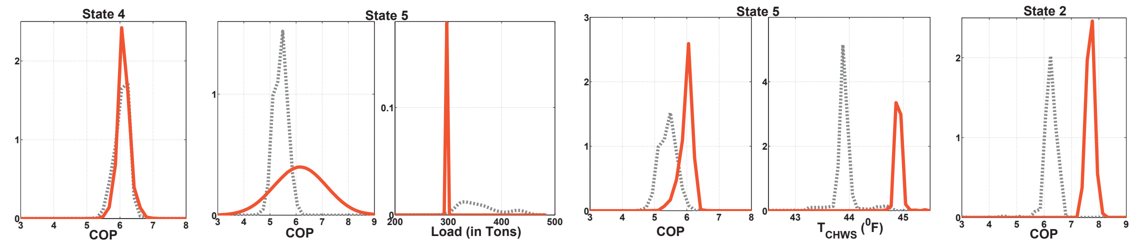

(a) Normal case (b) Temporal anomaly type 1 (c) Temporal anomaly type 2 (d) Spatial anomaly

Figure 7. (a-c) Performance Assessment of chiller 3 with respect to its past performance. The (gray) dotted curves correspond to the chiller's past performance, and the (orange) solid curves correspond to its current performance (d) Performance Assessment of chiller 3 with respect to its peers. The (gray) dotted curve correspond to chiller 3's performance, and the (orange) solid curve correspond to chiller 2's performance.

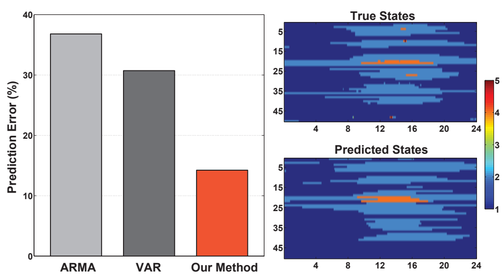

Figure 8. Prediction error using the three approaches. True states vs. Predicted states using our approach

rics such as power consumption and COP. Hence, predicting the future states of the chiller provides a wealth of information in terms of the input conditions as well as the future resource (power) requirements. This information can be extremely valuable to a building administrator in terms of chiller scheduling, and demand shaping.

One approach is to predict the two principal components using standard time-series prediction techniques (e.g., ARMA). The predicted principal component values can then be mapped to states, by assigning each predicted data point to the closest cluster in the projected space. The optimal values for the order of the AR, and the moving average parts can be chosen using the Akaike information criterion [1].

Alternatively, we propose to predict the chiller states directly, thereby reducing a multivariate time-series prediction problem to univariate time-series prediction. We use a random forest based ensemble classifier [1] to predict the chiller states based on the following seven input features: time of day, day of week, month of year, duration of the previous state, chiller state on previous day at same time, and weather forecast information such as temperature and humidity.

Figure 8 compares the performance of the three models in terms of their prediction error, i.e., the fraction of states misclassified. The results are for one day (24 hours) ahead prediction, averaged over 50 consecutive days. In each case, the training data corresponds to all the available past data. The figure shows that the classifier-based approach that directly predicts the states performs significantly better than the multivariate time-series prediction methods. Also shown are the true states in the 50 different test cases, along with the predicted states using the classifier-based approach.

4 Other Use Cases

4.1 Cooling Tower

In this section, we demonstrate the generality of our framework by replicating the above analysis for an ensemble of cooling towers. Cooling towers are heat removal devices that are used to transfer heat in the condenser water supply loop to the atmosphere. Hot water from the condenser of the chiller unit is sprayed from the top of the cooling tower. Dry air enters through the vertical faces of the cooling tower and flows past the sprayed water, transferring the heat from the water to the air. The cooled water that is collected at the bottom of the cooling tower is sent back to the condenser, while the warm, moist air is forced out into the atmosphere by huge blowers. Our site is equipped with three cooling towers that work in parallel. We treat the three cooling towers and the associated pumps as one physical entity in this analysis.

To select the features that determine the operating behavior of this entity, we again treat the entity as a black box, and determine its input and output parameters, shown in Figure 9(a). Note that the ambient temperature and humidity play a critical role for this entity, as the performance of the cooling towers depends on the temperature and the moisture in the air flowing in. Using domain knowledge, the above five features are reduced to the following four: T_{CWS}, $(T_{CWR} - T_{CWS})\dot{f}_{CWS}$, ambient temperature, and ambient humidity. Using PCA we further reduce three features that capture 95% of the variance in the feature data. Lastly, the projected data is partitioned into five clusters using k-means.

The output state sequence can then be used to estimate the *a priori* state probabilities and the transition probabilities, as described in Section 3.4. Figure 9(b) shows the state transition diagram for this entity, along with its mean COP value in each state. Note that this entity operates at low efficiency in state 1, high efficiency in states 2 and 4, and moderate efficiency in states 3 and 5. Given this state machine, we can perform assessment and anomaly detection for the cooling tower ensemble, similar to that for the chillers.

Figure 9(c) demonstrates a temporal anomaly, where the entity operated at a higher efficiency during a test period (orange solid curves) as compared to that of the training period (gray dotted curves), under identical input conditions. Identifying the cause for this anomaly, and improving the COP of the ensemble could potentially result in 8.7% savings in the

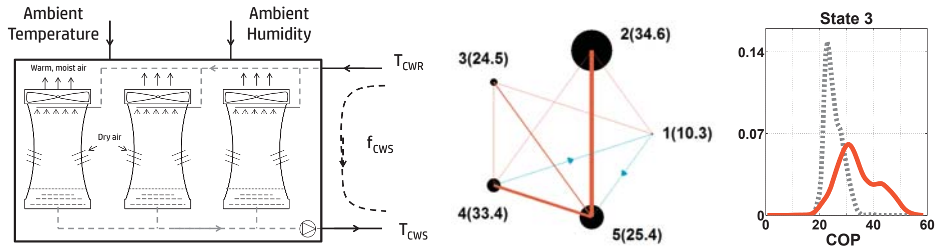

(a) Ensemble black box (b) State Transition Diagram (c) Temporal anomaly

Figure 9. (a) Features that determine the operating state of the ensemble of Cooling Towers (b) State Transition Diagram for the ensemble along with its mean COP value in each state (c) Performance Assessment of the ensemble with respect to its past performance.

power consumption of the ensemble in State 3.

4.2 Other Devices/Entities

In the future, we plan to extend the proposed state machine framework to other devices such as lighting, where the input features could correspond to occupancy information and ambient lighting, and to entities such as building aggregate loads, where the features could correspond to weather, time of day, workday/holiday, and occupancy information. However, there are challenges involved in modeling these devices, as it is difficult to measure some relevant input features. For example, the operating behavior of lights can depend on the behaviors of the occupants.

5 Related Work

Finite state machines are widely used for modeling systems. Hart used FSMs of appliances for energy disaggregation [3]. Parson *et al.* used similar state machines as prior models, again for energy disaggregation [8]. These studies only used states based on power levels.

There is a large body of work on anomaly detection. Chandola *et al.* provide a comprehensive survey on the techniques commonly used [2]. Katipamula *et al.* [4, 5] reviewed fault detection and diagnosis (FDD) in buildings. Zhou *et al.* describe a model-based FDD method for HVAC systems [13]. Regression models are built for performance indices related to the operation of HVAC sub-systems such as chillers, cooling towers, etc. The output from these models, built using data from normal operation, is compared with actual data to determine anomalies. Patnaik *et al.* model operation of an ensemble of chillers using state machines derived from chiller utilization data [9]. A hierarchical rule-based FDD method for HVAC is proposed by Schein *et al.* [10]. Most of these FDD methods require deep domain expertise. In comparison, our framework requires minimal domain knowledge, and applies to any entity. Some of these methods, e.g., [10], could be used in conjunction with our approach to determine possible causes of anomalies detected.

Comparing consumption of an entity with a peer group is commonly done. Kolter *et al.* create a predictive model for aggregate building consumption and enable users to compare their consumption with that of a similar population [6].

6 Conclusions and Future Work

This paper describes a data-driven framework for characterizing the resource consumption of building entities. A use case involving three chillers from a large commercial campus is used to demonstrate the importance of the framework. As an example, a peer assessment of two identical chillers revealed a 25% reduction in the power consumption for one of the chillers is possible.

We are extending our framework in several ways. First, we continue to evaluate our knowledge discovery techniques on an increasingly larger set of data. Second, we are implementing our framework on a cluster of servers (rather than one), to enable it to evaluate all appliances/devices/entities in a building or set of buildings in parallel. Finally, we are developing entity specific algorithms that can provide causal analysis for the observed anomalies, and can be implemented on top of our current framework.

7 References

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

[2] V. Chandola, A. Banerjee, and V. Kumar. Anomaly detection: A survey. *ACM Comput. Surv.*, 41(3):15:1–15:58, 2009.

[3] G. Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(2):1870–1891, 1992.

[4] S. Katipamula and M. R. Brambley. Methods for fault detection, diagnostics, and prognostics for building systems - A Review, Part I. *HVAC&R Research*, 11(1):3–25, 2005.

[5] S. Katipamula and M. R. Brambley. Methods for fault detection, diagnostics, and prognostics for building systems - A Review, Part II. *HVAC&R Research*, 11(2):169–187, 2005.

[6] J. Z. Kolter and J. Ferreira. A large-scale study on predicting and contextualizing building energy usage. In *AAAI*, 2011.

[7] M. J. Moran and H. N. Shapiro. *Fundamentals of Engineering Thermodynamics*. John Wiley and Sons, 1996.

[8] O. Parson, S. Ghosh, M. Weal, and A. Rogers. Non-intrusive load monitoring using prior models of general appliance types. In *AAAI*, 2012.

[9] D. Patnaik, M. Marwah, R. K. Sharma, and N. Ramakrishnan. Temporal data mining approaches for sustainable chiller management in data centers. *ACM Trans. Intell. Syst. Technol.*, 2(4):34:1–34:29, 2011.

[10] J. Schein and S. T. Bushby. A hierarchical rule-based fault detection and diagnostic method for hvac systems. *HVAC&R Research*, 12(1):111–126, 2006.

[11] B. W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, 1998.

[12] U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy. Energy efficiency trends in residential and commercial buildings, 2010.

[13] Q. Zhou, S. Wang, and Z. Ma. A model-based fault detection and diagnosis strategy for HVAC systems. *Int. J. Energy Res.*, 33(10):903–918, 2009.