



Discrete-Event Based Patient Flow Simulation of an Emergency Surgery Department

Downloaded from: <https://research.chalmers.se>, 2026-05-26 09:04 UTC

Citation for the original published paper (version of record):

Combrink, A., Johnson, D., Moldan, P. et al (2024). Discrete-Event Based Patient Flow Simulation of an Emergency Surgery Department. 10th 2024 International Conference on Control, Decision and Information Technologies, CoDIT 2024: 1243-1248.
<http://dx.doi.org/10.1109/CoDIT62066.2024.10708150>

N.B. When citing this work, cite the original published paper.

Discrete-Event Based Patient Flow Simulation of an Emergency Surgery Department

Alvin Combrink

*Department of Electrical Engineering
Chalmers University of Technology
Gothenburg, Sweden
combrink@chalmers.se*

Petr Moldan

*Department of Electrical Engineering
Chalmers University of Technology
Gothenburg, Sweden
petr.moldan@gmail.com*

David Johnson

*Department of Electrical Engineering
Chalmers University of Technology
Gothenburg, Sweden
davidjohnsson.96@hotmail.com*

Martin Fabian

*Department of Electrical Engineering
Chalmers University of Technology
Gothenburg, Sweden
fabian@chalmers.se*

Abstract—Increased demand for healthcare services is placing a significant strain on hospitals. Prolonged waiting times for patients are becoming commonplace, while healthcare staff are subjected to unsustainable workloads. Finding ways to increase patient flow through hospital departments is one crucial step toward efficient healthcare services.

In this work, a modelling framework is proposed to model patient flow through a healthcare department. Patient progression and resource allocation is simulated, offering insights into expected outcomes, bottlenecks, and other inefficiencies.

A discrete-event model of a hospital department is formulated and proposed to be used together with Monte Carlo simulations. Patient treatment is represented by a series of processes, each consisting of smaller tasks. Medical staff members are represented as resources with specific qualifications that decide what tasks they may execute. Resources are allocated dynamically to model department-specific procedures, therefore increasing the flexibility of the proposed framework and opening up modelling possibilities to different healthcare departments.

A real-world healthcare department is modelled and simulated using historic data and expert knowledge. In this way, the modelling flexibility of the framework is shown. Comparisons between simulation results and actual outcomes highlight the importance of establishing high-quality quantitative data collection in healthcare departments at an early stage to provide a stable foundation for operational modelling research. With accurate process times and resource usage data, the proposed framework has the potential to serve as an important support function, and ultimately contribute to a more sustainable and efficient healthcare.

Index Terms—Discrete-Event Modelling, Patient Flow, Emergency Department

I. INTRODUCTION

Healthcare providers worldwide are facing a significant strain [19]. Increased demand for healthcare services has surpassed the available supply in many regions, leading to prolonged waiting times for patients [19]. This in turn can exacerbate medical conditions and necessitate more complicated treatments. Moreover, the increased demand for healthcare

services is also impacting the efficiency of hospitals, resulting in unsustainable workloads [19] for personnel and resources.

Finding ways to improve patient flow through hospital departments is one way to contribute to a well functioning healthcare system. Understanding how various factors effect patient flow can help administrators make better decisions regarding patient scheduling, personnel scheduling, resource allocation, etc. The central aim of this work is to develop a simulation framework that is general enough to be applicable to a wide range of departments and that can be used to reveal insights regarding how treatment and resource usage may progress over a time horizon within a department.

Research on modelling and optimising patient flow is a multidisciplinary field. Many studies focus on specific sub-fields due to the complexity and scale of modern healthcare systems [9, 10]. While modelling patient flow is not new [3, 9, 15], it is still widely under-utilised within healthcare [15], and can still arguably be considered an unsolved problem in practice.

Previous studies have modelled entire hospitals on a strategic level [2, 8, 14], or less complex departments, such as elderly care [17] or emergency rooms [13, 20]. Generally, the research can be categorised as either operational or strategic in nature, with some studies proposing a third tactical category [9]. While the definition of strategic is not necessarily well-defined, strategic models generally focus on at least multi-department or hospital-level modelling, usually with time-frames measured in weeks. Highly relevant issues include predicting clinical pathways [2, 8, 14] and finding methods for forecasting demand for various treatments [2, 9], which are crucial problems to solve to achieve a higher utilisation of hospital resources. In contrast, operational models usually focus on subsystems, such as a department or a number of interdependent departments within a hospital [9], with objectives usually being to model operations over hours or days. Within the operational domain, which this work can be

considered to belong to, there exist multiple proposed patient flow modelling frameworks [9, 12].

The most prevalent models are based on discrete-event systems [9, 12]. Other options include agent-based models [7, 9] and queuing theory [2, 9]. However, the majority of models are not strictly confined to a single category but rather a combination.

Discrete-event based models are widely used in healthcare due to various favourable attributes. One crucial factor is the organisational structure of hospitals, which aligns well with the principles of discrete-event modelling. Most medical personnel have well-defined responsibilities [12]. Strict rules exist to define what tasks each specialised medical professional is allowed to perform, and which patients fall under that responsibility. While healthcare systems share some characteristics with other discrete-event systems, such as assembly lines, telecommunication, and computer systems, they differ from these technical systems due to the substantial presence of humans, both patients and caregivers.

Patients with similar diagnoses and planned medical procedures can have significantly different requirements. Thus, discrete-event models proposed for healthcare simulation typically include probabilistic elements [14, 16, 17, 20]. Incorporating probabilistic process times in a discrete-event simulation is a well established practice both for modelling patient flows [12, 16, 17, 20] and industrial applications [6].

Current proposed operational-level models tend to model units with linear patient flows such as elderly care [17, 18] with few, if any, shared resources; or emergency departments [16, 20], which generally have fewer types of shared resources and generally have different challenges, such as that the vast majority of patients are not diagnosed prior to arrival at the emergency department. Predicting length of stay in a surgery department is neither a novel nor a solved problem [4, 5].

The model proposed in this work represents patient treatment as a series of processes, each process representing a procedure that the patient undergoes. Each process consists of a number of tasks, each of which must be completed by exactly one healthcare staff member with a specific qualification for a certain amount of time. Tasks within a process are performed in parallel. Once all processes for a patient have been completed, the patient is considered treated and leaves the department. Healthcare personnel are modelled as resources with a set of qualifications that specify their eligibility to perform certain tasks. Each process has a specified available subset of the department's resources, which increases the modelling flexibility by allowing available resources to be set dynamically.

Given a schedule of patient arrival times and resource demands, as is commonly known beforehand in e.g. surgery departments, the formulated model can be used to simulate the progression of patient treatment and resource usage over time. Here, Monte Carlo simulations are performed. An advantage of performing Monte Carlo simulations is the intrinsic modelling freedom that it allows. As shown in Section III, the proposed model is flexible enough to relatively easily implement a

number of department-specific procedures, such as conditional resource availability and tasks that do not directly relate to patient care, and to also address limited knowledge of the processes themselves. This allows for different types of hospital departments to be modelled and simulated. However, such simulations are only as useful as the accuracy of both the simulation model's representation of the real system, and the parameters going into the model. Parameters may include processes' composition of tasks, durations, resource usage, etc. Past treatment data and expert knowledge is a source of insight into these parameters. By sampling from distributions of past process times and resource usage, Monte Carlo simulations may provide insights into expected patient treatment progression as well as their uncertainty.

This paper presents work [11] on modelling and simulating the Orthopaedic Surgery Department 1 at Sahlgrenska University Hospital, Mölndal, Sweden, demonstrating how the framework's flexibility can be utilised when faced with a number of rather department-specific aspects. Process times are sampled from historic data, given a patient's planned type of surgery and ASA-classification [1]. ASA-classification is a 6-point scale used by caregivers to assess a patient's physical health prior to surgery, here the ASA-classification is assumed to correlate with process times as well as resource intensity.

The outline of the paper is as follows: Section II presents the model formulation. A specific healthcare department is modelled using the proposed framework in Section III, followed by results in Section IV. Finally, discussions and concluding remarks are found in Section V.

II. MODEL FORMULATION

A. Preliminaries

The system represents a medical department where patients enter, are treated by medical staff members, and finally exit. A patient $\omega \in \Omega$ (with Ω the set of all patients) is associated with a time of arrival $t(\omega)$ at the department. Only at time $t \geq t(\omega)$ can the treatment of ω begin. Medical staff members are represented as *resources*, $r \in \mathcal{R}$, each associated with a set of medical qualifications $s(r) \subseteq \mathcal{S}$ and a priority value $\pi(r)$.

The treatment of a patient ω is modelled as an ordered set of $n_p^\omega \in \mathbb{N}$ *processes*, $\mathcal{P}(\omega) = \langle p^{\omega,1}, \dots, p^{\omega,n_p^\omega} \rangle$. These processes represent various medical procedures, such as a single process for an entire surgery or several processes for sequential parts of a surgery, and are executed in sequence. It is assumed that once all processes for a patient have been executed, the patient is done at the department and exits the system. A process $p^{\omega,i}$ is associated with a set of resources $\mathcal{R}(p^{\omega,i}) \subseteq \mathcal{R}$ and a priority value $\pi(p^{\omega,i})$.

A process $p^{\omega,i} = \{ \tau^{\omega,i,1}, \dots, \tau^{\omega,i,n_\tau^{\omega,i}} \}$ is a set of $n_\tau^{\omega,i}$ *tasks*. A task $\tau^{\omega,i,j}$ is associated with a medical qualification requirement $s(\tau^{\omega,i,j,k}) \in \mathcal{S}$, a duration $d(\tau^{\omega,i,j,k})$, and a requested resource $q(\tau^{\omega,i,j}) \in \mathcal{R}$. All tasks in a process execute in parallel, starting at the same time. However, they do not necessarily finish at the same time since tasks can

have different durations. A process is considered executed when all of its tasks have been executed. It is assumed that multiple tasks in a process do not have the same medical qualification requirement. Furthermore, time is discretised and task durations are assumed to be multiples of time-steps.

B. Resource Requesting and Booking

At some time t there can be any number of patients in the department. For each patient, there is a process that is either being executed or is waiting to be executed. Let the set $\mathcal{W} \subseteq \{p \in \mathcal{P}(\omega) \mid \omega \in \Omega\}$ contain all processes that are waiting to be executed. Initially $\mathcal{W} = \emptyset$. Process $p^{\omega,1}$ is added to \mathcal{W} at $t = t(\omega)$, $\forall \omega \in \Omega$, and process $p^{\omega,i}$ is added to \mathcal{W} when $p^{\omega,i-1}$ has finished executing.

At the same time as a process is added to \mathcal{W} , all of its tasks *request* a resource. Requesting a resource is not the same as *booking* a resource, which would mean that the task may use the resource and begin executing. Booking is done at a later phase.

A task $\tau^{\omega,i,j} \in p^{\omega,i}$ may only request a resource $r \in \mathcal{R}(p^{\omega,i})$ that contains the correct medical qualification, $s(\tau^{\omega,i,j}) \in \mathbf{s}(r)$. Let this set of resources be defined as

$$\mathcal{R}(\tau^{\omega,i,j}) = \{r \in \mathcal{R}(p^{\omega,i}) \mid s(\tau^{\omega,i,j}) \in \mathbf{s}(r)\}. \quad (1)$$

If a task requests a resource r , then $q(\tau^{\omega,i,j}) = r$. Resource r 's queue can be defined as the set of tasks requesting it,

$$\mathcal{Q}(r) = \{\tau \in p \mid p \in \mathcal{P}(\omega), \omega \in \Omega, q(\tau) = r\}. \quad (2)$$

Primarily, $r \in \mathcal{R}(\tau^{\omega,i,j})$ with the smallest queue is requested. If there are several resources with an empty queue, $\mathcal{Q}(r) = \emptyset$, the resource with the lowest priority is requested. Let the set of resources that $\tau^{\omega,i,j}$ may request and that have an empty queue be defined as

$$\mathcal{R}^\emptyset(\tau^{\omega,i,j}) = \{r \in \mathcal{R}(\tau^{\omega,i,j}) \mid \mathcal{Q}(r) = \emptyset\}. \quad (3)$$

The requested resource for $\tau^{\omega,i,j}$ is

$$q(\tau^{\omega,i,j}) = \begin{cases} \arg \min_{r \in \mathcal{R}^\emptyset(\tau^{\omega,i,j})} \pi(r) & \text{if } \mathcal{R}^\emptyset(\tau^{\omega,i,j}) \neq \emptyset \\ \arg \min_{r \in \mathcal{R}(\tau^{\omega,i,j})} |\mathcal{Q}(r)| & \text{otherwise.} \end{cases} \quad (4)$$

Once all resource requests have been made, resource booking begins. A resource can only be booked by one task at a time, and when that happens the task begins executing for as many time-steps as its duration, $d(\tau)$. When the task has been completed, the resource is released so that it can be booked by another task.

Booking is done for each process $p \in \mathcal{W}$ at a time, in order of priority and then according to first-in-first-out. All tasks in a process must book a resource before any of them may begin executing. This is because all tasks in a process begin executing at the same time, and each require a resource. To avoid deadlocks, if any task has requested an already booked resource and is therefore unable to book it at the current time, no task in the process will book their requested resources.

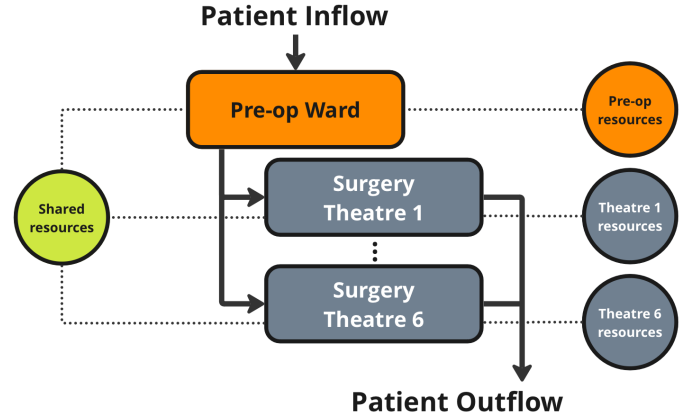


Fig. 1. A macro level overview of the patient flow in the department with corresponding resources.

More formally, for each process $p \in \mathcal{W}$, in descending order of first priority $\pi(p)$ and then time in \mathcal{W} , if none of the resources requested by the tasks $\tau \in p$ are booked, i.e. $\{q(\tau) \mid \tau \in p\} \cap \mathcal{B} = \emptyset$ ($\mathcal{B} \subseteq \mathcal{R}$ containing all booked resources), then all tasks may book their requested resources. Resources $r \in \{q(\tau) \mid \tau \in p\}$ are added to \mathcal{B} and p is removed from \mathcal{W} . The tasks $\tau \in p$ start, and execute for as long as their respective durations $d(\tau)$. After $d(\tau)$ time-steps have passed since task τ began executing, it is considered executed and resource $q(\tau)$ is removed from \mathcal{B} such that it can be booked by another task. Once all tasks $\tau \in p$ have executed, p is considered executed. Finally, if all processes $p \in \mathcal{P}(\omega)$ have executed, the treatment of patient ω is considered complete and ω leaves the system.

III. APPLYING THE MODEL TO A HOSPITAL DEPARTMENT

The Orthopaedic Surgery Department 1 at Sahlgrenska University Hospital, Mölndal, Sweden, consists of two wards: the *pre-op ward* and the *surgery ward*, the latter consisting of 6 *surgery theatres*, see Fig 1. At the pre-op ward, patients are prepared for surgery by undergoing a number of procedures. Once preparations are complete, the patient is moved to one of the available surgery theatres where an operation occurs. After surgery, the treatment of the patient is complete. The procedures that a patient undergoes at each of the wards are modelled as a series of processes. A patient ω 's set of processes $\mathcal{P}(\omega)$ therefore consists of a number of pre-op-related processes, followed by a number of surgery-related processes. A lack of knowledge regarding the exact procedures that a patient undergoes at each of the wards is addressed by using data collected by the department to estimate a patient's total time in each, as described in Section III-A. Given an estimate of the total time the patient spends at e.g. the pre-op ward, $t_{\text{pre-op}}^\omega$, a single process in $\mathcal{P}(\omega)$ can be created, containing the same number of tasks as required resources, and where the longest duration $d(\tau)$ is equal to $t_{\text{pre-op}}^\omega$. However, this would lock resources to the patient for the entire time they spend at the ward which is not accurate. In reality, staff members alternatingly tend to different patients. A solution to

TABLE I
NUMBER OF RESOURCES IN THE PRE-OP WARD, EACH SURGERY THEATRE,
AND THE SHARED RESOURCE POOL.

Resource specialisation	Pre-op	Surgery theatre i	Shared
NSSK	1	1	2
SSK	1		2
USSK	1	1	2
OSSK		1	2

TABLE II
THE SCHEDULED PATIENTS, EACH WITH AN OPERATION TYPE,
ASA-CLASSIFICATION, ASSIGNED THEATRE, AND THE NUMBER OF
HISTORICAL DATA POINTS WITH MATCHING OPERATION TYPE AND
ASA-CLASSIFICATION.

ID	Op-type	ASA	Theatre	Historical data
1:1	A	3	1	88
1:2	B	3	1	10
1:3	C	3	1	452
2:1	D	2	2	611
2:2	D	2	2	611
2:3	D	2	2	611
3:1	E	3	3	16
3:2	F	3	3	4
4:1	G	3	4	15
4:2	H	1	4	36
4:3	I	2	4	141
5:1	J	2	5	181
5:2	J	3	5	58
6:1	K	2	6	2
6:2	A	2	6	50
6:3	L	4	6	37

this, implemented here, is to create $t_{\text{pre-op}}^{\omega}/\delta$ processes, where $\delta = 1$ minute is a chosen time resolution. The same is done for surgery. This also allows for greater flexibility regarding resource demands, as described in Section III-B. No process prioritisation is used; all process priorities $\pi(p)$ are set to the same value.

The pre-op ward and each surgery theatre contain their own resources, which should be sufficient under normal circumstances. However, the department also has a pool of shared resources that can be used by either ward when needed, see Table I. In this case, each resource has only one specialisation. For the pre-op ward, the shared resources are only available during the morning rush. To model these resource constraints, a process p 's set of available resources $\mathcal{R}(p)$ is dynamically modified during runtime of the simulations to reflect this.

For this problem, 16 patients are scheduled for surgery and each have their own operation type, ASA-classification, assigned surgery theatre, see Table II, and a scheduled time to enter the department, seen by their respective start of the simulated pre-op phase in Fig 2. Surgery theatres are not modelled as resources. Instead, during the resource-booking phase, the availability of a patient's assigned surgery theatre is first checked. If a theatre is available, that theatre's resources becomes accessible to the process by populating its resource set $\mathcal{R}(p)$. Otherwise, $\mathcal{R}(p) = \emptyset$.

Finally, in addition to treating patients, the healthcare staff are also occupied with items not directly related to individual patients. These are modelled as *auxiliary tasks*; they are placed

TABLE III
PROBABILITY VALUES FOR ASA-CLASSIFICATIONS AND
RESOURCE INTENSITY-MODES.

	ASA-classification			
	1	2	3	4
Low	0.2	0.3	0.4	0.5
Medium	0.2	0.3	0.8	1.0
High	0.5	0.6	0.8	1.0

in the Ω set, with as many processes and tasks as required. The auxiliary tasks *prepare theatre*, *re-stock* and *take inventory* are added at every δ time-steps with probability 0.005, 0.07 and 0.005, respectively. Priority of these tasks are set to lower than patient tasks, except for prepare theatre.

A. Sampling Process-Times and Patient Processes

Since information regarding patient procedures at each of the wards is not available, they cannot be treated as separate processes. Instead, the time each patient spends at each ward is sampled from historic data containing these values, collected by the department. Data matching the patient's surgery type and ASA-classification is uniformly sampled, independently for the time spent at the pre-op and surgery wards, respectively. The number of historic data points matching each patient's parameters is seen in Table II, right column.

B. Sampling Resource-Demand for Tasks

Based on consultations with the healthcare staff at the department, it is assumed that a patient's ASA-classification is a determinant of how resource-intensive they will be, measured by the number of staff members needed to attend to the patient. Resource intensity is modelled here by the number of tasks τ in a process p ; each task requires one resource, and all tasks in a process must execute in parallel. Sampling of resource demand is done based on two parameters: a patient's ASA-classification and a resource intensity-mode, which can be low, medium or high. The resource intensity-mode is a simulation setting that can be useful for investigating a schedule's sensitivity to various levels of overall resource usage.

For each combination of ASA-classification and intensity-mode, a probability value σ based on the medical staff's expert knowledge is given, see Table III. A set of possibly required qualifications $\mathbf{q} \subseteq \mathcal{S}$ is given for each ward. Then, for each process $p \in \mathcal{P}(\omega)$, with probability $1 - \sigma$, a random subset $\bar{\mathbf{q}} \subseteq \mathbf{q}$ of medical specialisations is chosen and a task for each element is added to the process. Otherwise, \mathbf{q} is used instead. Therefore, the higher σ is, the higher the resource demand of patient ω will likely be.

IV. RESULTS

Simulations are run for the scheduling scenario described in Section III, which originates from a real scenario that took place at the department. Therefore, the simulated results are shown together with the actual outcome in Fig 2. Each patient's simulated pre-op start time is set to the scheduled arrival time, while for the actual outcomes, pre-op started when

the patient arrived at the department. The pre-op end time and the surgery start and end times shown are the average over the simulations, error bars show the variance.

Despite using the low intensity-mode, as was done here, the simulated total treatment times for most patients is longer than the actual outcomes. Additionally, the time variance is significant. It should be noted, however, that the surgery end time depends on the surgery start time, which in turn depends on the pre-op end time. Therefore, consideration should be taken that variance compounds for each successive time. However, the variance for the pre-op end time comes directly from the data and can still be considered significantly large.

V. DISCUSSION AND CONCLUSION

The primary purpose of this work was to develop a model capable of representing a hospital department in order to simulate patient flow, resource usage, and expected scheduling outcomes. The framework presented above is relatively simple, yet was used to model an emergency surgery department, including a number of department specific procedures. In this regard, it is conceivable to think that the framework could be used to model other departments and therefore provide insights that could further the work toward more efficient and sustainable healthcare.

The modelling approach was applied to a real hospital department, and results showed significant discrepancies between simulated and actual outcomes. This discrepancy only grows larger with the higher resource intensity modes. The reason for these discrepancies may be multiple, however, one strong contributing factor may be the lack of high-resolution, quantitative data regarding the composition and duration of processes. The specific hospital department addressed here does not have enough data relevant to the model to provide a foundation for gaining useful insights. For example, the underlying data used to sample process times consists of a start and end time for the pre-op and surgery processes, which include the time when resource constraints lead to delays in the treatment of those patients. Therefore, resource constraints in simulation that affect process times only compound onto the already extended times provided by historic data.

The simulated times, particularly for the pre-op end time which does not include any compounding effects, show that the method of sampling times from data matching operation type and ASA-classification yields high variance values. There are many potential reasons and solutions to this. The high variance could be due to the way the data is collected, the routines in the department, or that there are other parameters that better determine process times but were not found and used here. For example, it is possible that using other factors, such as which surgeon is performing the surgery, could result in more accurate process times. There are multiple ways for future work to address this: better predictive models could provide more accurate times with less variance; different operational routines at hospital departments could lead to lower variance

in the actual times, facilitating more accurate scheduling simulations.

Additionally, no data regarding which specific resources that were involved in the treatment of patients was available. Therefore, it is highly uncertain that the method of controlling resource intensity, as described in Section III-B, is accurate.

Besides using the model to simulate the patient flow, a potential avenue of further research could be to use the model for *falsification*, where the simulation model is automatically exposed to different input parameter values, and the simulation result is automatically assessed to guide the selection of new input parameter values with higher probability to break given requirement specifications. This could help to uncover faulty assumptions in the underlying system and parameters, as well as reveal, and therefore mitigate, potentially unwanted outcomes.

Based on these insights, this work highlights that modelling and simulation of a healthcare department can only be as accurate as the underlying knowledge of how the department fundamentally works, be it from data or experts. Therefore, for frameworks such as the one presented here to be applicable in healthcare, it is crucial that routines surrounding data and knowledge collection of patient flows and resource usage are established at an early stage.

REFERENCES

- [1] Amr E Abouleish, Marc L Leib, and Neal H Cohen. "ASA provides examples to each ASA physical status class". In: *ASA Monitor* 79.6 (2015), pp. 38–49.
- [2] Mor Armony et al. "On Patient Flow in Hospitals: A Data-Based Queueing-Science Perspective". In: *Stochastic Systems* 5.1 (2015), pp. 146–194. DOI: 10.1287/14-SSY153. eprint: <https://doi.org/10.1287/14-SSY153>. URL: <https://doi.org/10.1287/14-SSY153>.
- [3] Papiya Bhattacharjee and Pradip Kumar Ray. "Patient flow modelling and performance analysis of healthcare delivery processes in hospitals: A review and reflections". In: *Computers & Industrial Engineering* 78 (2014), pp. 299–312.
- [4] Oliver Buchholz et al. "Analyzing the relationship between physicians' experience and surgery duration". In: *Operations Research for Health Care* 36 (2023), p. 100377. ISSN: 2211-6923. DOI: <https://doi.org/10.1016/j.orhc.2022.100377>. URL: <https://www.sciencedirect.com/science/article/pii/S2211692322000388>.
- [5] Marinus JC Eijkemans et al. "Predicting the unpredictable: a new prediction model for operating room times using individual characteristics and the surgeon's estimate". In: *The Journal of the American Society of Anesthesiologists* 112.1 (2010), pp. 41–49.
- [6] George S Fishman. *Discrete-event simulation: modeling, programming, and analysis*. Vol. 537. Springer, 2001.
- [7] Marcia R. Friesen and Robert D. McLeod. "A Survey of Agent-Based Modeling of Hospital Environments". In: *IEEE Access* 2 (2014), pp. 227–233. DOI: 10.1109/ACCESS.2014.2313957.
- [8] Anastasia A. Funkner, Aleksey N. Yakovlev, and Sergey V. Kovalchuk. "Towards evolutionary discovery of typical clinical pathways in electronic health records". In: *Procedia Computer Science* 119 (2017). 6th International Young Scientist Conference on Computational Science, YSC 2017, 01-03 November 2017, Kotka, Finland, pp. 234–244. ISSN: 1877-0509. DOI: <https://doi.org/10.1016/j.procs.2017.11.181>.




Fig. 2. Simulated and actual outcomes for a scheduling scenario including 16 patients and 6 surgery theatres. Error bars indicate the variance of the simulated times at which they are centred: pre-op end time, surgery start time, or surgery end time.

URL: <https://www.sciencedirect.com/science/article/pii/S1877050917323918>.

[9] Francesca Guerriero and Rosita Guido. “Operational research in the management of the operating theatre: a survey”. In: *Health care management science* 14 (2011), pp. 89–114.

[10] Muhammet Gul and Ali Fuat Guneri. “A comprehensive review of emergency department simulation applications for normal and disaster conditions”. In: *Computers & Industrial Engineering* 83 (2015), pp. 327–344. ISSN: 0360-8352. DOI: <https://doi.org/10.1016/j.cie.2015.02.018>. URL: <https://www.sciencedirect.com/science/article/pii/S0360835215000972>.

[11] David Johnsson and Petr Moldan. “Patient Flow Simulation in an Orthopedic Emergency Surgery Department”. Master’s thesis. Electrical Engineering, Chalmers University of Technology, 2024. URL: <https://odr.chalmers.se/items/a5169198-ab7c-4d96-87c8-11448ba0c9c6>.

[12] Jonathan Karnon et al. “Modeling using discrete event simulation: a report of the ISPOR-SMDM Modeling Good Research Practices Task Force-4”. In: *Medical decision making* 32.5 (2012), pp. 701–711.

[13] Renata Konrad et al. “Modeling the impact of changing patient flow processes in an emergency department: Insights from a computer simulation study”. In: *Operations Research for Health Care* 2.4 (2013), pp. 66–74. ISSN: 2211-6923. DOI: <https://doi.org/10.1016/j.orhc.2013.04.001>. URL: <https://www.sciencedirect.com/science/article/pii/S2211692313000052>.

[14] Sergey V. Kovalchuk et al. “Simulation of patient flow in multiple healthcare units using process and data mining techniques for model identification”. In: *Journal of Biomedical Informatics* 82 (2018), pp. 128–142. ISSN: 1532-0464. DOI: <https://doi.org/10.1016/j.jbi.2018.05.004>. URL: <https://www.sciencedirect.com/science/article/pii/S153204641830087X>.

[15] Boris G Sobolev, Victor Sanchez, and Christos Vasilakis. “Systematic review of the use of computer simulation modeling of patient flow in surgical care”. In: *Journal of medical systems* 35 (2011), pp. 1–16.

[16] Jiacun Wang. “Patient Flow Modeling and Optimal Staffing for Emergency Departments: A Petri Net Approach”. In: *IEEE Transactions on Computational Social Systems* (2022), pp. 1–11. DOI: 10.1109/TCSS.2022.3186249.

[17] Sarah-Jane Whittaker, Karen Rudie, and James McLellan. “An Augmented Petri Net Model for Health-Care Protocols”. In: *IEEE Transactions on Automatic Control* 60.9 (2015), pp. 2362–2377. DOI: 10.1109/TAC.2015.2409932.

[18] Sarah-Jane Whittaker et al. “Choice-point nets: A discrete-event modelling technique for analyzing health care protocols”. In: *2009 47th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. 2009, pp. 652–659. DOI: 10.1109/ALLERTON.2009.5394920.

[19] “Why health-care services are in chaos everywhere”. In: *The Economist* (2023). URL: <https://www.economist.com/finance-and-economics/2023/01/15/why-health-care-services-are-in-chaos-everywhere>.

[20] Jiani Zhou, Jiacun Wang, and Jun Wang. “A simulation engine for stochastic timed Petri nets and application to emergency healthcare systems”. In: *IEEE/CAA Journal of Automatica Sinica* 6.4 (2019), pp. 969–980. DOI: 10.1109/JAS.2019.1911576.