


Routing of two elevators

Jasper de Boer

Supervisor: Wemke van der Weij

BMI Paper



Vrije Universiteit Amsterdam
Faculteit Exacte Wetenschappen
Business Mathematics and Informatics
De Boelelaan 1081a
1081 HV Amsterdam

Executive Summary

There are many approaches to model an elevator routing system. In this paper, I shortly describe some literature, wherein different manners to model the routing system are presented. A simulation approach, dynamic programming and polling systems are used to model the routing system. In the simulation approach zoning is often used. Zoning means that the floors are split into a number of zones, each consisting of a group of floors. Different zones are served by different elevators. Dynamic zoning can reduce passenger waiting times in high-rise buildings. But it is still unclear what the effect of zoning policies is in a small building.

Therefore I have modelled an elevator routing system in a small building. In the presented model two elevators are considered that are carrying passengers from the ground floor to higher floors in the building only (only up-traffic). Four different policies (three zoning policies and one without zoning) are compared on their performance (waiting times, journey times, average load, and average number of passengers waiting). Does zoning influence the performance of the use of the elevators; will the waiting times decrease?

Analyzing the results leads to the conclusion that using zoning policies in such a low building as is used in the model is not preferable above a policy without zoning. Although there is a gain in journey time in a situation with a high intensity of arriving passengers, there are too many drawbacks to use zoning policies in the elevator routing system in a building with a small number of floors.

Preface

My name is Jasper de Boer, born in 1984 and I am studying Business Mathematics and Informatics (BMI) at the Vrije Universiteit in Amsterdam. At the end of the Master BMI but before the internship, BMI students write a BMI paper. The main objective of this paper is to do some relevant literature research combining two of the three study disciplines covered by study BMI.

Annemieke van Goor, who coordinates the BMI Paper processes, helped me finding a nice supervisor. She lets me contact Wemke van der Weij, who came with an idea for the subject. I want to thank both for their efforts. And Wemke van der Weij in particular for her valuable comments and suggestions.

The paper is about the routing of two elevators. I have tried to model an elevator routing system. I have compared several different routing strategies with the goal to minimize the passenger waiting time and passenger journey time.

I hope the reader will enjoy this paper.

Amsterdam, July 2007
Jasper de Boer

Contents

Executive Summary	i
Preface	ii
Contents	iii
1 Introduction	1
1.1 Introduction	1
1.2 Literature	3
1.3 Problem	6
1.4 Structure of the paper	7
2 Model	8
2.1 Elevator system as discrete-event system	8
2.2 Model details	10
2.3 Policies	11
3 Implementation.....	16
4 Results	18
5 Conclusion.....	24
6 Further research.....	26
7 References	27

1 Introduction

1.1 Introduction

It is commonly known that an elevator is a transport device used to move goods or people vertically. Since there are buildings with multiple floors, the use of elevators is an essential part of daily life. Higher floors must be reachable for everyone, also for disabled passengers such as passengers in wheelchairs. The use of elevators was getting more and more important. And high-rise office buildings with 30 floors are not rare nowadays. In low buildings elevators are not always available and in high buildings they are often occupied. When an elevator is occupied, you have to wait in the lobby. Passengers experience it as annoying when they have to wait for a couple of minutes before they can use an elevator.

The following slogan gives a good expression of how passengers experience queuing: *“Waiting is frustrating, demoralizing, agonizing, aggravating, annoying, time consuming and incredibly expensive.”* This slogan from Federal Express (the overnight package delivery service) has become justly famous. [\[3\]](#)

Because of all this complaints of the passengers, it is interesting to please the passengers and help them reducing frustration. The goal of this is to keep the satisfaction of the users of the elevators (e.g. employees) high. Sasser et al. [\[6\]](#) provide good examples of managing the waiting times in some psychological way. They offer the example of ‘the well-known hotel group that received complaints from guests about excessive waiting times for elevators. After analyzing how elevator service might be improved, it was suggested that mirrors should be installed at places where passengers waited for elevators. The natural tendency of people to check their personal appearance substantially reduced complaints, although the actual waiting time for the elevators was unchanged.

This psychological effect can be interesting, but we focus on the performance of the routing of an elevator according to some performance measures, such as waiting time and journey time.

First we mention some categories of traffic. Traffic in an elevator consists of three components: incoming, outgoing and inter-floor components. Incoming passengers travel from an entrance floor to populated floors, outgoing passengers from populated floors to the entrance floor and inter-floor passengers between populated floors. The incoming traffic in peak situations is called up-peak traffic and outgoing traffic in these situations is called down-peak traffic. Office buildings typically have up-peak traffic in the morning when employees enter the building, intense two-way or inter-floor traffic during lunch time, and down-peak traffic when employees exit the building. This pattern is somewhat different in hospitals. In these kind of buildings is much more inter-floor traffic, because patients are moved from department to department when they have a surgery operation for example.

To handle all these kind of traffic, an elevator have to route in a certain way. There are many different ways to route an elevator. The simplest algorithm is the following [\[13\]](#):

- An elevator continues travelling in the same direction while there are remaining requests in that same direction.
- If there are no further requests in that direction, then the elevator stops and become idle, or change direction if there are requests in the opposite direction.

It is interesting to compare this simple algorithm with other algorithms and eventually improve the performance measures.

1.2 Literature

In the literature many papers are known about how to allocate an elevator. Usually the performance criteria are assessed in the up-peak traffic situation. The up-peak traffic is the peak traffic in the morning when all the employees have to move upwards to reach their work location. Since the morning up-peak gives a difficult traffic situation in office buildings, it is important to have an efficient routing scheme. Several papers use simulation to study the allocation of elevators. Shearn [7] considered for up-peak traffic the optimal limited stop scheduling for a single elevator. A building with 20 floors and five different policies is considered. The policies he considers are as follows:

- The elevator stops at every requested destination floor,
- The elevator stops at requested even numbered floors only (it is assumed that people requiring an odd numbered floor will travel to the floor above),
- The elevator stops at floors 3, 8, 13, 18 only, if requested, on every trip (these floors are optimal for the fixed stop algorithm in this case),
- The elevator stops at floors 3,8, 13, 18 on every trip,
- The elevator makes at most four stops, these being selected dynamically by the algorithm described in the previous section.

In Shearn's paper, besides the elevator, some stairs are considered. Walking up and down stairs incurring some costs. For every policy some performance measures are determined: costs of walking, journey time and waiting time. The last policy (at most four stops) outperforms the other policies. It has the shortest journey times and the costs of walking are also quite low. However, this model is not very realistic. The obliged use of the stairs is not very user-friendly and you will rarely find it in buildings.

Another interesting paper is published by Tervonen et al. [11]. In this paper ten alternative configurations are tested on their performance. The number of elevators, rated load (maximum number of persons in one elevator) and elevator speed are varying in the different configurations. There are six performance criteria:

- Costs (the ten alternatives are ranked from 1 to 10 for the costs of a particular configuration),
- Area (the shaft space plus waiting area space)
- Waiting time (WT),
- Journey time (JT),
- The percentage of WT's exceeding 60 seconds,
- The percentage of JT's exceeding 120 seconds.

The configuration with seven elevators with a rated load of thirteen persons and an elevator speed of 5 m/s is the best choice.

In the early 90s, some primitive zoning techniques for routing elevators were developed. Zoning means that the floors are split into a number of zones, each consisting of a group of floor. Different zones are served by different elevators, in order to optimize the performance of the elevators [1]. Zoning policies can be classified into static zoning and dynamic zoning. Static zoning refers to the permanent assignment of a group of elevators to serve a number of floors in a building. Temporary static zoning can be pre-scheduled during certain times of the

day. The period usually coincides with the peak traffic situation and the elevators are serving all floors outside this period. The merit of zoning is not controversial. Actually, it has been implemented in lots of commercial buildings. However, the existing control patterns of zoning are either pre-determined or fixed on a time-schedule basis, or in other words, they are not adaptable to the real-time traffic patterns. This shortfall initiates the idea of dynamic zoning. The concept of primitive dynamic zoning for up-peak only has been developed so that elevators make fewer stops and cars return to the lobby faster. More complex dynamic zoning consists up-peak traffic, down-peak traffic and inter-floor traffic, and the zoning is continuously changing with respect to the changing traffic patterns.

So and Yu [10] have used this complex dynamic zoning to improve elevator performance. From the first zoning techniques in 1990, they formed the foundation of a comprehensive dynamic zoning scheme. They prefer dynamic zoning above static zoning, because a static assignment scheme can only be effective for one specified traffic condition, such as up-peak, but may be ineffective for other traffic patterns such as down-peak.

The paper shows that dynamic zoning can improve an existing elevator system for three different traffic conditions normally encountered in commercial buildings i.e., up-peak, down-peak and heavy inter-floor situations by reducing the average passenger waiting/journey time. Dynamic zoning only behaves a little bit poorer for the travel time (journey time minus waiting time) during a down-peak condition. That may be due to the fact that during down-peak, elevators are normally fully loaded within a few floors near to the top of the building or the top of a zone. Then, the number of stops is greatly reduced even without zoning. In this way, dynamic zoning cannot produce its desirable effect. However, from a passenger point of view, the sum of waiting time and travel time should be of major concern. In this way, there is still an improvement of 7% by using dynamic zoning.

It is anticipated that dynamic zoning may be a normal supervisory control scheme for lifts in future years when our computational machines become faster and have a larger memory space. The gain in performance of dynamic zoning is increasing if the number of floors becomes larger.

Siikonen [9] developed an elevator control system based on artificial intelligence and historical data (the Traffic Master System 9000 (TMS9000)). A practical example where an old electronic control system was modernized with the TMS9000 control system showed an improvement of about 35-40 % in average hall call times (the time between a passenger gives a hall call and the arrival of the elevator responding the hall call). The paper claims that passenger waiting times correlate with call times when the traffic intensity stays below handling capacity, so improvement in passenger waiting times is about the same.

The stochastic simulation approach used in the papers described above is a manner to model the elevator routing problem. Shoham and Yechiali [8] have considered another approach to model an elevator system. An elevator can be seen as a queuing system consisting of N queues (channels) served by a single server which incurs switch-over periods when moving from one channel to another. Very often such applications are modelled as a polling system in which the server visits the channels in a periodic routine following a given polling table, or according to some probabilistic mechanism which allows designers to prioritize the different queues so as to affect and optimize overall system performance.

They concentrate on an important polling mechanism, the Elevator-type scheme: instead of moving cyclically through the channels, the server first scans the channels in one direction, i.e. in the order 1, 2, ..., N ('up' cycle) and then reverses its orientation and serves the channels in the opposite direction ('down' cycle). Then it changes direction again, and keeps moving in

this manner, scanning the channels back and forth. This type of service regime is encountered in many applications, such as in computer operating systems as an algorithm for scheduling hard disk requests.

Nikovski and Brand [4] presented another manner to model an elevator routing system. They used dynamic programming to model it. They detail an efficient scheduling algorithm based on dynamic programming for exact estimation and minimization of the expected waiting times of all known passengers in a group elevator system. Empirical comparison with a state-of-the-art scheduler in a very detailed discrete-event elevator bank simulator demonstrated that for a wide variety of buildings, ranging from 8 to 30 floors, and with 2 to 8 elevators, dynamic programming reduces waiting times by 30%-40% under very heavy traffic, and rarely under-performs the benchmark scheduler in light traffic.

Nowadays, there are some applications that use dynamic programming for the elevator routing system, for example the so called 'Destination floor control system' [12]. This system lets passengers designate the floor they want before the elevator car actually arrives. A kiosk then directs passengers to the elevator that will get them to their destinations with the shortest journey time. This results in an increasing handling capacity of the elevator by up to 30%. For buildings with heavy traffic at peak hours, this can mean dramatic reductions in lobby crowding and overall trip time. In fact, in certain applications, passenger journey times are often reduced by 25%.

1.3 Problem

In the previous section several papers about the elevator routing problem are discussed. There are many different approaches to model the elevator system: a simulation approach, modelling with dynamic programming and modelling as a polling system. Although there are some advanced systems nowadays, it is interesting to compare certain routing strategies for two elevators operating simultaneously. In the literature a lot is found about zoning. But it is still unclear what the effect of zoning policies in a small building is. How does this zoning influence the performance of the use of the elevators, will the waiting times decrease?

In the presented model two elevators are considered that are carrying passengers from the ground floor to higher floors in the building only; such situations are not widely studied. During the morning up-peak period this situation effectively occurs for a limited period. In the morning at offices, when all employees move to their work location, there is no or minimal inter-floor and down-traffic. Four different policies (three zoning policies and one without zoning) are compared on their performance. The following performance measures are considered:

- **The expected waiting time per passenger (EW)**
The expected waiting time per passenger is the time between the entering of the waiting area until the passengers enters the elevator, summed over all passengers and then divided by the total number of passengers considered.
- **The expected passenger journey time per passenger (ES)**
The expected passenger journey time per passenger is the time between the entering of the waiting area until the passengers leaves the elevator, summed over all passengers and then divided by the total number of passengers considered.
- **The average load of the elevators (N)**
The average load of the elevators is the average number of passengers (passengers in both elevators are summed) in the elevators during the simulation period.
- **The average number of passengers waiting (L)**
The average number of passengers waiting is the average number of passengers waiting in the lobby for an elevator during the simulation period.

The goal of this paper is to improve ES and EW by comparing several strategies. The other two performance measures are related to ES and EW. However, most important is that the passengers reach their work location fast and the passengers should have minimal complaints about waiting times.

1.4 Structure of the paper

The paper is organized as follows: in section 2 the model, that is used to solve the problem, is described. Section 3 gives insight to the way the model is implemented in a simulation environment. After the model is implemented, the results can be obtained. These are given in section 4. Analyzing the results leads to conclusions about the problem. These conclusions are discussed in section 5. The paper ends with some recommendations about further research in section 6 and the references in section 7.

2 Model

In this chapter the elevator problem is formulated as a mathematical model.

2.1 Elevator system as discrete-event system

The elevator system will be formulated as a discrete-event system. In this discrete-event system the events are the changes that are significant. These include requests for elevators at floors and the arrival of elevators at floors. A discrete-event systems (DES) consists of several key parts (Figure 1):

- Entities and their relations (logical statements)
- A simulation executive
- A central clock
- Random number generators
- Results collection and analysis

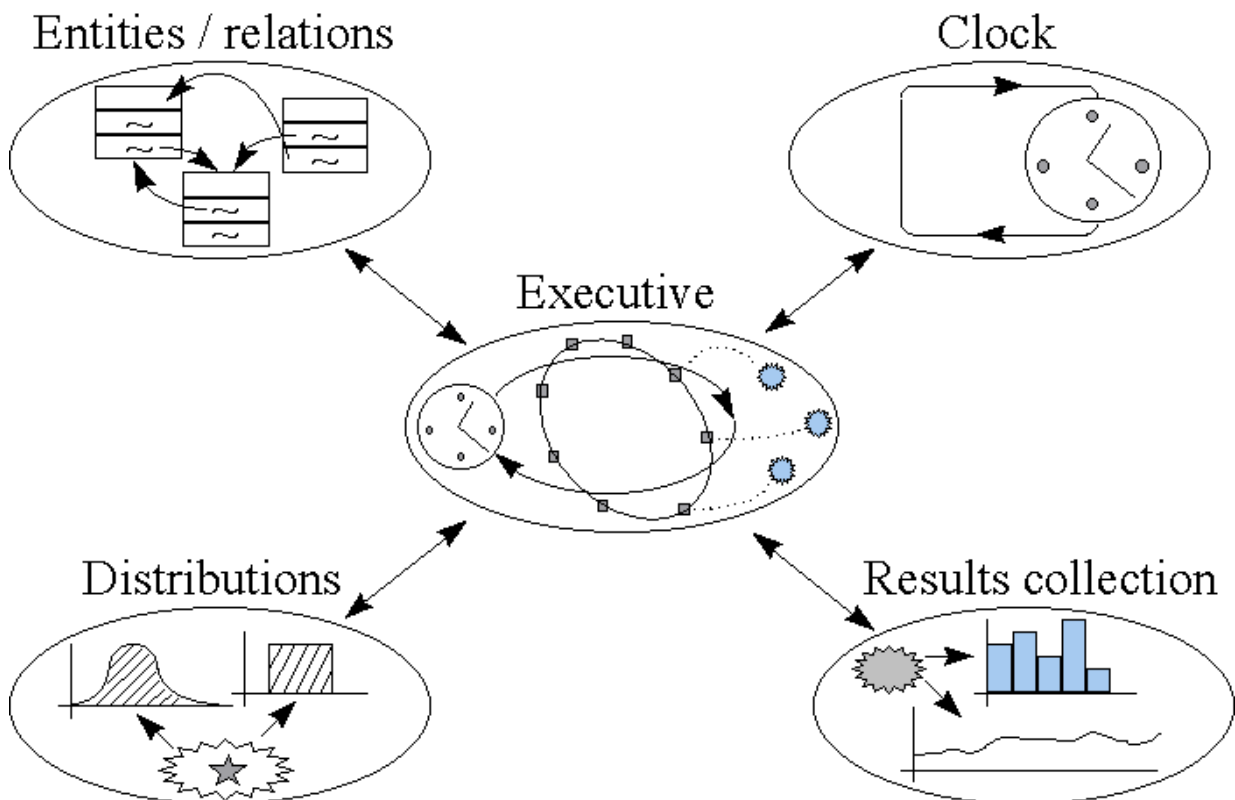


Figure 1. Structure of a discrete-event simulation system [5]

Entities are elements of a modelled system found in the real world. Entities can be either permanent or temporary. Temporary entities are entities which pass through the model, while permanent entities remain in the model throughout the simulation. The main objective of the simulation is to observe the behaviour of the temporary entities and collect information on them.

In the elevator system, the temporary entities are the passengers, who want to travel by elevator. The permanent entities are the elevators. The main objective of the simulation is to measure the performance of the passengers going through the elevator system. The relation between the entities is that the passengers use the elevators.

One of the central components of a discrete event simulation system is the simulation executive. It is responsible for controlling the logical relationships between the entities as well as the time advance. It provides the dynamic, time-based behaviour of the model.

However, in order to control the time advance, it is also necessary to have a clock in the system. This is, along with the simulation executive, one of the key parts of a DES. The central clock is used to keep track of time, and it is controlled by the simulation executive, which will advance the clock whenever necessary. There are two basic ways for controlling time advances. These approaches are:

- Next Event
- Time Slicing

In the Next Event case the model is advanced from the time of the present event to the time of the next event. It means that if nothing is going to happen in a certain period of time, the executive will move the model forward to the next event directly.

The Time Slicing mechanism differs somewhat from the Next Event mechanism. In the Time Slicing approach the model is forwarded in time at fixed intervals.

In our elevator system is chosen for a Next Event time advance.

All the elements mentioned above are key parts in a discrete event simulation. In addition, there are two other elements which are vital to any simulation system, including discrete event simulation: random number generators and results collection and analysis.

Random number generators provide stochastic behaviour for the model, defining a variation between certain ranges for every operation in the model. This makes it possible to mimic the operation of a real system being modelled very accurately.

In the elevator system the arrival process of the passengers at the ground floor lobby follows a Poisson process. And the destination of these passengers is uniformly distributed over the number of floors.

In order to get some meaningful information out of the simulation system, it is important to have the results collection and display features implemented in the system. These features provide meaningful analysis of the system being modelled.

In our elevator simulation system there are features implemented to obtain the performance measures.

The system is a stochastic simulation since there is a random component in the system, which is the arrival pattern of the passengers. (The behaviour of the elevators is deterministic, depending only on requests for service and scheduling policies. However, the model is driven by random input, in the form of the distribution of passengers. Thus the model decisions are seen as being deterministic, but the simulation itself as being stochastic, where we are interested in the use of elevators rather than the elevators themselves.)

2.2 Model details

There is chosen for a ground floor and four populated floors. More floors would make the model more and more complicated and four floors are enough to apply and compare some different policies. In the model only up-peak traffic is considered, so there is no down or inter-floor traffic. These last two categories of traffic would make the model too complex.

The passengers arrive at the ground floor approximately according to a Poisson process. This means that the inter-arrival times follow the exponential distribution, with the following probability density function: $\lambda e^{-\lambda x}$, where λ is the arrival rate. The destination of a passenger follows a uniform distribution. Every populated floor (1, 2, 3 and 4) has the same probability to be chosen.

The service times are assumed to be deterministic and only depend on the number of stops the elevator makes. The service time can be separated in three subparts:

- The doors open and passengers enter: 2 seconds,
- The elevator moves one floor up or down: 2.5 seconds,
- The elevator stops, passengers move out and doors close: 12.5 seconds.

These times are measured at the Vrije Universiteit in Amsterdam at Thursday, June 28. About fifteen observations were made to get an estimation of each of the service times.

There is no difference in travelling time between travelling two floors at once, or travelling two floors with a 12.5 second stop. It is both modelled as two times 2.5 seconds. In reality, travelling without stops will be faster because of some starting delay, but the model does not take this into account for modelling ease.

In the model, the maximum load of the elevators is infinite. This is not realistic; most systems have a load sensor in the elevator. The load sensor tells the computer how full the elevator is. If the elevator is near capacity, the computer would not make any more pick-up stops until some people have gotten off. Load sensors are also a good safety feature. If the elevator is overloaded, the computer will not close the doors until some of the weight is removed [2]. But in our model none of those features are implemented, because the number of passengers in an elevator will normally not be extremely high in the simulation runs.

A simulation is run till the performance measures converge. Because implementing of a stop function was difficult (some measures were calculated outside the simulation environment), the behaviour of the performance measures is analysed and in most cases after 250 000 time units (a time unit is equal to a second), the first two significant numbers of the performance measures become constant. So, actually the simulation is run for 250 000 time units.

2.3 Policies

The performance of two simultaneously working elevators is measured for four different policies. The two different elevators are denoted by elevator A and elevator B. The different policies are described below:

Policy 1 (no zoning)

Both elevators:

- The elevator stops at floor i , if the destination position of some passenger is i .
- The elevator serves the passengers in the order of direction: if there are passengers with destinations i and j and $i < j$, then the elevator stops at i before j .
- If all passengers in the elevator are served and the elevator is empty, the elevator changes direction and moves to the ground floor.

Policy 2 (zoning odd/even)

- Elevator A allows only passengers with an odd destination floor (1 or 3)
- Elevator B allows only passengers with an even destination floor (2 or 4)
- The elevator stops at floor i , if the destination position of a passenger is i .
- The elevator serves the passengers in the order of direction: if there are passengers with destinations i and j and $i < j$, then the elevator stops at i before j .
- If all passengers in the elevator are served and the elevator is empty, the elevator changes direction and moves to the ground floor.

Policy 3 (zoning low/high)

- Elevator A allows only passengers with an low destination floor (1 or 2)
- Elevator B allows only passengers with an high destination floor (3 or 4)
- The elevator stops at floor i , if the destination position of a passenger is i .
- The elevator serves the passengers in the order of direction: if there are passengers with destinations i and j and $i < j$, then the elevator stops at i before j .
- If all passengers in the elevator are served and the elevator is empty, the elevator changes direction and moves to the ground floor.

Policy 4 (nested zoning)

- Elevator A allows only passengers with a specific destination floor, namely the floor numbers 1 and 4.
- Elevator B allows only passengers with a specific destination floor, namely the floor numbers 2 and 3.
- The elevator stops at floor i , if the destination position of a passenger is i .
- The elevator serves the passengers in the order of direction: if there are passengers with destinations i and j and $i < j$, then the elevator stops at i before j .
- If all passengers in the elevator are served and the elevator is empty, the elevator changes direction and moves to the ground floor.

The policies are illustrated in more detail in the following schemes (Figure 2, 3, 4 and 5):

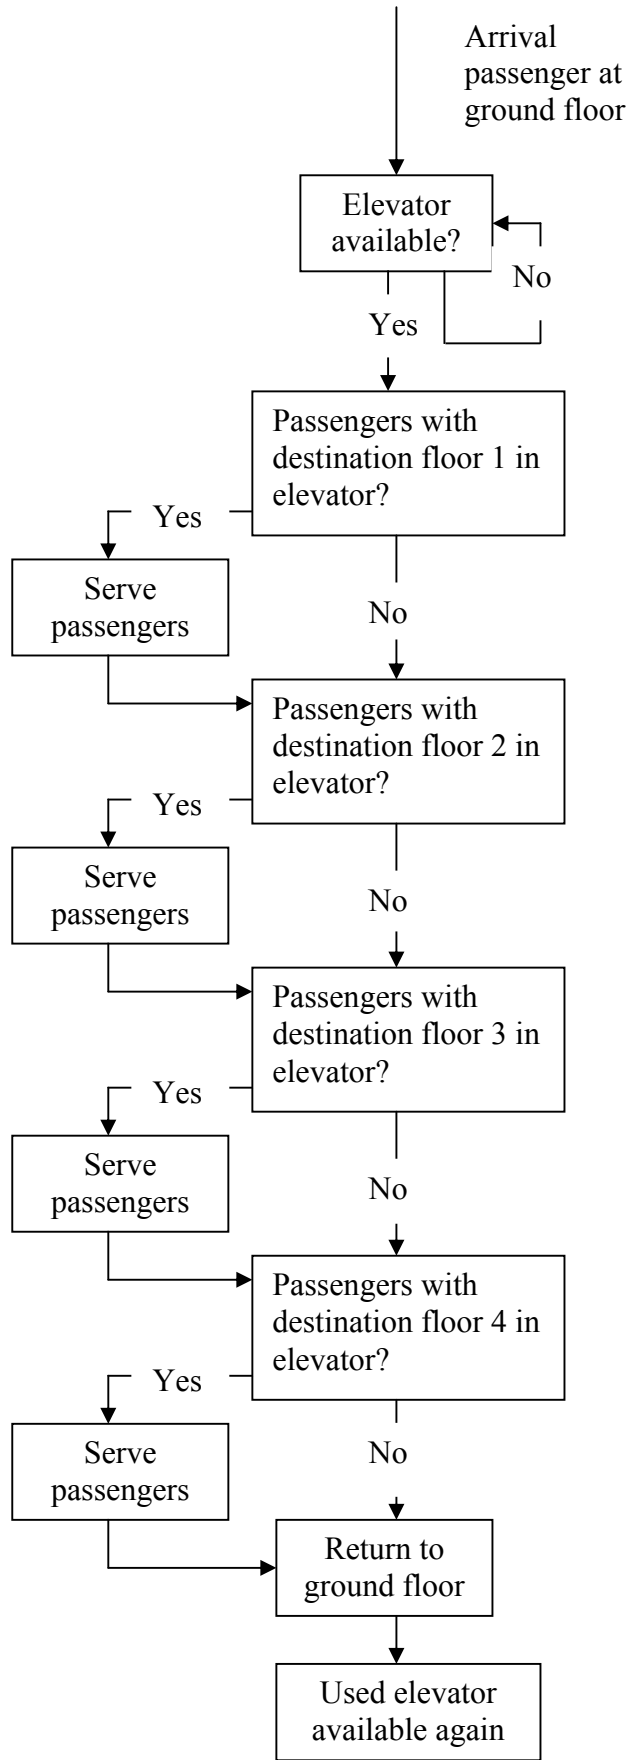


Figure 2. Policy 1 (no zoning)

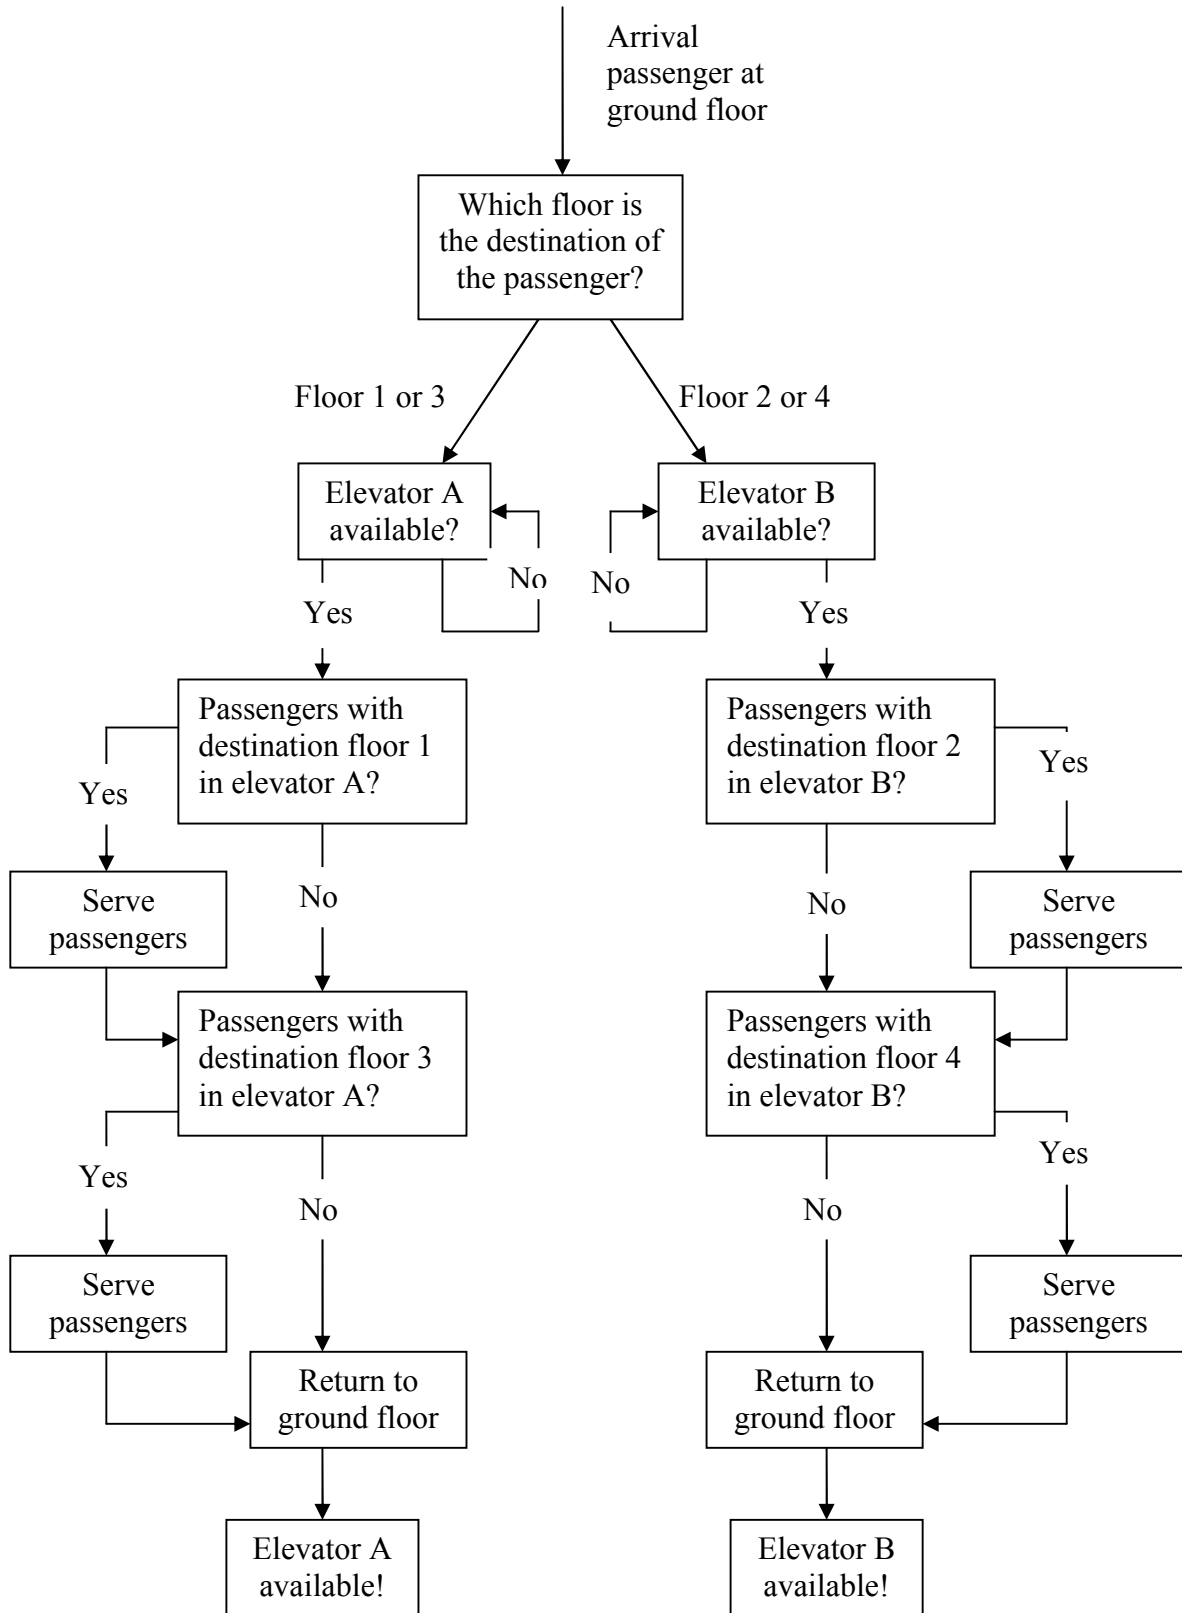


Figure 3. Policy 2 (zoning odd/even)

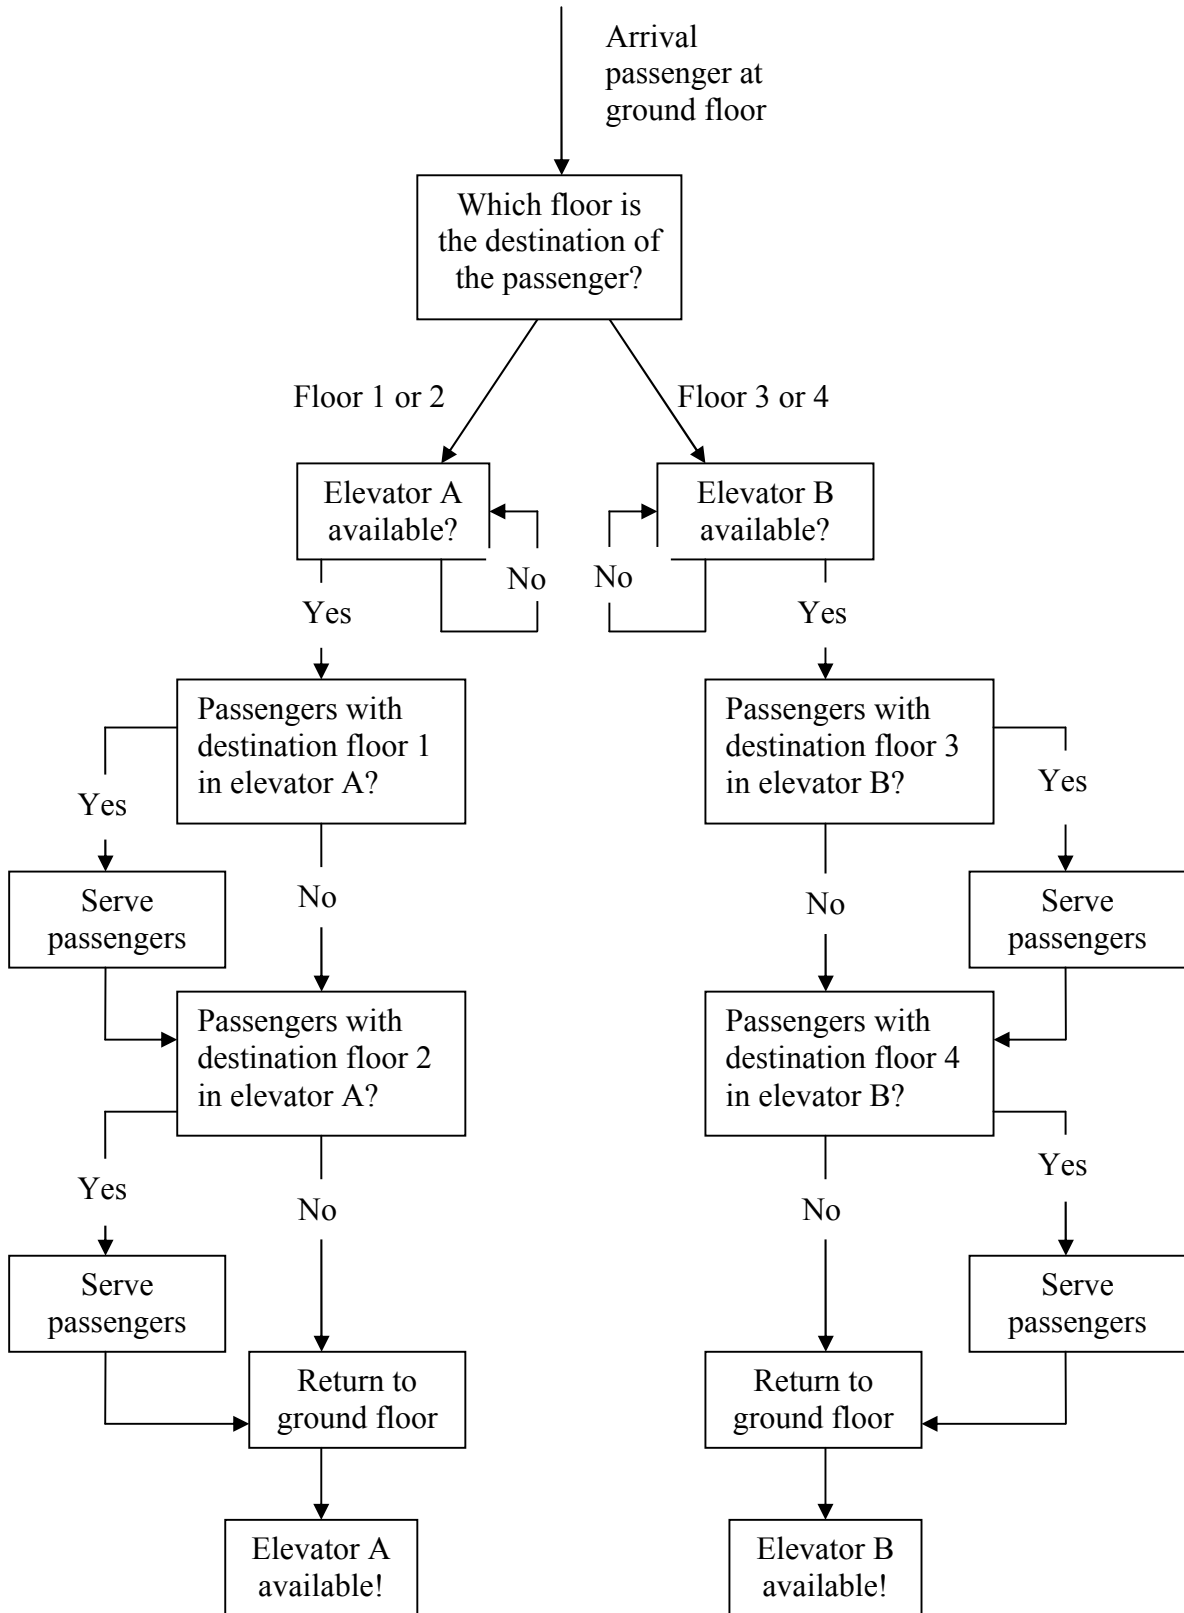


Figure 4. Policy 3 (zoning low/high)

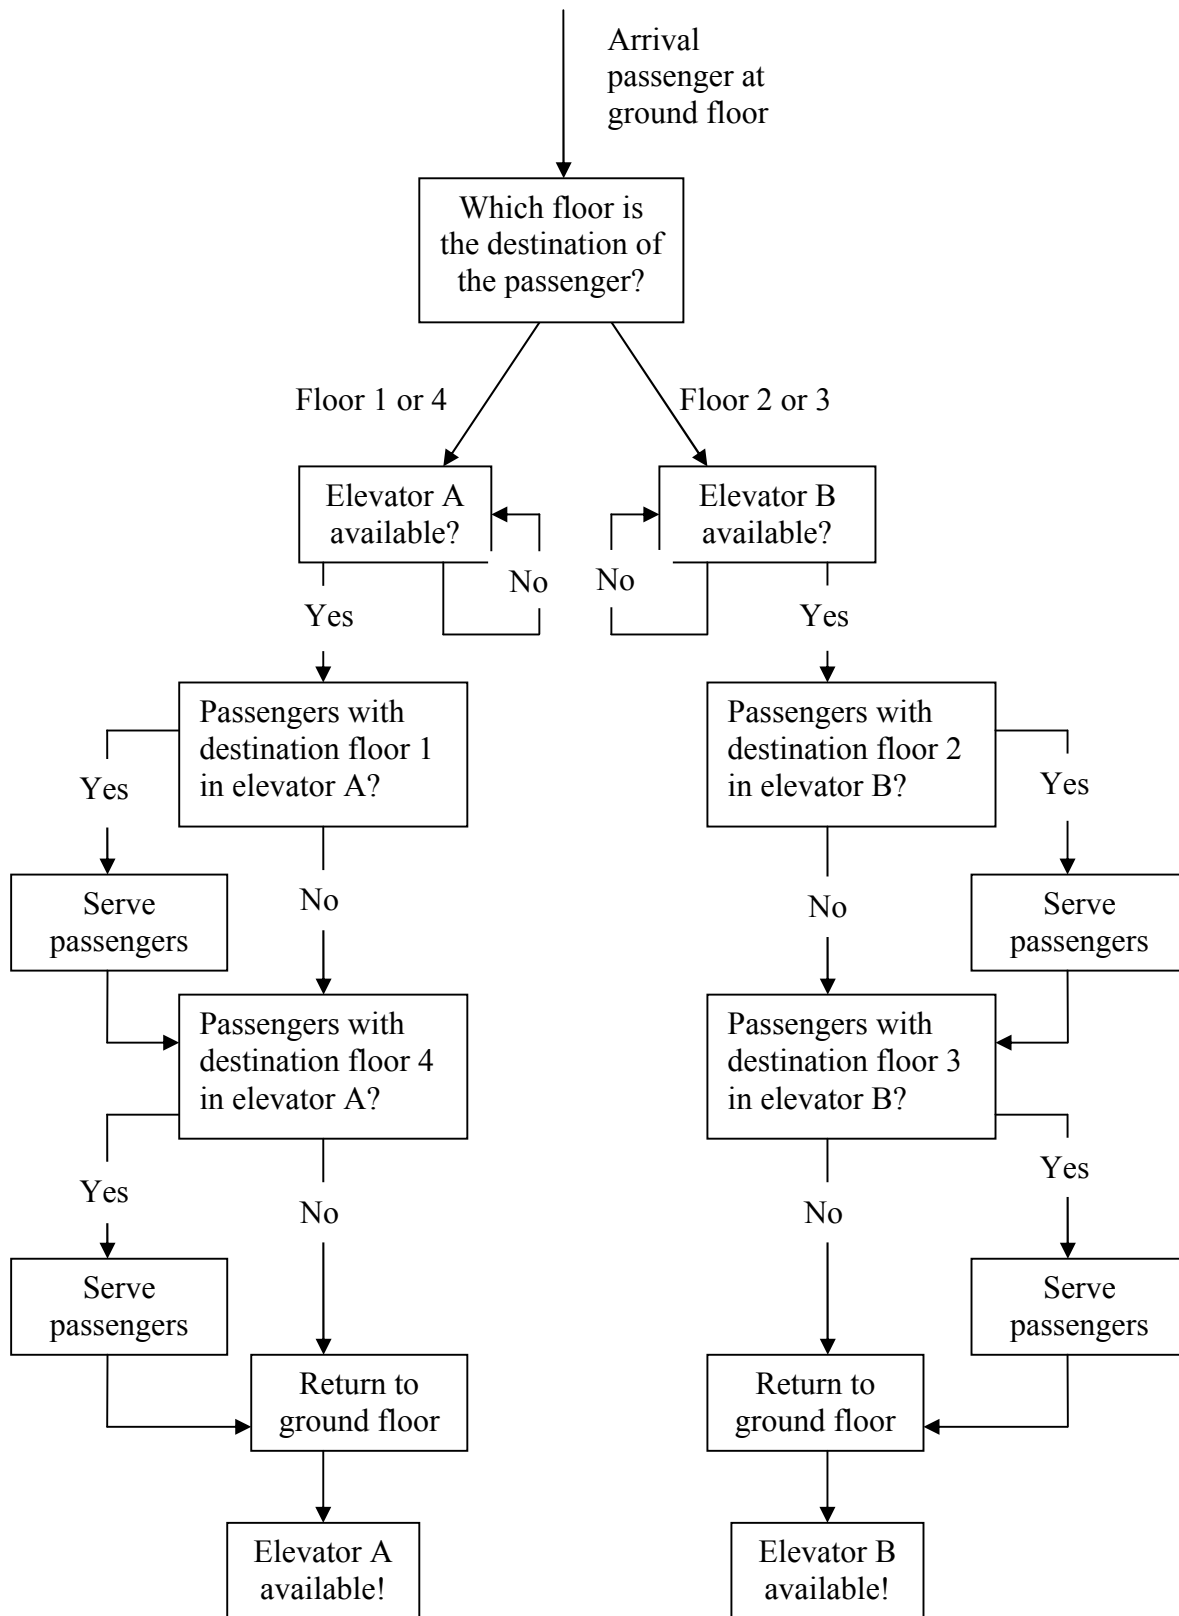


Figure 5. Policy 4 (nested zoning)

3 Implementation

The model as described in the previous section is implemented in a simulation environment called Extend (Version 5.0).

Extend (Imagine That, Inc.) is a simulation environment used to model, analyze, and optimize processes. It has a lot of features like libraries of components, hierarchies of models, linking with MS Office, and the ability to model continuous, discrete event, and hybrid systems. Extend has its own modeling language (ModL) which resembles C, and the ability to call code from other languages. It has specialized packages for Industrial Systems, Operations Research, and Continuous Process simulations.

The simulation program uses a generator that generates the arrivals of passengers at the ground floor lobby according to a Poisson process. An integer random number is drawn from a uniform distribution between the numbers 1 and 4. This number, which represents the destination floor, is assigned to the passenger. The passengers are grouped in four different queues, where each queue collects the passengers for one floor.

If an elevator is present at the ground lobby the passengers in the queues are served. After some time (the time the passenger needs to travel to its destination), the passenger leaves the system and the elevator becomes available. However, some delay time is needed, since the elevator needs to travel to the ground floor).

We let the simulation run till it converges (see section 2) and collect the results, which are the performance measures given in section 1.3.

Extend has an option called '*Show Animation*'. This option is very handy for validation of the model. The behavior of the passengers can be checked, so the programmer can easily see the model behaves in the way, the programmer wanted it to be defined.

A single run in the Extend simulation environment endures approximately 10 minutes. The two random components: the arrivals and the destination of the passengers use a certain seed in the simulation. The generator that generates the arrivals according to a Poisson process has got seed: 2 and the random number generator for the destination floors has got seed: 60.

Now that the model is implemented, it can be run with the four different policies described in section 2.3. But before the running, an arrival rate λ must be chosen. It is interesting to change this parameter and compare how the different policies react on the changes in arrival intensity. Three different arrival rates are chosen: high intensity ($\lambda = 0.4$), medium intensity ($\lambda = 0.133$ seconds) and low intensity ($\lambda = 0.04$ seconds).

A figure is made, to clarify how the model is implemented in the simulation environment, which is Figure 6.

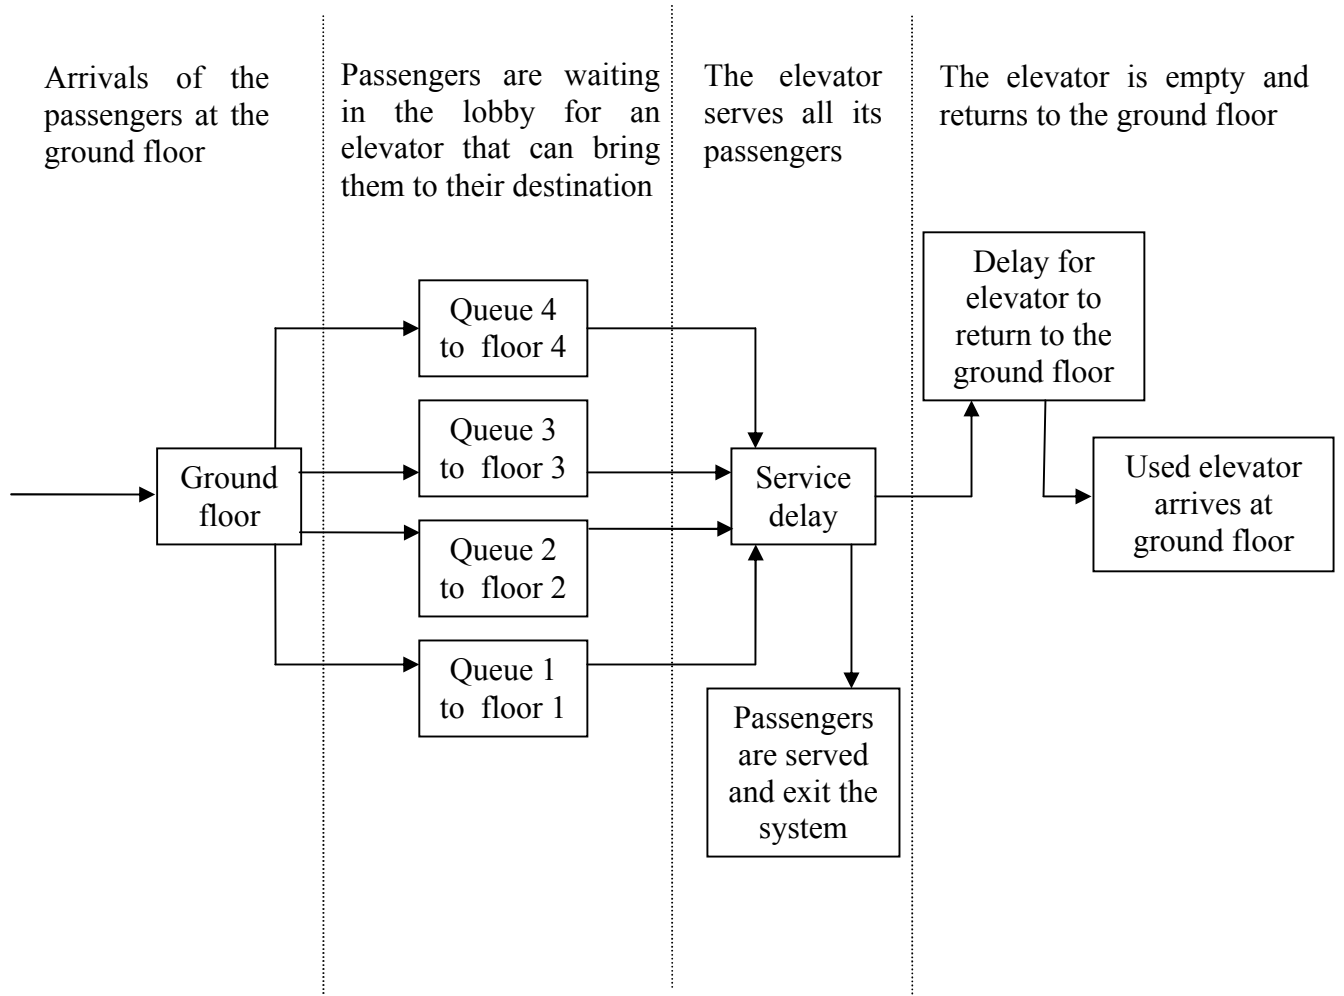


Figure 6. Scheme of the implementation

Figure 6 illustrates that first a passenger comes into the system. If there is not an elevator available that can serve the passenger, the elevator has to wait in a queue that belongs to its destination floor. If the elevator for that destination is available, the passenger is served. When all passengers are served and the elevator becomes empty, the elevator moves to the ground floor and after some time the elevator will be available for traveling.

4 Results

In this section the results of the model are evaluated. Runs are made with three different arrival rates: high intensity ($\lambda = 0.4$), medium intensity ($\lambda = 0.133$) and low intensity ($\lambda = 0.04$). The passengers arrive at the ground floor approximately according to a Poisson process. This means that the inter-arrival times follow the exponential distribution. An arrival rate λ of 0.4 seconds means that, on average, every second there are 0.4 passengers arriving to the ground floor lobby, i.e. on average the time between two arrivals is $1/0.4 = 2.5$ seconds.

The first performance measure that is considered is an important one: the expected passenger journey time (ES). The results of this performance measure are found in Table 1. The column 'Average' means the average over a high number of runs. The results of these runs are not the same, to give an impression of the spread of the results of the runs in the second column 'CI (95 %)' a 95 percent confidence interval is given.

	High Intensity ($\lambda=0.4$)		Medium Intensity ($\lambda=0.133$)		Low Intensity ($\lambda=0.04$)	
	Average	CI (95%)	Average	CI (95%)	Average	CI (95%)
Policy 1	44.316	[43.310; 45.322]	28.697	[27.478; 29.915]	11.366	[10.823; 11.909]
Policy 2	35.502	[34.993; 36.011]	27.247	[25.969; 28.525]	14.686	[13.902; 15.470]
Policy 3	34.035	[33.779; 34.291]	25.528	[25.165; 25.892]	14.554	[13.873; 15.235]
Policy 4	35.710	[35.450; 35.971]	27.322	[27.181; 27.464]	14.838	[14.594; 15.081]

Table 1. Expected passenger journey time (ES) in seconds

If the intensity of the passengers is very high ($\lambda=0.4$), policy 3 has the shortest ES. Policy 2 and 3 are not far away, they have slightly longer journey times. Policy 1 is significantly worse, than the others. If the intensity is medium, policy 3 is still the best. But policy 1 has a very high gain according to its performance with a high intensity of arrivals. When the intensity is very low, policy 1 even outperforms the other policies.

It is remarkable that all the policies behave in some logical way to each other (when the intensity changes), except policy 1. This policy has significantly shorter journey times when the arrival intensity is getting lower. This effect is shown in Figure 7. In the figure, the first bin, belonging to policy 1, is decreasing almost two times faster than the other ones.

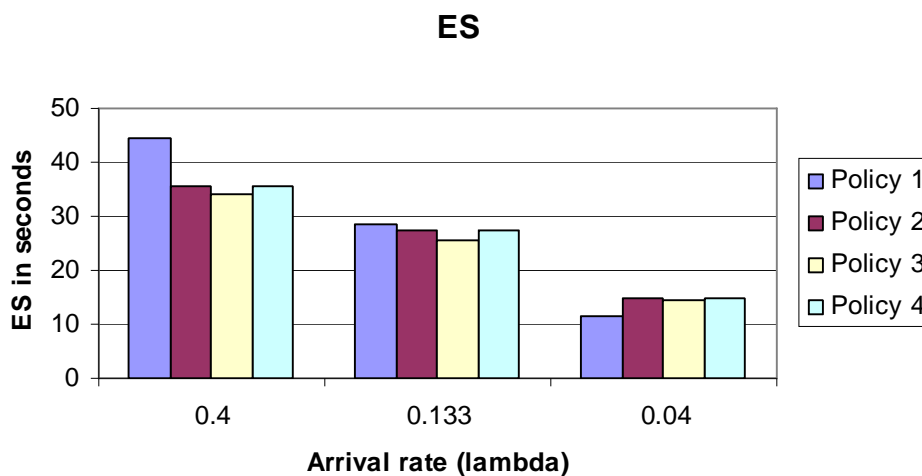


Figure 7. Expected journey time

The second performance measure that is discussed is the expected passenger waiting time. The results are given in Table 2.

	High Intensity ($\lambda=2,5$)		Medium Intensity ($\lambda=7,5$)		Low Intensity ($\lambda=0.04$)	
	Average	CI (95%)	Average	CI (95%)	Average	CI (95%)
Policy 1	20.809	[20.316; 22.106]	13.841	[12.709; 14.973]	4.346	[4.084; 4.608]
Policy 2	22.941	[22.781; 23.100]	16.868	[16.567; 17.169]	7.665	[7.407; 7.923]
Policy 3	21.623	[21.534; 21.712]	15.754	[15.475; 16.032]	7.470	[7.408; 7.533]
Policy 4	22.871	[22.480; 23.262]	17.066	[16.764; 17.368]	7.787	[7.689; 7.885]

Table 2. Expected passenger waiting time (EW) in seconds

Policy 1 has the shortest EW for each level of intensity. When the intensity decreases, the difference between the performance of policy 1 and the other policies is bigger. From the other three policies, policy 3 has the shortest EW. The following performance measure that is considered is the average load of the elevators.

	High Intensity ($\lambda=0.4$)		Medium Intensity ($\lambda=0.133$)		Low Intensity ($\lambda=0.04$)	
	Average	CI (95%)	Average	CI (95%)	Average	CI (95%)
Policy 1	8.732	[8.211; 9.253]	1.926	[1.827; 2.026]	0.276	[0.221; 0.332]
Policy 2	4.958	[4.781; 5.169]	1.295	[1.237; 1.353]	0.274	[0.224; 0.324]
Policy 3	4.927	[4.792; 5.061]	1.260	[1.188; 1.333]	0.271	[0.238; 0.304]
Policy 4	4.955	[4.847; 5.061]	1.299	[1.268; 1.330]	0.273	[0.254; 0.292]

Table 3. Average load of the elevators (N) in number of passengers

Table 3 shows that with a high intensity of arriving passengers, policy 1 has the highest average load of all policies, with a load between 8 and 9 passengers. The other policies perform almost the same with this high intensity and have on average a load of almost 5 passengers over the two elevators. When the arrival intensity is medium, the performance pattern of the policies is approximately the same. Policy 1 has still the highest average load, but the gain according to the performance in the situation with high intensity is somewhat less. If the average inter-arrival time of passengers is 25 seconds (arrival rate 0.04), the average loads are all at the same level. Figure 8 shows these results in a graphical way.

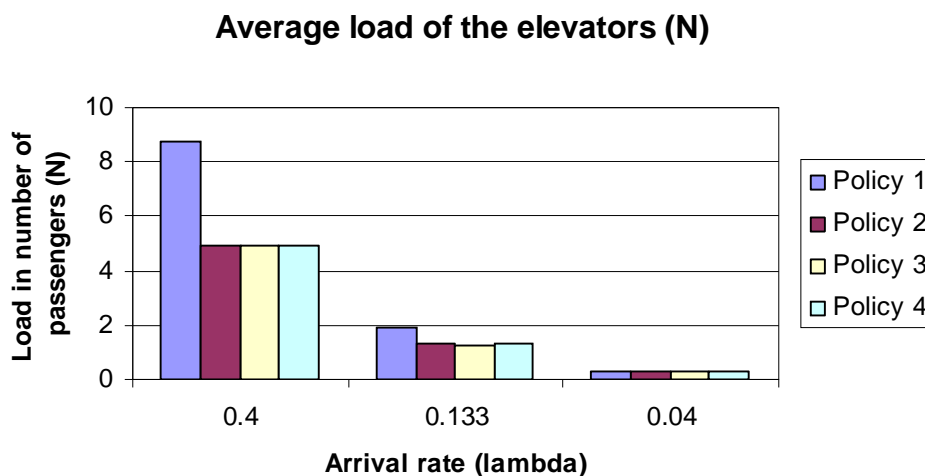


Figure 8. Average load of the elevators

The last performance measure that is discussed is the average number of passengers that are waiting at the ground floor lobby. The results of this performance measure for the three different arrival rates are given in Table 4.

	High Intensity ($\lambda=0.4$)		Medium Intensity ($\lambda=0.133$)		Low Intensity ($\lambda=0.04$)	
	Average	CI (95%)	Average	CI (95%)	Average	CI (95%)
Policy 1	7.349	[6.554; 8.145]	1.556	[1.371; 1.741]	0.097	[0.059; 0.135]
Policy 2	8.396	[8.082; 8.733]	1.944	[1.761; 2.127]	0.231	[0.160; 0.301]
Policy 3	7.868	[7.497; 8.238]	1.795	[1.691; 1.898]	0.224	[0.157; 0.291]
Policy 4	8.393	[8.303; 8.462]	1.970	[1.913; 2.027]	0.232	[0.205; 0.258]

Table 4. Average number of passengers waiting (L) in number of passengers

These results are somewhat similar as the results for the EW. When the elevators are routed according to policy 1, the littlest passengers are waiting at the ground floor lobby on average. Policy 3 has the best performance of the other policies again.

Policy 1 seems to have the best performance and reacts in a somewhat different way to a change of intensity than the other three policies. In the following part of this section, the results are discussed somewhat deeper and the results are tried to clarify the results with a decent argumentation.

In policy 1, both elevators serve all passengers, so they can stop on every floor. An advantage of this strategy is that passengers with a certain destination floor can use both elevators; this can reduce the waiting time. But a disadvantage is, that the elevator makes more stops and a stop of an elevator costs 12.5 seconds, five times the time that is needed for traveling one floor, this can cause longer expected passenger journey times. With this reasoning the results can be explained.

If the intensity of arrivals is high, there are high loads of passengers in the elevators. When there are a lot of passengers in the elevator, there is a huge probability the elevator has to stop frequently. A stop is a time consuming feature and because policy 1 serves all the floors with both elevators, it performs worse in a high intensity situation. But if there are not many passengers arriving, the elevators are often not in use. And with policy 1, one free elevator can always bring a passenger to its destination, which is not the case for the other three policies. If there is an elevator available for traveling in policy 2, 3 and 4, it can be an elevator that does not serve a certain floor the passenger wants to travel to. So in these policies the passengers sometimes have to wait even when there is an elevator available at the ground floor.

The final conclusion is that policy 1 behaves somewhat different to changes in the arrival rate parameter than the other policies because both elevators are serving all floors. The ES of this policy is long, when there are many passengers in the system, but if the intensity of arrivals is not that high, this policy benefits from the feature that passengers with all sorts of destination floors can use both elevators. This feature also causes shorter EW and L for policy 1. The loads are bigger in policy 1, because the elevator trip times are longer in this policy (in policy 1 an elevator can make four stops at maximum, in the other policies 2 stops is the maximum). If trip times are longer, there are more passengers waiting at the lobby and EW and L will increase. This clarifies the behavior of EW and L for policy 1 with different arrival rates. L and EW are shorter because in policy 1, both elevators serve all passengers with all kind of destination floors. But when the intensity is high, trips are longer and L and EW increase. It

can be concluded that policy 1 is getting better L and EW (according to the other policies), because both elevators can be uses by all passengers. But this effect shrinks when the intensity is getting higher, because the trips are getting longer.

To give another insight in the differences of performance between policy 1 and another policy, some graphs are shown in Figure 9. In these graphs the cumulative distribution function of ES and EW are given for policy 1 (low and medium intensity) and for policy 2 (low and medium intensity). The results are based on a single run in the Extend simulation environment.

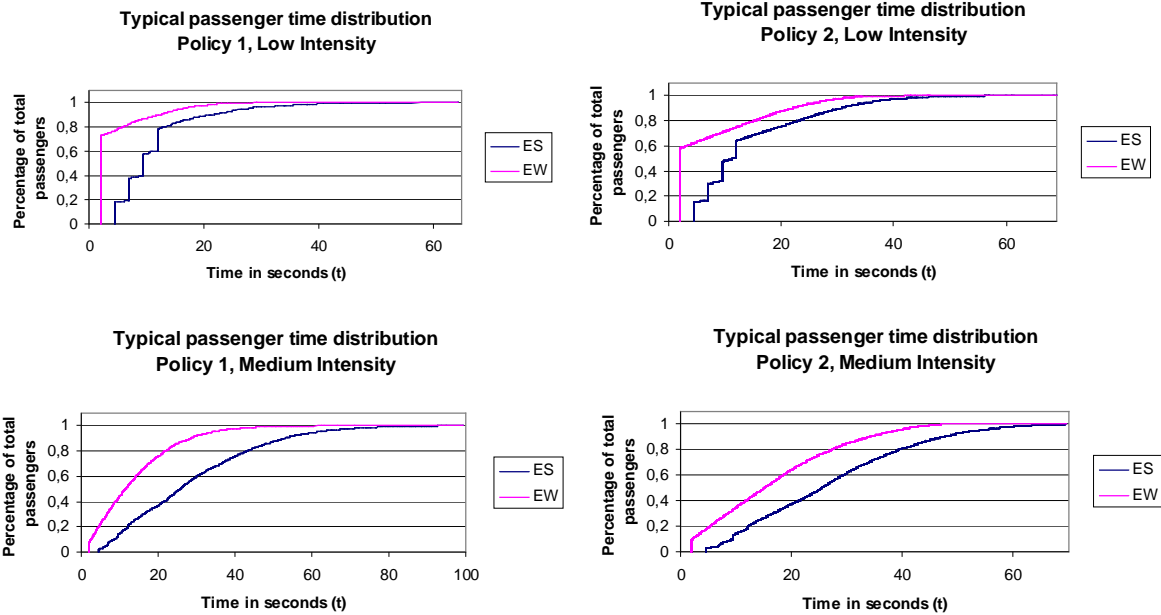


Figure 9. Distribution functions of ES and EW in four different situations

First have a look at the expected passenger waiting time. When the intensity of arriving passengers is low, there are a lot of passengers that only have to wait 2 seconds for the doors that open. The graphs display that for policy 1 this group is bigger than for policy 2. In policy 1 around 70 percent of all passengers do not have to wait before an elevator arrives, and in policy 2 this percentage is around 60 percent. Furthermore almost all passengers (97.5 %) have EW shorter than 20 seconds, in policy 2 this result is around 30 seconds. This is the consequence of the fact that in policy 1 both elevators can handle every passenger and in policy 2 destination floors can not be reached with every elevator.

In a situation with a medium arrival rate, the elevators are rarely idle. Expected waiting times of 2 seconds only cover around 10 percent of the passengers. Policy 1 still has shorter EW, 80 percent of all passengers have EW shorter than 20 seconds and in policy 2 a percentage of 80 percent is achieved when passengers with EW shorter than 25 seconds are taken into account.

Then have a look at the expected journey time. From previous results it is known that policy 1 performs best when the intensity is low, and the results of the performance measures are getting worse when the intensity is higher. Figure 9 also shows this effect. With a low intensity of arriving passengers, in policy 1 90 percent of the passengers have EW shorter than 22 seconds and in policy 2 this 90 percent is reached when passengers with EW till circa 31 seconds are take into account. Thus with a low intensity, policy 1 performs better.

But when the ES of a medium intensity situation are considered, policy 1 is getting worse. In this situation, 90 percent of the passengers in policy 1 have ES shorter than 53 seconds. In policy 2 this number of 90 percent is already reached at 48 seconds.

The same behavior can be seen, when the maximum EW and ES for policy 1 and 2 are considered (the same runs as discussed above). The rounded results are given in Table 5.

	Max ES		Max EW	
	Low Intensity	Medium Intensity	Low Intensity	Medium Intensity
Policy 1	64.5	99.4	29.2	71.1
Policy 2	69.0	69.5	46.5	47.0

Table 5. Maximum ES and EW

When the intensity of arrivals is low, policy 1 has the lowest maximum, but when the intensity is medium, policy 2 has a lower maximum. The two maximum EW at medium intensity are not very surprising.

In policy 1, the worst time a passenger can arrive is the moment that two elevators have just left and both carry passengers with all kind of destination floors (1, 2, 3 and 4). Then a passenger have to wait four times a elevator stop time (four times 12.5 seconds), plus the time for traveling four floors up and four floors down (eight times 2.5 seconds), plus the time for opening the doors (2 seconds). Summed, this gives a limit of 72 seconds. But the probability that passengers with four different destination floors arrive in a split second (which is needed to approach the limit of 72 seconds) after the first elevator moves up, is minimal. For this reason the maximum EW is somewhat shorter than the limit.

In policy 2, the worst time a passenger can arrive is the moment that the elevator that can bring the passenger to its destination is just moving upwards and that elevator is carrying passengers with destination floors 2 and 4. Then a passenger has to wait for two 12.5 second stops, and traveling four floors up and four floors down. After this the passenger has to wait another 2 seconds for opening the doors, this makes its EW amount 47 seconds. The maximum EW of policy 2 in the considered run approaches this limit.

Overall is concluded that policy 1 performs the best of all policies, except for ES (in a high and medium intensity situation), which is not an unimportant performance measure. The best zoning policy is policy 3. In this policy, elevator A serves the two low floors (1 and 2) and elevator B serves the higher floors (3 and 4). In this policy, elevator A has short ES, because the highest floor this elevator serves is floor 2. Due to this fact, for passengers that have low destination floors there is often an elevator available at the ground lobby. This can cause the slightly better performance than the other zoning policies. The other two zoning policies (2 and 4) are performing somewhat the same, the averages are in each others confidence intervals (most of the time). Elevator A and B of policy 2 performs approximately the same respectively to elevator B and A of policy 4. This is shown in table 6 for a medium intensity situation, where the average number of waiting passengers is displayed. For policy 1 only one value is given because both elevators are routed in the same way, so passengers do not wait for a specific elevator.

	Elevator A	Elevator B	Total L
Policy 1	1.56		1.56
Policy 2	0.86	1.07	1.94
Policy 3	0.66	1.13	1.80
Policy 4	1.07	0.90	1.97

Table 6. Average number of passengers waiting (L) per elevator ($\lambda=0.133$)

In the last part of this section there is look to reality. In real buildings there are some stairs besides the elevators. Some of the passengers travel up stairs instead of using the elevator. Although most passengers are commonly very lazy, so traveling by stairs does not occur very often. We now assume that 50 percent of the passengers that have to travel to the first floor walk one floor up stairs. All the other passengers always use the elevator. So in this situation, passengers that travel by an elevator that can stop at floor 1 will have a shorter EW, because the probability that another passenger wants out at floor 1 is somewhat smaller.

Table 7 shows the effect of using the stairs occasionally in a medium intensity situation. The gain on EW that is given in the table means the average reducing of EW of the passengers that travels by elevator. The gain is according to the situation that all passengers are using the elevators.

	Gain on EW
Policy 1	10,1 %
Policy 2	4,1 %
Policy 3	1,3 %
Policy 4	7,6 %

Table 7. Gain on EW, when using stairs occasionally ($\lambda=0.133$)

Policy 1 and 4 has the biggest gain on EW, when 50 percent of the passengers with destination floor 1 will take the stairs. This is not surprisingly, because in these policies, passengers that use elevators which stop at floor 1 have long ES. However, this example shows another benefit for using a no zoning policy in low buildings because in this more realistic type of model, policy 1 performs even better.

5 Conclusion

In the previous section the results are analysed and evaluated. The questions from section 1.3 can be answered. Which policy performs best at the four different performance measures? And what is the effect of zoning on these performance measures in a building with four floors?

Policy 1, the policy without zoning, performs the best at all the performance measures, except for ES in a high and medium intensity situation. But performance measure ES is one of the important performance measures, since the time that is needed for travelling and waiting is lost time. Reducing ES will please the passengers.

But is this bad performance on ES a reason to route elevators according to a zoning principle? It is unknown, what the real intensity of arrivals is at the ground floor lobby. If the intensity of arrivals is very low, policy 1 performs the best at all the performance measures, even ES. But such a low intensity is not very realistic. A medium intensity or high intensity situation is more realistic. In these situations the elevators are rarely idle. In the morning up-peak at real offices this is also the case.

Another remark is about EW. ES is an important performance measure, but so is EW. Policy 1 performs the best with all intensities on this performance measure. Unlike the passenger journey time is the most important measure according to the wasting of time issue, the passenger waiting time is the most important performance measure for the mood of a passenger. A passenger wants to be in the system and hates waiting, so in this perspective EW is more important than ES. In that case, policy 1 is recommended to use for the elevator routing system.

In reality there are some other advantages and disadvantages for using a zoning policy. A disadvantage of the zoning policies is the fact that a specific elevator does not serve all the floors. It might be possible that an arriving passenger enters a free elevator, and then concludes that it does not go to the wanted floor. So the passengers of the elevator system would need to be educated in its use, but since the passenger population in offices is relatively fixed this should not present a problem.

Another drawback of using a zoning policy for elevator routing is that it is only performing better in peak-situations. Outside these periods, there is light traffic and in such situations a policy without zoning is better. Therefore, if a zoning policy is used for the elevator routing system, it is only used in the peak traffic and in the other periods, the elevator must have another routing system. This will lead to more costs.

In our model it is suggested that there is no down-traffic and inter-floor traffic. In real situations, these two sorts of traffic are present, even in a morning up-peak situation. When one of the three zoning policies is implemented in the elevator routing system, some floors can not be reached from a certain higher floor than the ground floor. In policy 2 for example, elevator A only stops at floor 1 and 3, and elevator B only stops at floor 2 and 4. So when a passenger is located on floor 1 and the passenger wants to travel up to floor 4, the passenger can not use an elevator, because the elevator that stops at floor 4 do not stop at floor 1.

But in reality a passenger can reach floor 4 from floor 1, because in all buildings there are some stairs besides the elevators. However, this is not very user-friendly. There are

passengers that are not or barely able to walk up and down stairs. Nevertheless, if a passenger is able to use the stairs, it can reduce waiting time. In policy 2 for example, there can occur the following situation: a passenger wants to travel to floor 4 and the elevator that stops there, is busy, but the other elevator is available for use. Then the passenger can enter this other elevator, travel up to the nearest floor (floor 3) and walk the last floor up stairs.

When the use of stairs is taken into account, it is also showed that this leads to lower EW, because fewer passengers are travelling by elevator in a situation where some passengers walk up stairs. The gain of EW is the biggest for the policy 1, so in this more realistic model, a policy without zoning is performing even better.

Finally it is concluded that the benefit of using zoning policies in such a low building, as a building with four floors is, is not enough to prefer them above a policy without zoning. Although there is a gain of 10 seconds in ES in the high intensity situation, there are too many drawbacks for using zoning policies for the elevator routing system in a small building, with a small number of floors.

The effect of using zoning policies is probably bigger, when the number of floors becoming higher. If there are more floors and a policy without zoning is used, the elevators should have make on average even more stops, so the ES will become larger. This is also mentioned in the paper of So and Yu [\[10\]](#) about dynamic zoning.

6 Further research

This paper gave interesting insights in the effects of zoning in a building with a small number of floors on performance. However, it is advised to further investigate in extending this model. Hereby is thought of higher buildings, models that includes inter-floor traffic and down traffic, extension of the model with stairs. This might lead to an even better argumentation for the optimal policy. The modelling assumptions, e.g. estimated service times, are not certainly realistic. It is advised to measure these times in a more exact way or do some sensitivity analysis on the results, by changing the service times.

7 References

- 1 Chu, S.C.K., Lin, C.K.Y, Lam, S.S. 2003. In European Journal of Operational Research: Hospital lift system simulator: A performance evaluator–predictor.
- 2 Harris, T., 2007. How Elevators Work. <http://science.howstuffworks.com/elevator7.htm>.
- 3 Maister, D.H., 1985. The psychology of waiting lines. <http://davidmaister.com/pdf/PsychologyofWaitingLines751.pdf>
- 4 Nikovski, D., Brand, M., 2003. Decision-Theoretic Group Elevator Scheduling. <http://www.merl.com/papers/docs/TR2003-61.pdf>
- 5 Ruhonen, T, 2007. Improving the operation of an emergency department by using a simulation model. <http://dissertations.jyu.fi/studcomp/9789513928568.pdf>
- 6 Sasser, W.E., Olsen, J., Wyckoff, D.D., 1979. Management of Service Operations: Text, Cases and Readings. New York: Allyn and Bacon.
- 7 Shearn, D.C.S., 1983. Optimum limited stop scheduling for a single lift. Journal of the Operational Research Society 34 (1).
- 8 Shoham, R, Yechiali, U., 1992. Elevator-type polling systems. ACM SIGMETRICS Performance Evaluation Review Volume 20 , Issue 1
- 9 Siikonen, M.L., 1997. Planning and Control Models for Elevators in High-Rise Buildings. <http://www.sal.hut.fi/Publications/pdf-files/rsii97b.pdf>
- 10 So, A.T.P, Yu, J.K.L., 2001. Intelligent supervisory control for lifts: dynamic zoning. Sage journals online.
- 11 Tervonen, T., Hakonen, H., Lahdelma, R., 2006. Elevator planning with stochastic multicriteria acceptability analysis. www.tucs.fi/publications/attachment.php?fname=TR590.ps.gz
- 12 2005. Destination Dispatch. ThyssenKrupp Elevator Corporation. http://www.thyssenkruppelevator.com/images/brochure_images/pdf/TKE_DestDisp.pdf.
- 13 2007. Elevator. Wikipedia. <http://en.wikipedia.org/wiki/Elevator>.