A simulation model of a diagnostic radiology department

P. Ciaran O'KANE
Department of Business Studies, Queen's University, Belfast

Received September 1979
Revised October 1979

Unlike most of the component parts of a general hospital, which are designed to cater for patients with particular kinds of illnesses, the services of diagnostic radiology departments are utilised by almost every category of patient which enters the hospital system. Hence, efficient utilisation of X-ray facilities is a necessary condition for overall hospital efficiency. Because of the diversity of inputs and range of services provided a radiology department represents a very complex system. In order to assess the effect of various proposed operating policies, using a number of criteria, a simulation model (SIMRAD) of such a system has been constructed. The model has been tested in one particular hospital but it is believed that SIMRAD is sufficiently general to allow it to be utilised in other environments with the minimum of adaptation. This paper concentrates on a description of the system factors identified, the model structure, the input required and the output provided. The detail provided should allow the model to be applied by other workers in this field. A selection of experiments that have been performed is mentioned but detailed results will be given in another paper.

1. Introduction

Any hospital may be visualised as a macro-system comprising a large number of sub-systems which are quite complex both within themselves and in their interactions with each other. The overall performance of this macrosystem, and the impact of any suggested changes in its operation, can be measured using various criteria. However, since the main purpose of a hospital is the care of the sick a major consideration in any such evaluation must always be the effect on patients. Different parts of the system have greater impact on certain categories of patients than others; for example the number, type and efficiency of operating theatres will be of minimal concern to medical patients; similarly the care received by geriatric patients will not be a major consideration of maternity patients and so on. At the same time a few of the constituent sub-systems affect a wide cross-section of patients and among these must be counted such sub-systems as the pathology and the diagnostic radiography departments.

This paper gives details of a simulation model of a diagnostic X-ray department in a general hospital. While a considerable literature is now available on the development of OR models of various sections of the hospital system [1,2,3,4,5] relatively little has been published about details of models relating to radiology. Jeans et al. [6] have written about one such model while Fraser [7], Wilkinson [8] and Lev et al. [9] have also developed models of such departments. The model described in this paper is designed to build on this work to provide a planning tool which potential users will accept as a realistic simulation of the real-world system. It must be recognised that in a multi-organisation such as a hospital the success or failure of any proposal is judged by different people using different criteria, not all of which can be quantified. This model attempts to present to decision makers the consequences of any proposed course of action on a number of measurable factors. They are then at liberty to judge the efficiency of the proposal on the basis of those model results that are of immediate interest combined with their own judgement of what the impact will be on the important, non-quantifiable factors.

An attempt has also been made to allow for the effects of interactions between the modelled system and the rest of the hospital and also to include, to a limited degree, some of the human and personal factors which are involved in the administration of such a department. Although the model described here is based on the study of a particular department it has been written so that it can be applied to other hospital X-ray departments with the minimum of alteration.

The author would like to thank the radiologists, radiographers and other staff in the department in which this study was carried out for the help they provided so willingly to him during this project.

European Journal of Operational Research 6 (1981) 38–45

2. Description of system modelled

The radiography department in any general hospital is essentially a waiting line system with patients forming the inputs and one or more X-ray examinations being the service provided. The service facilities consist of examination rooms manned by trained radiographers. However, even a cursory examination of the operation of such a department demonstrates the existence of a number of factors which would invalidate any attempt to represent the system by a simple analytical model.

The inputs do not form a single stream, as patients come from a number of unrelated sources and thus generate a number of parallel and independent input streams each with its own characteristics and subject to different levels of control. For instance, the arrival of casualty patients which, because of hospital policy, is not subject to any control demonstrates a random, but quite stable, pattern. At the other end of the spectrum are patients with appointments whose rate of arrival is rigidly controlled. Between these extremes are patients from outpatient clinics and from the wards of the hospital itself. Arrivals of the former group are basically random but occur only within certain periods dictated by clinic time-tables, while the rate of arrival of the latter group can be controlled, to a limited extent, by the radiology department itself but still displays a high degree of randomness.

The main service required by arrivals is one or more X-ray examinations. In addition, some arrivals will require a subsidiary service facility in the form of changing cubicles. All patients enter a single waiting line. In theory the queue discipline is first-in–first-out but this order is more frequently broken than adhered to. Before a service can commence an examination room (and a cubicle if required) together with a radiographer must be available. Typically a trained radiographer can perform any examination. However, examination rooms are not all equiped in the same way so that only a subset of the rooms is suitable for any given patient's requirements. While the model can assign the first radiographer who is free to the patient at the head of the queue, this patient cannot be examined unless one of the rooms which is available suits his examination requirements. Therefore, a particular patient may be served out of turn if a room suitable for his requirements but not for those ahead of him in the waiting line, becomes available. The suitability of the room is a function of the examination required which in turn is a function of the source of the patient. It has also been observed in practice that a radiographer, if faced with a choice of suitable rooms, will have a 'favourite' room. A realistic model must take cognizance of this fact.

Another factor often ignored in models of this type is that the time at which a radiographer or a room become available does not necessarily, coincide with the end of an examination. When an examination is completed the patient can be dismissed from the system but frequently certain work has to be performed in the room before another patient can use it. Similarly, each radiographer has certain duties to perform in relation to the last patient she has examined such as checking films, dealing with records, etc., before she is free to deal with another patient. Only when these operations have been completed is the room and the radiographer free to accept another patient.

To enhance the realism of the model, provision has also to be made for radiographers' rest periods such as lunch and coffee breaks. It was evident, given all of these complications, that Monte-Carlo simulation was the most effective way of representing this system.

3. Input to the model

The model described below was developed after an extensive investigation of one particular department which appeared to be reasonably representative of many hospital X-ray departments. The detailed results of this investigation have been reported elsewhere [10,11] but the basic conclusion was that the following input data were the minimum required by the model if an acceptable level of realism was to be achieved:

(a) a description of the arrival pattern of patients,

(b) the distribution of examination requirements,

(c) the distribution of durations of each type of examination,

(d) the availability and suitability of examination rooms,

(e) the number of radiographers and a timetable giving their availability.

The values these data take vary from department to department. Hence before the model can be applied in any given situation an empirical investigation of the operations of that particular department will be necessary.

It was found, for instance, in the department in which the model was tested, that patients could be divided into eight main input streams or sources. These consisted of a stream for casualty patients, one for in-patients, another for patients with appointments, and five streams for patients from different categories of out-patient clinics. There were long periods, and even days, when no patients arrived from some of these sources. Thus the approach used was to identify to the computer the times between which each source could be considered to be open. For example, the casualty source was considered to be continuously open while clinic timetables dictated when out-patient sources were open. Even within these periods the rates of arrivals showed considerable variation. The most suitable method of describing the rate of arrival was to divide the periods during which any source was in operation into a number of 'sessions' so that a reasonably homogeneous group of arrivals was included in each session. It was then found possible to describe the distribution of arrivals by the negative exponential distribution, a different mean being used for each session and each source. In this way the model can be easily informed of changes in other parts of the hospital system.

The timetable of appointments is known in advance and can be input directly to the model but an allowance must be made for the unpunctuality of patients. This is supplied through an empirical distribution of the deviation of actual arrival times from the time of appointment. In the case of the particular department studied appointments were given only for certain types of examinations. Therefore, for appointees, the examinations required were pre-determined.

The probability of a patient from any other source requiring a given type of examination is a function of the source from which he comes. Definitions of the different types of examinations were drawn up in consulation with the staff of the department. It was found that 22 categories were sufficient to describe the complete work load of the unit.

The type of examination required by a patient from any given source is a random variable, the associated probabilities being estimated from an empirical distribution. Allowances are also made for the possibility that a patient will require more than one examination or a repeat of one or more examinations, the probabilities being derived from separate empirical distributions for each source. Further empirical investigations provide probability distributions for the duration of each type of examination which are also input to the model.

Information must also be provided as to which rooms are suitable for each type of examination and some device is necessary to reflect the 'popularity rating' of different rooms. For each examination type an array is read into the model which gives the numbers of the suitable rooms in order of popularity.

The total number of radiographers available is an exogenous variable and for each radiographer information is supplied about the times at which she is scheduled to have coffee breaks, lunch breaks, etc., and the duration of these breaks.

4. Structure of the model

The model, called SIMRAD, was written in FORTRAN IV and run on an ICL 1906S machine. Incremental time advance was used with the clock progressing in integer minutes from event to event. An event has been defined in the usual manner as being any set of conditions which may change the current state of any status variables. There are six possible events for each of which the relevant operations on the status variables have been defined:

(i) arrival of a patient,
(ii) the departure of a patient,
(iii) a change of session occurs within the period any given source is in operation or a new source comes into operation,
(iv) an examination room becomes available,
(v) a radiographer becomes available,
(vi) a radiographer's rest period commences.

An outline flow chart of the model is given in Fig. 1. The model, as constructed, allows any number of runs, each of any given length, to be made at a single call on the computer. Five days per week are simulated and the length of the working day can be varied at will. A summary of the operation of SIMRAD is given below.

Having read the input data described above and initialised status variables SIMRAD establishes its own initial conditions. No examinations commence before the set starting time but there may be patients waiting at that time whose examinations can commence as soon as the department comes into operation. In this particular department these patients tended to be those with appointments who had arrived early. The model extracts from the supplied file all appointments due on the current day, generates the actual time of arrival for each patient from the given punctuality distribution and places these in a special queue (APPQ).

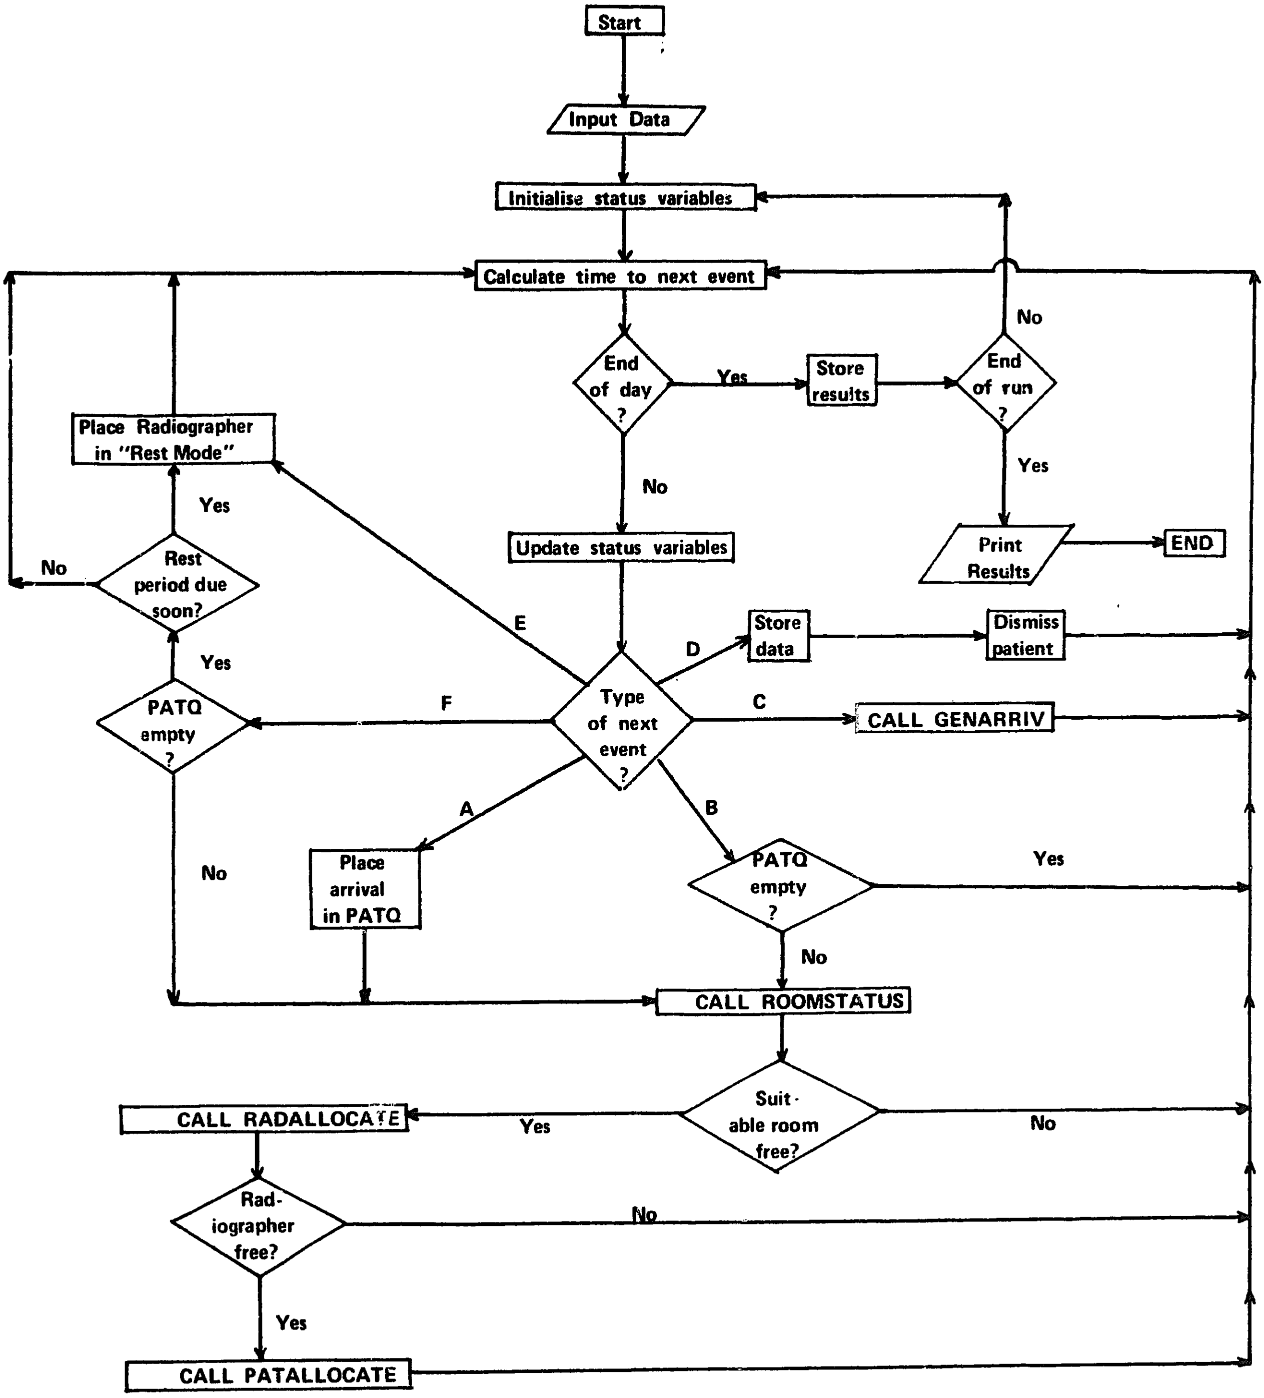

CODES FOR TYPES OF EVENT

A = Arrival of a patient:
B = Room becomes free:
C = End of session or
New source comes into operation:
D = End of examination:
E = Radiographers rest period due:
F = Radiographer becomes free.

Fig. 1. Outline of flow chart of SIMRAD.

This queue is scanned for any patients due to arrive before the opening time for that day and any such found are transferred to the active waiting line (PATQ), to await the opening of the system. This establishes the initial conditions. When the model opens the system any patients in PATQ are dealt with immediately. If PATQ is empty the system remains idle until the first arrival occurs.

The central control mechanism of the model is a routine called GENARRIV. It performs the following operations:

(i) opens and closes sources and changes sessions at the appropriate times as indicated by the input data,

(ii) records current status of each source.

(iii) records the time to next arrival from each source currently in operation,

(iv) generates new arrivals when appropriate.

An array, ALLSCEINT, records the time until the next arrival from each source; a value of zero in ALLSCEINT (K) indicates that a patient has just arrived from source K; a value of 1 000 indicates that no arrival was due from source K at the last call on GENARRIV while any other value gives the time remaining until the next arrival from source K, this having been generated previously. On entering GENARRIV the model scans ALLSCEINT until a value of 0 or 1 000 is found. The current status of the corresponding source is checked; if it is operational the current session is determined and the next arrival is generated using the parameters relevant to this source and session. If the generated arrival time is later than the scheduled time for a change of session the arrival is not permitted to take place and the source is flagged as being inoperative until the beginning of the next session. Otherwise the time until the arrival is entered in ALLSCEINT (K). When all of the cells of this array have been scanned control is returned to SIMRAD which will then have the following information available to it:

(i) the time until the next patient is due from all sources currently in operation,

(ii) the time when all sources not currently open will become operational.

SIMRAD orders patients in a special queue according to their time of arrival and also records the source from which each patient is due. The smallest element in this queue will be the next arrival unless a patient with an appointment, as indicated by APPQ, is due prior to this in which case his arrival takes precedence. The model then determines the time until the next occurrence of each of the other types of events, increments the time clock to the earliest of these and takes the action appropriate to the type of event.

When the next event is an arrival (event A, Fig. 1) the type, or types of examinations is generated from the distribution appropriate to the source of the patient. The examination times required are then generated from the distributions relevant to these types, allowance being made for the possibility of a repeat of any examination. The patient is placed in PATQ which has facilities for recording various items of data relating to him.

The model calls on a routine ROOMSTATUS which investigates the availability of a suitable room for a given patient. From the input data the room or rooms suitable for each type of examination together with the 'popularity rating' of these rooms is known. This routine checks the current availability of each room suitable for the first examination required in decreasing order of popularity. If no suitable room is free the patient remains in the waiting line. If a room is found the data is interrogated to find if a second examination is required; if so the routine checks if the previously found room is suitable for this second examination also. If this search is successful the number of the room is recorded and control is returned to the main programme. If it is not the routine resumes its check of other, less popular rooms suitable for the first examination and the process is continued until either a room suitable for both examination types is found or it is proved that no such room is available at that time. On leaving this routine the model knows the number, if any, of the most popular room which is both suitable for the current patient and immediately free.

Another routine, RADALLOCATE, is then called. This controls the allocation of radiographers to patients and rooms and ensures that staff are given scheduled rest periods. The model contains an array which holds data on each radiographer including the scheduled time and duration of her rest periods and her current status. By scanning this array the routine determines if any radiographer is free, and if so, which one. Before allocating this radiographer to a patient the position regarding her rest periods is checked. In general an attempt is made to ensure that each member of staff is allowed to take her rest period as close to the scheduled time as possible but there is a degree of flexibility available. Thus, if a radiographer is free and a rest period is due she is allowed to take it and hence is unavailable for the duration of that period. She may, however, be assigned to a patient even though this may cause the start of her rest period to be delayed. The only exception to this rule is where her allocation to a particular examination would cause an unacceptably excessive delay in the commencement of the rest period.

If the call to RADALLOCATE shows that a radiographer is free then she is allocated to carry out the required examination in the room already found to be suitable and free. This process is performed by a subroutine PATALLOCATE which also transfers the relevant patient from PATQ to the service queue (PATSERV). This routine also records the time at which both radiographer and room will become free allowing not only for the necessary examination time but also for the ancillary operations that are necessary after the patient has been examined.

When a room that has been engaged becomes free (event B, Fig. 1) the model changes the status of the room in question. If any patients are waiting in PATQ a call on ROOMSTATUS indicates if this room is suitable for any of these patients. If successful this call is followed by a call on RADALLOCATE followed by one on PATALLOCATE. A negative result at any stage in this process allows the model to proceed directly to the next event.

When a radiographer becomes free (event F, Fig. 1) a broadly similar process to the above is followed. In order to simulate reality more closely a radiographer is allowed to take her rest period early if no patients are waiting provided she becomes free within a certain specified period prior to her scheduled rest period. If her rest period is actually due or overdue (event E, Fig. 1) she is immediately placed in 'rest mode' and is flagged as unavailable until the scheduled end of that rest period. Otherwise, if any patients are waiting calls are made to ROOMSTATUS, RADALLOCATE and PATALLOCATE as above.

At the end of a session or at the opening of a new source (event C, Fig. 1) a call is made on GENARRIV to generate the first arrivals from that source or within that session.

When a service ends (event D, Fig. 1) the data relating to the relevant patient is accumulated, the patient is dismissed from the system and the room and radiographer are flagged as having completed the examination but still involved in the various ancillary tasks that are necessary.

5. Output of the model

The standard output from SIMRAD gives the following information for each run:

(a) mean and standard deviation of numbers of patients by source, day of week and week,

(b) maximum and minimum numbers of patients by day and by week.

(c) mean and standard deviation of examination times and waiting times by source, day of week and week,

(d) utilisation of examination rooms by individual room, day of week and overall,

(e) utilisation of staff by individual staff member, day of week and overall,

(f) number, mean and standard deviation of duration of each type of examination for single examinations with no repeats and with one repeat,

(g) maximum length of waiting line by day of week,

(h) mean length of waiting line by day of week, hour of day and overall,

(i) frequency distribution of queue length facing arrivals.

Options exist in the model which allow for more detailed output data if required. For instance, the model if requested, will provide a summary of results at the end of each day which include the number of the day, the time that the last examination on that day was completed, the total number of patients examined on that day, the maximum length of queue which developed during the day and the time at which it occurred.

A further option can be called which gives complete details of the value of all status variables at each time increment after a specified length of time has been simulated.

6. Validation of the model

The question as to how well this model represents reality cannot be fully answered at this stage. The ultimate test of any simulation model is, of course, the degree of accuracy with which it predicts the behaviour of the actual system. The ideal test would, therefore be to compare the actual effects of some change in the system to those forecast by the model but this test has not been possible to date. It would not have been valid to compare the results of the model with the data collected during the main survey since much of the survey data was used to derive a number of the exogenous variables fed into the model. It was possible, however, to get some idea of the usefulness of SIMRAD by comparing the results with those obtained in a pilot study which had been carried out over a four week period prior to the main survey. The results of the analysis of this data have not been included in the model and therefore they provide an independent measure of the performance of the system. The main drawback is that the information collected in the pilot study was less detailed than that in the main survey and therefore the accuracy of the more detailed predictions of the model could not be checked. Table 1 compares the results for total patient numbers. The simulated figures are from a run of 125 days. This table shows that in only one case (Mondays) are the estimated arrivals significantly different from the actual mean arrivals.

The comparison of the percentage of total patients from each of the four main sources given by the pilot study and by the model are shown in Table 2. The agreement is quite close except that the model underestimates the number of casualty patients. This may be explained by the fact that on one particular Monday during the pilot study almost three times more casualty patients attended than was usual and that this figure was 80% greater than the maximum observed in the detailed survey. While this may explain the variations in Table 2, it should also mean that the average figures in the pilot survey on a Monday should be greater than the forecast but Table 1 shows the opposite to be true. One possible explanation of this discrepancy is that the number of outpatients attending on Mondays during the pilot study was consistently less than during the main survey indicating that at least one outpatient clinic was not in operation during that period. The data was not sufficiently detailed to test this hypothesis further.

The overall conclusion was that while the model was not a perfect predictor it did replicate reality reasonably well and was sufficiently accurate to be used for experimental purposes.

Table 1
Comparison of simulated and actual numbers of patients

Day of week	Simulated values [a]		Actual [b]	$\frac{\overline{x}_1 - \overline{x}_2}{\sigma_1}$
	Mean ($\overline{x}_1$)	Std. Dev. (σ_1)	Mean ($\overline{x}_2$)	
Monday	121.5	8.71	103.3	+2.09
Tuesday	77.2	7.94	75.8	+0.18
Wednesday	87.9	9.27	97.5	−1.04
Thursday	94.0	9.69	100.5	−0.67
Friday	79.4	8.85	80.3	−0.10
Total	460.0	20.65	457.4	+0.13

[a] Based on 25 weeks.
[b] Based on 4 weeks.

Table 2
Comparison of simulated and actual percentages of patients from main sources

Source	Simulated %	Actual %
Outpatients	38.6	36.4
Inpatients	30.8	30.2
Casualty	19.9	24.0
Appointees	10.7	9.4

7. Discussion

This paper describes an attempt to model the operations of an important element of a hospital system. Efforts were made to produce a model which would be sufficiently realistic to be acceptable to the decision-makers and flexible enough to allow the impact of a variety of organisational changes to be assessed without being so detailed that excessive programming or running time would be required. A number of assumptions were made which simplified the structure of the model but were not thought to distort the realism of the results.

The most obvious of these were the exclusion of changing cubicles as explicit entities in the model and the assumption that, if multiple examinations were required, these would be limited to two, both of which would be performed on a single visit to an examination room. With respect to the first assumption, investigations in the hospital studied showed that more than two-thirds of patients did not use cubicles. It was also clear, in this department at least, that sufficient cubicle space was provided to allow a patient who required this facility to enter a cubicle while the previous patient still occupied the examination room. Hence time spent in a dressing cubicle is implicitly included in patient waiting time except in those cases where a patient arrives and a room is available immediately. The model admits such arrivals to the examination room at once but it was felt that the number of such patients was not large enough to justify the extra programming necessary to deal with such cases more accurately.

In certain cases when a patient requires more than one examination this will result in more than one visit to an examination room. Obviously this will require more time than is generated by the models which assumes all examinations are performed during a single visit. Again since only some 13% of patients had more than one examination and probably only a proportion of these made multiple visits it is not felt that the exclusion of this possibility from the model introduced any great distortion. It was also noted that only some 2% of all patients had more than two types of examination.

A further simplification incorporated in the model was to dismiss a patient from the system as soon as he vacates the examination room rather than to allow for the possibility that he may be asked to remain in the waiting area for a short period afterwards. The main distortion produced by this is that the total time that a simulated patient remains in the system may be less than that of a real patient. However, given that allowance has been made in the model for the possibility of repeat examinations and for the time that a radiographer spends checking films it was not felt that the complications that would be introduced by explicitly simulating this waiting period were justified by the resulting improvement in realism. Patients waiting after examination have minimal effect on the operations of the department.

This version of SIMRAD, written in FORTRAN IV, was run on an ICL 1906S machine. A simulation of 125 days, during which time some 11 500 patients were processed, required a mill time of 150 and occupied 28 *K* of core. This is considered to be quite satisfactory but recent experiments with a version of the model written in the simulation language SIMONE indicate that the above mill time requirements can be reduced quite considerably [12].

The model has been used for a number of experimental applications. These have included:

(a) an investigation of the effect on patient waiting time of changing the number of radiographers available.

(b) the determination of the 'best' number of radiographers;

(c) an investigation of the effect of introducing separate X-ray facilities for certain departments of the hospital,

(d) an investigation of the pattern of room usage on patient waiting time,

(e) an analysis of the effect of increases and decreases in the number of patients using the X-ray department,

(f) an analysis of the effect of changes in appointment timetables.

A detailed description of these, and other, experiments is currently being prepared for publication. Experience to data shows that a wide variety of types of experiments can be performed with only minimal alterations to the model. There has not been an opportunity to date of testing the author's belief that SIMRAD can be very easily adapted to model other radiology departments. The major problem in such an exercise lies in the generation of the necessary input data. Almost by definition these are peculiar to each individual radiology department and hence it is not considered advisable to attempt to derive common input parameters for different departments.

The experience of this project suggests that the effort spent in the collection of data constitutes an excellent investment, since the work results in a greatly increased understanding of a complex system.

References

[1] N.T.J. Bailey, A study of queues and appointment systems in hospital outpatient departments with special reference to waiting times, J. Roy. Statist. Soc. Ser. B 14 (1952).

[2] M.A. Feldstein, Improving the use of maternity beds, Operational Res. Quart. 16 (1965) 65–76.

[3] R.B. Fetter and J.D. Thompson, The simulation of hospital systems, Operations Res. 12 (1965) 689–711.

[4] F. Blewett, D.M. Grove, A. Massinas, J.M. Norman and K.H. Southern, Computer simulation model for multi-speciality wards, Operational Res. Quart. 23 (1972) 139.

[5] E.J. Rising, R. Baron and R. Averill, A systems analysis of a university health service out-patient clinic, Operational Res. 21 (1973) 1030–1047.

[6] W.D. Jeans, S.R. Berget and R. Gill, Computer simulation model of an X-ray department, B.M.J.1 (1972) 675–679.

[7] B.J. Fraser, The organisation of a radiology department, Unpublished Ph.D. Thesis, University of Reading (1969).

[8] E.W. Wilkinson, A simulation model of diagnostic radiology, Unpublished M.Sc. Thesis, University of Missouri (1966).

[9] B. Lev, B.S. Caltagirone and F.J. Shea, Patient flow and utilization of resources in a diagnostic radiology department, Investigative Radiology, (1972) 517–525.

[10] P.C. O'Kane, A statistical analysis of the operations of a hospital radiology department, Occasional paper, Department of Business Studies, Queen's University, Belfast (1974) 1–56.

[11] P.C. O'Kane and A.D. Cough, An evaluation of patient demand in a hospital radiology department, Hospital and Health Services Review (1978) 226–229.

[12] R.H. Perrot, A.K. Raja and P.C. O'Kane, A simulation experiment using two languages, to appear in Comput. J.