

Sustainable Operation and Management of Data Center Chillers using Temporal Data Mining

Debprakash Patnaik[†], Manish Marwah^{*}, Ratnesh Sharma^{*}, and Naren Ramakrishnan[†]

[†]Dept. of Computer Science, Virginia Tech, VA 24061, USA

^{*}HP Labs, Palo Alto, CA 94034, USA

{patnaik,naren}@vt.edu,{manish.marwah, ratnesh.sharma}@hp.com

ABSTRACT

Motivation: Data centers are a critical component of modern IT infrastructure but are also among the worst environmental offenders through their increasing energy usage and the resulting large carbon footprints. Efficient management of data centers, including power management, networking, and cooling infrastructure, is hence crucial to sustainability. In the absence of a “first-principles” approach to manage these complex components and their interactions, data-driven approaches have become attractive and tenable.

Results: We present a temporal data mining solution to model and optimize performance of data center chillers, a key component of the cooling infrastructure. It helps bridge raw, numeric, time-series information from sensor streams toward higher level characterizations of chiller behavior, suitable for a data center engineer. To aid in this transduction, temporal data streams are first encoded into a symbolic representation, next run-length encoded segments are mined to form frequent motifs in time series, and finally these metrics are evaluated by their contributions to sustainability. A key innovation in our application is the ability to intersperse “don’t care” transitions (e.g., transients) in continuous-valued time series data, an advantage we inherit by the application of frequent episode mining to symbolized representations of numeric time series. Our approach provides both qualitative and quantitative characterizations of the sensor streams to the data center engineer, to aid him in tuning chiller operating characteristics. This system is currently being prototyped for a data center managed by HP and experimental results from this application reveal the promise of our approach.

Categories and Subject Descriptors: H.2.8 [Database Management]: Database Applications - Data Mining; K.6.2 [Management of Computing and Information Systems]: Installation Management - Computing Equipment Management

General Terms: Algorithms, Experimentation, Measurement, Management, Reliability.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

KDD’09, June 28–July 1, 2009, Paris, France.

Copyright 2009 ACM 978-1-60558-495-9/09/06 ...\$5.00.

Keywords: data centers, chillers, clustering, motifs, frequent episodes, sustainability.

1. INTRODUCTION

Data centers have become a ubiquitous element of modern IT infrastructure, especially in the services sector that requires “always-on” capability. Practically every large IT organization hosts a data center, either in-house or outsourced to major vendors. Furthermore, the recent emergence of the software as a service (SaaS) paradigm or more generically cloud computing, coupled with emerging web-based business, social networking and media applications and services have led to a tremendous growth in the number, size, and power densities of data centers.

Recently, data centers have become an object of scorn for environmentalists, due to their increasing energy usage and the resulting large carbon footprints. For instance, a recent study [25] claims that a web search uses half the equivalent energy of boiling a kettle of water! The US EPA estimates that energy usage at data centers is experiencing successive doubling every five years, and that the annual electricity cost at these centers can amount to \$7.4 billion by 2011, with 10% of this amount being borne by the government. Furthermore, to meet the projected growth in demand would require the construction of 10 power plants! Sustainable data center research is hence becoming acknowledged as an important problem with the foremost of implications for future energy consumption and its environmental impact [14].

A key goal is to adaptively manage demands on power, cooling, and energy efficiency without compromising reliability and performance constraints. State-of-the-art data centers today provide capabilities for adaptive, configurable operation at the level of each subsystem such as power distribution, cooling, and compute server racks. The traditional approach to perform such adaptation is to conduct detailed computational fluid dynamics (CFD)-based modeling of air and temperature flows [19] through a data center. However, such methods are computationally intensive, and even for a small data center (say, about 3000 sq. ft.), it can easily take several hours for a model of reasonable accuracy to converge!

The computational infeasibility of such large-scale simulations coupled with the relative ease of gathering real-time data from sensors has led researchers to deploy sensor networks to track environmental data (e.g. temperature, humidity), operational state of systems and devices (e.g. utilization), and workload information (e.g. user requests). Since the sheer volume of such data precludes manual inspection, automated data mining and knowledge discovery

techniques are important to glean vital information. Furthermore, to close the loop between instrumentation and control/management leading to energy efficiency, what required are high-level capabilities to transduce the data streams into actionable knowledge for the data center administrator.

While data centers constitute a mix of computing elements, networking infrastructure, storage systems along with power management and cooling capabilities, all of which contribute to energy usage and sources of inefficiency, in this paper we primarily focus on the cooling infrastructure, especially chillers, as a case study for illustrating the importance of data mining solutions. In particular, we show how temporal data mining can bridge the gap between low-level, raw sensor streams, and the high-level operating regions needed for an engineer to efficiently manage the data center.

This paper makes the following contributions:

- We present a three-stage solution for mining data center chiller data involving (i) change point detection, (ii) motif mining, and (iii) sustainability evaluation. This design helps data center administrators to impart significant domain knowledge in the algorithms underlying each stage.
- Our notion of motifs can accommodate “don’t care” transitions (e.g., transients), an advantage we inherit by the application of frequent episode mining to symbolized representations of multivariate numeric data. They also help characterize sustainability metrics that serve as a handle for future tuning and management of data center operations.
- We present our algorithms in the context of managing the chiller subsystem of a data center managed by HP in Bangalore, India. To the best of our knowledge, our work is the first to use temporal data mining techniques to bear upon the cooling infrastructure of a production data center. Our KDD approach on this data center demonstrates that it is, in fact, possible to have sustainable operation as well as economical advantages, thus demonstrating that these are not conflicting objectives.

2. BACKGROUND

2.1 Data Centers

Figure 1 shows a data center consisting of IT equipment (servers, storage, networking) fitted in racks arranged as rows. A large data center could contain thousands of racks occupying several tens of thousands of square feet of space. Also shown in the figure are computer room air conditioning (CRAC) units that cool the exhaust hot air from the IT racks. Energy consumption in data center cooling comprises work done to distribute the cool air and to extract heat from the hot exhaust air. A refrigerated or chilled water cooling coil in a CRAC unit extracts the heat from the air and cools it within a range of 10°C to 18°C. The cooling infrastructure of a data center is shown in Figure 2.

Key elements of this infrastructure include CRAC units, plumbing and pumps for chilled water distribution, chiller units and cooling towers. Heat dissipated from IT equipment is extracted by CRAC units and transferred to the chilled water distribution system. Chillers extract heat from




Figure 1: Thermal map of a data center showing racks arranged in rows and CRAC units.




Figure 2: Typical cooling infrastructure of a data center.

the chilled water system and reject it to the environment through cooling towers or heat exchangers. In addition to the IT equipment, the data center cooling infrastructure can account for up to 50% of the total power demand [2]. The CRAC units provide two actuators that can be controlled. The variable frequency drive (VFD) controls the blower speed and the chilled water valve regulates the amount of chilled water flowing into a unit (between 0% and 100%). These built-in flexibilities allow the units to be adjusted according to the workload demand in the data center. The demand is detected via temperature sensors installed on the racks throughout a data center.

2.2 Data Center Chillers

The focus of this paper is on chiller units that receive warm water (at temperature, T_{in}) from the CRAC units, extract heat from it and recirculate the chilled water (at temperature, T_{out}) back to the CRAC units.




Figure 3: Operational state diagram of a chiller unit.

Chiller Architecture and Operation

Each chiller is composed of four basic components, namely, evaporator, multi-stage centrifugal compressor, economizer and water-cooled or air-cooled condenser. Liquid refrigerant is distributed along the length of the evaporator to absorb enough heat from the water returning from the data center and circulated through the evaporator tubes to vaporize. The gaseous refrigerant is then drawn into the first stage of compressor. Compressed gas passes from the multi-stage compressor into the condenser. Cooling tower water circulated through the condenser tubes absorbs heat from the refrigerant, causing it to condense. The liquid refrigerant then passes through an orifice plate into the economizer. Flashed gases enter the compressor while the liquid flows into the evaporator to complete the circuit.

Starting and stopping a chiller is a complex, multi-step process. Fig. 3 shows the operational state diagram of a typical chiller. On power-on, the chiller waits for the compressors to start, after a prescribed delay. On startup, the chiller utilization varies to match the cooling load. Based on chiller technology, chiller compressors can throttle in discrete stages or continuously. Feedback control is used to maintain the outlet temperature, T_{out} , close to a user-specified set-point temperature.

Terminology and Metrics

Here we define some terms used in context of a data center chiller unit.

IT cooling load. This is the amount of heat that is generated (and thus needs to be dissipated) at a data center. It is approximately equivalent to the power consumed by the equipment since almost all of it is dissipated as heat. It is commonly specified in kilowatts (KW).

COP. The coefficient of performance (COP) of a chiller unit indicates how efficiently the unit provides cooling, and, is defined as the ratio between the cooling provided and the power consumed, i.e.,

$$COP_i = \frac{L_i}{P_i} \quad (1)$$

where L_i is the cooling load on the i th chiller unit and P_i is the power consumed by it.

Chiller utilization. This is the percentage of the total capacity of a chiller unit that is in use. It depends on a variety of factors, mainly, the mass flow rate of water that passes through a chiller and the degree of cooling provided, that

is, the difference between the inlet and outlet temperatures ($T_{in} - T_{out}$). For a particular T_{out} , an administrator can control the utilization at a chiller through power capping or by changing the mass flow rate of water.

Chiller power consumption. This is simply the power consumed by a chiller unit. Although power meters that measure aggregate power consumption of data center infrastructure elements are usually available, meters that measure power consumed by an individual entity or a specific group (e.g. chillers) may not always be installed. In such cases, if the capacity of the unit and average COP are known, they, together with unit utilization, can be used to estimate power consumed.

$$P_i = \frac{U_i * C_i}{100 * COP_i} \quad (2)$$

where P_i is the power consumed, U_i the utilization, and C_i the capacity, all pertaining to the i th chiller unit.

Ensembles of Chiller Units

The number of chiller units required depend on the size and thermal density of a data center. While one unit may be sufficient for a small data center, several units operating as an ensemble may be required to satisfy the cooling demand of a large data center. Figure 4 shows an ensemble of chiller units that collectively provide cooling for a data center. Out of the five units shown, three are air-cooled while the remaining two are water-cooled. Also, to provide a highly available data center and ensure business continuity, sufficient spare capacity is usually provisioned to meet the cooling demand in the event of one or more units becoming unavailable as a result of failure or required maintenance.




Figure 4: Five chiller units work in tandem to provide cooling for a large data center.

Operational Challenges

Although operating curves for individual chiller units exist, no model is available for operation of an ensemble, especially one consisting of heterogeneous units. Additionally, shift and/or drift of response characteristics with time further complicate their management. The operational goals are to satisfy the cooling requirements while minimizing the total power consumption of the ensemble and maximizing the average lifespan of the units. While multiple factors impact the lifespan of a chiller unit, an important one is: rapid and large oscillations in utilization value. High amplitude and frequent variations in utilization due to varying load or some failure condition result in decreased lifespan, and, thus, need to be minimized.

Why are rapid oscillations bad? Frequent start and stop cycles lead to fatigue of mechanical parts due to high torque

requirements, and, deterioration of electrical circuitry due to high inrush current. Moreover, load fluctuations due to cycling can also lead to drop in power factor and potential penalties from the utility. In case of data centers with on-site generation, such fluctuations can lead to reliability issues at the generators as well. Downstream of chillers, pump performance and cooling tower efficiency can also be adversely affected. Typically chillers have an MTBF (mean time between failure) of 20,000 hours or more, which can reduce exponentially due to oscillations.

2.3 How can data mining help?

While administration of a single chiller unit is not complicated, configuring an ensemble of chillers for optimal performance is a challenging task, especially in the presence of a dynamically varying cooling load [3]. Typically, heuristics and rules-of-thumb are used to make decisions regarding:

- Which chiller unit(s) should be turned on/off, and when?
- What utilization range should a particular unit be operated at?
- How should the ensemble react to an increase or decrease in cooling demand?

Any guidance regarding questions posed above while maintaining performance and optimizing the above stated goals will be invaluable to a data center facilities administrator.

3. PRIOR WORK

3.1 Mining Systems and Installations

Many researchers have explored the use of data mining to optimize computer system design. For instance, modeling of rack-level temperature data specifically in relation to CRAC layout has been undertaken in [1, 23]. Optimization opportunities at multiple levels of smart center architecture have also been studied in [24]. More recent work [16] focuses on sensor data mining to identify anomalous and deviant behavior. Other related work includes the InteMon system from CMU [7, 6] that dynamically tracks correlations among multiple time series [18]. Interactive visualizations for system management have also been investigated [17]. Beyond these projects, most research at the intersection of data mining and high performance systems has focused on compute clusters, not necessarily data centers. For instance, failure prediction in IBM Blue Gene systems by analyzing event logs has been explored [10]. While these works constitute important strides in analytics, in order to support high-level knowledge discovery capabilities, we must raise the level of abstraction at which we study and infer patterns from sensor data streams.

3.2 Approaches to model time series data

There are numerous formalisms available to model time series data and good surveys are in [5, 12]. One broad class of representations conducts global decompositions of the time series (e.g., PCA, DFT, wavelets) and aims to use only the most significant components to model the series. Alternatively, piecewise representations are easier to compute and also lend themselves to a streaming mode of operation. Representations like SAX [13] first perform a piecewise aggregate approximation (the aggregate refers to the

notion of modeling the given, single, time series by a linear combination of multiple time series, each expressed as a box basis function) and symbolize the resulting representation so that techniques from discrete algorithms can be adapted toward querying, matching, and mining the time series. In our work, we seek to design representations for *multivariate* time series data (the works surveyed here focus on a single series) and also to enjoy the benefits of symbolization of series for ease and flexibility of analysis, as described below.

3.3 Finding motifs in time series

Motif mining is the task of finding approximately repeated subsequences in time series, and studied in various works, e.g., [4, 11, 19, 27]. Mining motifs in symbolized representations of the time series can draw upon the rich body of literature in bioinformatics, where motifs have been used to characterize regulatory regions in the genome. As the work closest to ours, we explicitly focus on the SAX representation, which also provides some significant advantages for mining motifs. First, a random projection algorithm is used to hash segments of the original time-series into a map. If two segments are hashed into the same bucket, they are considered as candidate motifs. In a refinement step all candidate motif subsequences are compared using a distance metric to find the set of motifs with the highest number of non-trivial matches. A contrasting framework, referred to as frequent episode discovery, is an event based framework that is most applicable to symbolic data that is not uniformly sampled [8, 9, 15, 20]. This enables the introduction of junk, or “don’t care” states, into the definition of what constitutes a frequent episode.

4. ALGORITHMS

Our primary goal is to link the multivariate, numeric, time series temperature data gathered from chiller units to high level sustainability characterizations. We decompose this goal into change point detection, motif mining, and sustainability characterization, thus using motifs as a crucial intermediate representation to aid in data reduction. The algorithms presented here, while being natural generalizations and combinations of prior work, have not been proposed before.

4.1 Change point detection

Prior work in multivariate change point detection (e.g., see [26] has posited statistical models for behavior within a segment and tracks changes in model parameters to denote qualitative change points. Our task is exacerbated by the lack of adequate models to characterize chiller behavior and also because of the varying interpretations that are attachable to multivariate data.

First, a motif or a trend can manifest in a single series or in multiple series. Second, even when it manifests in multiple series, the specificity with which it manifests can be fixed or variable. For instance, a motif can be ‘three chiller units show oscillatory behavior’ versus ‘*first three* chiller units show oscillatory behavior.’ Since we seek to mine motifs so that their presence/absence can be cross-correlated with the chiller design (e.g., air cooled versus water cooled), we seek motifs of the latter form.

A multivariate time series $T = \langle \mathbf{t}_1, \dots, \mathbf{t}_m \rangle$ is an ordered set of real-valued vectors of a particular variable. Each real-valued vector \mathbf{t}_i captures the utilizations across all the chiller

units. Although a streaming algorithm would be more suitable in the context of time series data, the current implementation is intended more as a diagnostic tool than for prediction. We first perform a k-means clustering on these vectors and use the cluster labels as symbols to encode the time series. Observe that the multivariate series is now encoded as a single symbol sequence. This sequence of cluster labels is then analyzed to detect the change points. See Fig. 5 for an illustration of this approach.

4.2 Motif mining

Already we have suitably raised the representation from multivariate numeric data. We raise the level of abstraction further by doing a run-length encoding of the symbol sequence and noting where transitions from one symbol to another occur. This gives us a sequence of events for input to serial episode mining as illustrated below.

Symbol Sequence : d d d b a c c d d d d c b
 \Downarrow
 Event Sequence : $\langle (d-b, 4), (b-a, 5), (a-c, 6), (c-d, 8), (d-c, 12), (c-b, 13) \rangle$

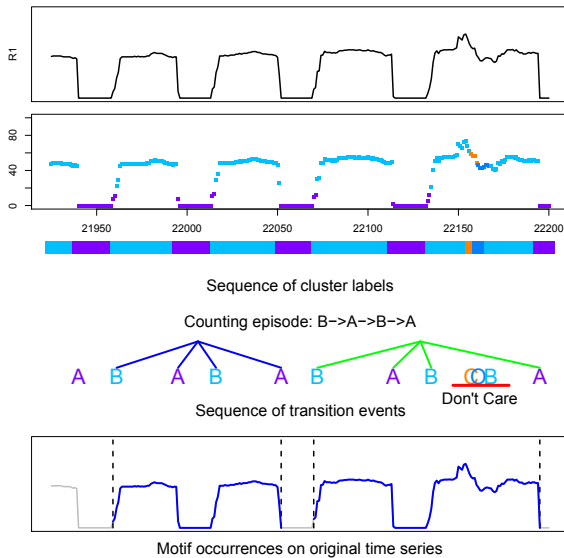


Figure 6: Illustration of motif mining in a single time-series using frequent episodes

Frequent episode mining is now conducted over this sequence of transitions. We adopt the framework of serial episodes with inter-event constraints. The structure of a serial episode α is given as:

$$\alpha = \langle E_1 \xrightarrow{(0, d_1]} E_2 \dots \xrightarrow{(0, d_{n-1}]} E_n \rangle \quad (3)$$

Here E_1, \dots, E_n are the event-types participating in the episode α and, for our domain, these event types are cluster symbol indices. Note that a serial episode requires a total order among the events. Each pair of event-types in α is associated with an inter-event constraint. For example the pair $E_1 \rightarrow E_2$, is associated with $(0, d_1]$ such that in an occurrence of α event E_2 occurs no later than time d_1 of event E_1 .

The mining process follows the level-wise procedure ala *Apriori*, i.e., candidate generation followed by counting. The candidate generation scheme is based on matching the $n-1$ size suffix of one n -node frequent episode with the $n-1$ size prefix of the another n -node frequent episode at a given level to generate candidates for the next level. The time complexity of the candidate generation process is $O(m^2n)$, where n is the size of each frequent episode in the given level, m is the number of frequent episodes in that level, since all pairs of frequent episodes need to be compared for a prefix-suffix match.

The algorithm for counting the set of candidates episodes is given in Algorithm 1. The count or frequency measure is based on non-overlapped occurrences [8]. Two occurrences of an episode are said to be non-overlapped if the events in one occurrence appear between the events in the other occurrence. This notion most naturally eliminates the problem of trivial matches highlighted in [19] where a match is found between two slightly shifted segments of the time series. Algorithm 1 takes as input the event-sequence and a set of candidate episodes and returns the set of frequent episodes for a given frequency threshold θ . The algorithm counts the maximum number of non-overlapped occurrences of each episode with the inter-event time constraint $(0, T]$. This approach also allows repeated symbols or events in the episodes.

The worst case time complexity of the counting algorithm is given by $O(lnm)$, where l is the number of events in the data sequence, m is number of candidate episodes and n is the size of the episode. The algorithm makes one pass of all the events in the event sequence and every time an event that belongs to an episode is seen the data structure s for the episode is updated. A hash map is used to efficiently locate only a subset of relevant episodes for each event seen in the event sequence. Since our method allows repeated symbols, in the case of such episodes the same event can update s structure at most n times. Therefore if the level-wise growth of candidates is sufficiently arrested by a suitable choice threshold, the algorithm scales linearly with data size.

Recall here that the events in mined frequent episodes correspond to transitions from one symbol to another. Our hypothesis here is that if motif occurrences are matched at transitions under an inter-transition gap constraint, then the corresponding time series subsequences will match under a suitable distance metric. In addition the episode mining framework allows for robustness to noise and scaling. The distance metric under which such motifs can be shown to be similar needs further investigation. None-the-less this technique is found to be very effective in unearthing similar time series subsequences in real data sets.

4.3 Sustainability Characterization

It is difficult (and subjective) to compare two motifs in terms of their sustainability impact by inspecting them visually. Therefore, it is necessary to quantify the sustainability of all motifs by computing a sustainability metric for them. This would enable quantitative comparisons between motifs; their categorization as 'good' or 'bad' from the sustainability metric point-of-view; and, furthermore, this information could be used to provide guidance to an administrator or a management system regarding the most 'sustainable' configurations of the chiller ensemble under a particular load.

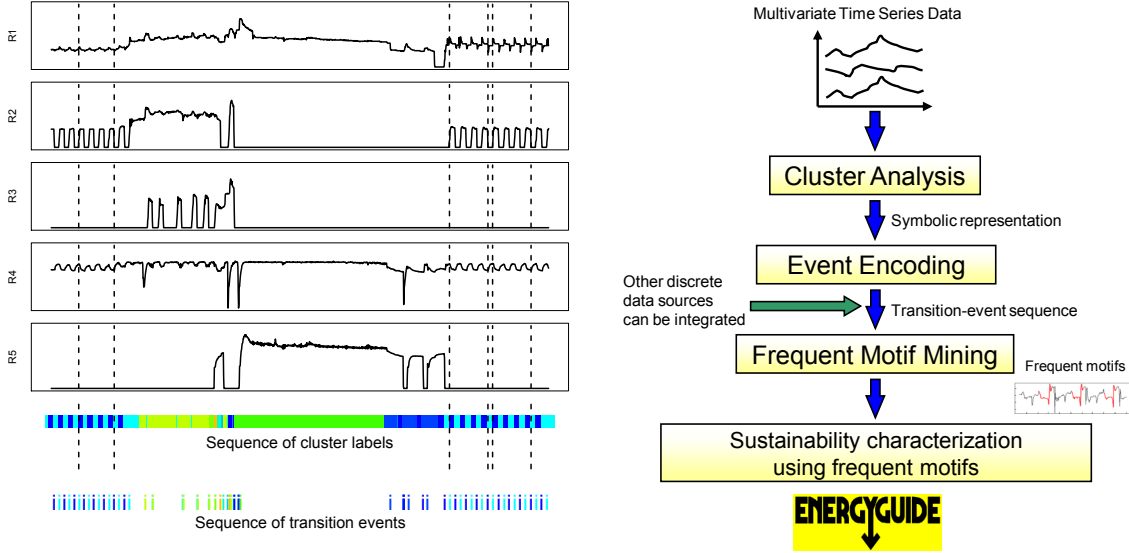


Figure 5: (left) Illustration of change detection in multi-variate time series data. (right) Our overall methodology of using frequent motifs for sustainability characterization.

There are several sustainability metrics, such as, power consumed, carbon footprint, and exergy loss [22]. Note that typically optimizing a sustainability metric, such as power consumed, also minimizes the total cost of operation.

Algorithm 1 Counting occurrences of serial episodes with inter-event time constraint $[0, T]$

Input: Candidate episodes $C = \{\alpha_1, \dots, \alpha_m\}$, where $\alpha_i = E_{\alpha_i(1)} \rightarrow \dots \rightarrow E_{\alpha_i(N)}$ is a N -node episode, Inter-event time constraint T and frequency threshold θ , Event sequence $S = \{(E_i, t_i)\}$.

Output: Frequent episodes $F: \alpha \in F$ if $\alpha.count \geq \theta$

```

/*Initialize*/
waits =  $\phi$ 
for all  $\alpha \in C$  do
   $\alpha.count = 0$ 
   $s =$  Array of size  $N$ , each cell initialized to  $-\infty$ 
  for  $i = 1$  to  $|\alpha|$  do
     $waits[E_{\alpha(i)}].append(\alpha, s, i)$ 
  for all  $(E_k, t_k) \in S$  do
    for all  $(\alpha, s, i) \in waits[E_k]$  do
      if  $(i = 1)$  or  $(t_k - s[i - 1] \leq T)$  then
        /*First event or Satisfies the time constraint*/
        if  $(i = |\alpha|)$  then
           $\alpha.count = \alpha.count + 1$ 
          Reinitialize all elements of  $s$  to  $-\infty$ 
        else
           $s[i] = t_k$ 
  Output  $F = \{\alpha : \alpha \in C \text{ such that } \alpha.count \geq \theta\}$ 

```

In this paper, we estimate two sustainability metrics for each motif: (1) the average COP of the motif; and, (2) a metric reflecting the frequency and amplitude of oscillations in utilization values. The average COP is calculated using Eqn. 1, where the load during motif i , averaged over all its occurrences, is used as L_i , and, the averaged power consumption of the motif as P_i . The power is estimated using

Eqn. 2 with a COP of 3.5 for the air-cooled and 6 for the water-cooled chillers. The COP of a motif quantifies the cooling effectiveness of the ensemble during that motif. In order to estimate the frequency of oscillations of a motif, we compute the number of mean-crossings, that is, the number of times the utilization crosses the mean value. This is very similar to number of zero-crossings that is commonly used in speech processing for estimation of frequency. This, together with standard deviation of a motif, allows oscillatory behavior to be compared.

5. EXPERIMENTAL RESULTS

5.1 Datasets

We applied our motif mining methodology to chiller data obtained from a large HP production data center covering 70,000 square feet with 2000 racks of IT equipment. Its cooling demand is met by an ensemble of five chiller units. The ensemble consists of two types of chillers – three are air-cooled and the remaining two are water-cooled. The data is spread over three different time periods: (1) July 2, 2008 to July 7, 2008; (2) November 27, 2008 to November 30, 2008; and (3) December 16, 2008 to December 26, 2008; totaling over 480 hours.

5.2 Motifs

In all, 22 motifs were discovered in the chiller utilization data whose qualitative properties are summarized in Table 3 (more on this later). From a quantitative point of view these motifs can be clustered into groups based on load. One such group (Group II) is depicted in Table 1 with other quantitative measures. Although each group has very similar load levels, the COP within a group varies with the motifs. In Group II, for instance, motif 8 has a COP of 4.87, while motif 5 is significantly more efficient at a COP of 5.4. This information provides key insights to an administrator regarding motifs that are more energy efficient compared to

Cluster ID	Power (in tons)	Avg Util (in %)	Dev.	#Samples	Color
4	944.38	185.34	187.88	681	Red
14	859.10	168.60	204.35	699	Orange
20	832.14	163.31	151.58	781	Light Orange
18	829.87	162.86	151.09	772	Yellow
12	768.14	150.75	47.68	1485	Light Yellow
19	706.24	138.60	166.98	1382	Yellow-Green
1	666.01	130.71	109.98	653	Light Green
5	659.68	129.47	115.78	1299	Green
8	597.20	117.20	40.71	3723	Light Green
6	583.26	114.47	143.87	667	Green
9	479.37	94.08	81.38	949	Light Green
16	464.30	91.12	164.41	562	Teal
17	456.84	89.66	71.88	1504	Cyan
15	431.86	84.75	80.54	975	Light Blue
11	395.65	77.65	215.82	1250	Blue
7	348.92	68.48	59.97	2330	Dark Blue
3	326.92	64.16	58.03	4313	Very Dark Blue
2	325.93	63.97	50.96	2280	Dark Purple
13	312.77	61.38	164.51	970	Purple
10	268.89	52.77	21.08	1388	Very Purple

Figure 7: Characterization of operating states of the ensemble of chillers using clustering

others. Furthermore, this information could be codified into rules for setting chiller ensemble configuration based on the current load.

Table 2 shows the most efficient and the least efficient motifs for Group II based on power consumption, determined by the motif’s COP value. Also shown are the potential power savings assuming that the least efficient motif could be transformed into the most efficient one.

Table 3 provides a qualitative description of all the motifs. They are grouped together manually based on similar utilization behavior of each of the chillers in the ensemble. Two factors are considered – the average utilization during the motif, and, the oscillatory behavior. The utilization level is labeled as low (L), medium (M) or high (H) based on the average utilization ranges of (0, 35), (35, 65) or (65, 100), respectively. The oscillatory behavior is marked (1) none (N) – which indicates steady values with minor variations; (2) small (S) – which indicates the values show oscillations that have a small amplitude; and finally, (3) large (L) – which indicates oscillations with a large amplitude. Furthermore, units that are not operating are marked ‘OFF’.

5.3 Comparing motifs 5 and 8

We investigate now in more detail the differences between motifs 5 and 8 of Group II. While both motifs 5 and 8 have three chillers turned on, they are of different types. In motif 8, all three operating chillers (C1, C2 and C3) are air-cooled. In motif 5, two air-cooled (C1 and C2) and one water-cooled chiller (C4) are running. The qualitative behavior of the chillers in the two motifs are show in Table 3. In motif 5, one chiller runs at high utilization (C4 at 66.5%), while the other two run at low utilizations (11.3% and 33.8%). In motif 8, one chiller runs at low utilization (17.63) while the other two operate at the medium range (49.1% and 44.3%). The amount of oscillatory behavior in the two motifs is about

the same as indicated by the average rate of mean-crossings and standard deviation in Table 1. In both the motifs, one air-cooled chiller show large oscillations.

5.4 Economical incentives for sustainable operation

Business growth is an important factor in planning for investment in infrastructure. Chillers, of the kind described in this paper, can cost upwards of \$150k and typically, depreciate over 15-20 years. In the absence of intelligent management, the annual operating cost, even at 50% load, can be equally high. Maximizing utilization and efficiency is crucial to achieve a favorable return on investment. Operating chillers among favorable motifs that satisfy these conditions will enable IT administrators to sustain growth in business without major capital or maintenance costs. It will not only prolong the useful life of the existing equipment but postpone expensive retrofits and further infrastructure investments.

The cost savings by operating in favorable motif regimes can be quite easily characterized by multiplying the kW savings by $0.11 \times 24 \times 365$ to obtain the annual savings in dollars (\$). Here, \$0.11 is assumed to be the cost per KWh (kilo watt hour) of electricity. In this case switching from motif 8 to motif 5 gives us a nearly \$40,000 in annual savings! Extrapolating this cost saving to other similar motifs gives us an idea of the utility of data mining algorithms in helping achieve cost effectiveness.

5.5 Carbon footprint calculation

Saving 1kWh of energy is equivalent to preventing release of 0.8 kg of carbon dioxide into the atmosphere. Based on the energy savings number, we can calculate the reduced carbon footprint.

$$\text{Saved Power} = 41\text{kW}$$

$$\text{Saved Energy (annually)} = 41\text{kW} \times 8760 \text{ hrs}$$

$$\begin{aligned} \text{Reduction in Carbon footprint} &= 41\text{kW} \times 8760 \times 0.8 \\ &= 287328 \text{ kg CO}_2 \end{aligned}$$

Observe that this is just the operational footprint; there is also an “embedded carbon footprint” of the chiller unit (as added in its manufacturing process). By maximizing operational life and utilization, we are managing this embedded carbon in the equipment as well. In other words, we are limiting the increase in embedded carbon in the environment while delivering the cooling required.

6. DISCUSSION

We have demonstrated a powerful approach to data mining for data centers that helps situate trends gathered from sensor streams in the context of sustainability metrics useful for the data center engineer. Our future work is in several categories. First, we seek to perform streaming analysis of the chiller data to provide real-time actionable input to manage the operational state of chillers. Second, we seek to explore the co-existence possibility of both short (i.e., lasting at most a few hours) and longer motifs (e.g., spanning half a day to possibly multiple days) and use this information to perform higher order minings of motifs [21], i.e., to detect patterns composed of motifs themselves. Third, we seek to use other sustainability metrics to characterize the

Table 2: The most and least efficient motif for each group and the potential power savings if the operational state of the chiller ensemble could be transformed from the least to the most efficient motif.

	Load (KW)		Most Efficient Motif	Least Efficient Motif	Potential Power Savings	
	Ave.	Std			KW	%
Group II	2089	35	5	8	41	9.83%

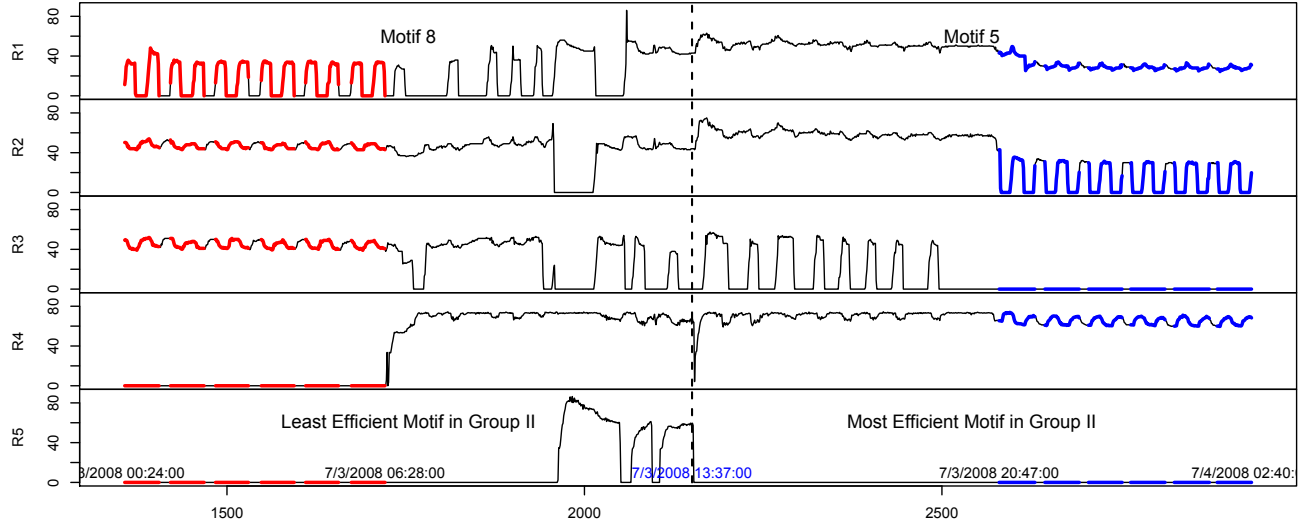


Figure 8: Two motifs with similar load but widely varying levels of efficiency. Observe how the least efficient motif transitions to the most efficient motif in the given time period.

Table 1: Summary of the quantitative measures associated with motifs. Shown here are statistics for motif in one of the load groups.

Motif	Load (KW)	Power (KW)	COP	MC	Std	Span (min)	N
Group II							
8	2028	417	4.87	.042	5.05	50	10
2	2073	400	5.18	.048	4.28	113	17
5	2076	384	5.4	.048	4.25	49	35
4	2083	387	5.38	.048	4.31	82	23
6	2119	422	5.02	.044	4.08	51	23
3	2121	414	5.13	.042	4.31	119	11
1	2123	395	5.38	.046	4.39	172	10

MC – number of mean-crossings per minute
 N – number of occurrences of the motif

motifs; for instance, through lifetime analysis to convert the loss in a chiller’s lifetime into a power number.

Eventually, we would like to use temporal data mining to uncover a complete model of the data center which can then be tuned/controlled/optimized, thus making data mining an integral part of the data center architecture. To achieve this, we need to model the transfer function in a way that encapsulates workload changes, manual steering of chiller operation, and other intermittent transients and recoverable faults. We propose to use our motifs as a handle to redescribe time series data so that significant portions

of time series progression can be captured via state transition diagrams. Control strategies can then be assessed by superimposing their likely trajectories over the state transition diagrams and quantifying their stability characterizations. This will enable us to use data mining not simply as a post-processing data interpretation phase but to actually form dynamic models of data center chillers and their control strategies.

7. REPEATABILITY

We are committed to repeatability of our results. Datasets and source codes used in this paper are available at: <http://neural-code.cs.vt.edu/chiller>

Acknowledgments

This work was supported in part by the Institute for Critical Technology and Applied Science (ICTAS), Virginia Tech. The authors also wish to thank Mohandas Mekanapurath and Vaibhav Bhatia for their support in data collection from a production data center environment.

8. REFERENCES

- [1] L. Bautista and R. Sharma. Analysis of environmental data in datacenters. Technical report, HP Labs, 2007.
- [2] C.L. Belady. In the data center, power and cooling costs more than the IT equipment it supports. *Electronics Cooling*, 13(1), Feb 2007.
- [3] T. Boucher et al. Viability of dynamic cooling control in a data center environment. *Electronic Packaging*, June 2006.

Table 3: Summary of the qualitative behavior of all the motifs. The ones showing similar behavior are grouped together.

Motifs	Number of Units ON	Chiller behavior									
		C1		C2		C3		C4		C5	
		Util.	Oscil.	Util.	Oscil.	Util.	Oscil.	Util.	Oscil.	Util.	Oscil.
7, 15	2	M	L	OFF		H	S	OFF		OFF	
13	2	L	L	OFF		M	S	OFF		OFF	
1,2,4,5,9,10	3	L	N	L	L	OFF		H	N	OFF	
3,6,11,12	3	OFF		L	N	L	L	H	N	OFF	
8,14	3	L	L	M	N	M	S	OFF		OFF	
16	3	L	S	M	N	M	S	OFF		OFF	
17	4	M	N	M	L	OFF		H	N	H	S
18	4	M	N	M	S	L	L	H	N	OFF	
19	4	M	S	L	L	OFF		H	N	H	L
20	4	M	N	L	L	L	L	H	N	OFF	
21	4	OFF		M	S	L	L	H	N	H	L
22	4	OFF		H	N	M	L	H	N	H	S

Utilization (Util.) – Low (L), Medium (M), High (H)
Oscillations (Oscil.) – None (N), Small (S), Large (L)

- [4] B. Chiu, E. Keogh, and S. Lonardi. Probabilistic discovery of time series motifs. In *Proc. KDD'03*, pages 493–498, New York, NY, USA, 2003. ACM.
- [5] H. Ding et al. Querying and mining of time series data: experimental comparison of representations and distance measures. *VLDB J.*, 1(2):1542–1552, 2008.
- [6] E. Hoke et al. Intemon: Continuous mining of sensor data in large-scale self-infrastructure. *SIGOPS Oper. Syst. Rev.*, 40(3):38–44, 2006.
- [7] E. Hoke, J. Sun, and C. Faloutsos. Intemon: Intelligent system monitoring on large clusters. In *Proc. VLDB'06*, pages 1239–1242. VLDB Endowment, 2006.
- [8] S. Laxman, P.S. Sastry, and K.P. Unnikrishnan. Discovering frequent episodes and learning hidden markov models: A formal connection. *IEEE TKDE*, 17(11):15051517, 2005.
- [9] S. Laxman, V. Tankasali, and R.W. White. Stream prediction using a generative model based on frequent episodes in event sequences. In *Proc. KDD'08*, page 453461, Las Vegas, USA, August 2008.
- [10] Y. Liang et al. Failure Prediction in IBM BlueGene/L Event Logs. In *Proc. ICDM'07*, pages 583–588, 2007.
- [11] J. Lin et al. Finding motifs in time series. In *Proceedings of the Second Workshop on Temporal Data Mining*, pages 53–68, 2002.
- [12] J. Lin et al. A symbolic representation of time series, with implications for streaming algorithms. In *Proc. DMKD'03*, pages 2–11. ACM, 2003.
- [13] J. Lin et al. Experiencing SAX: A Novel Symbolic Representation of Time Series. *Data Min. Knowl. Discov.*, 15(2):107–144, 2007.
- [14] S. Lohr. Demand for data puts engineers in spotlight. New York Times. Published June 17, 2008.
- [15] H. Mannila, H. Toivonen, and A.I. Verkamo. Discovery of frequent episodes in event sequences. *Data Mining and Knowledge Discovery*, 1(3):259–289, 1997.
- [16] M. Marwah et al. Stream mining of sensor data for anomalous behavior detection in data centers. Technical report, HP Labs, 2008.
- [17] P. McLachlan et al. Liverac: Interactive visual exploration of system management time-series data. In *Proc. CHI'08*, pages 1483–1492, New York, NY, USA, 2008. ACM.
- [18] S. Papadimitriou, J. Sun, and C. Faloutsos. Streaming Pattern Discovery in Multiple Time-series. In *Proc. VLDB'05*, pages 697–708. VLDB Endowment, 2005.
- [19] P. Patel et al. Mining Motifs in Massive Time Series Databases. In *Proc. ICDM'02*, page 370, 2002.
- [20] D. Patnaik, P.S. Sastry, and K.P. Unnikrishnan. Inferring neuronal network connectivity from spike data: A temporal data mining approach. *Scientific Programming*, 16(1):49–77, 2008.
- [21] J.F. Roddick et al. Higher order mining. *SIGKDD Explor. Newsl.*, 10(1):5–17, 2008.
- [22] A.J. Shah et al. Exergy analysis of data center thermal management systems. *Journal of Heat Transfer*, 130(2):021401, 2008.
- [23] R. Sharma et al. Application of exploratory data analysis (eda) techniques to temperature data in a conventional data center. In *Proceedings of ASME IPACK'07*, 2007.
- [24] R. Sharma et al. On building next generation data centers: Energy flow in the information technology stack. In *Proceeding of Compute 2008*, 2008.
- [25] A. Wissner-Gross. Revealed: The environmental impact of google searches. The Sunday Times. Published Jan 11, 2009.
- [26] X. Xuan and K. Murphy. Modeling Changing Dependency Structure in Multivariate Time Series. In *Proc. ICML'07*, pages 1055–1062, New York, NY, USA, 2007. ACM.
- [27] D. Yankov et al. Detecting Time Series Motifs under Uniform Scaling. In *Proc. KDD'07*, pages 844–853, New York, NY, USA, 2007. ACM.