

A GENERALIZED SIMULATION MODEL OF AN INTEGRATED EMERGENCY POST

Martijn Mes

Manon Bruens

University of Twente
Drienerlolaan 5

7522 NB Enschede, THE NETHERLANDS

Ziekenhuisgroep Twente
Zilvermew 1

7609 PP Almelo, THE NETHERLANDS

ABSTRACT

This paper discusses the development of a discrete-event simulation model for an integrated emergency post. This post is a collaboration between a general practitioners post and an emergency department within a hospital. We present a generalized and flexible simulation model, which can easily be adapted to several emergency departments as well as to other departments within the hospital, as we demonstrate with our application to the integrated emergency post. Here, generalization relates to the way we model patient flow, patient prioritization, resource allocation, and process handling. After presenting the modeling approach, we shortly describe the implemented and validated model of the integrated emergency post, and describe how it is currently being used by health care managers to analyze the effects of organizational interventions.

1 INTRODUCTION

We present a generic and flexible framework for discrete-event simulation (DES) modeling of emergency departments. Our work is motivated by a case study at an integrated emergency post (IEP) at Ziekenhuisgroep Twente, Almelo, the Netherlands. Here, the emergency department (ED) of the hospital collaborates with a general practitioners post (GP post) with the aim of (i) providing a single point of contact for patients with acute need of care and (ii) avoiding overcrowding of the ED by treating a part of the self-referrals at the GP post. Despite the reported potential benefits (see, e.g., Kool, Homberg, and Kamphuis 2008), there is no quantitative proof that the use of an IEP actually leads to such results. Therefore, we started a research project with the goal to form a scientific foundation on how to reach an optimal process design for the IEP, in which the right patient is provided with the right care by the right caregiver, with optimal use of resources and without unnecessary delays. As part of this research project, we here focus on the development of a flexible simulation model for the IEP.

Realistic modeling of an ED, and specifically an IEP, is a difficult and time-consuming task. Hence, the development of a flexible and reusable simulation model is desirable. The objective of this study is (i) to come up with a flexible modeling framework that can be applied to other EDs and (ii) the development of a validated simulation model for the IEP in Almelo. The simulation model will be used to test several organizational interventions for the IEP in Almelo, to provide insight into the benefits of introducing an IEP at various hospitals within the Netherlands, and to increase the understanding by health care managers of the dynamics and process changes involved in the introduction of an IEP.

Discrete-event simulation has proven itself as an effective tool to support decision making in health care to improve the allocation of resources and system performance. Comprehensive literature reviews on the use of simulation in health care are given by Jacobson, Hall, and Swisher (2006) and Gunal and Pidd (2010). For a vast majority of the simulation studies on EDs, the presentation is limited to a specific hospital and the results are often hard to generalize to other EDs, let alone, to other departments within the hospital. We mention a few exceptions to this statement. Ferrin, Miller, and McBroom (2007) use a semi-reusable product, called EDsim, to evaluate organizational interventions at a hospital in Arizona. Also aiming at reusability, Sinreich and Marmor (2004) present a simulation model that is flexible and intuitive

to use and that contains default settings that form a good starting point for modeling many EDs. Fletcher et al. (2007) report the use of a generic simulation model that has been used at various EDs within the UK. Facchin, Rizzato, and Romanin-Jacur (2010) present a flexible model of an ED based on common structural and functional characteristics of emergency rooms.

As stated by Hay, Valentin, and Bijlsma (2006), many simulation studies have taken a factory view of the ED, where patients come in like orders queued on a factory floor, often with fixed priority, and drive the process by requesting resources. As an alternative, they present a different way of modeling and studying ED processes. The study argues that with only a factory view of the world, low acuity patients will continually be starved from receiving services and many will not receive treatment at all. This issue has also been reported by Raunak et al. (2009), who describe a simulation architecture with an emphasis on modeling resources and processes. Using some early comparisons, they state that their framework offers modeling flexibility which is likely to be of value within the medical community.

The contribution of our work is the description of a generalized flexible simulation model. Here, generalization is not being sought in the similarities between different EDs, but merely on the way we could model the events in patient flow through an ED; specifically, related to patient prioritization, resource allocation, and process handling. The primary application of this model is an ED, but it can easily be extended to surrounding departments as we demonstrate with our application of the IEP.

The remainder of this paper is structured as follows. In Section 2, we provide a brief description of our case study at the IEP. In Section 3, we present our conceptual model, where we describe the patients, resources, processes, process dependencies, prioritization rules, and the structure of our simulation model. In Section 4, we describe the implemented simulation model. We close with conclusions in Section 5.

2 CASE DESCRIPTION

Outside regular office hours, people with an acute need for care are often confused where to go. They assess the seriousness of their demand of care and decide whether to go to a GP post or to the ED. At the ED these patients are called 'self-referrals'. Many of the self-referrals could have been helped by a GP, but since they already are at the hospital, they receive treatment at the ED. Capacity and cost restrictions create a growing need to bring down the number of self-referrals at the emergency department (Kool, Homberg, and Kamphuis 2008). From this need, the idea was created for a gatekeeper function: health care professionals who decide for the patient where to go to. The IEP is a practical example of such a gatekeepers function. This research takes place at the IEP located at Ziekenhuisgroep Twente (ZGT), in Almelo, the Netherlands. The IEP, which has been introduced at this hospital in 2010, consists of a collaboration between the ED of this hospital and the GP post of Almelo (CHPA) that was built directly next to the ED.

With the introduction of a IEP, patients can either come to the IEP as self-referral or call the IEP. During the call, a GP assistant assesses the situation and decides what the best course of action is for the patient: giving telephonic advice, sending a GP or ambulance to the patient, or asking the patient to come to the IEP for a consultation. Self-referrals are triaged by a GP assistant. When a patient needs more care, the GP assistant books the patient for a consult at the GP post, or directly refers the patient to the ED. This means that no self-referrals arrive at the ED. The goal of integrating the ED with a GP post is to reach higher efficiency of care, because patients can be helped at the right place by the right care giver. Higher efficiency is also expected because of reduction in overlap between the two departments. Furthermore, the use of one registration desk for acute care provides clarity for patients.

The goal of the project is to form a scientific foundation on how to reach an optimal process design for the IEP. An important prerequisite to reach this goal is the development of a flexible and valid simulation model for the IEP, which we describe in this paper.

3 SIMULATION MODELING

In this section, we present a flexible simulation model. Generally speaking, we distinguish three main components in discrete-event simulation: entities, resources, and processes. Entities are the moving parts that require service and resources act on these entities. Within the health care domain, the entities resemble the patients and resources include medical staff, operating rooms, hospital beds, medical equipment, etc. Both patients and resources are characterized by certain attributes, like urgency, age, gender, availability, and capacity. The processes are the services required by the entities, i.e., the treatments required by patients. Typically, patients require a sequence of treatments, which we denote by a care pathway. A treatment may require multiple resources, e.g., a treatment room, a nurse, and a medical specialist. After completion of a treatment, these resources may become available again, which could result in the simultaneous start of treatments for multiple patients.

Besides the basic components, patients, resources and processes, we also formulate (i) the dependencies between patients, resources, and processes, and (ii) the prioritization and appointment scheduling rules related to the processes. These concepts, together with realistic input data and the physical layout of the system, enable an accurate representation of reality. In the next subsections we present each of the components and illustrate them using our case study at the IEP.

3.1 Patients

Modeling patients involves modeling the arrival processes, patient attributes, and care pathways. Attributes of patients may include age, gender, diagnosis related group (DRG), triage urgencies, etc. The care pathway is treated in Section 3.3. Here we focus on the arrival process.

Arrivals are generally time dependent, which can be described by a stochastic arrival process with arrival rate $\lambda_{h,d,w}$, being the average number of patients arriving in hour h on the day of the week d in a week of the year w . For our problem, we observed the following. First, the relative difference in arrivals between days in a week does not differ significantly between the weeks. However, the absolute number of arrivals per week does differ significantly. Second, the relative difference in arrivals per hour does not differ significantly between the week days (Monday till Friday) but the absolute number of arrivals per week day differ (e.g., Friday is busier than the other week days). In addition, the arrival pattern on Saturday as well as on Sunday is significantly different. To cope with this, we decided to use the following procedure for calculating the arrival rates: $\lambda_{h,d,w} = \alpha_{h,d} \beta_d \gamma_w$, where $\alpha_{h,d}$ is the arrival rate of patients in hour h on day d , with $d = 1, \dots, 7$ (starting at Monday), but where $\alpha_{h,d} = \alpha_{h,d'}$ for $d, d' \in \{1, \dots, 5\}$; β_d is a day factor, which represents the fluctuation in the arrivals between week days, with $d \in \{1, \dots, 5\}$; and γ_w is a week factor, which represents seasonal effects, with $w \in \{1, \dots, 52\}$. The average of β_d over all week days and the average of γ_w over all weeks both equal one. This means that, on a particular day or week, they represent the deviation in expected number of arrivals on this day or week compared to the average number of arrivals on this day or week.

We now proceed as follows. At the beginning of each week w , we draw a week factor γ_w from a given distribution. Similarly, at the beginning of day d , we draw a day factor β_d . Each hour, the historic average number of patients in that hour on that day, given by $\alpha_{h,d'}$ is multiplied with the drawn factors, to get the parameter $\lambda_{h,d,w}$. We then generate Poisson arrivals with rate $\lambda_{h,d,w} = \alpha_{h,d} \beta_d \gamma_w$ using the thinning procedure (Law 2007). The assumption of Poisson arrivals seems to be valid on a given day. However, the resulting variance on a daily and weekly basis is not correct. To cope with this, we tune the variance of the distributions for β_d and γ_w in such a way that the daily and weekly variances match reality. The daily and monthly factors are distributed according to the normal and lognormal distributions.

3.2 Resources

We distinguish between three types of resources: staff, rooms, and equipment. For our model of the IEP, we separate staff into regular staff and external staff. The regular staff is included in the shift schedule. The

external staff is the staff that has to be called in, either coming from within the hospital or from externally. For the external staff members, there may be a delay in their arrival (see also the ‘delay tasks’ mentioned in Section 3.3) and there may be a delay upon exiting (within this exiting delay time, no additional delay would be required if we need this staff member for another task). We include the following staff types in our model: GP, GP assistant, triage assistant, ED nurse, medical resident, surgical resident, medical specialist, physician assistant, and diagnostic nurse. We consider the medical specialists and diagnostic nurses to be external staff since they are called to the ED when needed. For the regular staff members, we have to determine how many staff members of each type are present in each of the time intervals. We use input rosters that describe for each hour, how many staff members start and end their work shift.

Several rooms are taken into account in the model. Rooms directly related to the processing steps are two triage rooms, six GP rooms, eight ED rooms, two X-ray rooms, one CT room, and two plaster rooms. Other areas that are taken into account, mainly for animation purposes, are the call center, patient waiting areas at the GP post and the ED, and ‘waiting areas’ for staff members at the GP post and the ED. Most of the tasks take place in a particular room after which the patient goes to the waiting area. Except for treatment of a patient at the ED, where the ED room stays assigned to the patient until it leaves the IEP.

Regarding the medical equipment, we include portable ECG and ultrasound equipment, which means that tests performed with this equipment take place at the ED room.

3.3 Processes

We distinguish between three types of processes: regular tasks, parallel tasks, and delay tasks. A regular task is a treatment required by a patient. For most of the regular tasks, the patient and a staff member come together in a room. For a parallel task, the patient does not need to be present. As a result, the patient may undergo another task while a parallel task is in progress. An example of a parallel task is the reviewing of test results from an X-ray, ultrasound, or CT scan, for which the presence of the patient is not necessary. A delay task is a task that precedes a regular task and causes the patient to wait before the next treatment can be started. Examples of delay tasks are the required travel time of a staff member (e.g., a GP making a home visit), the cleaning of a room or medical equipment, etc.

The simulation process is driven by the regular tasks required by the patients. Note that the care pathway of patients is also defined in terms of regular tasks. As input for our simulation, we use the following specification of each regular task:

Task = [required resources, waiting location, delay task specification, parallel task specification, parallel check specification].

The required resources contain a list of resources required by this task (room, staff, equipment), where the use of, e.g., multiple staff members is possible. The waiting location states where the patient has to wait before starting this task. When the task can be started, and a room is filled in at required resources, the patient will be moved to this room. The delay specification states possible reasons for delays and is connected to one or more resources defined in the required resources. As an example, consider a home visit by a GP. The required resources for this task only involves a GP. Since a travel time is required to visit the patient, a delay task for this resource has to be started first. The parallel task specification consists of possible reasons to start a parallel task. If a parallel task is required after the regular task, it will be put on the task list after the regular task is finished. The parallel check specification states whether it is allowed for patients to have open parallel tasks before starting this regular task. As an example, the regular task ‘treatment at the ED’ requires the results of all diagnostic tests (reviewing a diagnostic test is considered as a parallel task).

For all three types of tasks we have a list with processing time distributions. However, we might have more distributions than tasks, since the processing times may not only depend on the task involved, but also on other factors, such as the staff member performing the task. For the regular and parallel tasks, we

define specifically which type of staff members are authorized to perform this task. This is not required for the delay task since they use the resources described by the corresponding regular task. For all tasks and types of staff members, we specify whether this staff member can do this tasks, or has certain restrictions (authorization can depend on the severity and kind of symptoms). Also, we have a similar list where we define priorities for these tasks in case there are multiple staff types present that are able to perform this task.

The care pathways of patients are defined in terms of regular tasks only. An overview of the different pathways through the IEP can be found in Figure 1. There are three ways in which patients can enter the IEP: (i) by calling the IEP, (ii) by coming to the IEP as a self-referral, or (iii) by coming to the ED through an external referral. In this last group, the patients are referred to the ED for example by the call center of the national alarm number or by another hospital. After entering the IEP, we have a sequence of treatments after which the patient can either (i) go home, (ii) be admitted into the hospital, or (iii) be transferred for admission in another hospital. For the simulation model, the destination of the patient is relevant since there might be a delay in admission (e.g., because there is no bed available at the ward) or in transfer (e.g., to wait for an ambulance). These patients stay in their ED room while waiting to move along and exit the IEP.

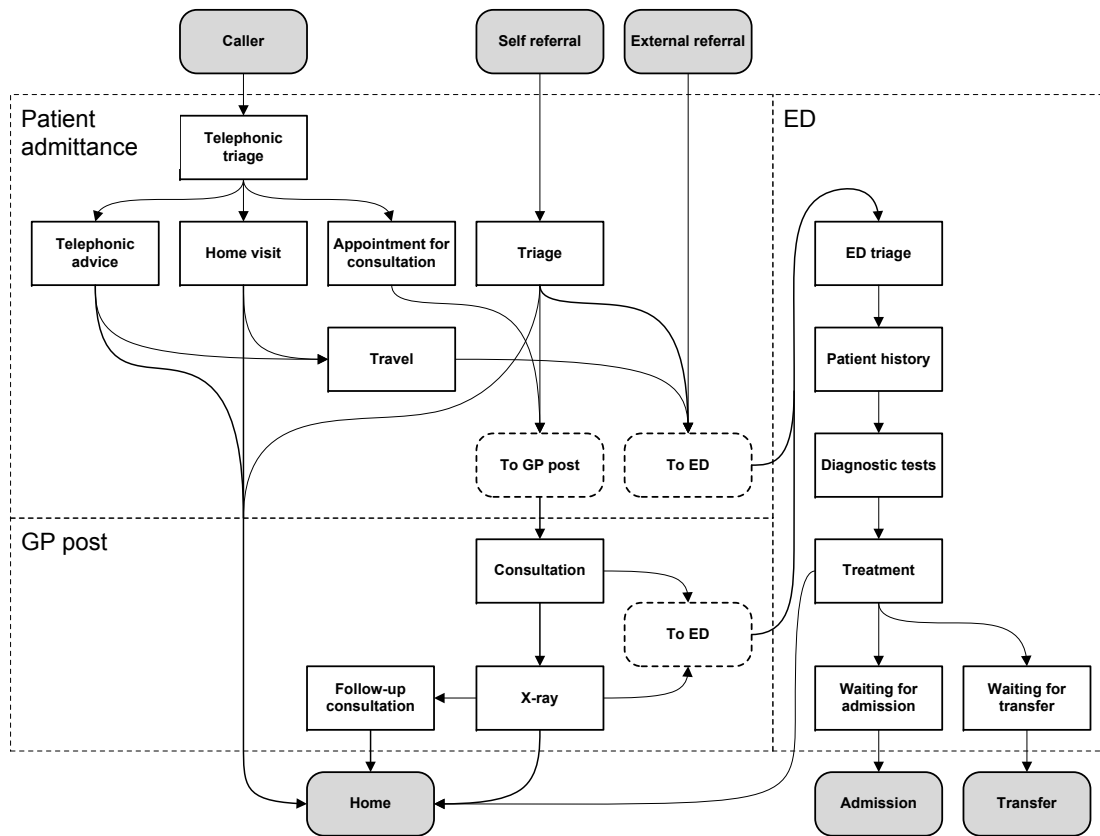


Figure 1: Flowchart of the processes of the IEP.

From Figure 1 we can see that the IEP process is split into three parts: patient admittance, which is the entry point of the IEP, the GP post, and the ED. The total number of paths going through the IEP is 31, of which three directly go from patient admittance to home, four through the GP post going home, eight paths go through the ED with three different exit destinations resulting in an additional 24 paths.

3.4 Dependencies

In practice there are many dependencies between patient arrivals, patient attributes, patient care pathway, process duration etc. Typical examples include the relation between time and urgency (e.g., patients arriving in the middle of the night tend to have a higher urgency) and the simulation group and required diagnostic tests (e.g., stroke patients are likely to receive a CT scan).

A common approach to handle dependencies between treatments is the use of care pathways. Here we generate a sequence of tasks using given input distributions. The disadvantage of this approach is that there might be many possible pathways and we have to define input distributions for each of them. For our model we make a few assumptions. First, we assume that the paths in the three areas as shown in Figure 1 are independent. Second, the diagnostic tests required at the ED are independent on the rest of the pathway. We realize that these assumptions might lead to a unrealistic pathway for an individual patient. However, on a group level these peculiarities will average out. Also, we do include indirect dependencies as we describe in the remainder of this section.

Upon creation of a patient (see Section 3.1), we follow a sequence of steps:

- Depending on the time of arrival, we determine the GP post urgency (GP post triage category).
- Depending on the GP post urgency, we determine the care pathway of the patient, for the processes within the part patient admittance.
- If the first part of the patients' pathway ends at the GP post, we determine the path through the GP post (independent of the path at patient admittance).
- If one of the paths end at the ED, we do the following:
 - We determine the ED urgency (ED triage category). If the patient went through the GP post, we determine this urgency based on the GP post urgency. Otherwise, the urgency depends on time.
 - Depending on time, we determine a so-called diagnosis group. These groups are defined by the eight most common 'Diagnosis Related Groups' (DRGs) at the IEP and two rest groups: surgical specialties and contemplative specialties.
 - Depending on the diagnosis group, we determine the required diagnostic tests and the duration of these tests (e.g., number of photos to take).
 - Depending on the ED urgency, we determine ED treatment requirements and duration (e.g., is an external medical specialist required and how long does it take to bring him to the ED).
 - Depending on the diagnosis group, we determine the way the patient will leave IEP.

Variables not included in the sequence of steps are assumed to be independent. Note that all these dependencies have been confirmed by our data analysis and by expert opinion.

3.5 Prioritization and Appointment Scheduling

In modeling the flow of patients, we might need to prioritize and schedule appointments. Prioritization is related to patients (e.g., which patient to treat first) and to resources (e.g., which staff member to assign to a specific task).

Tasks are prioritized as a function of (i) the triage categories, for which different systems are in use at the GP post and the ED, and (ii) the waiting time of the patient. Our prioritization function allows to increase the priority of a patient when its waiting time is close to its maximum waiting time. This maximum waiting time is derived from regulations that state for each triage category within how much time a patient must be seen by a staff member. After sorting the task list, we check for each task, starting with the first, if it can be started. Further, for tasks of patients with the same triage category, we give priority to a parallel task over a regular task. We do this because (i) the parallel tasks generally require limited processing time and resources and (ii) regular tasks might be delayed as a result of postponing parallel tasks. Note that the delay tasks are not included in the task list.

Our model also includes an appointment module. When a patient calls the IEP and he has to come to the IEP for a consultation at the GP post, he gets an appointment time, which depends on the availability of GPs. In our simulation we maintain a schedule for the GPs, which contains the working shifts of GPs, the appointments with patients, and reserved GP capacity for home visits that are not known yet. Every calling patient gets an appointment time at the first available time slot. Upon arrival at the IEP, the consultation is placed on the tasks list. Then, patients with higher urgency will be treated first automatically, independent on the appointment time.

3.6 Simulation Model Structure

Given the use of discrete-event simulation, we must identify events that trigger decisions and processes. The way we handle the events in our simulation is shown in Figure 2.

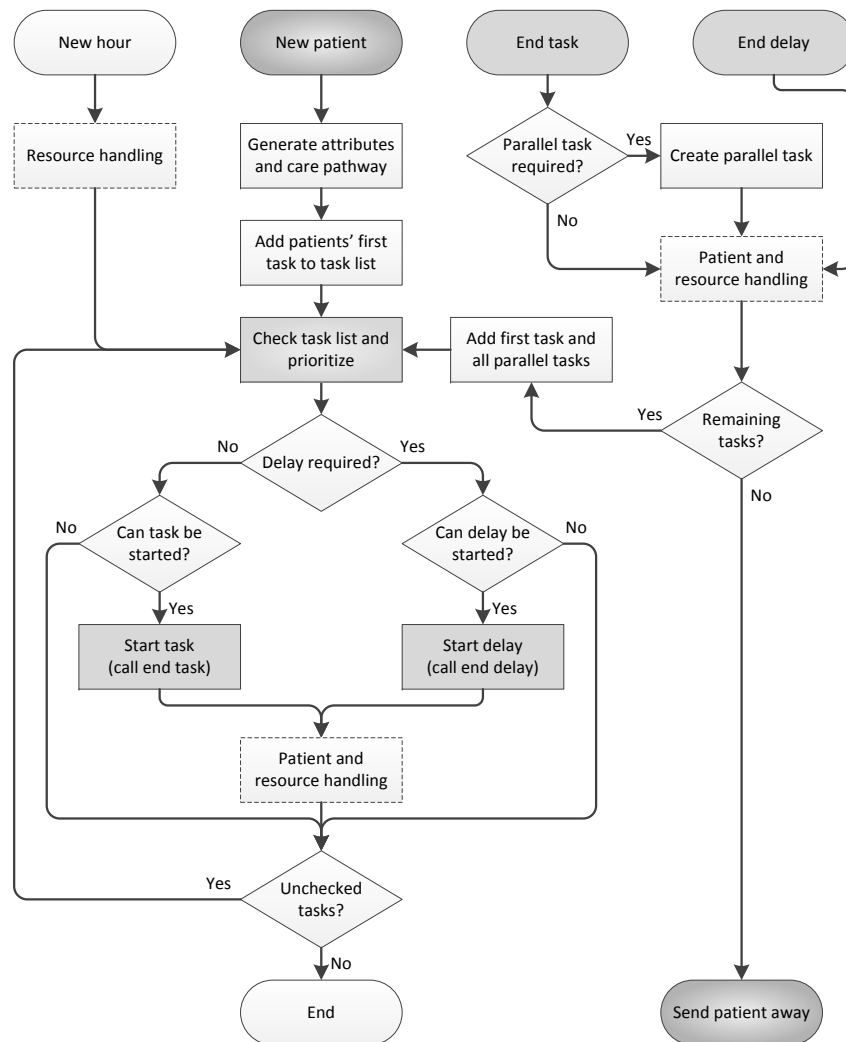


Figure 2: Simulation framework.

Upon the creation of a new patient, the patient attributes and its care pathway are generated. Even though we generate a full care pathway, we only add the first task to the task list. We then trigger a module that checks all open tasks and prioritizes these tasks. For each task, we also have to check whether a delay

is required. In case of a delay, one or more resources required by this task, will be claimed for a given amount of time before the regular task can be started. Upon the start of a task (one of the three types), we perform patient and resource handling. This involves updating counters, performance registration, and animation (i.e., moving patients, staff, and equipment).

When the processing time for the specific task ends, we again perform patient and resource handling, and check whether a parallel task is required (which we only do for regular tasks). Now, the resources required by the finished task become available. An exception here is the ED-room the patient is in, which in some situations stays assigned to the patient. The release of resources triggers the program to check whether new tasks can be started.

The last important trigger is the passing of time. For simplicity, we limit ourselves to hourly events. These hourly events may cause changes in the resource availability. Here we can think of the start of working time for staff members, opening of rooms, etc. At the moment the shift of a certain staff member ends, the staff member finishes the task he is currently performing and then leaves the IEP.

4 IMPLEMENTATION

We implemented our model in the discrete-event simulation software package Tecnomatix Plant Simulation. The basis for the simulation model is a map of the IEP in Almelo, which increases recognizability for the stakeholders and future users. The patients are shown as small figures that move over the map. Patients as well as staff members have different colors based on their characteristics. Colored markers are placed at the patients to indicate their urgency. Separate areas are created for patients and staff that are currently outside the IEP: 'waiting for an appointment', 'waiting for visit', 'being visited', 'in ambulance' and for staff members 'being off shift' and 'being underway'. A screen shot of the simulation model is shown in Figure 3.

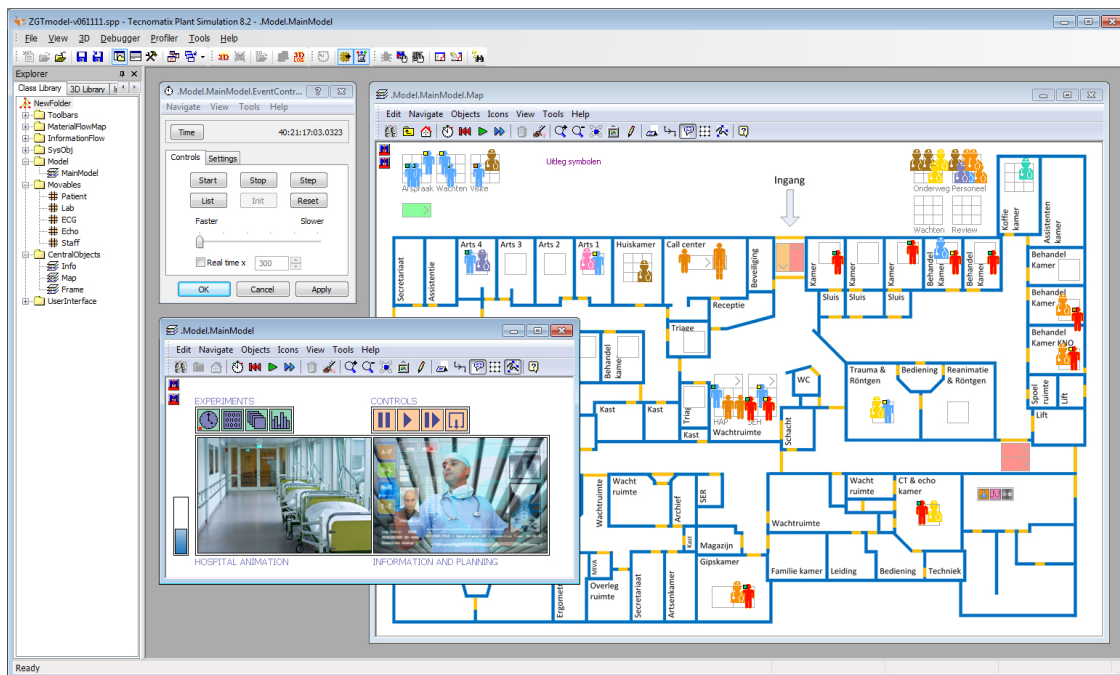


Figure 3: Screenshot of the simulation model.

Before the simulation model can be used, it has to be thoroughly verified and validated. Verification of our model was done using (i) code debugging, (ii) model reviewing, (iii) providing demonstrations to the stakeholders, (iv) comparing the generated outcomes (e.g., treatment durations) with the input distributions,

and (v) watching the animation of the patients moving through the IEP. From this analysis, we conclude that the current simulation is a good translation of the conceptual model.

Validation was done using the six validation techniques as mentioned by Law (2007). To provide an example of a quantitative validation technique, we compare the historical and the simulated patient arrival patterns. For the patient arrivals, we use a mixture of three distributions to generate the arrivals (see Section 3.1). Hence, the standard deviation in the number of arrivals per hour is not directly input for our simulation model. To validate this part, we perform a comparison between the historical information from the IEP and the results of four years simulation using the mean and variance of the number of arrivals per hour for the different days. As an illustration, Figure 4 shows the historic and simulated mean and standard deviation in the number of arrivals (callers and self-referrals) per hour on Saturdays. Here, we see a good match in mean number of arrivals per hour. The match in standard deviation is reasonable; in periods with relative high arrival rates we observe slightly lower variances, which is due to the characteristic of the Poisson distribution that the variance follows the mean.

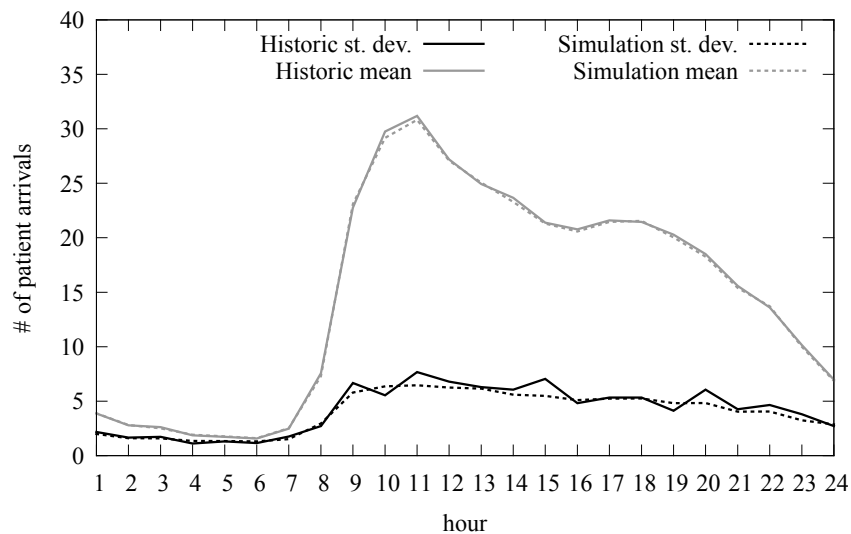


Figure 4: Deviation in mean and standard deviation of arriving patients per hour on Saturdays.

Similar comparisons were done with respect to weekdays, Sundays, the external referrals, and the waiting times for each treatment. From this analysis, we conclude that our simulation model matches reality accurately enough.

5 CONCLUSIONS AND DISCUSSION

In this paper we introduced a generalized flexible discrete-event simulation model for emergency departments. We illustrated our modeling approach with the design of a simulation model of an integrated emergency post (IEP), which is a collaboration between a general practitioners post (GP post) and an emergency department (ED). We showed that our simulation model accurately emulates the patient flow through the IEP and that it is ready to be used to analyze the effects of organizational interventions.

By now, two follow-up projects have been started in which the simulation model is used. The first project examines the benefits, for this and other hospitals, of integrating the ED and GP post into an IEP. The second project studies organizational interventions aimed at increasing the efficiency within the IEP using simulation optimization techniques.

ACKNOWLEDGMENTS

This research is part of a research project sponsored by ZonMw (project number 8271.1011). The participants of this project are ZGT, the 'Centrale Huisartsenpost Almelo' (CHPA), and the departments 'Health Technology and Services Research' and 'Industrial Engineering and Business Information Systems' of the University of Twente.

REFERENCES

- Facchin, P., E. Rizzato, and G. Romanin-Jacur. 2010. "Emergency department generalized flexible simulation model". In *Proceedings of the IEEE Workshop on Health Care Management*, 1–6. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Ferrin, D. M., M. J. Miller, and D. L. McBroom. 2007, December. "Maximizing hospital financial impact and emergency department throughput with simulation". In *Proceedings of the 2007 Winter Simulation Conference*, edited by S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, 1566–1573. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Fletcher, A., D. Halsall, D. Huxham, and D. Worthington. 2007. "The DH accident and emergency department model: A national generic model used locally". *Journal of the Operational Research Society* 58 (12): 1554–1562.
- Gunal, M., and M. Pidd. 2010. "Discrete event simulation for performance modelling in health care: a review of the literature". *Journal of Simulation* 4:42–51.
- Hay, A. M., E. C. Valentin, and R. A. Bijlsma. 2006, December. "Modeling emergency care in hospitals: a paradox - the patient should not drive the process". In *Proceedings of the 2006 Winter Simulation Conference*, edited by L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, 439–445. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Jacobson, S. H., S. N. Hall, and J. R. Swisher. 2006. "Discrete-Event Simulation of Health Care Systems". In *Patient Flow: Reducing Delay in Healthcare Delivery*, edited by F. S. Hillier and R. W. Hall, Volume 91 of *International Series in Operations Research & Management Science*, 211–252. Springer US.
- Kool, R., D. Homberg, and H. Kamphuis. 2008. "Towards integration of general practitioners posts and accident and emergency departments: a case study of two integrated emergency posts in the Netherlands". *BMC Health Services Research* 8 (225).
- Law, A. 2007. *Simulation modeling and analysis*. McGraw-Hill series in industrial engineering and management science. McGraw-Hill.
- Raunak, M., L. Osterweil, A. Wise, L. Clarke, and P. Henneman. 2009. "Simulating patient flow through an Emergency Department using process-driven discrete event simulation". In *Proceedings of the 2009 ICSE Workshop on Software Engineering in Health Care, SEHC '09*, 73–83. Washington, DC, USA: IEEE Computer Society.
- Sinreich, D., and Y. N. Marmor. 2004, December. "A simple and intuitive simulation tool for analyzing emergency department operations". In *Proceedings of the 2004 Winter Simulation Conference*, edited by R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, 1994–2002. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHIES

MARTIJN MES is an assistant professor within the department Industrial Engineering and Business Information Systems at the University of Twente, the Netherlands. He holds a MSc in Applied Mathematics (2002) and a PhD in Industrial Engineering and Management at the University of Twente (2008). His research interests are health logistics, transportation, multi-agent systems, stochastic optimization, discrete-event simulation, and simulation optimization. His email address is m.r.k.mes@utwente.nl.

Mes and Bruens

MANON BRUENS is a researcher at the Integrated Emergency Post at Ziekenhuisgroep Twente (Hospital Group Twente) in Almelo, the Netherlands. Besides research, she also works as policy advisor in this hospital. She has a BSc in Industrial Engineering and Management and a MSc in Health Sciences. Her research interests are health care and operations research techniques. Her email address is m.bruens@zgt.nl.