Aalto University

MS-E2177 - Seminar on Case Studies in Operations Research D

Simulating and Optimizing a Multi-Car Shaft Elevator System
Final Report

Case Group MLS Lift
Mikael Tunturi (Project Manager)
Daud Abucar
Yunseon Lee
Emma Tonteri

June 13, 2025

Contents

1 Introduction

1.1 Background

This is a project done in collaboration with MLS Lift Consulting, for simulating and optimizing a multi-car elevator system.

MLS Lift Consulting is a company founded in 2021 by CEO Marja-Liisa Siikonen, who is our contact person. The company offers consultation on people flow in buildings for customers such as architects and construction planners.

Modern-day buildings are increasing in height: the tallest building in 2004, Taipei 101 in Taiwan, was 508 meters tall. But just five years later, in 2009, Burj Khalifa in the United Arab Emirates, stood at an impressive height of 828 meters. Although a large part of the height is so-called vanity height (meaning the height above the highest usable floor), the height trend is pointing upwards. And with taller buildings comes more people.

Tall buildings mean high people flow and a high need for elevators. The height of modern-day buildings imposes restrictions on elevators: a single elevator cannot span the whole height of a really tall building, as the weight of the ropes of the elevator would be too heavy. This leads to elevator groups serving different sections of the building and to an increased surface area required for the elevators. Because of this, the floor plans of modern-day buildings are quite inefficient: up to half of the usable core space is consumed by elevators [1].

In order to tackle this inefficiency, multi-car elevator (MCE) systems have been designed. In MCE systems, multiple cars move in the same shaft. Some designs include moving in horizontal direction as well, in addition to the traditional vertical. These systems require thorough design and planning, as the risk of collisions is present and the safety of elevators is non-negotiable. The development of MCEs is still mostly theoretical, as currently there are no MCEs in use in commercial buildings, apart from the old Paternoster design.

Paternoster is an elevator where the cars move in two connected and adjacent shafts. The cars do not have doors and they do not stop at any floor, meaning that the user has to step in and out of the moving car. Paternosters are efficient and have been common especially in Western Europe, but the construction of new Paternosters has been banned due to safety reasons [2]: they pose a risk to old and disabled users, and there have been multiple fatal accidents.

The aim of the project is to examine the efficiency of MCE systems, by simulating and optimizing an MCE system using a traditional Paternoster elevator in three different steps, which will be discussed in detail in the following section. In addition to developing the model, the project has contained a visit to the Stockmann department store in Helsinki city centre and to the KONE main building in Keilaniemi, Espoo.

In Section 2 we will go through literature review, and in Section 3 the methods are presented. Results are stated in Section 4 and discussed in Section 5, while Section 6 sums up our conclusions. Finally, self-assessment is done in Section 7, followed by references.

1.2 Objectives

As stated in the introduction, the aim was to implement three different simulators. In the first step, a traditional Paternoster elevator is simulated. In the second simulation doors are added to the elevator, and it has to stop for a passenger to be able to board the car. In addition, every car can accommodate multiple passengers instead of one. In the third and final simulation, the cars can move independently in the elevator shaft, answering to calls on their own.

During the course, the main objectives for our project changed from simulating three multi-car elevator systems in Python to simulating just two. The advanced simulation in the project report was eliminated. In both simulators, three key parameter metrics, average waiting and journey time, and the number of passengers handled, are used to track and evaluate their performance. It is assumed that passengers can only enter from the lobby, and not on other floors.

In a conventional paternoster, the three key parameter metrics are calculated, not optimized. This simulator does not consider stopping or the opening and closing times of the doors. As a result, when passengers get off at their destinations, it does not impact the round-trip time (RTT) of the continuously moving elevator. It also means that passengers leaving the elevator will not affect the waiting times of passengers in the lobby. Therefore, adjusting the number of passengers handled in a continuously moving elevator to reduce waiting times is meaningless. As a result, no optimization can be applied to these three key parameters as there is no correlation between them.

In Simulator 2 with doors, the RTT of elevators increases with the number of destinations. This is because each additional destination adds extra time for opening and closing the doors. Consequently, when elevators stop at multiple destinations, the waiting time for passengers in the lobby increases. To address this issue, adjusting the number of passengers handled can help reduce waiting times. As a result, optimizing the number of passengers can lead to improvements in both average waiting time and RTT while maintaining the elevator HC as high as possible.

After the optimization step, both simulators would be run in up-peak morning traffic and the results would be visualized. To be specific, we use the default value for C for simulator 1 and the optimized value for C for simulator 2 when running both simulators to obtain the results.

Glossary

C Car capacity.

HC Handling capacity.

H Average highest reversal floor.

L Number of elevators.

M Average number of passengers inside the elevators.

N Number of floors.

RTT Round-trip time, see t_{rt}.

S Average probable number of stops.

t_p Average one-way passenger transfer time.

t_{rt} Round-trip time of the system.

t_s Time consumed in stopping.

t_v Time to travel between two standard pitch adjacent floors at a rated speed.

2 Literature Review

This section examines earlier research and theoretical foundations that inform the design and optimization of multi-car elevator systems.

Two of the three key parameter metrics, the average waiting time and the number of passengers handled, are used to determine the efficiency of elevator systems [3]. In [3], these two values are adjusted by the five parameters in [3], that is,

1. number of cars,
2. passengers per car and its moving speed,
3. proper control systems,
4. the layout of the elevators,

5. scheduling of transporting the passengers to their destinations.

In this project, parameters 1 and 4 would remain the same, while parameters 2, 3 and 5 would be adjusted to optimize the number of passengers handled for Simulator 2. In order to adjust parameters 2 and 5, we aim to analyze passenger traffic, where we look at their distribution and how they move to their destinations during a single workday.

Up-Peak Traffic

In particular, we focus on analyzing up-peak passenger traffic, defined as the movement of passengers from the lobby to their destinations in [3], especially during the morning. The reason being that this is when the handling capacity of the system, HC, is predicted to have reached its maximum. As a result, we assume that if Simulator 2 can be optimized to handle morning up-peak passenger traffic, it can be easily adapted to other situations as well.

In this project, HC is the number of passengers handled in five minutes. This is calculated using the formula in [3], given as

$$HC = \frac{0.8 \times 300 \times L \times C}{t_{rt}}, \tag{1}$$

where L is the number of elevators, t_{rt} is the round-trip time of the system, and C is the car capacity, given the same as in Section 1.2. The number 0.8, is the optimal utilization factor of elevator systems.

Round-Trip Time

According to [4], the round-trip time is the average period of time for a single elevator car to travel during up-peak traffic a full cycle so that the main terminal (lobby) is reached again after serving the registered car calls.

The formula for the most common round-trip time (RTT) is given as follows [1]:

$$t_{rt} = 2Ht_v + (S+1)t_s + 2Mt_p, \tag{2}$$

where H is the highest reversal floor, S is the probable number of stops, t_v is the time to travel between two standard pitch adjacent floors at rated speed (s), t_s is the time consumed in stopping (s), M is the average number of passengers inside the elevators, and t_p is the average one-way passenger transfer time (s).

Here, S is calculated using the formula in [1], given as

$$S = N\left[1 - (1 - \frac{1}{N})^M\right], \tag{3}$$

where N is the number of floors.

For continuously running paternosters, equation (2) simplifies to

$$t_{rt} = 2Ht_v. \tag{4}$$

Handling Capacity and Interval

According to [4], the handling capacity is the maximum sustainable number of passengers that a single lift or a lift group can transport in a specified period for a specific traffic mix under specified loading constraints.

The interval, on the other hand, is the average time between successive car departures from the main entrance floor in up-peak. The interval is calculated using the formula in [1], given as

$$t_{int} = \frac{t_{rt}}{L}, \tag{5}$$

where L is the number of elevators and t_{rt} is the round-trip time.

Destination Control vs. Conventional Control

According to [1], the conventional control system has a disadvantage during up-peak periods that all waiting passengers step to the first stopping car and passengers cannot be guided to other cars. This leads to decreased efficiency and increased passenger journey times.

In contrast, the destination control system (DCS) guides passengers to certain cars, resulting in fewer stops and improved handling capacity. In a DCS, passengers with the same destination floor are allocated into the same car. [1]

In a widely used implementation of the destination control system, passengers specify their desired floor before entering the elevator. The system then assigns the optimal elevator car to the passenger, who is directed to the designated elevator. There is no need for car calls to be made inside the elevator itself. [1]

Destination calls can be made using keypads located on the entrance floor or other floors. A display is essential because it immediately shows the identifier of the car that will serve the call, allowing passengers to know which car to enter based on their destination. [1]

Destination calls can also be made using the access card which has the passenger home floor coded in it. System-wise, the passenger call is automatically moved from the access control to elevator control, which then allocates the given call to the best car available. [1]

3 Simulation Design & Methods

This section describes the technical setup, tools, and methodology used to simulate and analyze elevator systems. Because we aim to optimize the HC value in Section 2 for simulator 2, using the results for simulator 1, both simulators will run for 5 minutes. For both simulators, passengers are assumed to follow a Poisson distribution [5]. This means that we assume that passengers arrive independently. In both of our simulators, this is set to λ passengers per 5 minutes.

Using this arrival rate, we can calculate the system utilization factor, ρ. This shows how much the system is occupied by the passengers. This is calculated using the formula

$$\rho = \frac{\lambda}{L \times \mu}. \tag{6}$$

where L is the number of elevators, λ is the passenger arrival rate per second, and μ is the passenger handling rate for each elevator per floor interval in [5]. Here, we set $L = 22$.

The waiting time (WT) is the amount of time passengers wait at the lobby. The journey time (JT) is the sum of the WT, transit time of passengers inside the car, and the opening and closing time of the doors for each elevator. In paternosters, the transfer time per passenger to enter or exit the car is 1 s [1]. We aim to simulate two cases: one case where the opening and the closing time of the doors is zero, and the other where they are nonzero.

In both simulators, elevators move at a speed of 0.5 m/s between the 10 floors of a single shaft, since only up-peak traffic is considered. When they move inside a high building that is expected to be at least as tall as the Petronas Tower in Kuala Lumpur, of 88 floors with 450 m [6]. Because this project is carried out on a smaller scale, we set the height of the shaft to approximately 1/10 of the size of the high building, which makes it 50 m, and 10 floors plus the entrance floor.

Then, the distance between adjacent floors becomes 5 m. Then, the time it takes for an elevator to travel between floors becomes

$$\frac{5 \text{ m}}{0.5 \text{ m/s}} = 10\ s.$$

As we have 22 elevators moving between 2 shafts of 50 m from equation (4), t_{rt} for the continuously moving paternoster elevator is calculated as

$$t_{rt} = \frac{2 \times 50 \text{ m}}{0.5 \text{ m/s}} = 200\ s.$$

As each elevator's passenger processing rate is set to 0.8 passengers / second in [5], μ is set to $(0.8 \times 1)/200\ s = 0.8/200\ s$. In simulator 1, where optimizing the HC value is not necessary, using the formula for HC given in the literature review, we have handled

$$HC1 = \frac{0.8 \times 300 \times 22 \times 1}{200} \approx 26 \ \text{persons} \ / \ 5 \ \text{minutes} \tag{7}$$

during the simulation run.

We test three different scenarios (low traffic, medium traffic, and high traffic). For low traffic, we set $\lambda = 20$ persons / 5minutes. For medium traffic, we set $\lambda = 40$ persons / 5minutes. For high traffic, we set $\lambda = 60$ persons / 5minutes. These are equivalent to setting cases with low, medium, and high traffic as $\frac{20}{300}$ persons / second, $\frac{40}{300}$ persons / second, and $\frac{60}{300}$ persons / second respectively.

We get the corresponding ρ values for these three different scenarios (low traffic, medium traffic, and high traffic) as follows:

$$\rho_1 = \frac{\lambda}{L \times \mu} = \frac{\frac{20}{300}}{22 \times \frac{0.8}{200}} = 0.7576,$$

$$\rho_2 = \frac{\lambda}{L \times \mu} = \frac{\frac{40}{300}}{22 \times \frac{0.8}{200}} = 1.5152,$$

$$\rho_3 = \frac{\lambda}{L \times \mu} = \frac{\frac{60}{300}}{22 \times \frac{0.8}{200}} = 2.2727.$$

In the simulators, we assume that one round trip in the morning up-peak passenger traffic consists of the following steps in [6], given as

1. leaving the previous elevator,
2. passengers making calls to their destinations,
3. calls being canceled when the elevator stops to the floor and opens doors,
4. passengers entering the car,
5. elevators closing the doors and leaving the floor.

Other cars stop at upper floors to unload the passengers at their destination floors simultaneously with the car at the entrance floor. No extra time to the round-trip time is caused.

In Simulator 1, we assume that the elevators are continuously moving inside the shaft, which allows steps 2 and 3 in the passenger journey to be meaningless, as they would not stop for the passengers, regardless of their actions. C is set to one, as it would be difficult for passengers to step safely inside a continuously moving vehicle at the same time. This is the case where the opening and closing times of the doors are zero.

In Simulator 2, we assume that after step 2, the elevator would first stop in the lobby to pick up the passengers and then stop at their destinations to drop off the passengers. Hence, the closing and opening times of the doors for each elevator are added when computing the t_{rt} value. As denoted in [1], we will set them to be 2 seconds. Hence, we need to add 4 seconds per elevator to t_{rt}. Moreover, we also need to add the transfer time per passenger. This means that we need to add 1 second for each passenger that moves out of the elevator to t_{rt}.

As we have the door mechanism, it is now safe for passengers to enter the elevator at the same time. Hence, we set $C \geq 1$, to be specific, as $C = 5$. From equation (2), this makes t_{rt} become $200\ \ s + 22 \times (4 + 1 \times 5)\ \ s = 398\ \ s$. Now that the values for C and t_{rt} have changed, μ is now $(0.8 \times 5)/398 = 4/398\ s$. Using equation (1), we have at least handled

$$HC2 = \frac{0.8 \times 300 \times 22 \times 5}{398} \approx 66 \text{ persons / 5 minutes} \tag{8}$$

during the simulation run, as a result.

Hence, the corresponding ρ values for these three different scenarios (low traffic, medium traffic, and high traffic) change as follows:

$$\rho_1 = \frac{\lambda}{L \times \mu} = \frac{\frac{20}{300}}{22 \times \frac{4}{398}} = 0.3015,$$

$$\rho_2 = \frac{\lambda}{L \times \mu} = \frac{\frac{40}{300}}{22 \times \frac{4}{398}} = 0.6030,$$

$$\rho_3 = \frac{\lambda}{L \times \mu} = \frac{\frac{60}{300}}{22 \times \frac{4}{398}} = 0.9045.$$

The ρ values are always smaller than 1 for cases when the value for λ is smaller than the value for HC and vice versa for both simulators. After assigning the passengers to the right elevators, we use the sum of the total number of passengers and the number of active elevators as a metric to track the performance of simulator 2 [7].

To find out the car capacity of Simulator 2 so that the same amount of traffic can be handled as with Simulator 1, the following process can be used.

Create both simulators without optimization. After running both simulators, compare the handling capacity of the system. After comparison, calculate the smallest car capacity, C, which would allow Simulator 2 to achieve the same handling capacity as Simulator 1. Use this in Simulator 2 in the next round. Repeat the steps until a sufficiently small C is found. Another extra step would be to assign passengers to the right elevators according to their destination [1].

4 Results

In Figure 1, representing Simulator 1, passengers board the continuously moving elevator cars. The left photo shows the movement of the Paternoster-like simulator, while the right photo illustrates the process of passengers leaving their elevator cars and entering their destination floors.

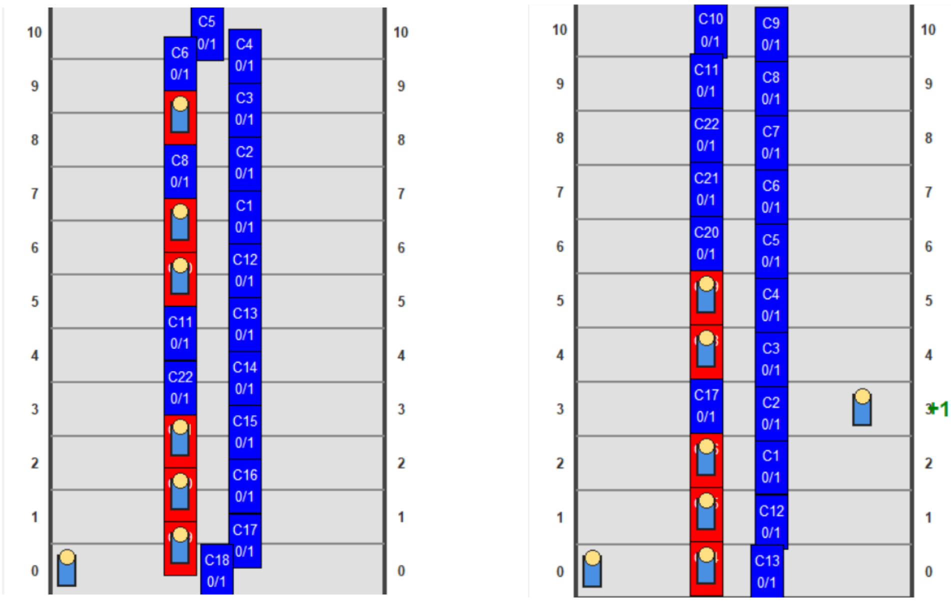

Figure 1: Photos refer to Simulator 1.

In Figures 2-4, the average waiting (AWT) and transit times (TT) are presented over simulated time, for Simulator 1, having an individual figure for each passenger arrival rate: Figure 2 (20 persons per 5 minutes) in which the calculated HC of Simulator 1 is relatively close to the passenger arrival rate λ, Figure 3 (40 persons per 5 minutes), and Figure 4 (60 persons per 5 minutes). In Figures 3 and 4, the calculated HC of Simulator 1 is significantly lower than the corresponding passenger arrival rates.

In Figure 2, AWT should remain approximately constant for a continuously moving paternoster. However, AWT line slightly increases over time. For TT in Figure 2, there should not be a rapid increase if passengers leave the elevator on upper floors.

For Simulator 1 with the calculated handling capacity (HC) of approximately 26 persons per 5 minutes, Figure 3 shows that increasing the passenger arrival rate λ from 20 to 40 persons per 5 minutes significantly increases the average waiting times. A similar effect can be seen in Figure 4 in which the passenger arrival rate has been increased to 60 persons per 5 minutes.

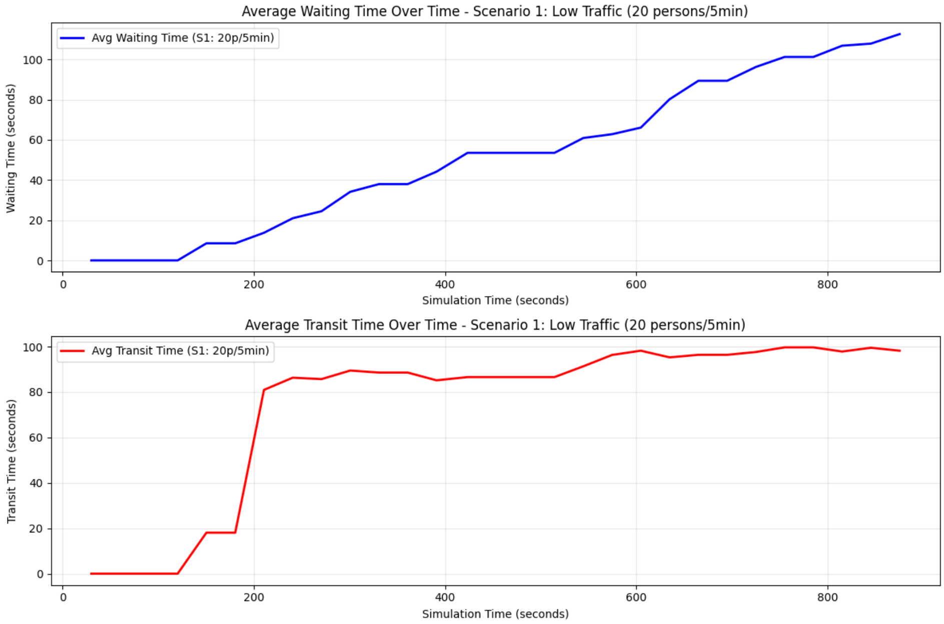

Figure 2: Average waiting times (AWT) and transit times (TT) over simulated time in Simulator 1, when $\lambda = 20$ persons / 5 minutes.

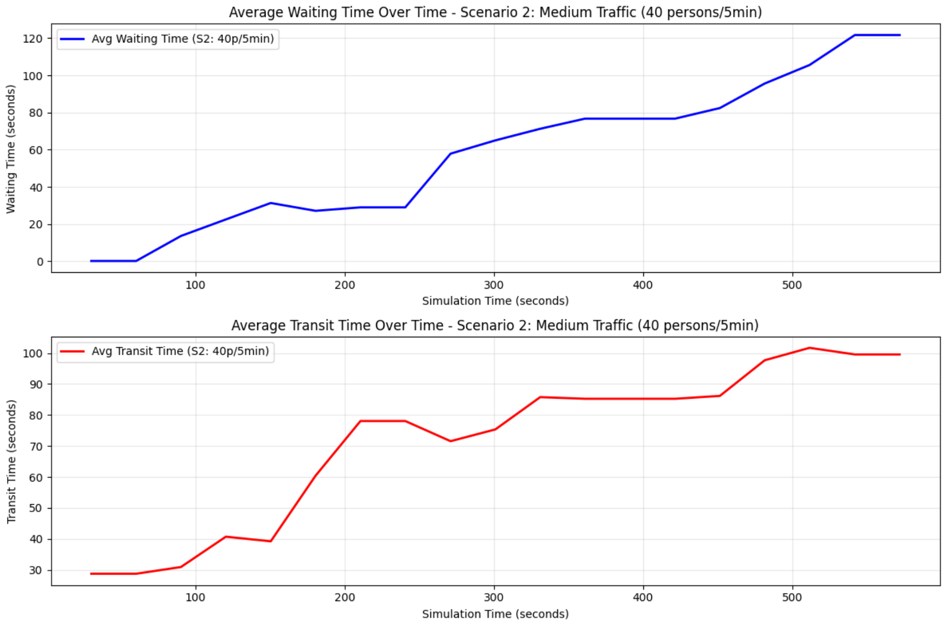

Figure 3: Average waiting times (AWT) and transit times (TT) over simulated time in Simulator 1, when $\lambda = 40$ persons / 5 minutes.

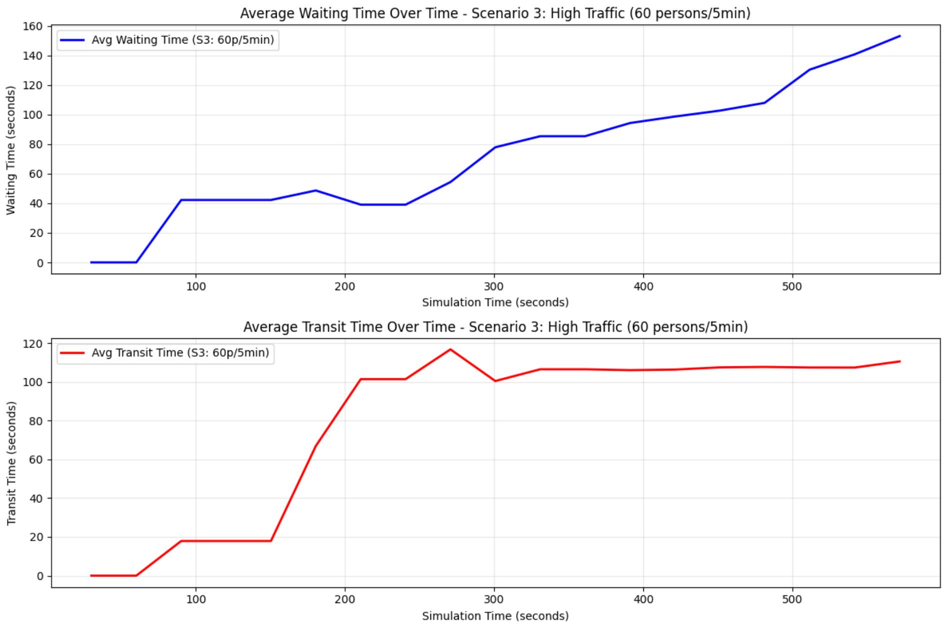

Figure 4: Average waiting times (AWT) and transit times (TT) over simulated time in Simulator 1, when $\lambda = 60$ persons / 5 minutes.

In Figure 5, representing Simulator 2, which stops, opens the doors, takes the passengers in, closes the door, and continues moving up, until it stops again, and lets the passengers step out of the elevator car. The left photo shows the Paternoster-like movement of the simulator between the stops, while the right photo illustrates the stopped system while one passenger leaves their car on the third floor, which is their destination floor.

In Figures 6-8, the average waiting (AWT) and transit times (TT) are presented over simulated time, for Simulator 2, having an individual figure for each passenger arrival rate, similarly to Simulator 1. In Figure 6 with the passenger arrival rate $\lambda = 20$ persons per 5 minutes, the calculated HC of Simulator 2 is significantly higher than λ, which leads to low average waiting times below 20 seconds. In Figure 7 with the passenger arrival rate $\lambda =$ 40 persons per 5 minutes), the calculated HC of Simulator 2 is still clearly higher than the increased λ. In Figure 8 with the passenger arrival rate $\lambda =$ 60 persons per 5 minutes, the calculated HC of Simulator 1 is close to the λ, but it is still slightly higher than the increased λ.

For Simulator 2 with the calculated handling capacity (HC) of approximately 66 persons per 5 minutes, all the three passenger arrival rates λ are lower than HC so the average waiting times stay relatively short. That is discussed more in Section 5.

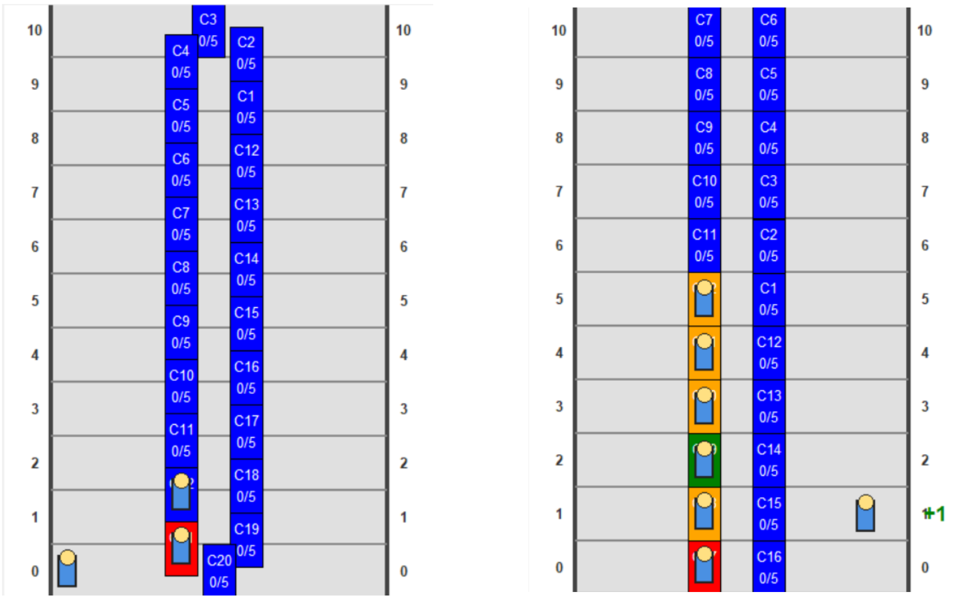

Figure 5: Photos refer to Simulator 2.

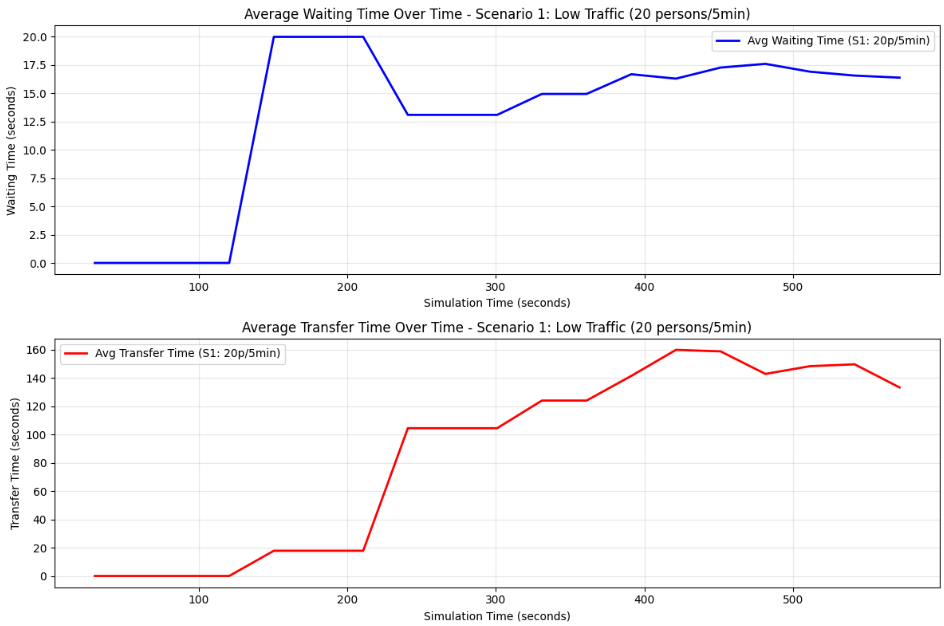

Figure 6: Average waiting times (AWT) and transit times (TT) over simulated time in Simulator 2, when $\lambda = 20$ persons / 5 minutes.

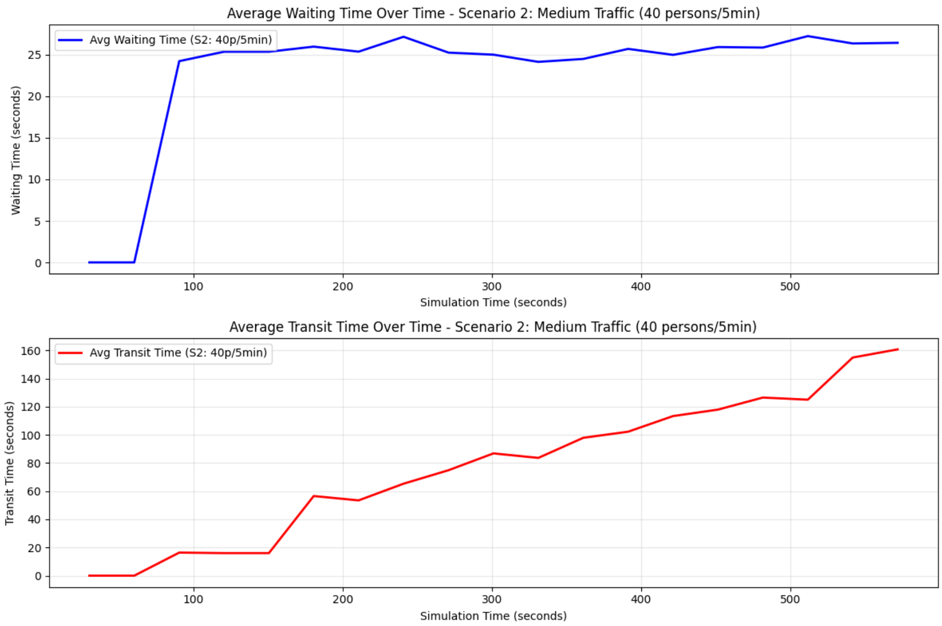

Figure 7: Average waiting times (AWT) and transit times (TT) over simulated time in Simulator 2, when $\lambda = 40$ persons / 5 minutes.

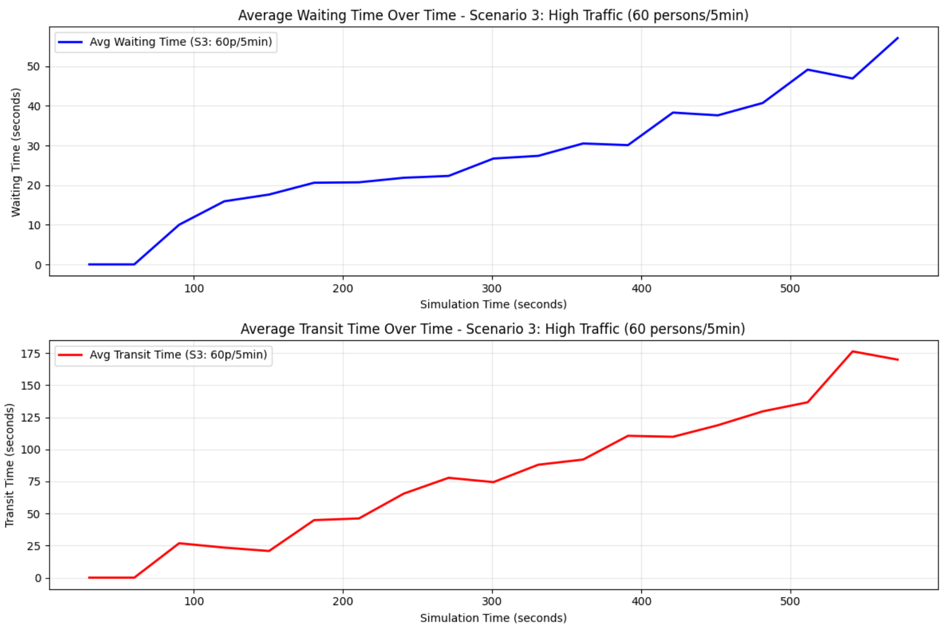

Figure 8: Average waiting times (AWT) and transit times (TT) over simulated time in Simulator 2, when $\lambda = 60$ persons / 5 minutes.

While consistent axis scaling across Figures 2–4 and 6–8 would have improved the clarity of AWT and TT comparisons, this was not fully implemented due to time constraints. Nevertheless, the presented figures and tables still allow for a sufficient comparison between the performance of Simulators 1 and 2.

Table 1 summarizes the calculated up-peak RTT, average waiting times (WT), and average transit times (TT) for both Simulator 1 and Simulator 2, in all three scenarios with different traffic profiles. When light traffic, the arrival rate is set to 20 persons per 5 minutes. When medium traffic, the arrival rate is set to 40 persons per 5 minutes. And when high traffic, the arrival rate is set to 60 persons per 5 minutes. The handling capacity (HC) is set to the constant maximum capacity that an elevator can transport passengers during 5 minutes. HC values are rounded to the closest integer below the decimal value. RTT values are calculated with the equations inserted into the Simulator scripts, and the results are shown in the Python GUI.

In Table 1, waiting time (WT) denotes the time passengers need to wait before boarding the elevators. Transit time (TT) denotes the time passengers need to spend inside the elevator car from the moment they board the car to the moment they arrive at their destination floor. Those values are averages over time and they are updated in real time. Both AWT and TT should be approximately constant for Simulator 1. However, for Simulator 2, AWT and TT increase when the passenger arrival rate λ increases.

Table 1: Simulators 1 and 2 in all three traffic profile scenarios.

No	Car size	Calculated HC (persons / 5 minutes)	Arrival rate λ (persons / 5 minutes)	Calculated RTT (s)	Avg WT (s)	Avg TT (s)
1	1	26	20	200	50	90
1	1	26	40	200	75	85
1	1	26	60	200	110	105
2	5	66	20	398	15	140
2	5	66	40	398	25	130
2	5	66	60	398	40	160

Tables 2-4 present the comparisons across different traffic profiles to support the analysis of which simulator performs better under varying passenger arrival rates.

Table 2: Simulators 1 and 2 at $\lambda = 20$ persons per 5 minutes.

No	Car size	Calculated HC (persons / 5 minutes)	Arrival rate λ (persons / 5 minutes)	Calculated RTT (s)	Avg WT (s)	Avg TT (s)
1	1	26	20	200	50	90
2	5	66	20	398	15	140

Table 3: Simulators 1 and 2 at $\lambda = 40$ persons per 5 minutes.

No	Car size	Calculated HC (persons / 5 minutes)	Arrival rate λ (persons / 5 minutes)	Calculated RTT (s)	Avg WT (s)	Avg TT (s)
1	1	26	40	200	75	85
2	5	66	40	398	25	130

Table 4: Simulators 1 and 2 at $\lambda = 60$ persons per 5 minutes.

No	Car size	Calculated HC (persons / 5 minutes)	Arrival rate λ (persons / 5 minutes)	Calculated RTT (s)	Avg WT (s)	Avg TT (s)
1	1	26	60	200	110	105
2	5	66	60	398	40	160

5 Discussion

This project was conducted in collaboration with MLS Lift Consulting to simulate and compare two elevator systems: a continuously moving Paternoster (Simulator 1) and a stop-based system with door mechanics (Simulator 2). The key emphasis was on evaluating performance during the morning peak period in a 10-story building with both simulators and comparing the handling capacity of the system (HC), the round-trip time (RTT), and the passenger waiting times.

Round-trip time (RTT) played a key role. A shorter RTT typically results in higher HC, as observed when optimizing parameters in Simulator 2. This reinforces the importance of reducing unnecessary stops and improving boarding efficiency.

Simulator 1, the continuously moving elevator system, showed naturally with one-person cars higher handling capacity (HC) due to its uninterrupted functioning. However, it does not stop or have doors, which makes it unsafe and unsuitable for modern buildings. Simulator 2, on the other hand, initially had lower HC due to door operations and stopping, but this could be improved by increasing the car capacity to 5.

In our simulations, the calculated HC of Simulator 1 was approximately 26 persons per 5 minutes, while Simulator 2 reached approximately 66 persons per 5 minutes. These values remained constant across scenarios and were compared with different arrival rates λ of 20, 40, and 60 persons per 5 minutes to assess system performance under light, medium, and heavy traffic.

The plots in Figures 2–4 showed that for Simulator 1, the average waiting times (AWT) began to increase significantly once λ exceeded the calculated HC. This behavior was especially evident in Figures 3 and 4. These results suggest that average waiting times stay relatively optimal when the handling capacity HC of the elevator is more than the arrival rate of passengers λ, as follows: HC $> \lambda$. However, if the arrival rate λ exceeds the handling capacity HC, then the waiting times begin to increase significantly, as seen in Simulator 1 with arrival rates of 40 persons per 5 minutes and 60 persons per 5 minutes.

The same figures also showed a slight increase in AWT even for $\lambda = 20$ persons per 5 minutes, even though in theory, a continuously moving system like Simulator 1 should maintain a relatively stable AWT. This minor deviation may be due to cumulative queue effects or implementation delays in our simulation model.

Simulator 2 has a high handling capacity, so there was not similar effect even though the arrival rates were increased. Sure, the waiting times increased, but not relatively much. Figures 6–8 confirmed this trend: AWT remained low for Simulator 2 even as λ increased from 20 to 60 persons per 5 minutes. Even under high traffic, AWT stayed below 40 seconds. However, the average transit time (TT) increased gradually with higher arrival rates in Simulator 2, likely due to additional boarding and exiting delays as more passengers shared each car.

On the other hand, when comparing transit times (TT), Simulator 1 has clearly lower transit times than Simulator 2. That is due to the continuous movement of the paternoster elevator. This difference was visible in the simulation plots, where TT for Simulator 1 remained relatively stable, while Simulator 2 showed a progressive increase in TT under heavier traffic.

It is worth noting that the consistency of axis scaling across Figures 2–4 and 6–8 would have made comparisons easier. While this was not fully implemented due to time constraints, the trend differences between simulators were still clear enough to draw valid conclusions.

Depending on the building type and its passenger traffic profile, the choice between systems can vary. Based on these results, Simulator 2 becomes preferable due to its car capacity in high traffic scenarios, and it also offers more comfort, safety, and accessibility. However, Simulator 1 is more preferable when the traffic is very light so that stopping is not necessary and the continuous movement of the paternoster has a competitive edge over the stopping-based elevator system.

For high-throughput use cases where safety is less critical, such as restricted industrial settings, a continuously moving system like Simulator 1 could be a viable option. But for public or commercial buildings with varied users, a well-optimized stop-based system like Simulator 2 is more realistic and preferable for future use.

From the numerical results and simulation experience, we can recommend prioritizing solutions that balance performance and user safety. Even though Paternoster systems are theoretically efficient, their limitations in safety and accessibility outweigh the benefits in most contexts. Modern destination control systems and multi-car elevator (MCE) designs can achieve similar or better throughput with safer and more user-friendly operation.

Finally, Table 1 provided a clear numerical summary of RTT, AWT, and TT under each traffic scenario. These results supported the visual plots and enabled direct average performance comparison between the two simulators. Tables 2–4 further clarified the differences between Simulator 1 and Simulator 2 specifically under each traffic profile.

There is still room for improvement in our simulator models. We assumed simple passenger behavior (Poisson arrivals, morning up-peak), which does not reflect all real-life scenarios at different times. Future extensions could include more complex traffic profiles or other realistic peak behaviors.

6 Conclusions

This report presents the results of a project carried out for MLS Lift Consulting as a part of the course *MS-E2177 - Seminar on Case Studies in Operations Research D*. A multi-car elevator system simulation was built using Python in two steps: first, simulating a conventional Paternoster, and second, modifying it to include doors and a stopping mechanism.

A literature review was conducted to understand key concepts and previous research related to elevator systems, especially under up-peak traffic conditions. We focused on important performance metrics such as round-trip time (RTT), handling capacity (HC), and interval, and explored how these are affected by factors like elevator speed, number of passengers, and number of stops. The review also compared conventional control systems (CCS) with destination control systems (DCS), which improve efficiency by grouping passengers heading to the same floor. These insights from the literature shaped our simulation design and optimization methods.

In the initial two phases, we successfully developed and tested two elevator simulations under various parameters, analyzing performance and exploring optimization strategies. These simulations highlighted how system modeling and performance evaluation can support better system-level decision-making. Special attention was given to the clarity and functionality of the Python simulation, which now provides reliable figures to support traffic analysis and elevator planning.

Unfortunately, due to time constraints and technical challenges, we were unable to complete the third and final simulation. It offered important lessons in scope management, prioritization, and the realities of working within project constraints, which are highly relevant skills in a professional setting.

We are grateful for the opportunity to collaborate with MLS Lift Consulting, and we hope that our contributions are useful for the company. The detailed comparison between two elevator systems, supported by clear visuals and well-structured analysis, offers practical insight for similar real-world cases. The simulations and findings can serve as a foundation for future development or as a reference for similar tasks. The insights gained through our analysis may inform better decision-making in the design or management of elevator systems.

Working on this project also strengthened our teamwork, communication, and problem-solving skills, which are key aspects of real-world operations research. Additionally, the process of transforming theoretical models into working simulation tools gave us valuable hands-on experience with applied operations research and systems design.

Future work could explore more complex passenger behavior or integrate more advanced optimization algorithms.

7 Self Assessment

How closely did the actual implementation of the project follow the initial project plan? Were there any major departures and, if so, what?

The actual implementation of the project deviated somewhat from the initial project plan, as we completed two of the three planned steps. The third simulation, in which the elevators would move independently, could not be implemented. This setback was communicated with the client and we agreed to focus on the first two simulations to a high standard, which allowed us to provide meaningful analysis and reliable results.

In what regard was the project successful? In what regard was it less so?

The project was successful in completing the two simulations in a thorough and informative manner. The simulations provided clear visualizations, well-structured analysis, and a robust Python tool that supports traffic analysis and elevator planning. Not completing the third simulation was a disappointment for our group. However, the experience offered important lessons in scope management, prioritization, and the realities of working within project constraints, which are crucial skills in professional practice. Additionally, our visits to Stockmann and KONE enriched our understanding of elevator systems, including Paternoster operation, destination control, safety, and maintenance. Working collaboratively throughout the project also strengthened our teamwork, communication, and problem-solving skills.

What could have been done better, in hindsight?

In hindsight, we could have improved significantly in both scheduling and communication. Our scheduling was too lenient in the beginning of the course, leading to a cramped schedule towards the end and ultimately the incomplete third simulation. Communication both within the group and with the company could also have been improved. Proactively seeking guidance and feedback earlier would likely have increased the chances of delivering the final simulation and enhanced the overall project flow.

References

[1] Marja-Liisa Siikonen. *People Flow in Buildings*. Wiley-Blackwell, 2021. ISBN: 9781119545569.

[2] *ISO 8100-1:2019 - Lifts for the transport of persons and goods — Part 1: Safety rules for the construction and installation of passenger and goods lifts*. Available from https://www.iso.org/standard/72168.html. 2019.

[3] B.E. A. Bhatia. *Overview of Building Elevator Systems*. Last accessed 13 May 2023. 2020. URL: `https://www.pdhonline.com/`.

[4] *ISO 8100-32:2020 - Lifts for the transport of persons and goods — Planning and selection of passenger lifts to be installed in office, hotel and residential buildings*. Available from https://www.iso.org/standard/72746.html. 2020.

[5] Lutfi Al-Sharif, Ahmad Abu Alqumsan, Weam Ghanem, Islam Tayeh, and Areej Jarrar. "Modelling of Elevator Traffic Systems Using Queuing Theory". In: *Proceedings of the 4th Symposium on Lift and Escalator Technologies*. Sept. 2014. DOI: `10.13140/RG.2.1.1392.0165`.

[6] Marja-Liisa Siikonen. "Planning and Control Models for Elevators in High-Rise Buildings". PhD thesis. HUT, 1997.

[7] Micah Allen, Davide Poggiali, Kirstie Whitaker, Tom Rhys Marshall, Jordy van Langen, and Rogier A Kievit. "Reinforcement Learning for Elevator Control". In: *Welcome Open Research* 4 (2021), p. 63.