Predicting Peak-Demand Days in the Ontario Peak Reduction Program for Large Consumers

Yuheng Helen Jiang, Ryan Levman, Lukasz Golab and Jatin Nathwani
University of Waterloo, Canada
{y29jiang,rslevman,lgolab,nathwani}@uwaterloo.ca

ABSTRACT

In this paper, we propose a heuristic algorithm for day-ahead prediction of the top K days having the highest peak hourly demand for electricity over a given year. This problem, which arises in the context of critical peak pricing in Ontario, Canada, is difficult because we may have to wait till the end of the year to find out which K days ended up being the peak days. Our solution is to leverage short-term load forecasts and call tomorrow a peak day if it has sufficiently high probability of being a peak day in the time window covered by the forecast. Using Ontario demand data from 2007 till 2013, we show that our algorithm may need to call about $2K$ peak days to ensure that most if not all of the actual K peak days are included.

1. INTRODUCTION

Reducing peak electricity consumption is an important problem, which has led to a variety of peak pricing schemes in many jurisdictions. In this paper, we analyze the Five-Coincident-Peaks (5CP) program that affects large industrial and commercial consumers (whose monthly peak exceeds 5 megawatts) in the province of Ontario, Canada. In this program, large consumers pay heavy surcharges for the electricity they consumed during the five days with the highest peak hourly demand [1]. For some customers, these surcharges are higher than their volumetric charges [3].

Ontario's 5CP program is different to, e.g., California's Critical Peak Pricing [2], in which utilities choose which days will be peak-pricing days according to some criteria, and they notify the participating customers one day in advance. In 5CP, Ontario's Electricity System Operator (IESO) waits till the end of the year and applies the surcharge to each large consumer based on its contribution to the load (at the peak hour) on the actual five peak days of this year. Thus, without the benefit of hindsight, it is difficult for consumers to know when to curtail load in order to avoid surcharges.

We propose a practical algorithm that, at the end of every day, predicts whether tomorrow will be one of the five peak days of the current year, given only the publicly-available information such as short-term and long-term load forecasts and historical load statistics. We define the *precision* of such an algorithm as the fraction of days identified by it that are in fact peak days, and *recall* as the fraction of actual peak days that were identified as such by the algorithm. For example, suppose that the actual peak days for some year were June 1, July 2, July 3, July 25 and August 8. Suppose that during the course of this year, the algorithm predicts the following *six* days as being peak days: June 1, July 1, July 2, July 3, July 20 and August 8. Its precision is $\frac{4}{6}$ and recall is $\frac{4}{5}$.

Obtaining perfect recall is easy: we predict that each day will be a peak day. Of course, precision will be very low and there will be many false alarms, causing customers to curtail operations unnecessarily and lose business. Ideally, we should achieve high precision (few false alarms) and high recall (few missed alarms).

The IESO publishes 12-month load forecasts, but they are not accurate because Ontario's peak demand is strongly correlated with daily high temperatures, especially in the summer when the daily peak is caused by high air-conditioning use in the afternoon. The 14-day short-term forecasts are quite accurate, and the proposed algorithm uses these as described below.

2. OUR SOLUTION

At the beginning of the year, we are given the actual peak hourly demand for each day in the past year and the IESO 12-month long-term forecast for the current year. At the end of each day in the current year, we will be given the actual peak demand for that day, the 14-day short-term forecast from the IESO and the weather forecast for tomorrow. At the end of each day, the algorithm will compute a probability that tomorrow will be one of the five peak days between the beginning of the year and 14 days from today. If this probability exceeds a threshold τ (which will be defined shortly), the algorithm calls tomorrow as one of the five peak days so that consumers may react accordingly. Since Ontario has been a summer-peaking province since 2005, rather than running the algorithm for an entire year, we only run it from May 1 to September 30.

Throughout the year, we maintain another threshold, τ_D, which is a lower bound for a peak day, i.e., any day whose peak demand forecast is below τ_D will never be called as a peak day. For the initial value of τ_D, we use the maximum peak demand from the IESO long-term forecast, minus a Load-Forecast-Uncertainty (LFU) value of 1600 megawatts. The LFU value is published by the IESO and is related to the uncertainty of the long-term forecast. We also apply two filtering criteria based on domain knowledge: Saturdays and Sundays will never be called as peak days, and nor will days whose weather forecast is not *extreme*, which we define as 30 degrees Celcius or higher. These filtering criteria are meant to avoid false positives.

Let D_i be the actual peak hourly demand on day i and let $\hat{D}_{iL}$ be the estimated peak hourly demand on day $i+L$ as of day i based on the IESO short-term forecast.

e-Energy'14, June 11–13, 2014, Cambridge, UK.
ACM 978-1-4503-2819-7/14/06.
http://dx.doi.org/10.1145/2602044.2602076.

Table 1: Computing ranking probabilities. Define $\theta(x, y)$ as follows: $\theta(x, y) = P(Rank_{future} = x) \times P(Rank_{past} = y)$

Final Ranking	Formula
$P(Rank_{overall} = 1)$	θ(1,1)
$P(Rank_{overall} = 2)$	θ(1,2) + θ(2,1)
$P(Rank_{overall} = 3)$	θ(1,3) + θ(2,2) + θ(3,1)
$P(Rank_{overall} = 4)$	θ(1,4) + θ(2,3) + θ(3,2) + θ(4,1)
$P(Rank_{overall} = 5)$	θ(1,5) + θ(2,4) + θ(3,3) + θ(4,2) + θ(5,1)

Calculating Probabilities

To compute the ranking probabilities over short-term forecasts, we need to identify their distributions. According to the chi square goodness-of-fit test, we verified that the residuals of the short-term forecasts, computed by $\hat{D}_{iL} - D_i$, from 2006 till 2013 are normally distributed with a mean of zero and some standard deviation that depends on L. Thus, every short-term forecast is a random variable with a mean equal to the forecast value and a standard deviation computed from historical data.

We need the following three probabilities. $P(Rank_{future} = j)$ is the probability of $\hat{D}_{i1}$ ranking jth among the 14 days for which we have a short-term forecast, i.e., among $\hat{D}_{i1}$ to $\hat{D}_{i,14}$. $P(Rank_{past} = j)$ is the probability of $\hat{D}_{i1}$ ranking jth compared to the peak demand on the days we have seen so far, i.e., D_1 to D_i. $P(Rank_{overall} = j)$ is the probability of $\hat{D}_{i1}$ ranking jth within the days we have seen plus those for which we have a short-term forecast, i.e., D_1 to D_i and $\hat{D}_{i1}$ to $\hat{D}_{i,14}$.

Assuming that the short-term forecasts for different days are independent, $P(Rank_{future} = j)$ and $P(Rank_{past} = j)$ are easy to compute. For example:

$$P(Rank_{future} = 1) = \prod_{j=1}^{14} P(\hat{D}_{i1} \geq \hat{D}_{ij}) \quad (1)$$

Since we assumed that the residuals of the short-term forecast are normally distributed, we compute $P(\hat{D}_{i1} \geq \hat{D}_{ij})$ using the probability density function for a normal distribution. $P(Rank_{overall} = j)$ is a bit more complex. Table 1 shows how to compute it for j between one and five. For example, $\hat{D}_{i1}$ can rank second overall under two conditions: either it ranks first in the past and second in the short-term future, or it ranks second in the past and first in the short-term future.

The Algorithm

Figure 1 gives the pseudocode. If the day-ahead peak demand forecast exceeds the lower bound τ_D and tomorrow is a weekday and the weather forecast is extreme (line 2), then we compute the required probabilities (lines 4-6) and we check if tomorrow has a high probability of ranking fifth or higher. If this probability exceeds the threshold τ_p, we predict that tomorrow will be a peak day.

We then check if the lower bound τ_D should be adjusted (lines 9-12). If tomorrow's weather forecast is not extreme, but the demand forecast exceeds the lower bound, then we should raise the lower bound to tomorrow's peak demand forecast (lines 9-10). On the other hand, if tomorrow's weather is expected to be extreme but the lower bound on the peak demand is not going to be exceeded, then we should lower the lower bound (lines 11-12).

To choose a value for τ_p, we use the following data-driven approach. Using the actual and estimated demand data from the previous year, we compute $P(Rank_{overall} \leq 5)$ for each day in the past year. We then check this probability for the actual 5CP days and choose τ_p to be the minimum of these. For example, in 2012, each actual 5CP day had $P(Rank_{overall} \leq 5)$ above 0.1, so for 2013 we can set $\tau_p = 0.1$.

1. FOR i=May 1 to Sep 29
2.   IF $\hat{D}_{i1} \geq \tau_D$ and weekday and extreme weather forecast
3.   FOR j=1 to 5
4.   Compute $P(Rank_{future} = j)$
5.   Compute $P(Rank_{past} = j)$
6.   Compute $P(Rank_{overall} = j)$ based on Table 1
7.   IF $P(Rank_{overall} \leq 5) \geq \tau_p$
8.   Predict "tomorrow will be a peak day"
9.   ELSE IF $\hat{D}_{i1} \geq \tau_D$ and not extreme weather
10. $\tau_D = \hat{D}_{i1}$
11.   ELSE IF $\hat{D}_{i1} < \tau_D$ and extreme weather
12. $\tau_D = \hat{D}_{i1}$

Figure 1: Proposed algorithm for the 5CP problem.

Experimental Evaluation

We implemented the proposed algorithm in Matlab and computed its precision and recall on Ontario's demand data from 2007 till 2013. The historical load, short-term and long-term forecast data were downloaded from the IESO Website at `ieso.ca`. Figure 2 shows the results: for each year, we plot the precision, recall and the number of called peak days. The average recall across the seven years is 0.94 and it varied from 0.8 to 1, i.e., the peak days called by the algorithm included four or all five of the actual peak days for that year. The average precision is 0.55 and it varied from 0.4 to 0.7. The number of called peak days varied from 7 to 11. Thus, to identify four or five of the actual five peak days, we may need to call a total of about 10 days as peak days.

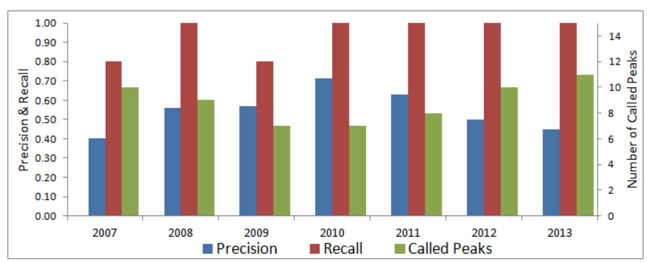

Figure 2: Precision, recall and the number of called peak days from 2007 to 2013.

3. REFERENCES

[1] The 5CP program. `http://ieso-public.sharepoint.com/Pages/Participate/Settlements/Global-Adjustment-for-Class-A.aspx`

[2] J. Bode, C. Churchwell, S. George, and F. Sullivan. 2012 California statewide non-residential critical peak pricing evaluation. Technical report, Freeman, Sullivan & Co, 2013.

[3] J. Spears. Ontario power fee sets new record. Accessed on 18 October 2013, at `http://www.thestar.com/business/economy/2013/09/03/ontario_power_fee_sets_new_record.html`.