PERGAMON

Socio-Economic Planning Sciences 33 (1999) 25–38

Bed allocation techniques based on census data

Sophie D. Lapierre[a,]*, David Goldsman[b], Roger Cochran[c], Janice DuBow[d]

[a]CRT and Department of Mathematics and Industrial Engineering, École Polytechnique, C.P. 6079 succ. Centre-ville, Montréal, QC, Canada H2V 1N2

[b]School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0205, U.S.A.

[c]Cochran/Merrick Associates, 839 Southern Shore Drive, Peachtree City, GA 30269, U.S.A.

[d]Northside Hospital, 1000 Johnson Ferry Rd, Atlanta, GA 30319, U.S.A.

Abstract

This paper considers the problem of allocating a number of beds to the different medical and surgical specialties in a hospital. This task is complicated by the fact that patterns of patient arrivals differ among specialties, that the scheduling of medical procedures varies over a week, and that the demand for various medical services can show seasonality. We develop a time series model using hourly census data to make good decisions regarding the size of each unit while minimizing data collection and modeling efforts. This model is simple enough to be used widely and, we suggest, represents an improvement versus methods currently employed by hospitals. Results are presented from an application to data collected at Northside Hospital, Atlanta, GA.

1. Introduction

Bed allocation concerns the permanent number of beds assigned to the different medical and surgical specialties in a hospital. An appropriate bed allocation is important for cost-efficiency of hospitals. Too few beds assigned to a specialty may lead to customer assignments to units with inappropriate equipment and inadequate staff training. This, in turn, would likely result in lower quality of care. Further, too many beds increases costs through an underutilization of resources.

* Corresponding author. E-mail:sophiel@crt.umontreal.ca.

PII: S0038-0121(98)00008-1

A fundamental principle for appropriate bed allocation is that beds needed by a specialty should be determined such that sufficient specialty beds are available "most of the time" [3, 4, 10]. This basic concept has been used for the design of public systems such as sewage, electricity, and phone networks since it would typically cost too much to satisfy every possible contingency. Hospitals can sometimes cope with higher than usual loads by "bumping" customers to other medical or surgical units, setting up temporary beds, or postponing elective surgeries or treatments. Of course, these overflow situations are not desirable as they put stress on staff and can affect the quality of services [12]. Administrators should therefore allocate beds to serve demands for each medical and surgical service such that equally important services have a similar allocation level.

Bed allocation is a macro-level type decision, where the relevant question to be answered is: "How many beds are needed to support a medical or surgical unit under its current admission, treatment, and discharge practices?" Most work on bed capacity planning has used a simulation approach. Fetter and Thompson [4] first used simulation models to determine capacity needs. Esogbue and Singh [3], also using simulation, analyzed the repartition of a fixed number of beds among different specialties within a unit. On the other hand, Hershey et al. [9] used a queueing model, while Kao and Tung [10] used a combination of simulation, queueing theory, and statistical analysis to determine bed requirements.

Subsequent work has extended bed allocation models to address efficiency questions such as "How well can we use those beds?" Harris [8] thus assessed operating timetables along with the resulting bed requirements, while Dumas [2] compared how different admission and transfer policies affect bed space. McFadden [14] studied new obstetric patient movements and how they would improve overall operational efficiency. Vassilacopoulos [16] incorporated the dimension of the length of a hospital's waiting list while Romanin-Jacur and Facchin [15] tied efficiency of bed use to staff load. The bed efficiency questions examined by these extensions are mainly of intra-departmental concern, where answers to them are not necessary in order to address hospital-wide bed allocation decisions. Kao and Tung [10], however, suggest that if both issues are to be considered, bed allocation decisions should be made only after departmental efficiency questions have been dealt with.

All models cited above make several simplifying assumptions about patient arrival rates and length of stay. Most of them thus assume that arrivals are Poisson, so that the arrival rate does not change according to the hour of the day and day of the week. Kao and Tung [10] recognize this simplification, but show that it does not affect the sensitivity of their model. Hancock et al. [7] show that a significant proportion of admissions is elective (i.e. planned admissions) and that bed requirements are lower than what models based on the Poisson assumption would suggest. More recent simulation models, such as those developed by Romanin-Jacur and Facchin [15] and Lowery and Martin [12], use several arrival rates per day and break down arrival distributions into different patient groups.

Simulation models have become highly complex as health care problems have increased in sophistication. The simplifying assumptions that were appropriate when the health care system was relatively *stable* and lengths-of-stay were long, are simply not valid anymore. The average length-of-stay in hospitals has been reduced drastically; e.g. following a long period of stability, the average hospital stay at Northside Hospital after a normal vaginal delivery has been steadily declining, from 3.8 days in 1976 to 1.7 days in 1995. Also, many surgical procedures that used to require a hospital stay, such as cataract removal and mastectomies, are now normally performed on an outpatient basis. In addition, closure of hospitals has brought new customers to the remaining hospitals. Useful bed allocation models in this new context now require data that are recent and that can be collected and analyzed quickly. The most recent simulation models, however, require what we view as excessive development time. This paper thus introduces a model that attempts to reduce the difficulty of data collection and model formulation, while providing "good" solutions.

Our model was developed to assist Northside Hospital, a 450-bed regional hospital in Atlanta, GA, in making hospital-wide bed allocation decisions with an emphasis on obstetric services. We begin by discussing the information typically available in hospitals to help health care managers with their decisions regarding bed allocation. Next, we present our model using hourly census data, where census refers to the number of patients in the different units at a given time. We then examine how hourly census data are extracted from the existing information system for use in our model. Finally, we discuss results from the application of our model to Northside Hospital.

2. Existing information and models

Hospital information systems have been built around an organizational model where patients are admitted in the early afternoon and discharged early the next morning. In such an organizational environment, the peak census occurs at midnight and the length-of-stay is counted in days, with the first and last days of the stay considered full days. Therefore, hospital information systems typically keep "census-at-midnight" data obtained from their computerized admissions system, and "patient-days" per specialty obtained from the cumulative number of days patients have spent at the hospital during a particular month (or year). Given today's changing health care environment, where length-of-stay is diminishing, data kept by information systems are likely insufficient for making effective decisions regarding bed allocation.

As a result of conditions noted above, the assumption that occupancy peaks occur at midnight may not always apply. Fig. 1 plots the data from the average census-at-midnight per day of the week of an obstetric unit at Northside Hospital. It appears that some days are busier than others, but this does not reveal the "real" picture. As we see in the figure, the midnight census "hides" higher census peaks that can occur at different times of the day. Our study showed that a discrepancy of up to 20% exists between the average census-at-midnight and the average census at some other time during the same day for this particular unit. A further complication in extrapolating midnight census data to bed allocation decisions is that the size and timing of the discrepancy varies from unit to unit. Thus, our suggestion that census-at-midnight data represent insufficient information on which to make bed allocation decisions.

Similarly, patient-days per month also represent an insufficent database. Because of different practice and discharge patterns, two different services with the same number of patient-days per month might thus require a different unit size. For example, planning is simplified in departments such as optometry, where elective surgery accounts for a large proportion of the

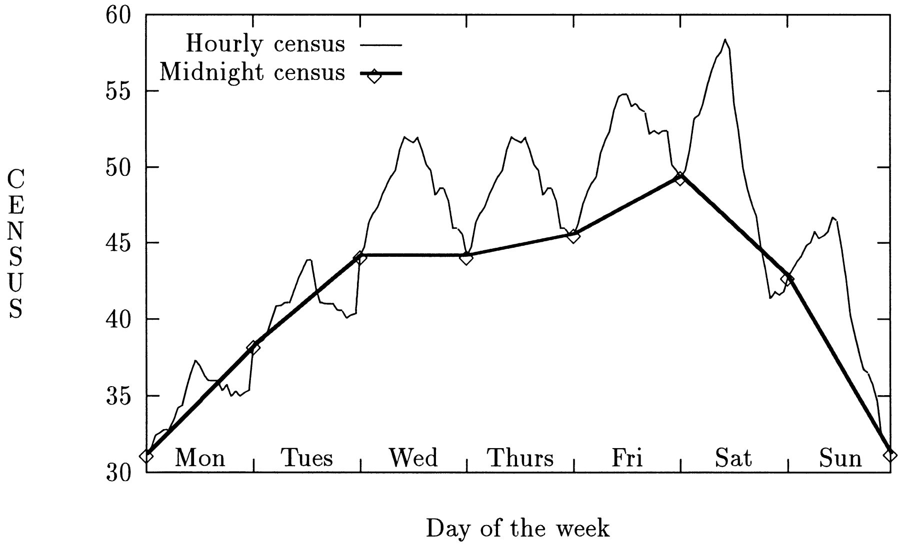

Fig. 1. Perinatal unit—average census.

cases; because admissions and discharges can be better controlled, fewer beds are required for a given number of patient-days. By contrast, the difficulty of forecasting obstetric deliveries means that more beds would be needed for the same number of patient-days to ensure beds are available should there be a surge in admissions. If admissions and discharges were truly random, patient-days per month would provide enough data to use queueing theory to decide on bed allocation. But even in an obstetric unit, admissions and discharges are not truly random; discharges seldom occur at night and admissions for induced deliveries will generally be during the daytime on weekdays.

Patient-days information must therefore be complemented with information on admission and discharge patterns. These patterns, however, often depend on the exact procedure to be performed. For instance, planned Cesarean surgeries are scheduled mainly for Tuesdays and Wednesdays and generally require a 3–4-day stay. However, distinction between a planned and an emergency Cesarean is not reported to the information system. With hospitals performing up to 3000 different procedures, this lack of information makes discerning patterns in admission and discharge extremely difficult. The only way to pinpoint exact procedures (and, by extension, admission and discharge patterns) is to specifically extract the data from patients' medical files, a very labor-intensive task. Because of the need for data beyond patient-days data, standard simulation models are thus not an attractive option.

Given these circumstances, we have modeled hourly census data, and used patient-days yearly reports to validate our census data. Our model has been developed to accurately define the number of beds a medical or surgical specialty needs even if admission patterns change according to the day of the week and/or hour of the day. Importantly, this model must be combined with long-term trends, as admission patterns and lengths-of-stay will likely change over long periods of time.

3. The model

In presenting our method, it is assumed that we have in hand hourly census data from the busiest times of the year for a given medical or surgical specialty. Hourly census data are obviously *autocorrelated* because the census for a given day and hour depends on the census from the previous hour. However, the census from one week to another can be regarded as roughly independent. This independence stems from the fact that the average length of stay in all units of Northside Hospital is less than one week. We can consequently look at data on a weekly period, breaking down data sets per day of the week. Fig. 2 shows census data of an obstetric unit for five different Sundays. We can see from the figure that the daily curves have similar shapes. The model we introduce here has generated the curves given in Fig. 3. (Notice that Figs. 2 and 3 are similar-looking.) The basic idea is that, for each day of the week, we can separate the variation into a weekly variation and an hourly variation. Mathematically, the model is as follows:

$$B_{i,j} = \mu_{i,j} + A_i + \varepsilon_{i,j} \quad i = 1, \dots, 7;\ j = 0, \dots, 23 \tag{1}$$

where:

$B_{i,j}$ is the number of beds required on day i, hour j. (See Fig. 4 for a graphical representation of each component of $B_{i,j}$.)

$\mu_{i,j}$ is the mean number of beds required on day i and hour j which we estimate with the sample mean $\bar{X}_{i,j} \equiv \frac{1}{n}\sum_{k=1}^{n} B_{i,j,k}$ [see Eq. (4)].

A_i introduces variation from week to week; it is the average difference between the real and average census curves; $A_i \sim N\ (0,\ S_i^2)$, i.e.A_i follows a normal distribution with mean 0 and variance S_i^2, where $S_i^2 \equiv \frac{1}{24}\sum_{j=0}^{23} S_{i,j}^2$ and $S_{i,j}^2$ [see Eqs. (5) and (6)] is the sample variance on day i, hour j.

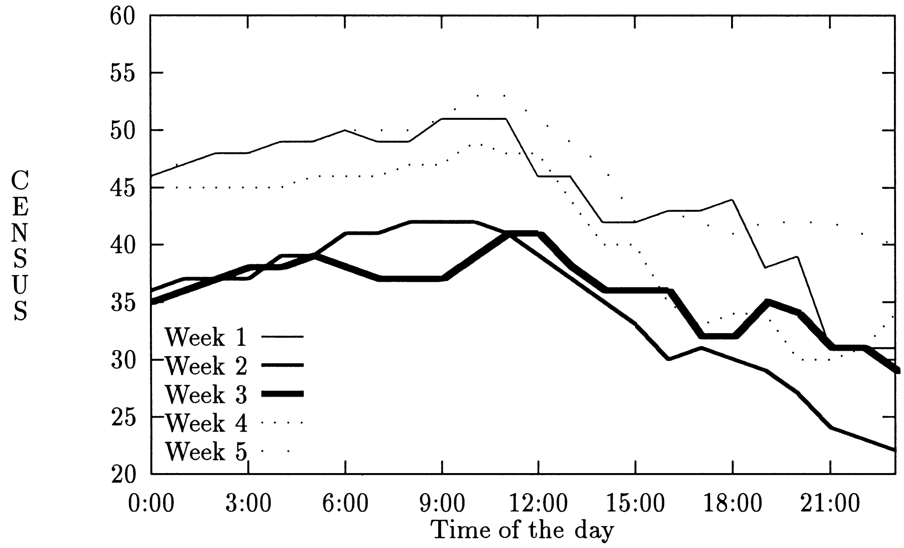

Fig. 2. Obstetric unit—real data for five Sundays.

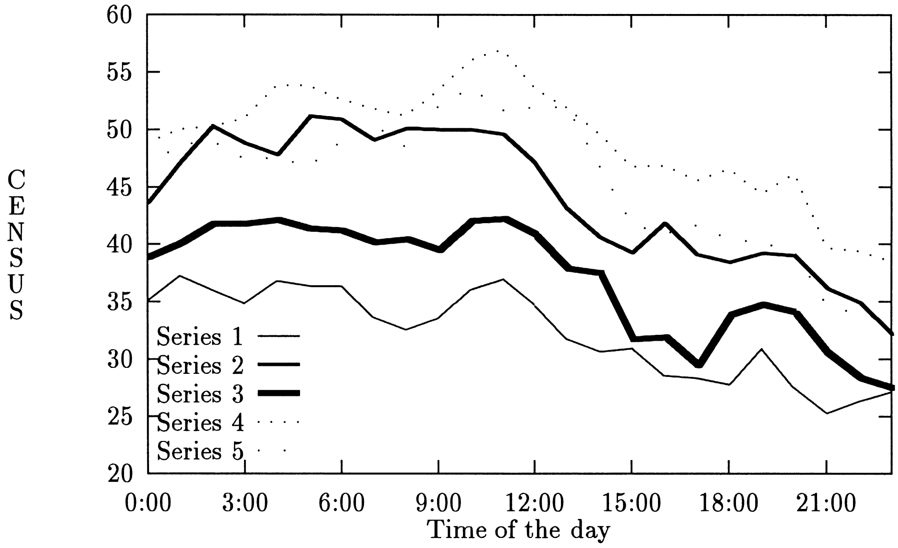

Fig. 3. Obstetric unit—simulated time series data for five Sundays.

$\varepsilon_{i,j}$ introduces hourly variation; it is an autoregressive noise factor, to be discussed below.

We model here the $\varepsilon_{i,j}$'s as a *first-order autoregressive (AR(1))* process with coefficient ϕ_i, i.e.

$$\varepsilon_{i,j} = \varphi_i \varepsilon_{i,j-1} + z_{i,j},$$

where

$$z_{i,j} \sim N(0, \sigma_{z_i}^2).$$

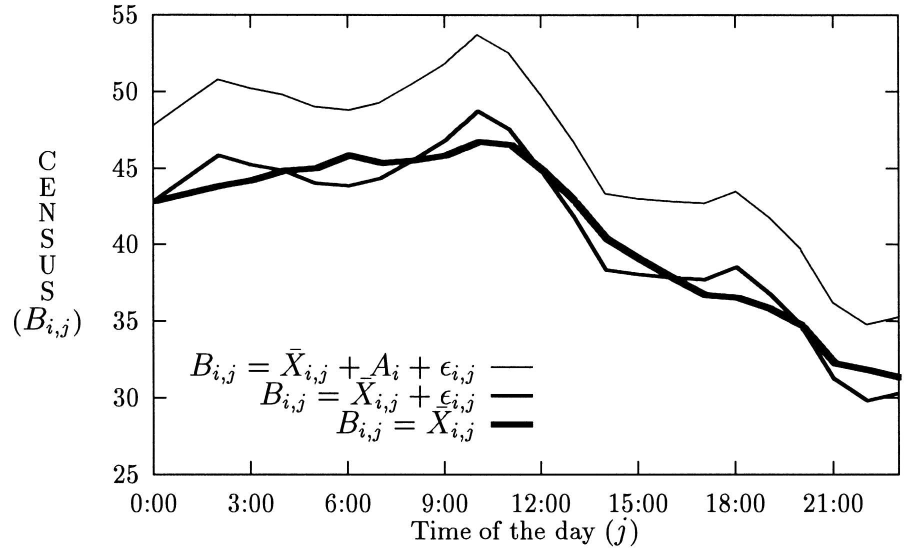

Fig. 4. Sunday—effect of each model component.

Such AR(1) models are often used in modeling elementary autocorrelated time series. The $\varepsilon_{i,j}$'s have the following properties:

$$\varepsilon_{i,j} \sim N\left(0, \frac{\sigma_{z_i}^2}{1-\varphi_i^2}\right)$$

$$\mathrm{Cov}(\varepsilon_{i,1}, \varepsilon_{i,1+k}) = |\varphi_i|^k.$$

From Fuller [5], we can estimate the ϕ's and σ_z^2's in the following way:

$$\hat{\varphi}_i = \frac{\sum_{j=0}^{n-2} \varepsilon_{i,j}\,\varepsilon_{i,j+1}}{\sum_{j=0}^{n-1} \varepsilon_{i,j}^2} \tag{2}$$

$$\hat{\sigma}_{z_i}^2 = \frac{1}{n-3}\sum_{j=0}^{n-2}\left(\varepsilon_{i,j+1} - \hat{\varphi}_i\,\varepsilon_{i,j}\right)^2. \tag{3}$$

3.1. Computation of parameters

We now transform (2) and (3) to account for the Northside Hospital data.

Assume that we have n weeks of data available. We break down our actual data set as follows:

$B_{i,j,k}$ $\quad i = 1, \ldots, 7$ (day of the week)

$j = 0, \ldots, 23$ (hour of the day)

$k = 1, \ldots, n$ (week of study).

We now compute, for all i and j,

$$\hat{\mu}_{i,j} = \bar{X}_{i,j} = \frac{1}{n}\sum_{k=1}^{n} B_{i,j,k}, \tag{4}$$

$$S_{i,j}^2 = \frac{1}{n-1}\sum_{k=1}^{n}(B_{i,j,k} - \bar{X}_{i,j})^2, \tag{5}$$

and

$$S_i^2 = \frac{1}{24}\sum_{j=0}^{23} S_{i,j}^2. \tag{6}$$

Eq. (4) averages data shown in Fig. 2 for each hour. Eq. (5) estimates the variance of the census for each hour and day of the week. Finally, Eq. (6) suggests that we estimate the variance of A_i from the average sample variance of the ith day of the week.

Thus far, we have estimated all model parameters except those of the underlying AR(1) model. We determine the parameters for each day of data [Eqs. (7)–(10)]; this is followed by averaging these estimated parameters across the n replicates [Eqs. (11) and (12)]:

$$A_{i,0,k} \equiv B_{i,0,k} - \bar{X}_{i,0}, \tag{7}$$

$$\varepsilon_{i,j,k} \equiv B_{i,j,k} - \bar{X}_{i,j} - A_{i,0,k}, \tag{8}$$

$$\hat{\varphi}_{i,j} \equiv \frac{\sum_{j=0}^{22} \varepsilon_{i,j,k}\ \varepsilon_{i,j+1,k}}{\sum_{j=0}^{23} \varepsilon_{i,j,k}^2}, \tag{9}$$

$$\hat{\sigma}_{z_i,k}^2 \equiv \frac{1}{21}\sum_{j=0}^{22}(\varepsilon_{i,j+1,k} - \hat{\varphi}_{i,k}\ \varepsilon_{i,j,k})^2, \tag{10}$$

$$\hat{\varphi}_i \equiv \frac{1}{n}\sum_{k=1}^{n} \hat{\varphi}_{i,k}, \tag{11}$$

$$\hat{\sigma}_{z_i}^2 \equiv \frac{1}{n}\sum_{k=1}^{n} \hat{\sigma}_{z_i k}^2. \tag{12}$$

Notice that Eq. (8) uses only the weekly variation at time 0:00, $A_{i,0,k}$, as we wish to interfere as little as possible in our modeling of the noise, $\varepsilon_{i,j+1,k}$. An alternative would have been to compute weekly variation for each hour and then take the average [in order to generate the parameter needed for Eqs. (2) and (3)]. However, since there is no clear cut difference between weekly variation and noise modeled by the AR(1), it is difficult to determine if this would add significantly to the accuracy of the model. Eqs. (11) and (12) are therefore viewed as but approximations of Eqs. (2) and (3).

3.2. Simulation

Now that the model and its parameters have been established, we use this information to simulate the hourly census. For every day i ($i = 1,\ldots,7$), we redo the following procedure n times, n being some large number of replications (say, at least 1000):

- Generate $A_i \sim S_i\ N(0,1)$ (Several techniques exist for generating a standard normal random variable. We refer the reader to Law and Kelton [11].)

- Generate independent $Z_{i,j}$'s$\sim\hat{\sigma}_{z_i}$ $N(0,1)$ $j = 0,\ldots,23$.
- Compute $\varepsilon_{i,j} = \hat{e}_i\, \varepsilon_{i,j} + Z_{i,j}$ $j = 0,\ldots,23$ (take $\varepsilon_{i,0} = 0$).
- Compute $B_{i,j} = \bar{X}_{i,j} + A_i + \varepsilon_{i,j}$ $j = 0,\ldots,23$.

For bed allocation, we are interested in knowing the census frequency distribution function in order to determine the number of beds that should be allocated to cover "most of the demand". Beds can be allocated based on the scenario of the busiest day, the busiest week, or a combination of the busiest day and week. These three ways to make bed allocation decisions are formally:

1. Sufficient beds should be allocated to satisfy the daily demand at the $(1-\alpha_d)$ level $(0\leqslant \alpha_d \leqslant 1)$, where

$$\alpha_d = \frac{\sum_{k=1}^{n}\sum_{j=0}^{23} d_{i,j,k}}{n * 24}$$

and

$$d_{i,j,k} = \begin{cases} 1 \text{ if } B_{i,j,k} > \{\text{Number of beds allocated}\} \\ 0 \text{ otherwise } (k \text{ is the simulation run number}). \end{cases}$$

2. Sufficient beds should be allocated to satisfy the weekly demand at the $(1-\alpha_w)$ level $(0\leqslant \alpha_w \leqslant 1)$, where

$$\alpha_w = \frac{\sum_{i=1}^{7}\sum_{k=1}^{n}\sum_{j=0}^{23} d_{i,j,k}}{n * 7 * 24}.$$

3. Sufficient beds should be allocated to satisfy the weekly demand at the $(1-\alpha_w)$ level, *and* sufficient beds should be allocated to satisfy the daily demand at the $(1-\alpha_d)$ level $(0\leqslant \alpha_w, \alpha_d \leqslant 1)$.

4. Data

The quality of the data is very important in seeking valid results. Indeed, it might even be more important than the mathematical model itself. In the following discussion, we examine the problems of obtaining and validating data.

4.1. Collection of data

We employed two methods to obtain an hourly census. The first was to regenerate census data from discharged patients' records. This method required extensive programming but had the advantage of enabling access to data over a long period of time. We encountered problems due to the fact that some patients' records were not retrieved from the system; the resulting census was thus lower than the real one. Since patient numbers were based on diagnosis, the lower numbers were likely due to inaccurate entry of the diagnosis into the admissions information system. We then settled on a second method that involved saving each unit's census data every hour. This method has the advantage of being simple to implement. The disadvantage is that it requires special data collection (since, currently, hospitals only record the midnight census). Further, analysis of the data must obviously wait until the period of data collection is complete. For the current study, we collected six weeks of hourly census data, a period considered to be representative of normal activity. We decided that accurate data over a period of six weeks was more desirable than unreliable data sets spanning a longer time period. Additionally, such data sets might reflect trends that no longer exist.

4.2. Validation of data

The quality of data brings up the issue of validation of data. In this regard, we used a basic result from queueing theory, Little's formula [6], which relates that over a relatively long period of time, the average number of new customers per day multiplied by the average time a customer spends in the service system equals the average number of people in the service system. For our application, we can use Little's formula to compare our study data to the data collected by the administrative system in the following way:

$$\left(\sum_{i=1}^{7}\sum_{j=0}^{23}\bar{X}_{i,j}\right)/(7*24) - (\text{Average new customers per day}) \times (\text{Average length-of-stay}) \quad (13)$$

$$= \frac{\text{Customers per year}}{365} \times \frac{\text{Patient-days per year}}{\text{Customers per year}}. \quad (14)$$

Differences between both sides of the equation can occur for several reasons:

- **Seasonality:** study data might have been taken during a busy or quiet season. Plots of past patient-days per month for a year or more will show if there is seasonality or not. Our obstetric unit data did not show any significant seasonality despite the fact that the staff perceived summertime to be busier.
- **Trend:** business loads and practice patterns can change over time. A "smoothing average" can show trends without seasonality effects [13]. Our data showed that obstetric services had a declining census for the last two years but a constant number of deliveries per year. This declining census is the result of health insurance pressure to reduce hospital stay after deliveries.
- **Other reasons:** several other factors can cause discrepencies between the study data, administrative data, and reality. For example, if data are extracted by speciality, they will include days spent in emergency rooms, operating rooms, and intensive care units. On the other hand, if data are extracted by medical and surgical units, they will not include patients who had to be located on another floor due to bed availability problems. An additional

interesting fact is that a significant difference in total patient-days per year exists between information extracted from medical records and information extracted from administrative data. This is generally due to data entry errors and lack of medical expertise on the part of administrative personnel. This problem should be overcome in the future as newer generations of hospital information systems better integrate medical and administrative functions.

5. Results

We analyzed the six weeks of hourly census per unit for the entire hospital. Specifically, we computed the average hourly census weekly curve as shown in Fig. 1 for each medical and surgical unit, and then applied the time series model, but only for the obstetric services unit. We now discuss our results at these two levels.

5.1. Hospital-wide results

Bed re-allocation was affected in several units. The average hourly census and standard deviation of these curves were enough for Northside Hospital administrators to move some specialties to units with fewer beds and other specialties to units with more beds. Since floor units come with an already defined number of beds, the reassignment of units only needs to be "relatively" accurate.

Not all suggestions made by the model could be implemented. In certain cases, our results showed that the number of beds should be modified, but no changes were made because these units had only recently been overhauled and the hospitals' administration did not wish to disrupt them too frequently. In other units, the number of beds was not cut since demand was expected to increase with the recruitment of new physicians. Nevertheless, the hourly census per day of the week gave administrators more precise information than previously.

A question of interest here is the effect on bed requirements of combining two specialties in the same unit, i.e. creating "pool" beds for two or more specialties. Pooling resources usually reduces variation in demand; however, because of weekly variance, pooling two services that have their peak censuses at exactly the same day and time will not reduce bed requirements as much as pooling two units that have their peak censuses at different days of the week and/or times. Our model can determine if pooling two units is able to reduce bed requirements. Data analysis showed that most units at Northside Hospital had the same census patterns, i.e. the censuses peaked on Thursday and Friday at mid-day, thus limiting the benefits of pooling resources.

5.2. Obstetric services results

Women's Services at Northside Hospital was considering the partitioning of a fixed number of beds among those needed by women before delivery (antepartum), during delivery, and after delivery (perinatal). The number of *birthing room* (delivery/perinatal) beds also needed to be considered. Although these different types of beds are destined for the same customers, the needs of the women differ at each stage of the delivery process. They are best served by being located in the "correct" unit at each stage. Because of state regulations, Northside Hospital cannot increase the total number of obstetric beds without state approval but can change the subspeciality allocation. Due to the recent trend of decreasing hospital lengths-of-stay after birth, Northside Hospital felt that a reallocation of beds could increase the delivery per year capacity of its very busy obstetric services. The time series model was thus developed to provide a good estimate of each subspecialty in an effort to support a new reallocation of beds.

Our time series model was applied to the perinatal unit data in order to define how many perinatal beds could be transferred to the antepartum units and delivery rooms. Women's Services agreed that perinatal bed needs should be satisfied at the $\alpha_w = 0.02$ level. Our model generated the curve given in Fig. 5 and showed that 67 beds were needed to satisfy the perinatal needs at the $\alpha_w = 0.02$ level, i.e. 22 beds above the average daily census of 45. Since the perinatal unit had more than 67 beds at the time, this confirmed that the additional beds could be re-allocated to the other subspecialties.

Women's Services' administrators were pleased by the results coming from the time series models and used these results to plan the new layout of the obstetric services area. After studying the results and graphs, administrators asked us to generate different census frequency curves with our simulation models, like the one shown in Fig. 6. These, they felt, would help them in scheduling staff and in sizing permanent and float team nursing staffs.

6. Conclusion

Hourly census data appear to provide sufficient information for decisions regarding bed allocations to different medical and surgical units within a hospital. We have added a time

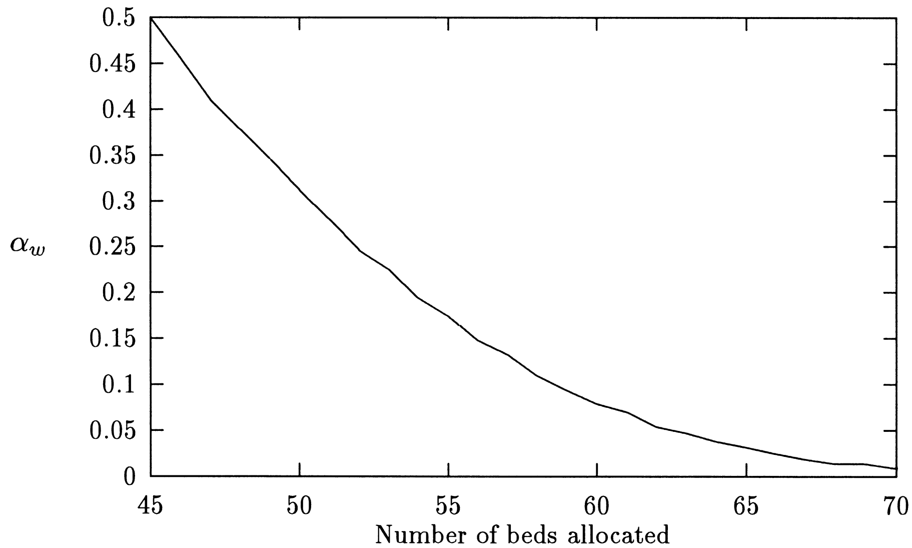

Fig. 5. Perinatal unit—beds allocated vs α_w.

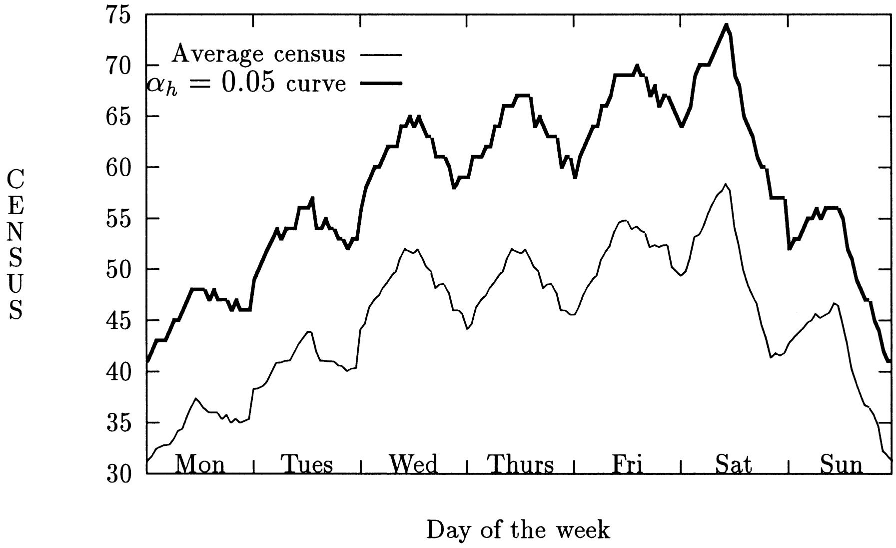

Fig. 6. Perinatal unit—census curve with $\alpha_h = 0.05$.

series AR(1) model to the average census curve in order to simulate the census. With such a model, we can predict the probability that the census will be higher than a given number of beds available. Health administrators can then use this model to allocate beds such that there is an acceptable level of overflow situations, and/or they can make their allocations such that medical and surgical services have similar occurrences of bed shortages.

While hospitals typically keep only midnight census data, our models require some work to retrieve the hourly census. A special data collection thus had to be organized in order to obtain six weeks of hourly census values. We validated and adjusted our data to reflect seasonality and census trends. Little's formula [6], a basic result from queueing theory, proved useful in validating these data.

Our model was applied to data from Northside Hospital, a 450-bed facility located in Atlanta, GA. Most of our efforts were directed towards helping the obstetric services to better partition their beds between units that focus on different stages of delivery. We also looked at the requirements of all other specialties in the hospital. The model's recommendations served as the basis for a new, hospital-wide reallocation of beds. Administrators had based their decisions primarily on bed requirements, but also had to consider physical constraints of the individual units as well as a specialty's place in the strategic growth of the hospital.

With a valid simulation model now in hand, we are in a position to investigate several of other interesting problems; these form the foundation for future research.

- **Selection among alternatives:** hospitals are often faced with the problem of choosing among a number of alternatives. For instance, how should beds be allocated among the subunits of the pediatrics unit? It is natural to ask which of these alternative allocations is the best, where "best" is defined according to some criterion of goodness (possibly determined by hospital administrators). Such problems fall into the realm of *ranking and selection*, for which a wide body of literature exists (cf. [1]). An open research area is that of applying ranking and selection procedures in the context of hospital-care problems.

- **Staff forecast:** to efficiently utilize nursing staff given the highly fluctuating demand that hospitals deal with, mechanisms such as float teams (nurses with a fixed schedule but a flexible unit assignment), temporary nurses, and overtime are utilized. Decisions relating to scheduling have to be made from within a few hours to a few days before the shift occurs. A good forecasting model can thus assist in the making of more effective staffing decisions earlier on in the process. In this regard, an interesting research activity would be to determine the accuracy of forecasts based on our proposed simulation model.

Acknowledgements

We acknowledge the precious collaboration of the Nursing Services from Northside Hospital. We would like to thank the Editor-in-Chief, Barnett Parker, and two anonymous referees for their helpful comments and suggestions.

References

[1] Bechhofer RE, Santner TJ, Goldsman D. Design and analysis of experiments for statistical selection, screening, and multiple comparisons. New York: Wiley, 1995.
[2] Dumas MB. Simulation modeling for hospital bed planning. Simulation 1984;43:69–78.
[3] Esogbue AO, Singh JA. A stochastic model for an optimal bed distribution in a hospital ward. Operations Research 1976;24:884–98.
[4] Fetter RB, Thompson JD. The simulation of hospital systems. Operations Research 1965;13:689–711.
[5] Fuller WA. Introduction to statistical time series. New York: Wiley, 1976.
[6] Gross D, Harris CM. Fundamentals of queueing theory. 2nd ed. ed. New York: Wiley, 1985.
[7] Hancock WM, Martin JB, Storer RH. Simulation-based occupancy recommendations for adult medical/surgical units using admissions scheduling systems. Inquiry 1978;15:25–32.
[8] Harris RA. Hospital bed requirements planning. European Journal of Operational Research 1986;25:121–6.
[9] Hershey JC, Weiss EN, Cohen MA. A stochastic service network model with application to hospital facilities. Operations Research 1981;29:1–22.
[10] Kao EPC, Tung GC. Bed allocation in a public health delivery system. Management Science 1981;27:507–20.
[11] Law AM, Kelton WD. Simulation modeling and analysis. 2nd ed. New York: McGraw-Hill, 1991.
[12] Lowery JC, Martin JB. Design and validation of a critical care simulation model. Journal of the Society for Health Systems, 1992;3:15–36.
[13] Makridakis SG, Wheelwright SC. Forecasting, methods and applications. 5th ed. ed. New York: Wiley, 1989.
[14] McFadden KL. Hospital policy changes in obstetric patient movement. International Journal of Operations and Production Management 1996;16:28–41.
[15] Romanin-Jacur G, Facchin P. Optimal planning for a pediatric semi-intensive care unit via simulation. European Journal of Operational Research 1987;29:192–8.
[16] Vassilacopoulos G. A simulation model for bed allocation to hospital inppatient departments. Simulation 1985;45:233–41.