Forecasting Uncertain Hotel Room Demand

Mihir Rajopadhye, Mounir Ben Ghalia† and Paul P. Wang

Department of Electrical and Computer Engineering

Duke University

Durham, North Carolina 27708-0291

Timothy Baker and Craig V. Eister

Bass Hotels and Resorts

Atlanta, Georgia 30346-2149

Abstract

Economic systems are characterized by increasing uncertainty in their dynamics. This increasing uncertainty is likely to incur bad decisions that can be costly in financial terms. This makes forecasting of uncertain economic variables an instrumental activity in any organization. This paper takes the hotel industry as a practical application of forecasting using the Holt-Winters method. The problem here is to forecast the uncertain demand for rooms at a hotel for each arrival day. Forecasting is part of hotel revenue management system whose objective is to maximize the revenue by making decisions regarding when to make rooms available for customers and at what price. The forecast approach discussed in this paper is based on quantitative models and does not incorporate management expertise. Even though, forecast results are found to be satisfactory for certain days, this is not the case for other arrival days. It is believed that human judgment is important when dealing with external events that may affect the variables being forecasted. Actual data from a hotel are used to illustrate the forecasting mechanism.

Key Words: Forecasting; Uncertain variables; Economic revenue management systems; Hotel industry; Holt-Winters approach.

† *Corresponding Author*

Duke University

Department of Electrical and Computer Engineering

Durham, North Carolina 27708-0291, U.S.A.

TEL: +1 (919) 660-5233 FAX: +1 (919) 660-5293

E-Mail: mbg@ee.duke.edu

1 Introduction

Forecasting room demand is a very important part of modern day hotel revenue management systems. The objective of these systems is to maximize the revenue under the constraint of fixed room capacity. To this end, most hotels have implemented some form of inventory optimization and controls. These optimization routines are carried out over several days prior to the arrival day, so an estimate of the demand for rooms for that particular target day is required to carry out the optimization. This paper deals with the problem of forecasting unconstrained hotel room demand. Unconstrained room demand is the number of rooms that can be rented if there are no capacity or pricing constraints. Room allocation and optimization are separate issues and are not addressed in this paper.

Optimization of the inventory is very important to the yield management system. The optimization problem involves selling the right type of room to the right customer at the right price, with the objective of maximizing the revenue. This is not a simple task, considering that there are multiple room types which can be sold at different rates to customers requesting multiple night stay. This makes forecasting an important issue, since a better forecast would result in improved inventory optimization, and consequently, increased revenue. Indeed, forecasting and optimization are among the primary components of the yield management system [1], and both components are vital for the performance of the system.

A lot of the work done on hotel revenue management systems deals with the optimization problem [2, 3, 4]. However, to the best of our knowledge, there has been little or no published work on the room demand forecasting aspect. In this paper, we show how a particular forecasting procedure can be applied to the hotel room demand problem.

Several methods have been used for the purpose of forecasting data in a variety of business applications [5]. Different methods vary in the manner in which the historical data is modeled. Regressional methods seek to explain the data with one or more input variables, and the model relates the data to the inputs with a set of coefficients [5]. An application of this method is found in [5], which uses a linear regression model to fit the monthly maintenance expense data of a manufacturing plant. Another method involves fitting a structural time series model to the data. Such a model is set up in terms of components which have a direct interpretation [6]. A very popular forecast method is the Box-Jenkins approach to time series modeling and forecasting [7]. This approach does not assume that successive observations and errors are independent. Future values are forecasted as a function of past observations and errors. An example of such a model can be found in [5], in which the viscosity of a product in a chemical process is modeled using an AR(2) model. However, these models are very complicated and difficult to implement. Smoothing methods, on the other hand, are simple and give equivalent performance with the right choice of model. Smoothing procedures discount past observations in predicting future data, but the manner in which past data is discounted is *ad hoc* [6].

The exponential smoothing procedure is a simple method to forecast future data based on past observations [8]. In this method, previous observations are discounted such that recent observations are given more weights and observations further in the past are given less weight. The weights decrease by a constant ratio, and thus lie on an exponential curve. This method, however, can be used only for non-seasonal time series showing no trend. Due to this, certain adaptations are required in order to use it for time series that arise in real problems.

A more general variation of the simple exponential smoothing procedure is the Holt-Winters method [9]. The latter considers the local linear trend and seasonality in the data. The trend represents the direction in which the time series is moving, while the seasonality explains the effects of different seasons in the data. This method owes its popularity to the fact that it is very simple to implement and is comparable with any other univariate forecasting procedure in terms of accuracy [10]. Also, the components of the forecast (viz. mean, trend and seasonality) lend themselves to an easy interpretation. These components are discussed in detail in the next section. In this paper, we apply the Holt-Winters procedure to forecast unconstrained room demand for an actual hotel. Data collected from an actual hotel is used in the initialization of the forecast components.

The objective of this paper is to apply and evaluate the Holt-Winters procedure to the forecast of hotel room demand . The forecast generated is based only on the hard data in the form of historical data and current booking activity. In this paper, no human input is accounted for in the forecast mechanism. The current study is part of an ongoing research aiming at developing a robust forecast system where both hard data and human input are combined. The motivation for this research comes from the fact that a good forecast would greatly enhance management decision making.

The remainder of the paper is organized as follows: We first present the Holt-Winters forecasting procedure in Section 2. Both, additive and multiplicative models are presented, and model selection criteria are discussed. Section 3 deals with the room demand forecasting. We discuss the reservation characterization and formulate the forecasting problem. Analysis of historical data and simulation of reservation requests are also included in this section, followed by the forecast algorithm. Simulation results are presented in Section 4. Finally, in Section 5, we present the conclusions and suggestions for further research work.

2 Holt-Winters Forecasting Method

The Holt-Winters method is an extension of the exponentially weighted moving average (EWMA) procedure [6]. The EWMA algorithm forecasts future values based on past observations, and places more weight on recent observations. In the Holt-Winters method, the forecast components are updated in a similar fashion, i.e. more weights are placed on recent values. The distinctive feature of the Holt-Winters procedure is that it incorporates linear trend and seasonality into the simple exponential smoothing algorithm [6].

The Holt-Winters models the time series with three components: mean, local trend and seasonality. Depending on how the seasonal variation is included in the model, there are two versions of the Holt-Winters forecast procedure: the additive model and the multiplicative model. The additive model assumes that the forecast at time t, y_t, is given by

$$y_t = (\text{mean at } t-1 + \text{local trend}) + (\text{seasonality}) + \text{error} \tag{1}$$

while the multiplicative model assumes the forecast is given by

$$y_t = (\text{mean at } t-1 + \text{local trend}) \times (\text{seasonality}) + \text{error} \tag{2}$$

Each component of the forecast is described below.

2.1 Mean

The mean component, denoted by m_t, of the model gives the level of the time series at time instant t. It is the base component of the time series which is modified by the trend and seasonality effects to give the final value. For a constant, non-seasonal process, the mean is taken as the forecast of future observations.

2.2 Trend

A majority of the time series do not fluctuate about a constant level, but exhibit shifts in either the upward or downward directions. This effect is modeled by the trend component, denoted by b_t, and it gives the general direction in which the series is progressing. The trend can be classified as global or local, linear or non-linear, etc. The Holt-Winters procedure models a local, linear trend.

2.3 Seasonality

Time series generally show seasonal variations i.e. there is a periodic shift in the level of the series. This is especially true in case of hotel room demand, which have distinct periods of high and low demand depending on the type of hotel property, time of year etc. Seasonal effects are cyclic i.e. it repeats itself after a fixed interval of time. The Holt-Winters method models a finite number of seasonal variations, and the seasonal component is denoted by c_t.

The seasonality component determines the version of the Holt-Winters forecast procedure. The additive version, in which the seasonal component is added to the base and trend components, is used when the amplitude of the seasonal variation is independent of the level of the time series. The use of the multiplicative version is appropriate when the amplitude of the seasonal variation is proportional to the level of the time series. In the multiplicative version, the seasonal effect is multiplied to the base and trend components.

Additive Model

The time series is represented by the model

$$y_t = m_t + b_t t + c_t + \epsilon_t \tag{3}$$

where ϵ_t is the random error component with mean 0 and variance σ^2. We assume the length of a season to be s periods. The equations for updating the corresponding components are

$$\begin{aligned} \hat{m}_t &= \alpha(y_t - \hat{c}_{t-s}) + (1-\alpha)(\hat{m}_{t-1} + \hat{b}_{t-1}) & (4) \\ \hat{b}_t &= \beta(\hat{m}_t - \hat{m}_{t-1}) + (1-\beta)\hat{b}_{t-1} & (5) \\ \hat{c}_t &= \gamma(y_t - \hat{m}_t) + (1-\gamma)\hat{c}_{t-s} & (6) \end{aligned}$$

where α, β and γ are the smoothing constants for the base, trend and seasonal components respectively, and $\hat{m}_t$, $\hat{b}_t$ and $\hat{c}_t$ are the estimates of the base, trend and seasonal components respectively at time t. The forecast for any future time $\tau = 1, 2, \ldots$ is given by

$$\hat{y}_{t+\tau} = \hat{m}_t + \hat{b}_t \tau + \hat{c}_{t+\tau-s} \tag{7}$$

In the above equation, we use the estimate of the seasonal component at time $t+\tau$ computed s periods ago.

Multiplicative Model

The multiplicative model assumes the time series to be of the form

$$y_t = (m_t + b_t t)c_t + \epsilon_t \tag{8}$$

where, as above, ϵ_t is the random error component. The update equations are

$$\begin{aligned} \hat{m}_t &= \alpha y_t / \hat{c}_{t-s} + (1-\alpha)(\hat{m}_{t-1} + \hat{b}_{t-1}) & (9) \\ \hat{b}_t &= \beta(\hat{m}_t - \hat{m}_{t-1}) + (1-\beta)\hat{b}_{t-1} & (10) \\ \hat{c}_t &= \gamma y_t / \hat{m}_t + (1-\gamma)\hat{c}_{t-s} & (11) \end{aligned}$$

where α, β and γ have the same meaning as in the additive model. The forecast for any future time $\tau = 1, 2, \ldots$ is given by

$$\hat{y}_{t+\tau} = (\hat{m}_t + \hat{b}_t \tau)\hat{c}_{t+\tau-s} \tag{12}$$

As seen from equations 1 through 10, it is quite straightforward to implement the Holt-Winters method (either version) on a digital computer. We have used the multiplicative version to forecast room demand, based on the assumption that the seasonal effects are proportional in size to the local mean. Initially, the smoothing constants are assigned values arbitrarily, and α is optimized before each forecast run to minimize the most recent forecast error.

3 Forecasting Unconstrained Room Demand

3.1 Characterizing Reservation Requests

A reservation request is characterized by three quantities: the arrival day, market segment and the length of stay. A request for room reservation always specifies a particular arrival date. Availability and/or price considerations may cause the customer to change the arrival day, but since we are interested in forecasting unconstrained demand, we assume that there is no change in the requested arrival day of each reservation. Also, hotels have different types of rooms (e.g. suites, double rooms, economy etc.) and each of these is sold at different rates. Additionally, the same type of room may be sold at different rates to different customers, depending on promotional packages, concessional rates for government employees, etc. Thus, each reservation request is characterized by the market segment or rate category requested. Lastly, the demand for a room is also characterized by the number of nights the reservation is requested for. A reservation request may be for one night, two nights, or several nights. The length of stay is an important factor as it has a direct impact on the daily revenue, capacity, etc.

3.2 Characterizing Cancelation Requests

Like reservations, cancelations of some existing reservations are also characterized along similar lines. However, one must realize two important differences between the two. Firstly,only an existing reservation can cancel. In this sense, cancelations depend on the outstanding reservations. The other difference is that it is important to know when the cancelation occurred, i.e. how far before the arrival date was the reservation canceled. As seen later in section 3.4, this affects the short term demand forecast.

The forecast problem can then be defined as estimating the number of net reservations (reservations - cancelations) that will be received for each future arrival date per rate class for each

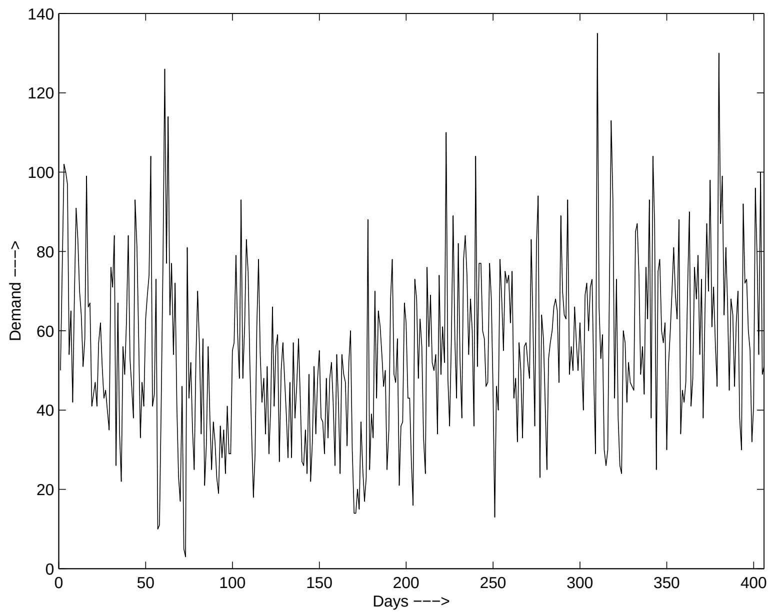

Figure 1: Plot of actual unconstrained net demand of rooms

possible length of stay. In this paper, we assume that the reservations are distributed over different stay durations according to historical data, and we focus our attention on forecasting the net demand for each future arrival day for each market segment. The simulations are carried out for one rate class. Extending this work to forecast unconstrained demand for several rate classes is straightforward.

3.3 Analysis of Historical Data

We have used data from an actual hotel for the initialization and testing of the forecast algorithm. Reservation data for 58 weeks (406 days) was used for this purpose. The property from which the data was obtained is a business/convention center property. The methodology developed here is general, and can be applied to any type of hotel (eg. business traveler property, leisure traveler property). Figure 1 shows the time series plot of the actual unconstrained net demand for the hotel property.

From the graph, we observe the following points:

- A distinct seasonal effect is seen in the plot. The average level of the series decreases between 75 and 200 days as compared to the other portions of the time series.

- While no distinct global trend is seen in the plot, it is possible to identify local linear trends in the data.

- Large and sudden increases or decreases in demand are seen at different parts of the plot. These may correspond to certain events taking place in the hotel or city, and these are classified as non-random effects.

The seasonal effect is specially important in the hotel revenue management problem. Hotels generally have distinct high demand and low demand seasons, and different room allocation strategies are used in these different seasons. As an example, hotels may work with only a part of the room inventory during low demand season to reduce the overheads associated with excess rooms. Consequently, it is very important to be able to estimate the room demand based on the current season.

3.4 Simulating Reservation Requests

To evaluate the performance of the forecast, we need to simulate the process of receipt of requests for hotel rooms. The output of the forecast is the number of customers in each market segment that will actually show-up on any particular arrival day. Thus, we need to simulate the build-up of net demand (reservations minus cancelations) for each market segment and each arrival day in the simulation period. Thus, the problem is that of generating random reservation and cancelation requests based on historical data.

Reservation requests are generated using a Poisson distribution, while a binomial distribution is used to model cancelation requests. Modeling reservations and cancelations using Poisson distribution and binomial distribution respectively are very common in the literature and have been used by several researchers [2, 3, 4]. In [2], a Poisson distributed random variable is used to simulate different types of reservation requests, and a binomial distribution for modeling no-shows and cancelations. Optimal strategies for renting hotel rooms taking multiple days stay into account are studied in [3], and a truncated Poisson process is used to describe the arrival process. Dynamic operating rules for motel reservations with the objective of maximization of profit are developed in [4], and a binomial distribution is used for cancelations and no-shows in the simulation.

A Poisson distribution is intuitively appealing since the requests are received primarily for individual rooms rather than groups. The mean of the Poisson distribution, which represents the rate at which reservation requests are received, is obtained from historical data. Since the rate at which reservations are requested vary depending on how close or far the arrival day is from the simulation day, we must use time-varying reservation rates to model the process more accurately. To accomplish this, we divide the booking horizon into a number of booking periods, and each booking period has its reservation request and cancelation rate. This concept of booking periods is illustrated in Figure 2 below. In the figure, r_i and c_i represent the reservation and cancelation rates, respectively, in booking period $t = i$. Note that different booking periods contain different number of days. This is because booking periods closer to the arrival day receive greater number of requests, and so they have fewer days in the period. Conversely, booking periods further away from the arrival day have more number of days since they record small number of requests. The different reservation and cancelation rates are also obtained from historical data for that particular arrival day and market segment. This implementation gives a more realistic approximation to the actual build-up curve.

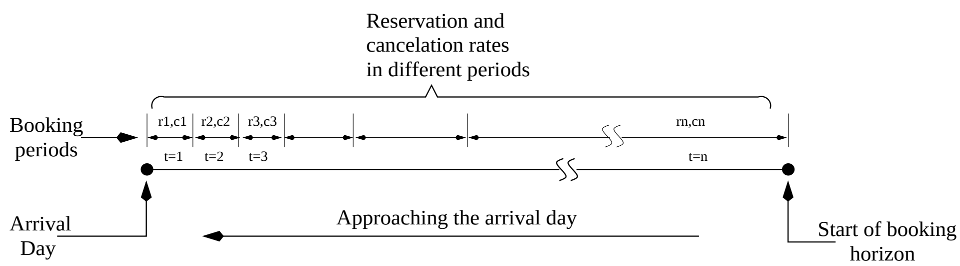

Figure 2: Booking Horizon and Booking Periods

The probability function for generating random reservation requests according to the Poisson law with time varying rates is given by

$$f(x) = \begin{cases} \frac{e^{-\lambda_t}(\lambda_t)^x}{x!} & \text{for } x = 0, 1, 2, \ldots, \\ 0 & \text{otherwise} \end{cases} \quad (13)$$

where,

$$\lambda_t = \frac{\text{number of requests received in booking period } t}{\text{number of days in booking period } t} \quad (14)$$

In the above equations, $f(x)$ is the probability that the random variable (number of requests received per day) takes on values from 0,1,2, ... The λ_t's are the different request rates in periods $t = 1, \ldots, n$, where $t = 1$ is the booking interval closest to the arrival day, $t = n$ is the interval farthest from the arrival day and n is the total number of booking periods.

A reservation may also be canceled prior to the arrival day, and we seek to model these cancelations as well. Cancelations have to be taken into account to be able to compensate for the loss in potential revenue due to canceled reservations. We model the cancelation with a binomial random variable. For each reservation, a binomial random variate is generated, with parameter equal to the probability of the reservation cancelling. This probability is obtained from historical data, and varies according to how close or far the arrival day is from the simulation day. A binomial distribution was chosen as there are only two possible outcomes: either the reservation will cancel or it will not. By generating such a random cancelation variate for each reservation, we can calculate the net reservations for each arrival day/market segment combination. We can then calculate the percentage of these net reservations that actually show-up on the arrival day by multiplying the net reservations and the

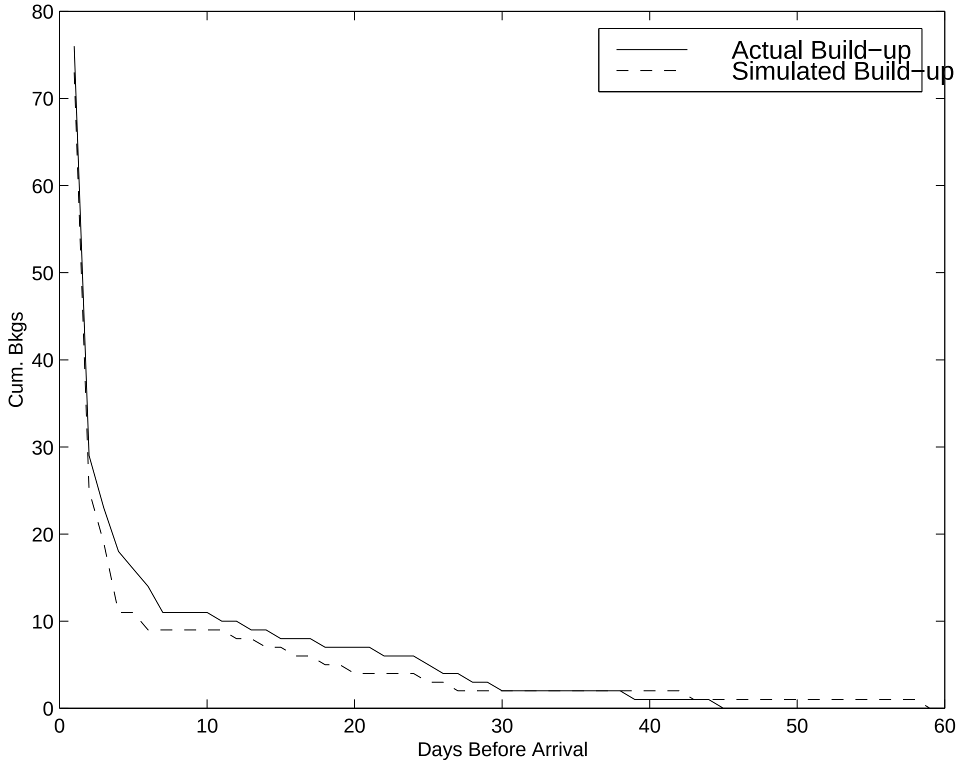

Figure 3: Actual and simulated build-up curves

show-up rate for that arrival day/market segment combination. The show-up rate is also obtained from historical data.

The probability function of a binomial distribution used to generate random cancelations is given by

$$f(x) = \binom{n}{x} p^x (1-p)^{n-x} \tag{15}$$

where x is the random variable, n is the number of Bernoulli trials and p is the probability that reservation will cancel. Here n, which represents each reservation, equals 1, and $f(x)$ is the probability that the random variable x takes values 0 or 1. The random variable itself represents the event that the reservation will cancel, with $x = 1$ meaning that the reservation will cancel and $x = 0$ meaning it will not.

Actual and simulated build-up curves are shown in Figure 3. The ordinate in the two figures correspond to the cumulative bookings that showed up on the arrival day. As seen from the two figures, we are able to obtain a realistic approximation of the build-up process using the approach described above.

3.5 Long Term and Short Term Forecasts

The forecasted value of demand is comprised of two components: the long term and the short term forecasts. As the name indicates, the long term forecast estimates the final demand for the different arrival dates/market segment combinations well in advance of the arrival dates. The short term forecast on the other hand, estimates the final demand only after the hotel property starts receiving booking for an arrival day. Typically, most of the advance booking requests are received during the 60 days before the arrival day; hence the name short term forecast. The final forecast is a weighted combination of the long term and the short term forecasts.

The long term forecast is the dominant component of the final forecast when the arrival day is far away from the processing day. Conversely, as the arrival day approaches, the short term forecast becomes the dominant component, since the short term forecast depends on the actual booking rate. The corresponding weights associated with each of the two forecasts also change as the arrival day approaches. Initially, when the arrival day is far away, the long term forecast weight is near unity, while the short term forecast weight is almost zero. As the arrival day gets closer, the short term forecast weight increases and the long term forecast weight decreases. Eventually, when the arrival day is very close, the short term forecast weight will approach unity while the long term forecast weight will be near zero. Each component of the forecast and the forecast weights are described below.

We wish to point out that this approach of using two forecasts (short term and long term) and then calculating a combined estimate is well known in the literature [11]. Notably, estimation of load on electric power systems use this methodology very frequently. Procedures for electric load estimation using two forecasted are discussed in [12]. This paper uses two models - a stochastic load model based on historical data and a weather load model which takes into account the effect of weather variables on the load pattern. The forecasts from both models are then combined optimally to give the final forecast.

3.5.1 Long Term Forecast

The long term forecast estimates the demand for future arrival dates based on historical data. The forecast may be made as much as a year ahead of the arrival date, and so it is called the long term forecast. We use the Holt-Winters process (described in Section 2) for this forecast.

3.5.2 Short Term Forecast

The short term forecast is an estimate of unconstrained net demand for future dates based on the actual advance booking activity. This is in contrast to the long term forecast, which estimates the demand entirely on the basis of the historical demand pattern. In estimating the demand, the short term forecast uses the current reservation held, the cancelation rates for these held reservations and the number of reservations that were turned down. In an equation form, it is expressed as

$$\begin{aligned}\text{S.T. forecast} &= (\text{net reservations held} \times \text{cancelation rate }) + \\ &\quad \text{net reservations turned down}\end{aligned} \tag{16}$$

The terms on the right side of Eqn. 16 are self explanatory. The net reservations held correspond to the number of reservations for an arrival date that are outstanding at the time of calculating the short term forecast. This quantity (net reservations held) when multiplied by the cancelation rate gives the number of reservations that will actually show up on the arrival date. The second term on the right side of Eqn. 16 represents an estimate of the number of reservations that were turned down or denied due to capacity or pricing constraints. If these requests were accepted, then the demand would be higher by an equivalent amount. This amount is included in the short term forecast since we want to forecast the unconstrained demand.

3.5.3 Forecast Weights

The final forecast consists of combining the long term and the short term forecasts to produce a single composite forecast. This is achieved by taking a weighted sum of the two forecasts. The objective is to give the long term forecast a higher weight than the short term forecast when the processing day is far away from the arrival day. Alternately, the short term forecast is given a higher weight when the processing day is close to the arrival day. The weights are normalized, i.e., their sum is unity.

Initially, the weights are set according to the number of bookings and number of forecasted demand. During the execution of the forecast program, the weights are continually updated. The update factor depends on the mean square error (MSE) between each of the forecast and the actual value. The new weights are given by

$$\text{new weight} = \frac{\text{ST forecast MSE}}{\text{ST forecast MSE} + \text{LT forecast MSE}} \tag{17}$$

The weights are themselves updated by taking a weighted average of the new weights and the old weights. The updated weights are given by

$$\text{updated weight} = \eta \times \text{old weight} + (1 - \eta) \times \text{new weight} \tag{18}$$

The parameter $0 < \eta < 1$ is fixed arbitrarily, depending on how much importance we wish to give to the new weights. We use $\eta = 0.9$ in the simulations.

3.5.4 Combined Forecast

After having calculated the short term and the long term forecasts and the forecasts weights, the final combined forecast can be calculated as

$$\text{final forecast} = \text{lt-weight} \times \text{LT forecast} + \text{st-weight} \times \text{ST forecast} \tag{19}$$

where

$$\begin{aligned} \text{lt-weight} &= \text{updated weight} \\ \text{st-weight} &= 1 - \text{lt-weight} \end{aligned}$$

The final forecast is calculated for each arrival day.

3.6 Forecast Algorithm

Having seen the different components of the forecast, we can now study the actual procedure used. The forecast algorithm can be easily understood with the help of the flowchart shown in Figure 4.

The flowchart shows how the two forecast components are calculated and combined to give the final forecast. Initialization of the long term component involves setting the values of the mean, trend and seasonal components, and is performed only once, at the beginning of the forecast process. One year (52 weeks) worth of data is used for the initialization. All the other components of the long term forecast are updated nightly. The next step in the long term forecast is to find the optimal value of the smoothing parameter α (refer Section 2) that minimizes the current forecast error. The current forecast error is the difference between the forecasted demand for the current day and the actual demand for the same day. We also use the root square forecast error to detect non-random effects in the data. Depending on whether such effects are present, the increments with which α is changed is decided. To ensure the stability of the Holt-Winters procedure, upper and lower bounds (0.5 and 0.025 respectively) are imposed on the value of α. We start with initial values of $\alpha = 0.2, \beta = 0.05$ and $\gamma = 0.1$. Once the values of the smoothing constants are fixed, the individual components of the long term forecast are computed for all future arrival days. These individual components are then used to generate the final long term forecast.

The initialization for the short term forecast comprises of determining the cancelation rates from the historical data. This is required since the actual cancelation rate for a given arrival day will be known only after the arrival day has passed. We assume that the cancelation rates do not change from year to year. The short term forecast uses the reservations held, cancelation rate and the turndowns to estimate the net demand. Each day, the number of reservations and turndowns are observed and these together with the cancelation rate are used in the generation of the short term forecast.

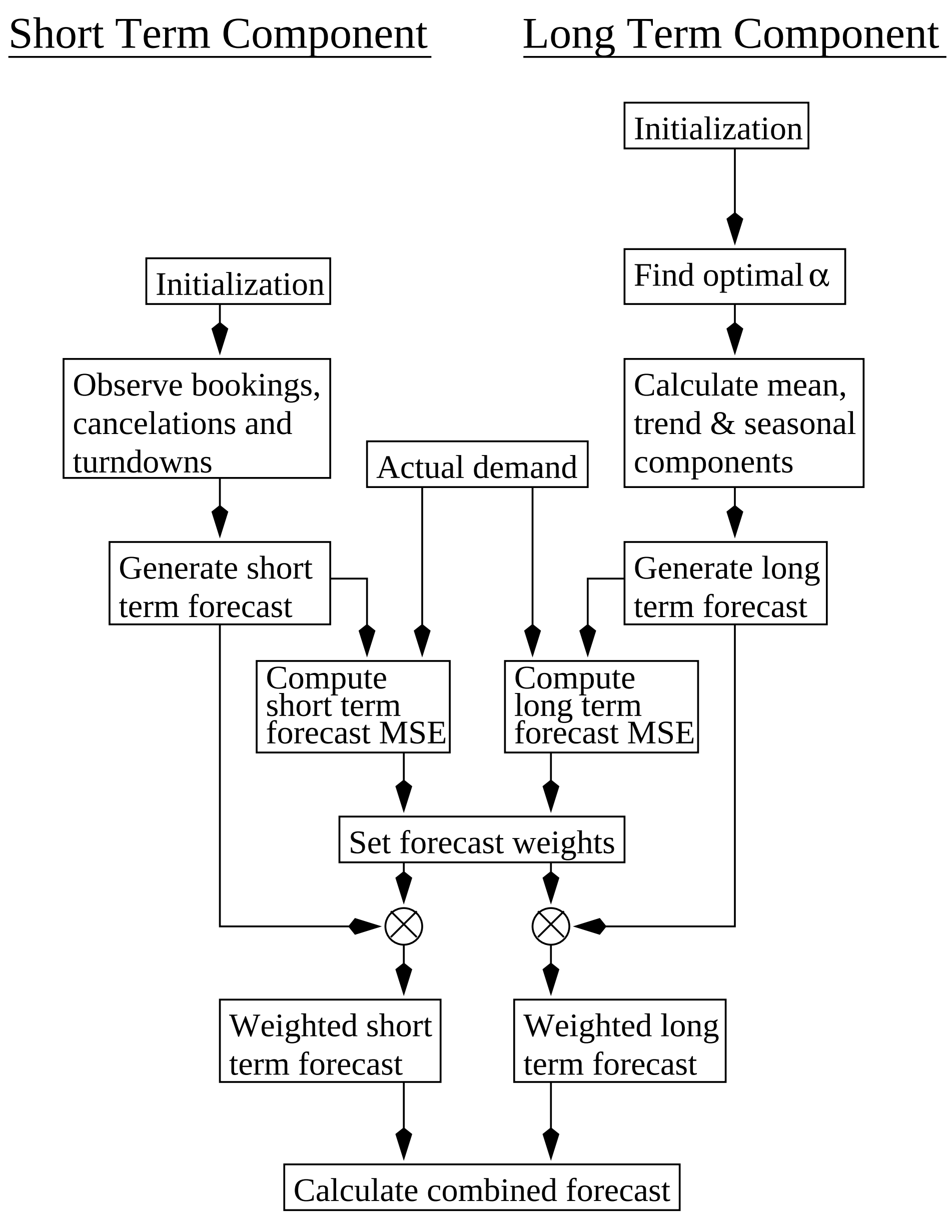

Figure 4: Flowchart of Forecast Algorithm

The long term and short term forecasts generated are then each compared with the actual demand for the current day, and the current forecast mean square errors are calculated for each forecast. The forecast weights, which are given initial values during the long term forecast initialization, are updated based on these errors as given in equations 17 and 18. The weighted short term and long term forecasts are then summed up to give the final composite forecast.

4 Simulation Results

As discussed in Section 3.2, 58 weeks of data from an actual hotel property were used for initialization and simulation purposes. Data from the first 52 weeks are used for initialization of the forecast parameters, and data from the following 6 weeks are used in obtaining the reservation and cancelation request rates to simulate random reservation requests as well as cancelation requests.

To see the effects of the long term and the short term forecasts for a particular arrival day distinctly, we should start forecasting the demand for that arrival day well in advance. Since the simulation period is only 6 weeks, we can expect to observe both components of the forecast for the last two weeks of the simulation period. We choose one weekday and one weekend day in the last week of the simulation period to test the forecast algorithm. This choice was made because the hotel property from which the data was obtained caters primarily to business travelers, and so different booking patterns can be expected depending on the day of the week.

Figure 5 shows the actual buildup of reservations, the combined forecast and its components for a weekday (Test Day 1) in the last week of the simulation period. The forecast of demand for a particular arrival day is done every night prior to the arrival day. For this example, the forecast starts 38 nights before the arrival day.

In Figure 5, the combined forecast and its two components give estimates of the demand upto one day before the arrival day. The arrival day corresponds to the zero coordinate on the x-axis. The final forecasted values lie on the vertical line drawn at one day before the arrival day. The point at which the actual demand curve meets the y-axis gives the value of the number of reservations that showed up on the arrival day.

As seen from the figure, the combined forecast follows the long term forecast closely when the arrival day is far from the day the forecast is made. The number of actual reservations initially is very small, and so the short term forecast does not give a very accurate estimate. As the arrival day approaches, the rate of reservation requests increases, and the short term forecast gives better estimates. The combined forecast, being a weighted sum of the two forecasts, lies in between the two estimates. Eventually, the combined forecast gives a satisfactory estimate of the final demand.

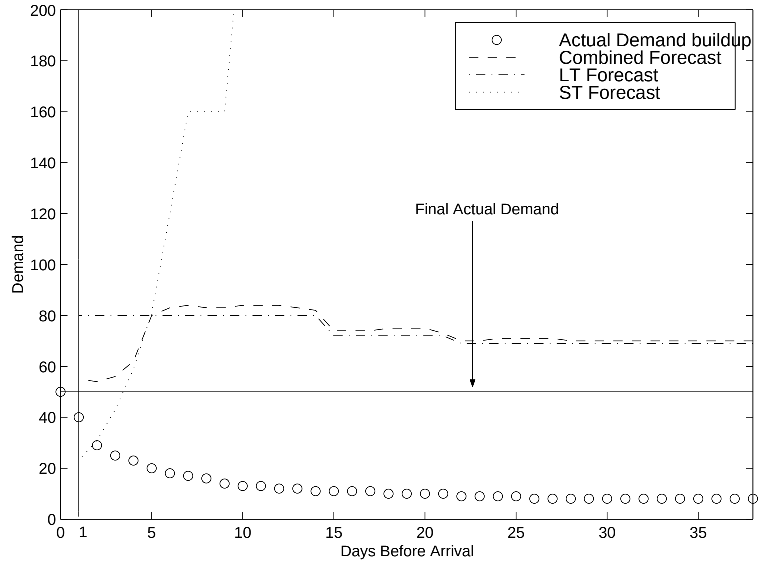

Figure 5: Actual and Forecasted Unconstrained Demand for Test Day 1

Figure 6 shows a similar graph for Test Day 2 in the last week of the simulation period. Similar to the previous case, the forecast for this arrival day starts 40 nights before the arrival day.

The results obtained in the second test case are similar to those obtained in the first. The combined forecast follows the long term forecast initially, and as the arrival day approaches, it takes into account the effect of the short term forecast. We obtain a good estimate of the unconstrained room demand in this case as well.

Figure 7 shows a similar plot for another day in the last week. In this case however, there is a large deviation of the forecasted value from the actual value. The reasons for this behavior of the forecast are discussed later in the section. The forecsat for this day started 36 nights prior to the arrival day.

From the above results, we see that neither the long term nor the short term forecast can give a reasonable estimate of the demand individually, but a weighted sum appears to be effective in some cases. This should not be surprising, since the long term forecast gives an estimate on the basis of historical demand. The short term forecast, on the other hand, depend on actual reservations and turndowns and this makes the task of forecasting difficult when the arrival day is very far away from the processing day. This provides the motivation for the use of a long term - short term forecast methodology.

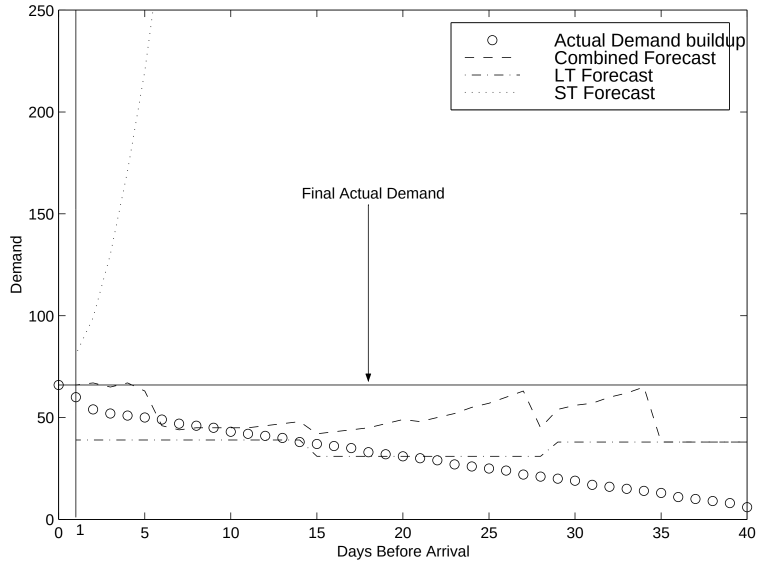

Figure 6: Actual and Forecasted Demand for Test Day 2

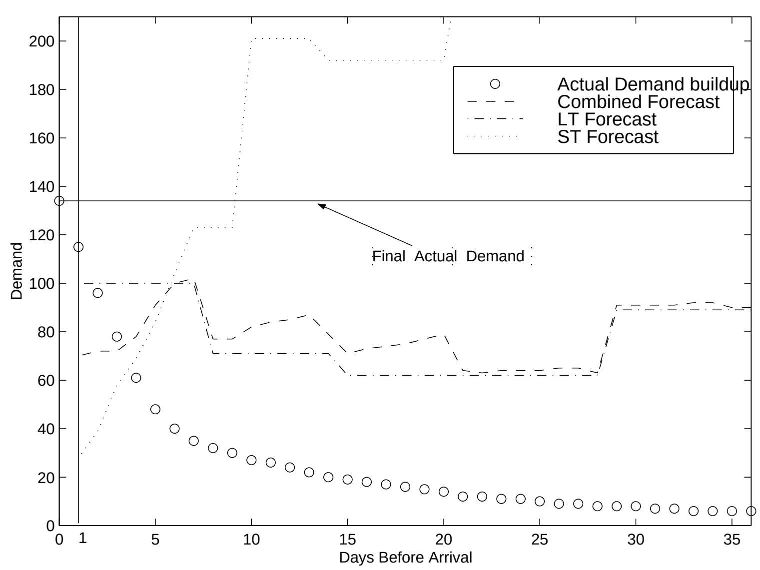

Figure 7: Actual and Forecasted Demand for Test Day 3

Monte Carlo simulations have been used to compare the actual and forecasted demand. This was necessitated by the fact that the reservation and cancelation requests are generated randomly, and it is essential that we average over several realizations of these requests to obtain a realistic approximation to the reservation/cancelation process. In this paper, the actual and forecasted demand have been averaged over 100 runs for each arrival day.

The mean absolute deviation (MAD) is used as a measure of performance of the forecast algorithm. This is calculated over the last 2 weeks of the simulation period. The mean absolute deviation is given by

$$\text{MAD} = \frac{\sum |\text{actual demand} - \text{forecasted demand}|}{\text{number of days}} \tag{20}$$

A plot of the mean absolute deviation is shown in Figure 8.

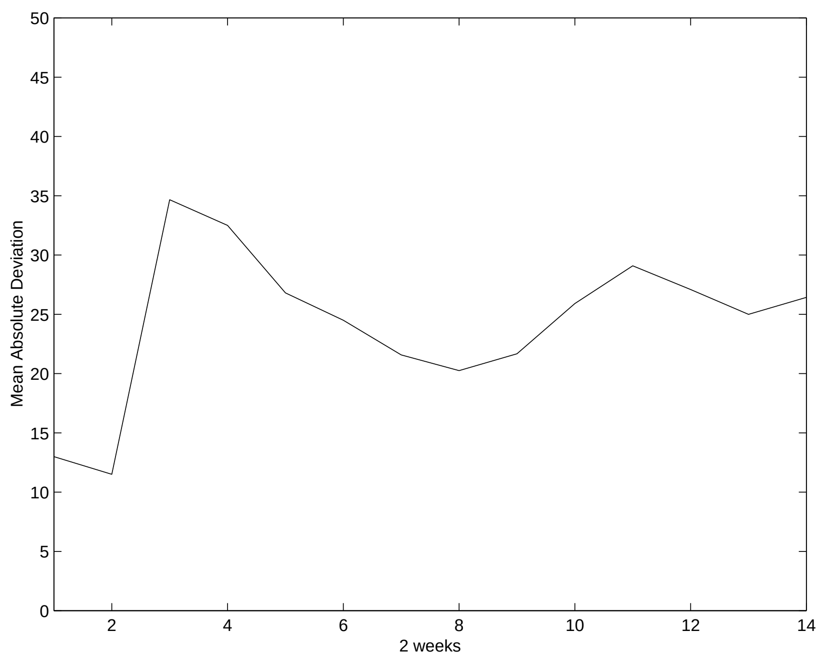

Figure 8: Mean Absolute Deviation for final 2 weeks of simulation period

Another measure of performance is the mean absolute percentage error (MAPE). MAPE is computed by averaging the absolute error between the actual and forecasted demand expressed as a percentage of actual demand over time. Figure 9 shows a plot of the mean absolute percentage error.

$$\text{MAPE} = \frac{\sum \frac{|\text{actual demand} - \text{forecasted demand}|}{\text{actual demand}}}{\text{number of days}} \tag{21}$$

From the various graphs, we notice that while a good forecast is obtained from some days, there are occasions on which the algorithm does not perform satisfactorily. The primary

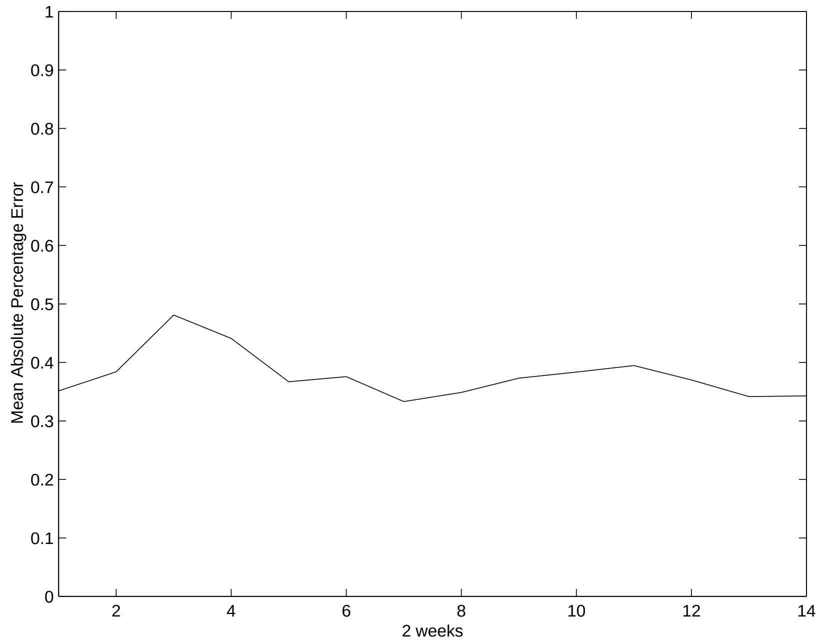

Figure 9: Mean Absolute Percentage Error for final 2 weeks of simulation period

reason for this is that the forecast is based entirely on a quantitative model and hard data. It does not take into account external/non-random effects which may have influenced the demand, but known to the hotel manager. Also, the algorithm is unable to distinguish non-recurring events based only on historical data. This provides a good motivation for including human knowledge in the forecast procedure.

5 Discussion and Future Research

This paper discussed the Holt-Winters forecasting procedure and its application to forecasting hotel room demand. The paramaters of the Holt-Winters model are initialized using historical data obtained from an actual hotel. The Holt-Winters forecast approach was used to compute the long term forecast of room demand. The short term forecast was computed based on actual booking activity. The final forecast is a weighted sum of the long term and the short term forecasts, and the forecast weights are decided by the mean square error between the forecasted and the actual values. Random reservation and cancelation requests are simulated using Poisson and binomial distributions, respectively. The parameters for the distributions were obtained using historical data. The simulations are carried out for a single rate class. Extending the algorithm for multiple rate classes is straightforward.

A lot of the difficulties encountered in testing the algorithm were data related. Having 58 weeks of data is barely sufficient to initialize and test the forecast algorithm. Due to this, it is not possible to update the seasonal component and observe its effect. Additionally, it is possible that the data contained non-random effects (which may affect the demand) or unexplained crests or troughs in demand and this caused the algorithm to give erratic demand values on some occasions.

There are several issues which are still to be addressed and accounted for in order to obtain a truly accurate forecast algorithm. One of the most important objectives of this project is to be able to incorporate expert knowledge into the system forecast. It has been observed that hotel managers are able to give very accurate forecasts within a period of 2-3 weeks of the arrival day. This is because they are able to supplement their experience with their knowledge of events, demand patterns, etc. We seek to represent this knowledge into the forecast algorithm. The resulting algorithm is not intended to replace the hotel managers; instead it would take the expert knowledge as an input and combine it with a mathematical forecast to give the best possible forecast.

The role of the forecast should also be seen in the proper perspective. As mentioned in the beginning of the paper, we feel that a good forecast will result in better inventory optimization and management. Thus, the objective is to be able to obtain a good forecast consistently for all future days, rather than having an exact prediction of the demand on some days. This logically takes us to the use of fuzzy logic in the forecast algorithm. These ideas are being investigated and they certainly present a new approach of tackling the problem. Work along this line of thought will be reported in a future paper.

References

[1] Cross, R. G., *Revenue Management*, Broadway Books, New York, 1997.

[2] Bitran, G. R., and Gilbert, S. M., *Managing Hotel Reservations with Uncertain Arrivals*, Opns. Res., Vol. 44, No. 1, (Jan. - Feb. 1996) pp. 35-49.

[3] Bitran, G. R., and Mondschein, S. V., *An Application of Yield Management to the Hotel Industry Considering Multiple Day Stays*, Opns. Res., Vol. 43, No. 3, (May - Jun. 1995) pp. 427-443.

[4] Ladany, S. P., *Dynamic Operating Rules for Motel Reservations*, Dec. Sci., Vol 7, (1976) pp.829-840.

[5] Montgomery, D. C., and Johnson, L. A., *Forecasting and Time Series Analysis*, McGraw-Hill, New York, 1976.

[6] Harvey, A. C., *Time Series Models*, MIT Press, Cambridge, MA, 2nd ed., 1993.

[7] Box, G. E. P., and Jenkins, G. M., *Time Series Analysis, Forecasting and Control*, Holden-Day, San Francisco, Rev. ed, 1976.

[8] Chatfield, C., *The Analysis of Time Series: An Introduction* Chapman and Hall, London, 2nd ed., 1980.

[9] Winters, P. R., *Forecasting Sales by Exponentially Weighted Moving Averages*, Mgmt. Sci., Vol 6, No 3, (1960) pp 324-342.

[10] Chatfield, C., *The Holt-Winters Forecasting Procedure*, Applied Statistics., Vol 27, No 3, (1978) pp.264-279.

[11] Sage, A. P., *Optimum Systems Control*, Prentice-Hall, Englewood Cliffs, NJ, 1968.

[12] Gupta, P.C., and Yamada, K., *Adaptive Short-Term Forecasting of Hourly Loads using Weather Information*, IEEE Tran. Power Appar. Syst., Vol. PAS-91, Sept./Oct. 1972, pp. 2085-2094.

[13] Harvey, A. C., *Forecasting, Structural Time Series Models and the Kalman filter*, Cambridge University Press, Cambridge, 1989.