Tactical taxibot planning at Amsterdam Airport Schiphol under uncertainty

Thesis Paper

C.F. van Winkel
4538226

Tactical taxibot planning at Amsterdam Airport Schiphol under uncertainty

Thesis Paper

by C.F. van Winkel

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on May 9, 2023 at 09:30.

Student number:	4538226	
Project duration:	May 2022 – May 2023	
Supervisors:	Ir. P. C. Roling,	TU Delft, supervisor
	R. K. Kotey,	Royal Schiphol Group, supervisor

TU Delft

Tactical Taxibot Planning at Amsterdam Airport Schiphol under Uncertainty

C. F. van Winkel
Aerospace Engineering, Delft University of Technology

Amsterdam Airport Schiphol aims to have a fully autonomous taxibot fleet that can taxi all incoming and outgoing flights by 2050. To support the transition towards this future state, wherein the demand for taxibots will initially be higher than the supply and operational concepts will continuously evolve, a flexible taxibot allocation model is developed. This model can optimize for different objectives: fuel savings, schedule robustness in terms of flight delays, a healthy working environment, and fairness towards participating airlines. The model is proven to be versatile and applicable to many phases towards the 2050 vision. However, this also makes the model computationally complex, which should be investigated in further research.

I. Introduction

Amsterdam Airport Schiphol (AAS) has multiple pillars in their 2050 vision. One of these pillars is Quality of Life, which also contains the goal to be the most sustainable airport by 2030 [1]. The intention is to improve the quality of the living environment, both globally and locally. One of the local mitigation options for both noise and emissions is the use of taxibots. Taxibots are trucks that can tow an aircraft from their gate to the runway and vice versa, instead of having the aircraft taxi on their engines. This procedure uses 50-65% less fuel and emissions when using diesel taxibots[2]. In addition, as the aircraft engines can start up later (outside the bay), this procedure greatly improves the working environment for ground handlers, who work in the bays.

With approximately 1500 flights per day, a challenging question for the development of taxibot operations is how to allocate taxibots to flights. There are multiple objectives that can be optimized for. For example: sustainability, on-time performance, and fairness to the participating airlines. Next to this, there are many constraints for the planning, such as the compatible aircraft types, potential charging time of electric taxibots, runway configuration, et cetera. All this, while the flight schedule is also dynamic; one cannot tell a day in advance what time exactly an aircraft will land or be ready for departure.

The sustainable taxiing project is an ongoing project, which means that it will expand in the coming years. The taxibot fleet will change (from diesel to electric, from manned to autonomous), what aircraft are compatible with taxibots will change, runways to operate on will expand, etc. The mission allocation system should be able to adapt to these changes, as the project expands. Next to that, the system could be used to help choose what direction the project goes by giving valuable insights in terms of the effects of design choices.

The objective of this research is as follows:

> To create and assess a mathematical model that can assign a given fleet of taxibots to incoming and outgoing flights, while optimizing for different objectives, taking into account the relevant uncertainties in the operation.

This objective captures a number of requirements that Schiphol has. Schiphol will have a fixed fleet of taxibots available, therefore a given fleet must be assigned to flights. The different objectives that Schiphol would like to optimize for are sustainability and fairness to participating airlines, while being able to tweak the amount of risk the planning is able to take.

II. The Modified Electric Vehicle Scheduling Problem

To generate a taxibot allocation, a mixed integer linear programming model is defined. To achieve the research objective, three modifications are made to this model; one to allow for steering on robustness, one for steering on outbound movements and one for steering towards fairness.

A. Electric Vehicle Scheduling Problem Introduction

It was concluded that an Electric Vehicle Scheduling Problem (EVSP) best represents the problem at hand. In an EVSP, a number of (electric) vehicles are assigned a set of tasks, with a fixed start and end time. The vehicles must drive from task to task and can visit charging stations along the way. At the beginning and end of the day, a vehicle must be positioned at one of multiple depots. The full EVSP model defined for this research is described in Appendix A.

B. Model adaptations

Once the EVSP model is defined, adaptations to achieve the research objective can be made.

1. Robustness

To make the model robust means to take the uncertainties in the operations into account. For this research, it was decided to only incorporate uncertainty in arrival/departure time, as this was found to be the most relevant uncertainty. A list of other identified uncertainties is provided in Appendix B.

To gain a deeper understanding of the uncertainty in arrival and departure times, a delay distribution study was performed over 2019 AAS data. If flight specific delay distributions can be made, consecutive taxibot trips can be allocated accordingly. For example, if a flight regularly arrives early, it would not be wise to allocate that taxibot trip directly after another taxibot trip of a flight that regularly arrives late. The full delay study can be found in Appendix C.

The delay distributions for each flight are used to find the probability of conflict between two flights. This requires the probability that the taxibot has completed the first trip and that the taxibot does not need to be underway to the second trip at each moment in time. For this, the delay distributions of two consecutive flights are taken, and the later flight's distribution is inverted. The first flight's distribution is now the probability that the flight has started its taxibot trip, the second flight's distribution is the probability that the flight does not yet need a taxibot. The first flight's distribution is shifted later by the amount of time the taxibot trip takes, the second flight's distribution is shifted earlier by the time needed to drive from the first trip to the second trip. These distributions are shown in Figure 1.

The conflict probability between two flights i and j is computed through Equation 1. For the conflict probability of a flight followed by a charging task, the delay distribution of the flight is shifted only by the time required to drive from the flight to the charger. Then, the the probability of not having completed the flight and deadhead trip yet at the intended time of charging is defined as FP_{ict}. The probability of having to have started a trip to flight i from charger c at a certain time is computed similarly and denoted as SP_{ict}.

$$CP_{ij} = 1 - P(\text{no conflict})_{ij} = 1 - P(\text{completed})_i \times P(\text{not started})_j \tag{1}$$

It was chosen to incorporate arrival/departure delay uncertainties though stochastic constraints, as these are not computationally heavy as compared to stochastic programming or robust optimization [3]. Additionally, it allows for the user to steer the maximum allowed risk to be taken. However, as these are constraints, it does not minimize or evenly divide conflict probability under the user specified maximum conflict probability. Four sets of constraints must be added to the model. Constraints 2 ensure that if a taxibot k completes both task i and j, the conflict probability between those two consecutive flights CP_{ij} (taking into account taxi time and deadhead time), cannot exceed the user specified maximum conflict probability CP_{max}. Constraints 3

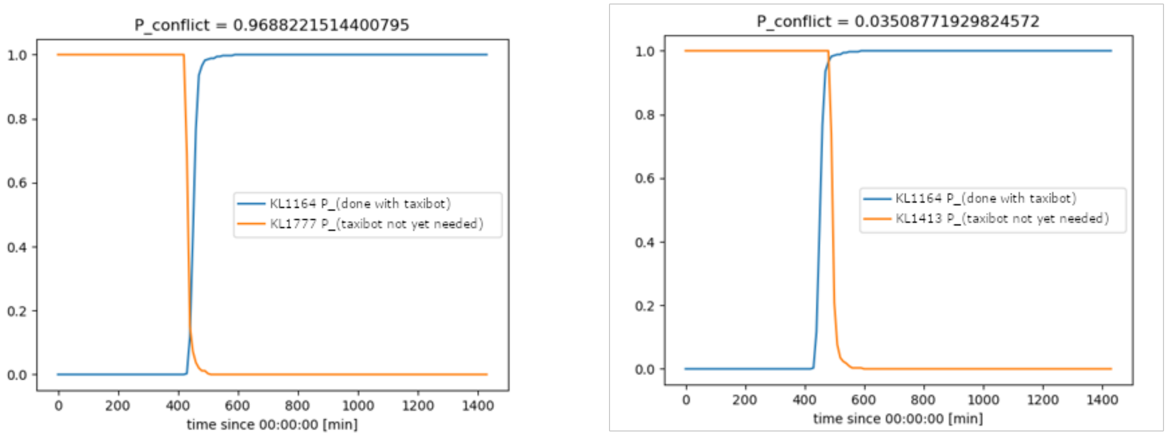

(a) Two flights with high conflict probability

(b) Two flights with low conflict probability

Figure 1. Probability distributions for different flights

are similar, except that it takes into account that the model might cheat constraints 2 by going to a charger between flights i and j. Therefore, Constraints 3 ensure that if a taxibot completes flight i, heads to a charger and then continues from that charger to flight j, the conflict probability between flights i and j still does not exceed CP_{max}.

$$x_{ijk} \times CP_{ij} \leqslant CP_{max} \quad ,\forall i, j \in \mathbf{N}, k \in \mathbf{M} \tag{2}$$

$$-M\left(2 - \sum_{c\in\mathbf{C}} x_{ick} - \sum_{c\in\mathbf{C}} x_{cjk}\right) + CP_{ij} \leqslant CP_{max} \quad ,\forall i, j \in \mathbf{N}, k \in \mathbf{M} \tag{3}$$

Constraints 2 and 3 prevent flight conflicts, however, conflicts could also occur between a charging task and a flight. It could be considered desirable to also prevent charging too close to the beginning/end of a taxibot trip. Constraints 4 and 5 ensure this. For example, in Constraints 4, if a flight is headed from taxibotting flight i to a charger c, then the taxibot should not be charging at that charger during time t if the probability of having finished taxibotting and driving to the charger (FP_{ict}) is greater than the maximum allowed conflict probability. Constraints 5 are similar, except adapted for the case when a taxibot is heading to a flight from a charger.

$$-M(1 - x_{ick}) + y_{ckt} \times FP_{ict} \leqslant CP_{max} \quad ,\forall i \in \mathbf{N}, k \in \mathbf{M}, c \in \mathbf{C}, t \in \mathbf{T} \tag{4}$$

$$-M(1 - x_{cik}) + y_{ckt} \times SP_{ict} \leqslant CP_{max} \quad ,\forall i \in \mathbf{N}, k \in \mathbf{M}, c \in \mathbf{C}, t \in \mathbf{T} \tag{5}$$

2. Healthy Working Environment

The main reason for using taxibots is sustainability. However, sustainability is a broad term, and therefore optimizing for fuel savings is not the only way to optimize for sustainability. Schiphol also desires a sustainable working environment for its employees. Using taxibots can aid in improving the working environment for employees in the bays. Outbound movements are expected to result in the biggest gain in terms of a healthy environment. This, because the engines are running considerably longer in the bay when departing than when arriving and because the aircraft must come into motion out of a standstill. Additionally, when leaving the bay, the exhaust of the aircraft is often directed towards the people working in the bay. This means that optimizing for a healthy working environment primarily means maximizing the number of outbound taxibot movements. Optimizing for outbound movements also allows for airlines to take the taxibot trip into account when fueling a departing aircraft.

To optimize for the number of outbound flights, the objective of the model is adjusted. An outbound weight factor is added, a value between 0 and 1, outlining to what extent it is desirable to focus on fuel savings ($W_o \to 0$) or on the number of outbound movements ($W_o \to 1$). Each flight is given an attribute O_i, which is 1 if the flight i is an outbound flight, and 0 otherwise. The outbound flights element of the objective is also multiplied by the average fuel savings of a flight on the day being optimized (F_{av}), to balance the two parts of the objective. Although this would not result in an exact balance between the two, it should still result in a better balance. The complete new objective is given in Equation 6

$$\max \left[Z = \sum_{i \in \mathbf{N}} \sum_{j \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D}} \sum_{k \in \mathbf{M}} \left(((1 - W_o) \times W_{ik} + W_o \times F_{av} \times O_i) \times x_{ijk} \right) \right] \tag{6}$$

3. Fairness

Especially during the trial period, Schiphol wants to have a fair distribution of taxibot trips among the participating airlines: KLM, Transavia, TUI fly and Corendon. As these airlines are investing their time, equipment and money on this project, it is considered fair to give all an equal amount of time with the taxibots. Additionally, different airlines can provide different perspectives that could lead to valuable learnings throughout the project.

To incorporate fairness into the model, the question as to what fairness is must first be answered. After talks with Schiphol, it was concluded that the fairness towards airlines is important in terms of fuel savings and exposure. However, dividing the flights as 'evenly' as possible is difficult with four airlines of different size. Due to the low number of flights for some participating airlines, it was decided that a proportional division is too arbitrary; if 1 out of 2 Corendon flights would be taxibotted on a specific day, also half of the KLM flights should be taxibotted, which would be impossible for a limited amount of taxibots. As, at this moment, the planning is only made one day ahead, a longer horizon is not an option either. Therefore it was decided absolute number division is a better method.

Incorporating this fairness as a constraint was considered to be most applicable, as this allows sustainability to remain the primary objective. The smallest airline is identified on the chosen day, and the other airlines are measured against that through Constraints 7 and 8.

$$\sum_{i \in \mathbf{N}} \sum_{j \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D}} \sum_{k \in \mathbf{M}} \left((x_{ijk} \times \text{smallest}_i) - 1 \right) \times FF \leqslant \sum_{i \in \mathbf{N}} \sum_{j \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D}} \sum_{k \in \mathbf{M}} \left(x_{ijk} \times \text{airline}_i \right) \quad , \forall \text{airline} \in \text{airlines} \setminus \text{smallest} \tag{7}$$

$$\frac{\sum_{i \in \mathbf{N}} \sum_{j \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D}} \sum_{k \in \mathbf{M}} \left(x_{ijk} \times \text{smallest}_i \right) + 1}{FF} \geqslant \sum_{i \in \mathbf{N}} \sum_{j \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D}} \sum_{k \in \mathbf{M}} \left(x_{ijk} \times \text{airline}_i \right) \quad , \forall \text{airline} \in \text{airlines} \setminus \text{smallest} \tag{8}$$

Constraints 7 ensure that airlines larger than the smallest airline will not be taxibotted significantly less than the smallest airline. Conversely, Constraints 8 ensure that the larger airlines are not taxibotted significantly more than the smallest airline. The fairness factor (FF) is a value between 0 and 1, denoting the importance of fairness. When the fairness factor is 0, fairness is not of importance and both sets of constraints are omitted.

The +1 in the numerator of Constraints 8 ensures that if the smallest airline does not fly on the given day at all, other flights can still be taxibotted. However, the number of taxibot movements for the other airlines is still limited.

III. Results & Discussion

The experiments specified before were run, and the results are analyzed in this section.

A. Solution Evaluation

Before being able to analyze the results, some evaluation methods must be defined.

1. End of Day Schedule Evaluation

To evaluate the performance of the schedule, it is compared to the actual (end of day) arrival and departure times. This is done by creating a list of planned flights for each taxibot. The algorithm works as follows. The list of scheduled tasks for a single taxibot is sorted according to their actual required starting time. Charging tasks are also included in this list, and are fragmented in single tasks the length of a time slot. The algorithm goes through this task list from beginning until the end, appending the feasible tasks to the actual task list. If the next task is a charging task, it will only append the charging tasks up until the moment that the taxibot must leave for the next feasible flight that was scheduled to be taxibotted. When appending the tasks, the taxibot's state of charge is monitored. If there is not enough battery left to go back to the depot after the flight, the flight will not be taxibotted. This should prevent taxibots from stranding with an empty battery in the field. This is done for each taxibot individually; one taxibot does not take any tasks that were scheduled for another taxibot.

Figure 2 shows an example of what the taxibot schedule was at the beginning of the day (grey bars, only flights scheduled to be taxibotted are shown), compared to what tasks were executed (green) and what tasks were missed (red). Fuel savings that were scheduled but not executed in the actual schedule can be used to quantify the performance of the schedule. One could also look at the percentage of flights that could not be taxibotted. This process could validate the use of scheduling the taxibots.

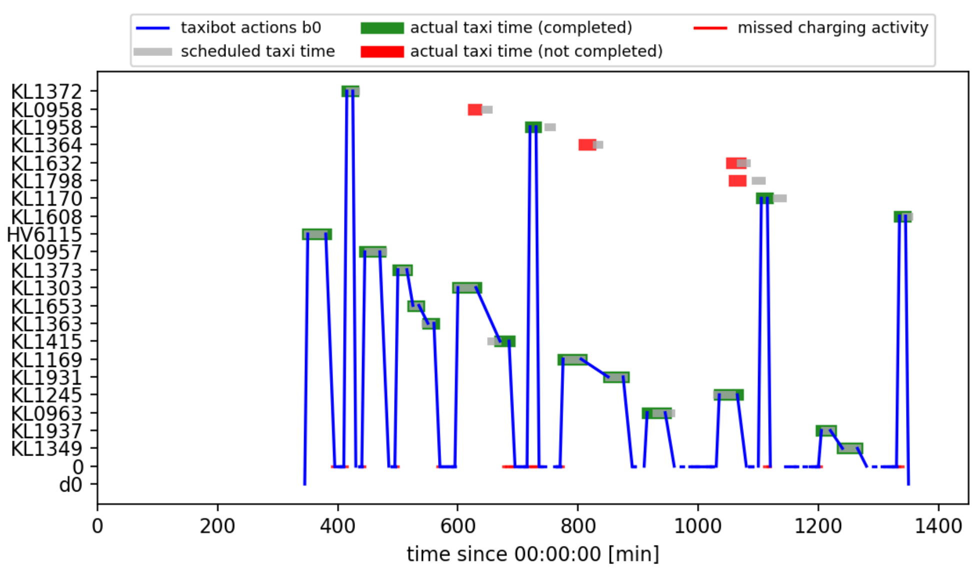

Figure 2. Example of the taxibot schedule compared to the actually executed taxibot trips.

2. Ad Hoc Baseline

In order to assess the true effects of planning, an ad hoc schedule is also made. This ad hoc schedule is made using the true arrival and departure times of flights. Each taxibot is assigned the first compatible flight that can be taxibotted, taking into account the time required to drive to the starting location. Should a taxibot not have enough energy left in its battery to complete this flight, it is assigned to the nearest available charger, where it will charge to the maximum allowed battery percentage (as specified by the user) before continuing taxibotting tasks. If there is no charger available at the time of running empty, the taxibot must wait until a charger is available, and no taxibotting tasks can be completed in the meantime.

Having to wait for chargers to become available is hypothesized as one of the major disadvantages of ad hoc task allocation. It is expected that, especially for electric taxibots with smaller battery capacities, the ad hoc schedule results in a poorer performance for this reason. Depending on how the taxibots will be deployed, airlines will not be able to anticipate being taxibotted (in terms of additional taxi time and fuel), which is not desirable. Additionally, fairness towards participating airlines cannot be guaranteed or steered upon, which is one of the desires from Schiphol. Another disadvantage is that this method does not optimize for fuel savings; taking the next available aircraft does not allow for choosing a more polluting aircraft. Despite these disadvantages, the ad hoc schedule does not need to keep uncertainties into account. This makes this way of scheduling computationally efficient. This could also mean that on days with many delays, more aircraft can be taxibotted, as the ad hoc schedule does not require waiting for the next scheduled flight.

B. Robustness

As flight delays vary from day to day, and these delays can cause large differences when comparing a schedule to the actuals, it was chosen to look at multiple days. For this experiment, the entire month of March 2022 is taken. The (retrospective) runway configuration data is only available from 2022 onward, so time in 2022 had to be chosen. Additionally, due to the operational difficulties at Schiphol during the May 2022 holidays and summer, passenger limits were imposed. March 2022 is the last month in 2022 without these disruptions and limitations. Of course, weather can also cause disruptions. Looking at the weather for March 2022 in Figure 3, no extreme weather was observed. Also the wind speeds were moderate in March 2022 [4]. Should a day still be more disturbed than others, this should be visible in the results of the experiment.

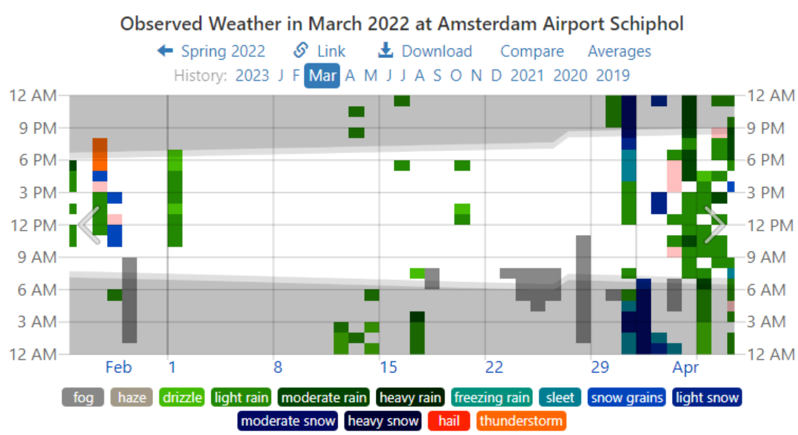

Figure 3. Hourly observed weather in March 2022. Should two types of weather be observed simultaneously, the most severe weather is shown [4].

For each day in March 2022, the model is optimized for 6 different maximum conflict probabilities (0.0, 0.2, 0.4, 0.6, 0.8, 1.0). Note that outliers are removed from the flight distributions, so 0% overlap probability is possible. All runways are taken into account, to avoid differences due to runway usage. Two identical electric taxibots are available for taxibotting, only being able to taxibot 73W ICAO subtype aircraft (Boeing 737-700 Winglets). This results in the model optimizing for 40-70 flights per day. There is only one charger available, so the two taxibots cannot charge simultaneously.

After running all instances, the scheduled to be taxibotted flights could be used to find which of those scheduled taxibot trips could actually be completed on the day. The ad hoc algorithm is also shown in these results for comparison. Please note that the ad hoc algorithm has no maximum allowed conflict probability, as it is not planning ahead.

Figure 4a shows a box and whisker plot for each allowed conflict probability over the 31 days of March 2022. It shows the total fuel savings due to taxibotting for the ad hoc, scheduled, and actual solutions. As can be seen, there is an outlier for each conflict probability. This outlier is March 1st; it does not perform very well in this metric. On March 1st, the Polderbaan was not used during the day due to adjustments being made to the runway lighting [1]. Flights from the Polderbaan need to taxi longer and therefore result in a larger fuel savings when taxibotted. Although March 1st has an average number of compatible flights (46 flights on March 1st compared to 48.6 flights on average), the total potential fuel savings from these 46 flights is only 2592 kg as compared to 4768 kg on average in March. This confirms that the Polderbaan is a good choice to start taxibot operations with.

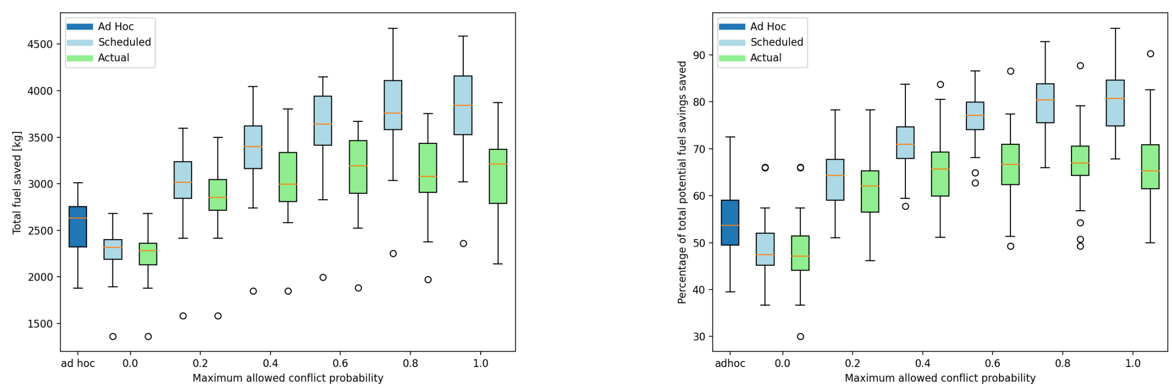

(a) Fuel savings boxplot for March 2022

(b) Average percentage of potential fuel savings

Figure 4. Fuel savings results for different allowed conflict probabilities

In order to come to better conclusions in terms of uncertainty incorporation, the fuel savings are plotted as a percentage of the total potential fuel savings. This is shown in Figure 4b. As can be seen and expected, the less risk that is taken, the lower the scheduled fuel savings. The scheduled fuel savings increase asymptotically with increasing the maximum allowed conflict probability; increasing the risks taken leads to less and less additional fuel savings. However, when looking at the actually realised fuel savings, this asymptotic behavior is amplified. Allowing only 40% conflict probability leads to very similar actual fuel savings as compared to allowing 100% conflict probability.

Comparing the results from Figure 4b with the ad hoc algorithm, it can be seen that ad hoc only performs better than allowing 0% conflict probability. As compared to other boxes in the boxplot, the ad hoc's spread is also larger, meaning that it would be difficult to accurately predict ad hoc fuel savings for other days.

Schedule reliability is also important, especially for outbound flights (for fuelling reasons). Figure 5 shows the percentage of scheduled taxibot flights that could be taxibotted in the actual schedule. Ad hoc is not taken into account in this case, as it does not schedule flights.

[1]https://bezoekbas.nl/project/28-02-2022-polderbaan-overdag-buiten-gebruik-op-1-maart/

It can be seen from Figure 5 that allowing 0% conflict probability creates the most reliable schedule; most days all flights that were scheduled to be taxibotted could also be taxibotted. This percentage drops quickly with increasing allowed conflict probabilities. It follows a similar asymptotic trend as in Figure 4b. However, as the difference between scheduled and actual fuel savings increases with allowed conflict probability, the reliability of the schedules as shown in Figure 5 shows the asymptotic behavior less.

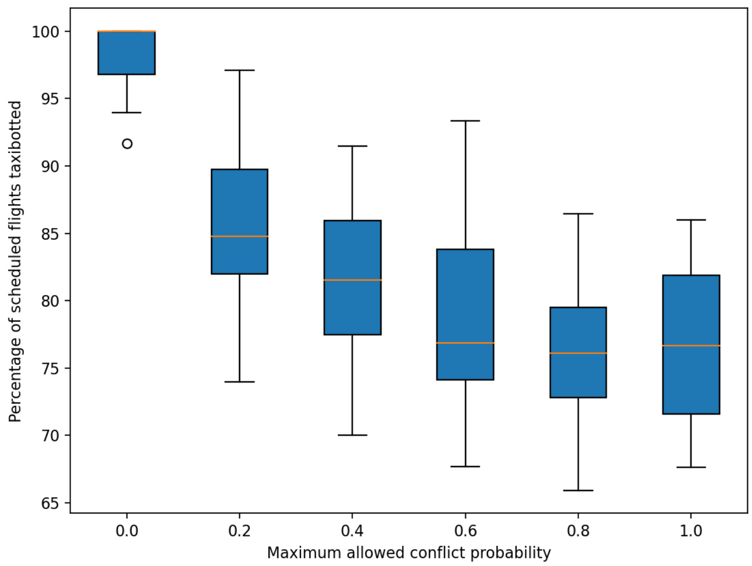

Figure 5. Percentage of scheduled taxibot trips that could be performed for different conflict probabilities

It was concluded that in the majority of the cases the ad hoc algorithm scores worse than the 'actual' schedule. It may therefore also be interesting to know how much worse the ad hoc algorithm scores on average. The percentage difference between the actual fuel savings and the other solutions is shown in Table 1. It is clear that the ad hoc algorithm only performs better when very little risk is taken in the schedule. In all other cases, on average, the ad hoc schedule performs significantly worse than the actual execution of the schedule. The table also shows that the difference between the scheduled fuel savings and the actual fuel savings increases when taking more risk, while the difference with the total potential fuel savings does not change much when more than 40% conflict probability is allowed.

Table 1. Average percentage difference with actual fuel savings for different maximum conflict probabilities

CP_{max}	**0.0** [%]	**0.2**[%]	**0.4**[%]	**0.6**[%]	**0.8**[%]	**1.0**[%]
Scheduled	+2.37	+4.30	+9.51	+16.31	+20.76	+22.43
Ad hoc	+15.37	-10.20	-16.30	-17.38	-17.99	-17.60
Potential	+114.96	+66.43	+55.13	+53.33	+52.19	+53.17

1. Recommendations

The effects of adding the stochastic constraints are clear in the results; the lower the maximum allowed conflict probability, the better the schedule could actually be executed. This means that stochastic constraints are an effective method of achieving robustness.

Using Figures 4b and 5, it can be concluded that allowing 40% conflict probability results in a favorable balance between maximizing fuel savings and reliability of the schedule. Nonetheless, should reliability of schedule be a larger factor, less risk should be taken.

It is recommended to investigate the uncertainties further. It could be interesting to investigate different maximum allowed conflict probabilities for different planning horizons and moments. For example, it is hypothesized that more risk can be taken when (re)scheduling each hour, as more information becomes available. Also, more research should be done in terms of the other uncertainties in the operation, to validate the choice of only including flight delays in the model.

It was also concluded that the Polderbaan plays a large factor in fuel savings; when the Polderbaan is in use, aircraft with origin or destination Polderbaan are likely to be taxibotted. This also means that when the Polderbaan is not in use, but it is predicted (from weather) that the Polderbaan will be in use again soon, that the taxibots should be as available as possible. This would mean ensuring that the taxibots have enough energy left in their batteries to taxibot the more polluting taxi trips ahead. In the current model, this is already done, as the actual (retrospective) runway configurations are used. However, in the future, when runway configurations are not known beforehand, a predicted runway configuration change could have a large impact on the schedule.

C. Healthy Working Environment

A similar base scenario is used as for the experiments regarding robustness: all runways available for taxibotting, all airlines participate, all 73W subtype aircraft are compatible with the 2 available taxibots, and there is one charger located in a bay. However, it is assumed that the runway configuration during the day has significant impact on the results, especially regarding the Polderbaan. Therefore, it was decided to run the experiments on three different types of days: days when the Polderbaan was primarily used for inbound flights (3rd, 9th, and 12th of March 2022), days when the Polderbaan was primarily used for outbound flights (2nd, 6th and 31st of March 2022), and days that the Polderbaan was not used at all (October 10th to 13th 2022, as the Polderbaan was under maintenance on those days [2]). The experiment is run for six different outbound weights: 0, 0.2, 0.4, 0.6, 0.8, and 1.

The results for the three different types of runway usages and the different outbound weights are shown in Figures 6a-6c. It is visible that the different types of days respond differently to the outbound weight.

On days that the Polderbaan is mainly used as inbound runway, changing the outbound weight has the most effect, as seen in Figure 6a. On these types of days with an outbound weight factor of 0, many flights coming from the Polderbaan would be taxibotted, as those are the biggest fuel savers. However, when more outbound flights must be taxibotted due to a higher outbound weight factor, other choices must be made. Nonetheless, the percentage of potential fuel savings saved does not drop as drastically. When the outbound weight factor is changed from 0.2 to 0.4, the percentage of outbound flights taxibotted increases by nearly 12%, whereas the percentage potential fuel savings saved drops less than 3%. This shows that on days that the Polderbaan is mainly in use for inbound flights, the outbound weight factor can easily be used to adjust the outcome to suit the current needs.

On days that the Polderbaan is primarily used as outbound runway, changing the outbound weight factor does not change much, as seen in Figure 6b. This confirms that outbound movements for the Polderbaan are a win-win situation; the largest fuel savers are taxibotted, as well as as many outbound movements as possible. Note the different scale on the percentage outbound axis compared to Figure 6a; changing the outbound weight only makes a 1.4% difference. It can be concluded that changing the outbound weight factor on days that the Polderbaan is in use as an outbound runway makes an insignificant difference.

On days that the Polderbaan is not in use at all, the difference is small as well, as seen in Figure 6c. However, it can be seen that both the percentage of outbound flights taxibotted and percentage of total potential

[2]https://www.schiphol.nl/nl/schiphol-als-buur/pagina/baanonderhoud-en-werkzaamheden/

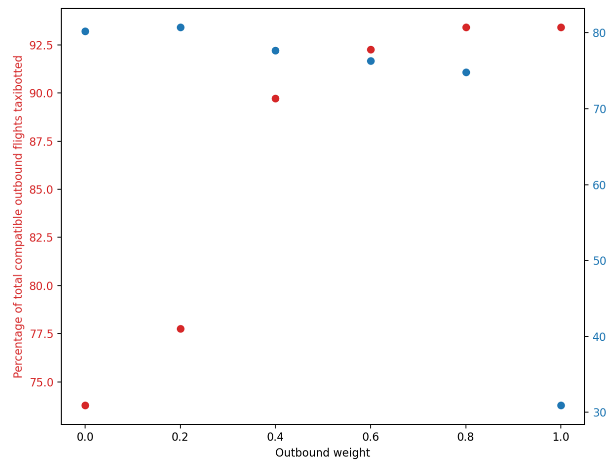

(a) Results for days with Polderbaan as inbound runway

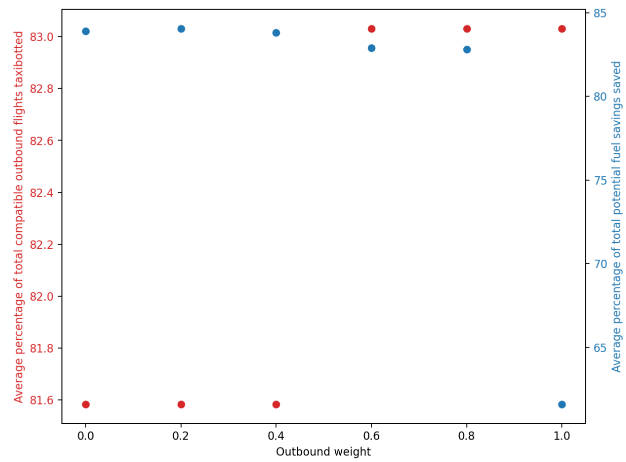

(b) Results for days with Polderbaan as outbound runway

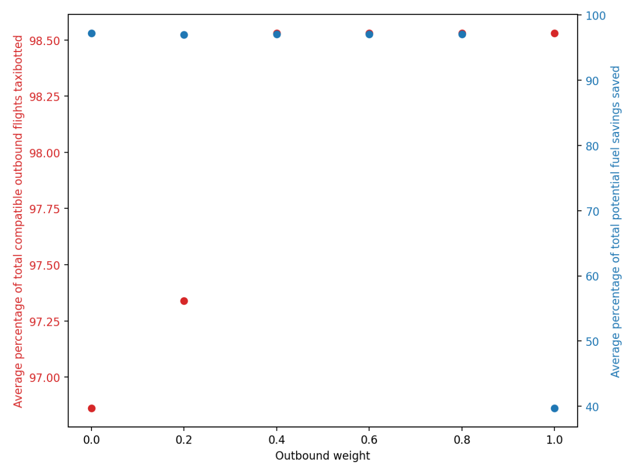

(c) Results for days without Polderbaan in use

Figure 6. Average percentage of outbound flights and fuel savings for different types of runway configurations

fuel savings saved are higher than on the other types of days. Nearly all flights can be taxibotted when the Polderbaan is not in use. This is likely due to the long driving distances to the Polderbaan. If the Polderbaan is not in use, the fuel saved per flight may be lower, but more flights can be taxibotted and platform employees should have a healthier working environment. For this selection of flights and runways, no true conclusion can be drawn, as too many flights can be taxibotted to see the true effects. The experiment would have to be repeated with a larger number of flights, or fewer taxibots, to reach conclusions.

For all three types of days, focusing fully on outbound flights ($W_o = 1$) yields undesirable results. The percentage of outbound flights taxibotted does not change, however the percentage of potential fuel savings saved drops drastically, for all three types of days. This is because the model has no incentive to increase the fuel savings by the objective function in Equation 6 when $W_o = 1$. For a healthy working environment, a good balance between outbound movements and total fuel savings is desired, such that the number of outbound movements is high, while also optimizing for fuel savings (and therefore emissions).

One might notice that the average percentage fuel savings for $W_o = 0.2$ are slightly higher than when $W_o = 0$ in Figures 6a and 6b. This is due to the time limit set for the optimization process, as well as the 5% MIP gap termination criterion. This time limit is necessary for the completion of the experiments in a reasonable amount of time. This does, however, mean that solutions are not necessarily optimal and that, in this case, the results for $W_o = 0.2$ are closer to the optimal solution than the results for $W_o = 0$. However, by definition of the objective function 6, should the optimizer be given enough time the results for $W_o = 0$ would yield higher/equal fuel savings as for $W_o = 0.2$. Nonetheless, it was considered that the differences in this case are minimal and that the conclusions drawn are still valid. It is recommended to further validate the results of these experiments for more days to allow for better statistical averages to be made.

D. Fairness

For this experiment it is required to select only the participating airlines. To select the compatible aircraft types, the fleet of the four airlines is compared. Corendon has the most limited fleet, with only the 737-800 Winglets (73H) and 737-800MAX (7M8) aircraft. Fortunately, the other three airlines all have the 73H aircraft type in their fleet as well. To increase the chances of having a Corendon and/or TUIfly flight in the set, also the 7M8 aircraft type is selected. As the fairness division is most relevant in the trial phase of the sustainable taxiing project, only taxibot trips to/from the Polderbaan and Zwanenburgbaan are selected. It was verified that those runways were in use on March 2nd to 15th 2022, so those dates were selected for this experiment. No uncertainty is taken into account, there are two taxibots and one charger. The fairness factor is varied between 0 and 1 in steps of 0.2.

The results for the different fairness factors for the fourteen chosen days are shown in Figure 7. The boxplot in Figure 7a shows that in general, a decrease in fuel savings is expected with higher fairness factors.

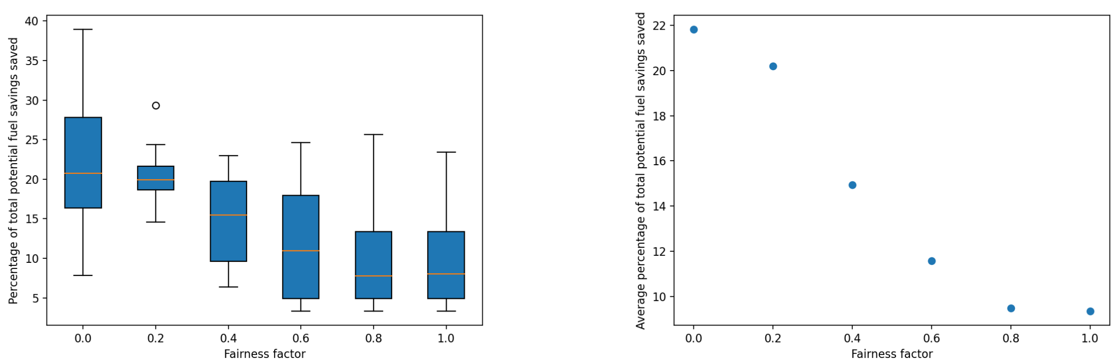

(a) Fuel savings vs fairness factor boxplot

(b) Average fuel savings vs fairness factor

Figure 7. Results of fuel savings vs fairness factor experiments

Figure 7a also shows one of the large drawbacks of the model; the size of the model increases drastically with the number of flights, therefore increasing solving time. In Figure 7a, it can be seen that there are days that perform better in terms of fuel savings when the fairness factor is high. By definition of Constraints 7 and 8, this should not be possible. In order to keep the research within a reasonable timeframe, a time limit of 50 minutes was set to solving the model. However, the MIP gaps for some instances were still much larger than the goal of a 5% MIP gap, explaining the irregularity. This happened in these instances and not in previous experiments, as the number of flights taken into consideration in these instances is roughly twice as large as the experiments done before.

When looking at average percentage fuel savings in Figure 7b instead of the boxplots, some conclusions can still be drawn. As expected, the fuel savings drop when increasing the fairness factor. However, there is little to no difference between a fairness factor of 0.8 and 1.0. This has to do with Constraints 8; should the smallest airline have two flights taxibotted in a day, the left hand side (the upper bound for other airlines), becomes 3.75 for a fairness factor of 0.8, and 3 for a fairness factor of 1.0. This means that this change in fairness factor does not add an additional flight for other airlines. Only on the 5th of March in the days tested did the smallest airline have more than 2 flights; it had 4. This one day did not have much of an impact on the average percentage fuel savings in Figure 7b, though it can be seen from the high percentage savings whiskers in Figure 7a.

1. Recommendations

It would be recommended to simplify the model to allow for larger amounts of flights to be investigated. However, each simplification comes at the cost of accuracy. If the model cannot be simplified, it would be recommended to attempt solving the instances again, without a time limit. Though this will likely take a long time, it would likely give better results.

Additionally, the implementation of fairness should be assessed. For this research it was chosen to look at absolute numbers, but perhaps a percentage wise distribution is more applicable. Similarly, for this research it was chosen that airlines that do not fly on a given day are still considered in the fairness constraints. This could also be altered, and would give vastly different results. However, this would likely result in an even larger spread in the aggregated results, as the spread in the number of flights per airline would be even larger.

Although the added constraints can aid in dividing the taxibot trips in a fair manner, it is questionable whether it is the best implementation of fairness. As stakeholders are a vital to keep the project running, more research must be done as to how this fairness can be incorporated into the project outside of taxiing. For example, it can be deemed reasonable to collect all costs and benefits of sustainable taxiing in one place and then distribute these fairly.

Not driving the taxibot for fairness misses the point of the sustainable taxiing project. This conclusion would likely be different if the participating airlines were more similar in size, however, this is not the case for Schiphol.

E. Other applications

The developed model also has the capability to be used for the design and development of the sustainable taxiing project. As the electric taxi bots do not exist yet, some design parameters for the taxibots could be experimented with using the model. The results of those experiments can then, for example, be used to set requirements to the developer of the electric taxibots.

Similarly, the model could aid in designing infrastructure for the project. The charging locations of electric taxibots are not yet known, just like the grid requirements or charging speeds.

The model also takes taxibot size and compatibility into account. That means that the model could also aid in deciding next acquisition steps. If there is demand and budget for an additional taxibot, the model could aid in deciding whether that should be a large or a small taxibot.

Experiments were performed in terms of battery capacity and charger locations. The results of these experiments can be found in Appendix D.

IV. Conclusion

Overall, it can be concluded that the research objective was accomplished. The model is able to assign a given fleet of taxibots to incoming and outgoing flights, with different objectives in mind. The model can assign taxibots based on robustness, a healthy working environment, and fairness. Nevertheless, there are improvements possible and many more questions to be answered.

First of all, the model defined for this thesis is complex. It is accurate with few simplifications, however, this comes at a computational cost. Especially for larger instances, the exact solver used was unable to solve the model within reasonable time. This also led to some ambiguous results, especially in the experiments on fairness. Both the model and the solver should be considered for improvement.

Generally speaking, many factors can be investigated using this model. In this research, only one variable was changed at a time to see the effects of that one variable. However, a combination of variables changing could also have a large impact on the results. It is recommended to explore this model further, to learn more about all the relationships between inputs.

References

[1] "Vision 2050 Storyline," , 7 2020.

[2] Bresser, D., and Prent, S., "Sustainable taxiing: Taxibot trial," , 2020. URL `https://www.schiphol.nl/en/innovation/blog/sustainable-taxiing-taxibot-trial/`.

[3] van Schaijk, O. R., and Visser, H. G., "Robust flight-to-gate assignment using flight presence probabilities," *Transportation Planning and Technology*, Vol. 40, No. 8, 2017, pp. 928–945. doi: 10.1080/03081060.2017.1355887.

[4] WeatherSpark, " March 2022 Weather History at Amsterdam Airport Schiphol," , 3 2022.

[5] "ICAO Aircraft Engine Emission Databank," , 7 2021.

[6] Zoutendijk, M., and Mitici, M., "Probabilistic Flight Delay Predictions Using Machine Learning and Applications to the Flight-to-Gate Assignment Problem," *Aerospace*, Vol. 8, No. 6, 2021, p. 152. doi: 10.3390/aerospace8060152.

[7] Manthey, N., "DPD Switzerland gets electric truck with 680-kWh battery," , 6 2020.

A. Full EVSP Model

1. Variables and Objective

Before all the constraints are explained, the different sets, variables and parameters must be defined. Then, the objective is defined.

Sets:

$\mathbf{C}$	set of chargers
$\mathbf{C}_{loc}$	set of chargers at a certain charger location
$\mathbf{D}$	set of starting/ending depots
$\mathbf{L}$	subset of flights, dependent on the constraint
Locs	set of charger locations
$\mathbf{M}_c$	set of combustion vehicles
$\mathbf{M}_e$	set of electric vehicles
$\mathbf{M}$	set of all vehicles, i.e. $\mathbf{M} = \mathbf{M}_c \cup \mathbf{M}_e$
$\mathbf{N}$	set of taxibottable flights
$\mathbf{T}$	set of time slots

Decision Variables:

x_{ijk}	1 if vehicle k is assigned to trip j after it completes trip i, 0 otherwise
y_{ckt}	1 if (electric) vehicle k is being charged at charger c at time interval t, 0 otherwise
p_{kt}	charging power level [kW] applied to electric vehicle k during time slot t

Parameters:

D	duration of one time slot [min]
B_k	battery capacity of EV k [kWh]
SOC_k	state of charge of EV k, with bounds $[SOC_k^{min}, SOC_k^{max}]$ [%]
G_t	electricity grid capacity available during time slot t [kW]
$CAP_{C_{loc}}$	number of taxibots that can charge at a charger location
S_i	starting time slot of taxibot trip i
F_i	finishing time slot of taxibot trip i
W_{ik}	fuel savings of tour i when being taxibotted by vehicle k [kg]
E_i	energy required for vehicle k to do tour i [kWh]
TDH_{ijk}	number of time slots required for vehicle k to dead-head from tour i to tour j
EDH_{ijk}	energy required for vehicle k to dead-head from tour i to tour j [khW]
size_k	the size of a taxibot or flight [small/large]
M	a sufficiently large number (used mainly for "if" constraints)

The primary objective of this model is to maximize the overall fuel savings. This is denoted by Equation 9. However, constraints 10-32 are needed to ensure feasibility of the plan made. Each constraint will be elaborated upon.

$$\max \left[Z = \sum_{i \in \mathbf{N}} \sum_{j \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D}} \sum_{k \in \mathbf{M}} W_{ik} \times x_{ijk} \right] \tag{9}$$

2. Assignment Constraints
The assignment constraints take care of the basic aspects of the scheduling problem; they ensure that the assignment of taxibots to flights are realistic.

$$\sum_{j \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D}} \sum_{k \in \mathbf{M}} x_{ijk} \leqslant 1 \quad \forall i \in \mathbf{N} \tag{10}$$

Constraints 10 ensure that each trip i is handled at most once in the schedule. It does not matter what trip follows trip i or what vehicle k handles the trip.

$$\sum_{i \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D}} x_{ihk} - \sum_{j \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D}} x_{hjk} = 0 \quad \forall h \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D}, \forall k \in \mathbf{M} \tag{11}$$

Constraints 11 ensure the chaining of operations. If a taxibot heads to a task h ($x_{ihk} = 1$), it must also do a task afterwards ($x_{hjk} = 1$).

$$S_j \geqslant F_i + TDH_{ijk} - M\left(1 - x_{ijk}\right) \quad \forall i, j \in \mathbf{N}, \quad \forall k \in \mathbf{M} \tag{12}$$

Constraints 12 ensure the timing of operations. If trip i and j are completed by the same taxibot ($x_{ijk} = 1$), then the start time of the latter trip must be later than the finish time of trip i, plus the time it takes to travel from the end of trip i to the beginning of trip j.

$$-M\left(1 - \sum_{j \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D}} x_{ijk}\right) + \sum_{l \in \mathbf{L}_3} \sum_{j \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D} \setminus i} x_{ljk} + \sum_{l \in \mathbf{L}_3} \sum_{j \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D} \setminus i} x_{jlk} \leqslant 0 \quad \forall i \in \mathbf{N}, \forall k \in \mathbf{M} \tag{13}$$

Constraints 13 ensure that there is no overlap between trips assigned to the same vehicle. In this constraint, $\mathbf{L}_3$ is a set of trips with overlap with trip i, including the time that is required to deadhead to/from trip i. It states that if trip i is completed by a taxibot, regardless of the trip that follows ($\sum_{j \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D}} x_{ijk} = 1$), then any trip in $\mathbf{L}_3$ cannot be completed by that taxibot.

$$-M\left(1 - x_{ijk}\right) + \sum_{l \in \mathbf{L}_4} \sum_{m \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D}} x_{lmk} \leqslant 0 \quad \forall i, j \in \mathbf{N}, \forall k \in \mathbf{M} \tag{14}$$

$$-M\left(2 - x_{ick} - x_{cjk}\right) + \sum_{l \in \mathbf{L}_4} \sum_{m \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D}} x_{lmk} + \sum_{c_1 \in \mathbf{C} \setminus c} \sum_{t=F_i}^{S_j} y_{c_1 kt} \leqslant 0 \quad \forall i, j \in \mathbf{N}, \forall c \in \mathbf{C}, \forall k \in \mathbf{M} \tag{15}$$

Constraints 14 and 15 ensure that each taxibot can only be at one place at a time. In these constraints, $\mathbf{L}_4$ is the set of trips that happen between the end of trip i and the start of trip j. They state that if a taxibot is scheduled to do trip i, followed by trip j (with a charger trip in between in the case of constraints 15), it cannot be assigned any trip between the end of trip i and the start of trip j. These constraints are therefore complementary to constraints 11.

3. Charging Constraints
The charging constraints are necessary for electric vehicles. They ensure that battery percentages are within reason and that charging is done at the correct chargers at the correct times.

$$S_j \geqslant F_i + \sum_{F_i}^{S_j} y_{ckt} + TDH_{ick} + TDH_{cjk} - M\left(2 - \sum_{c \in \mathbf{C}} x_{ick} - \sum_{c \in \mathbf{C}} x_{cjk}\right) \quad \forall i, j \in \mathbb{N}, \forall c \in \mathbf{C}, \forall k \in \mathbf{M}_e \tag{16}$$

Constraints 16 are for when a taxibot goes to charge. Constraints 16 ensure that if a trip j is assigned to a taxibot k, but the taxibot charged before going there, then the taxibot must have had enough time to go from the trip before charging (i) to the charger (c), charge for an amount of time, and drive to trip j.

$$SOC_k^0 + \frac{100}{B_k}\left(\sum_{t \leqslant S_i} D \times p_{kt} - \sum_{l \in \mathbf{L}_1} \sum_{j \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D} \setminus i} E_{lk} \times x_{ijk} - \sum_{l \in \mathbf{L}_1 \cup \mathbf{C} \cup \mathbf{D}} \sum_{j \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D} \setminus i} EDH_{ljk} \times x_{ijk}\right) \leqslant SOC_k^{max} \quad \forall i \in \mathbf{N}, k \in \mathbf{M}_e \tag{17}$$

$$SOC_k^0 + \frac{100}{B_k}\left(\sum_{t \leqslant F_i} D \times p_{kt} - \sum_{l \in \mathbf{L}_2} \sum_{j \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D}} E_{lk} \times x_{ljk} - \sum_{l \in \mathbf{L}_2 \cup \mathbf{C} \cup \mathbf{D}} \sum_{j \in \mathbf{N} \cup \mathbf{C} \cup \mathbf{D}} EDH_{ljk} \times x_{ljk}\right) \geq SOC^{min} \quad \forall i \in \mathbf{N}, k \in \mathbf{M}_e \tag{18}$$

Constraints 29 and 18 are the battery constraints. Constraints 29 ensure that at the start of trip i the the initial state of charge plus the amount charged up until the start of trip i, minus the energy drained doing the trips before trip i, minus the energy drained deadheading to/from all the trips before trip i, is never more than the user-specified maximum state of charge. Similarly, Constraints 18 ensure that this value is never below the minimum specified state of charge. It must be noted that $\mathbf{L}_1$ contains all the trips up to trip i and $\mathbf{L}_2$ contains all the trips up to *and including* trip i (to ensure that the taxibot cannot be low on battery at the end of the trip).

$$\sum_{t=S_i}^{F_j} \sum_{c \in \mathbf{C}} y_{ckt} - M\left(1 - x_{ijk}\right) \leqslant 0 \quad \forall i, j \in \mathbf{N}, k \in \mathbf{M} \tag{19}$$

$$\sum_{t=S_i - TDH_{ci}}^{F_i} \sum_{c_1 \in \mathbf{C}} y_{c_1 kt} - M\left(1 - x_{cik}\right) \leqslant 0 \quad \forall i \in \mathbf{N}, c \in \mathbf{C} \cup \mathbf{D}, k \in \mathbf{M} \tag{20}$$

$$\sum_{t=S_i}^{F_i + TDH_{ic}} \sum_{c_1 \in \mathbf{C}} y_{c_1 kt} - M\left(1 - x_{ick}\right) \leqslant 0 \quad \forall i \in \mathbf{N}, c \in \mathbf{C}, k \in \mathbf{M} \tag{21}$$

Constraints 19-21 all have to do with when a taxibot is allowed to charge. Constraints 19 ensure that if a taxibot completes two consecutive trips i and j, it will not charge between the start of trip i and the end of trip j. Constraints 20 and 21 ensure that if a taxibot has just left a charger ($x_{cik} = 1$) or is moving to a charger ($x_{ick} = 1$), it is not allowed to charge at any charger c_1 while on those trips.

4. Depot Constraints

The depot is the starting and ending location of a taxibot. Some rules regarding depots are specified in Constraints 23-26.

$$\sum_{d \in \mathbf{D}} \sum_{j \in \mathbf{N}} x_{djk} = 1 \quad \forall k \in \mathbf{M} \tag{22}$$

$$\sum_{i \in \mathbf{N}} \sum_{d \in \mathbf{D}} x_{idk} = 1 \quad \forall k \in \mathbf{M} \tag{23}$$

Constraints 23 and 22 ensure that a taxibot leaves the depot exactly once and it enters the depot exactly once. This simply ensures the beginning and ending locations of the taxibot.

$$-M\left(1-\sum_{d\in\mathbf{D}} x_{idk}\right)+\sum_{l\in\mathbf{L}_6}\sum_{j\in\mathbf{N}\cup\mathbf{C}} x_{ljk}+\sum_{c\in\mathbf{C}}\sum_{t\geqslant S_i} y_{ckt}\leqslant 0 \quad \forall i\in\mathbf{N},\forall k\in\mathbf{M} \tag{24}$$

$$-M\left(1-\sum_{d\in\mathbf{D}} x_{dik}\right)+\sum_{l\in\mathbf{L}_5}\sum_{j\in\mathbf{N}\cup\mathbf{C}} x_{ljk}+\sum_{c\in\mathbf{C}}\sum_{t\leqslant S_i} y_{ckt}\leqslant 0 \quad \forall i\in\mathbf{N},\forall k\in\mathbf{M} \tag{25}$$

Constraints 24 and 25 are regarding the starting and ending locations of a taxibot. Depots are locations where a taxibot can be parked, or where it can charge. The taxibots are meant to be located at a depot at both the start and the end of the planning horizon. This means that no trips can be handled before a taxibot has left the depot, and similarly, it cannot handle any more trips after it has entered the depot. These rules are specified in constraints 25 and 24 respectively. $\mathbf{L}_5$ is the set of trips starting before trip i, $\mathbf{L}_6$ is the set of trips starting after trip i.

$$\sum_{k\in\mathbf{M}} x_{ijk}=0 \quad \forall i,j\in\mathbf{C}\cup\mathbf{D} \tag{26}$$

Constraints 26 forbid movement between depots and chargers. Although these constraints are not completely necessary (nothing is wrong with moving between depots/chargers), it is unnecessary movement and adding these constraints simplifies the formulation of the model in some other constraints.

5. Infrastructure Constraints

Lastly, some constraints are needed to make solutions feasible in terms of infrastructure.

$$\sum_{k\in\mathbf{M}_e} p_{kt}\leqslant g_t \quad \forall t\in\mathbf{T} \tag{27}$$

$$\sum_{c\in\mathbf{C}} P_c^{\min}\times y_{ckt}\leqslant p_{kt} \quad \forall t\in\mathbf{T},\forall k\in\mathbf{M}_e \tag{28}$$

$$p_{kt}\leqslant\sum_{c\in\mathbf{C}}\left(P_c^{\max}\times y_{ckt}\right) \quad \forall t\in\mathbf{T},\forall k\in\mathbf{M}_e \tag{29}$$

Constraints 27 ensure that the total grid power used at a certain time should not exceed the grid capacity. Constraints 28 and 29 ensure that the taxibot cannot charge faster or slower than the maximum and minimum charger speed.

$$x_{ick}=0 \quad \forall i\in\mathbf{N},\forall k\in\mathbf{M},\forall c\in\mathbf{C}\backslash\mathbf{C}(i) \tag{30}$$

Constraints 30 ensure that a taxibot can only charge at compatible chargers. Due to the nature of the model (each trip can only be completed once, a charging trip has no specific start and end time), it was not possible to treat a charging trip the same as a flight taxibotting trip. Therefore, to ensure that all other constraints can be met, each flight is assigned a unique set of chargers. This means that the model knows $n_{chargers}\times n_{flights}$ chargers. Constraints 30 ensure that after handling flight trip i, a taxibot cannot go to a charger that is not associated with that flight.

$$\sum_{j\in\mathbf{N}\cup\mathbf{C}\cup\mathbf{D}} x_{ijk}=0 \quad \forall i\in\mathbf{N},\forall k\in\mathbf{M}\text{ if size}_i\neq\text{size}_k \tag{31}$$

In order to ensure that only compatible aircraft are taxibotted, constraints 31 are added. It ensures that a flight i is not taxibotted by taxibot k if flight i and taxibot k are not compatible in size. This type of constraint could also be added for other forms of compatibility.

$$\sum_{k \in \mathbf{M}_e} \sum_{c \in \mathbf{C}_{loc}} y_{ckt} \leqslant CAP_{C_{loc}} \quad \forall \mathbf{C}_{loc} \in \mathbf{Locs}, \forall t \in \mathbf{T} \tag{32}$$

Constraints 32 are added to ensure that the number of taxibots charging at a charging location does not exceed the number of chargers at that location. For example, if a charging location has 3 chargers, no more than 3 taxibots can charge there simultaneously.

B. Present Uncertainties

- **Landing time** One of the most uncertain parameters is the landing time. For this research, the landing time (and not the in block time) is relevant, as a taxi-bot would pick up the aircraft from a location near the runway. Landing times are less uncertain as time progresses, and their uncertainty with time varies greatly with region of departure.
- **Off-block time** Off-block time is also rather uncertain. The whole turnaround process varies a lot in the time that is required, especially also due to the late arrival of some passengers.
- **Taxi time and route** According to pilots (who were spoken to privately), even when the runway configuration and gate is known, the pilot does not know (and cannot predict) the taxi route until they are told the route when they land. Similarly, other traffic can influence the taxi time of an aircraft. However, for each runway configuration, some taxi routes can be eliminated (one cannot usually cross an active runway). For simplicity, it is assumed that there is no other traffic on the taxiways, and a fixed taxi time for each runway-gate combination is taken. Although it is known that taxi times vary even for the same route, this variation is assumed to be insignificant compared to the variation in arrival/departure time.
- **Runway configuration** Runway configurations can change throughout the day, and Schiphol does not know them beforehand. Nonetheless, this uncertainty is considered to be difficult to take into account one day ahead. For example, a change from Kaagbaan outbound to Polderbaan outbound has such a large impact on taxi times that this exceeds the purpose of robustness (replanning would be necessary).
- **Runway assignment** When two runways are used for inbound or outbound flights, there are no clear rules as to what aircraft is assigned to what runway. According to a private conversation with an air traffic controller, it is likely that there are some patterns to be found, but those patterns are dependent on each runway configuration and flight destination. It was suggested to simply assume an East-West assignment; flights with origin/destination to the West of Schiphol are assigned the most Western active runway and vice versa. Incorporating this as one of the uncertainties would, just like runway configuration, go beyond the purpose of robustness. Therefore, the actual, retrospective runway assignment is used for this project.
- **Gate assignment** Initial gate assignment is done one day ahead. However, as conflicts emerge gates may change. A change in gate may result in a change in taxi time. However, as a gate change is very unpredictable, and the change in taxi time is insignificant compared to potential delays, these potential changes in gate assignment are neglected.
- **Weather** Weather can change from moment to moment and can, in turn, change delay patterns and runway configurations. Similarly, it is likely that the taxibot's performance is impacted by different weather conditions. However, as this change in performance is yet to be investigated, no reasonable estimate can be made for different weather conditions.
- **Exact battery/fuel use** Batteries will degrade with time and therefore the number of missions that are possible with a charge can vary within a fleet of taxibots. Similarly, with weather changes, traffic, and taxibot driver patterns, exact battery or fuel use of a taxibot during a mission can vary. However, as

these factors are also still unknown and the goal is to be fully autonomous at some point, this uncertainty is not included in this research.

- **(un)Loading time** (un)Loading time will likely vary, especially with less experienced drivers and different weather conditions. However, this uncertainty is expected to be insignificant compared to arrival/departure time.
- **Driver actions** Drivers may drive different from one another, they may be sick, or they may go on strike. These uncertainties are not included, as the goal is to have autonomous taxibots in the future.
- **Taxibot unscheduled maintenance/break-downs** When making a schedule ahead of time, the breakdown or unscheduled maintenance of a taxibot cannot be taken into account. Should this happen, the taxibot schedule should be replanned, which is considered out of the scope of this research.
- **Exact engine type** An individual aircraft type can have different engines associated with it. As the engines under one specific aircraft type perform similarly when it comes to emissions [5], one individual engine is assumed per aircraft sub-type. This means that all B737-800 will have the same engine type, but the B737-700 will have another engine type. However, engine warm-up times also vary from engine to engine (and is also dependent on weather and flight schedule).

C. Delay Study AAS

A full year of airline traffic was analyzed in order make more accurate presence probability distributions for the specific case of Schiphol. The data from 2019 was analyzed, as that was the last 'regular' year (no COVID-19). With a full year worth of data all factors that could be of influence can be analyzed.

The data was taken from the Central Information System Schiphol (CISS). The features that could be of interest were extracted from CISS and made managable for processing. From CISS, the following parameters are extracted for each flight: flight number, scheduled in/off block time, airline, gate assigned, the main handler, aircraft registration, aircraft category, aircraft type, actual landing/take-off time, actual in/off block time, the airline's "bestimated" in/off block time (most recent airline update), destination/origin airport, and in/outbound runway.

Using this data, the following parameters were added: origin/destination region (Europe, Asia, North America, Latin America, Africa, and the Middle East), deviation of the scheduled in/off block time, time of day (in blocks of 2 hours), month, and day of the week. All these parameters together should be enough to deduce a delay distribution for any flight.

In order to analyze which parameters are most of influence, the different values a parameter can take should be plotted. It was chosen to work with a cumulative distribution, such that the conflict probabilities between two flights can easily be computed. Outliers on the outermost 1% (in tardiness and earliness) of the distribution are not taken into account.

The features of each flight that were chosen to be analyzed are: airline, destination/origin region, main handler, aircraft category and type, time of day, month, and day of week. An example of the plots made is shown in Figures 8 and 9, which is specific for the 10 largest airlines at AAS (in no particular order). In an ideal case, the presence probability for a flight would be a step function; 1 after the scheduled arrival time, 0 before the scheduled arrival time. This allows for planning without conflicts. If the presence probability is much more spread out, as the case is for EJU airline (light blue), much more buffer time must be planned for potential delay (or earliness). This means that for planning, flights ideally arrive exactly on time, as opposed to early (which is usually better for passengers).

As can be seen from Figure 8, the different airlines have a large variation in their distributions of arrival time. Where TFL has a low probability (~20%) of arriving on time (or early), nearly 60% of the DAL flights arrive before their scheduled time. When it comes to departure time (Figure 9), the variation is less pronounced, but there still is a roughly 20% spread when it comes to departing on time. Note that only the largest 10 airlines are featured in these figures for clarity; the spread is much larger when taking the other (large) airlines into account as well.

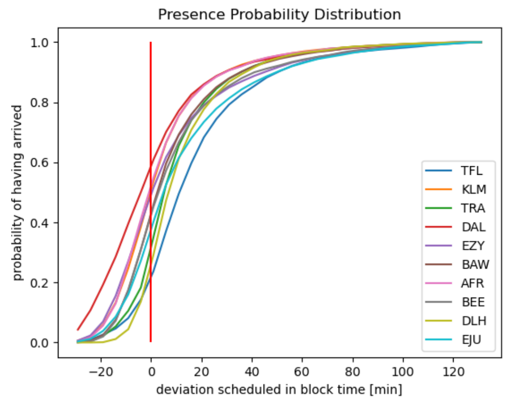

Figure 8. In block presence probability for different airlines

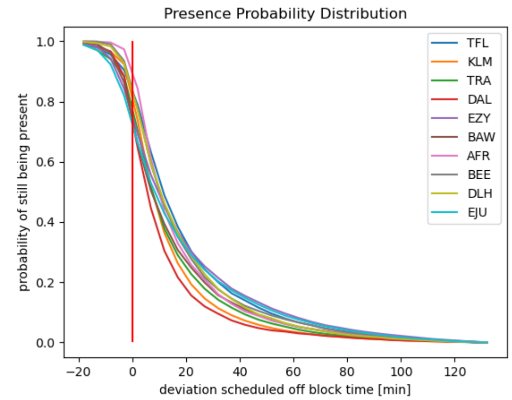

Figure 9. Off block presence probability for different airlines

After visual inspection of the aforementioned features and plots, it was concluded that the following features have the most significant spread in distributions: airline, aircraft type, origin/destination region, time of day, and month.

1. Individual Flight Analysis

In order to investigate the effect of each feature on the distribution of a single flight number (with all the same features), the different distributions are also plotted in one figure. Three different types of flights are shown in Figure 10: short haul budget airline, long haul flag carrier based in the Netherlands, medium haul flag carrier (not based in the Netherlands). All of these flights were scheduled to fly (nearly) daily in 2019, so each has a decent number of samples to make a distribution.

A number of things stand out from the different distributions. Looking at the THY flight, in terms of both arrival and departure it performs 'worse' than all of its features; the distribution is more spread out and there are more late arrivals/departures than each individual feature would suggest. Of course, a real distribution is not simply a sum of its components, but it is still notable. In general, looking at these three flights alone, one can already conclude that there is no clear relation between the true flight's distribution and its individual features.

Another noticeable relationship is the close relationship between the overall AAS distribution and the European region distribution. This is not strange, as a large majority (84%) of flights arriving at or departing from AAS are flying within Europe. However, this does mean that the 'region' feature has little to no added value for flights flying within Europe. For those two flights in Figure 10, it seems that the time of day is most representative of the delay distribution.

One thing to note is that none of these distributions take seasonality (month) into account. It was concluded that seasonality does have an effect on the distributions. This means that it would be desirable to take it into account. However, looking at a specific flight number for one month only would yield an insignificant amount of flights to make a reasonable conclusion. Therefore, seasonality will not be taken into account for this project, as assigning a weight to seasonality in order to affect the distributions has no scientific basis.

As an extensive investigation into these delay distributions is not within the scope of this thesis, some simplifying choices must be made. As it would be preferable for the model to make choices based off of the historic data, it was decided that flights that flew at least 100 times in 2019 take on the distribution of that

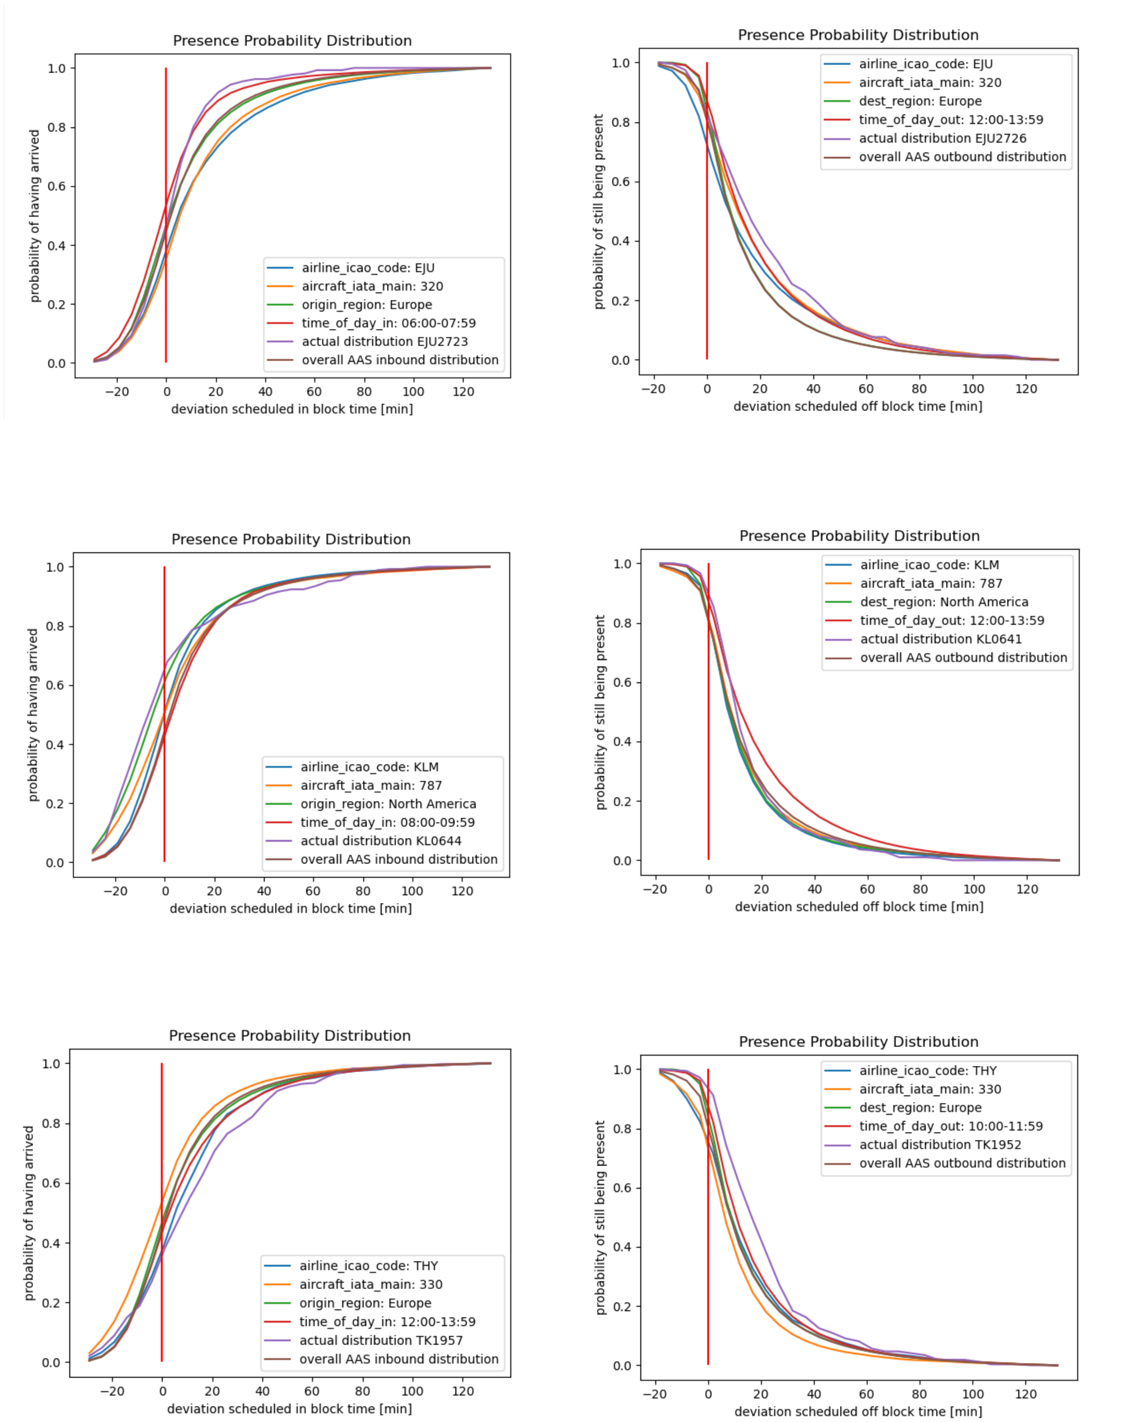

Figure 10. Presence probability distributions for different types of flights.

flight in 2019. This means that an individual flight only makes up 1% of the distribution. Roughly 46% of the outbound flights and 44% of the inbound flights meet this criterion.

Flights that did not fly 100 times in 2019 should also have a distribution. These flights take on the distribution of the region they fly to/from. Choosing region ensures that each new flight that may be added will have a distribution, despite no historical data being available. As each flight can be assigned to a region with a large amount of flights going there, there should be no problem with getting a distribution for yet non-existing flights. Had airline been chosen, some airlines may have an insignificant number of flights going to AAS, or a new airline may start flying at AAS, meaning that not all new flights would have a distribution. As it was noticed that the overall AAS inbound distribution is biased towards European flights, using region will be more accurate for intercontinental flights. This also means that the European flights will receive a rather generic distribution if they did not fly regularly in 2019, but choosing another feature would be arbitrary and inconsistent with the intercontinental flights.

It was a wish from Schiphol to look more specifically into the flights and airlines that are participating in the TaxiBot trials (Corendon Dutch Airlines, KLM, Transavia, and TUI fly). However, especially the holiday airlines (Corendon, Transavia, and TUI fly), do not fly to destinations regularly enough to make a meaningful distribution. KLM does, but those flights are taken into account through the 100 flights/year criterion.

Of course, much more extensive research can be done on these distributions. Seasonality should be added (also by analyzing more years), and a machine learning algorithm could help predict the delay distributions, even for new flights. In [6], two probabilistic forecasting algorithms are applied for flights at Rotterdam The Hague Airport. Doing such research for Schiphol could aid the accuracy of this planning algorithm, as well as the gate planning.

D. Using the Model for Design

To keep this thesis within reason, it was chosen to further investigate battery capacity and charger locations only.

A. Battery Capacity

One of the research questions was: *What are the effects of using electric taxibots versus diesel taxibots?*. To investigate this, the optimization is run for different battery capacities, after which it can be compared to a diesel taxibot (or an electric taxibot with infinite battery capacity). As it was noticed that the largest downfall of driving the taxibot ad hoc was the unplanned charging, the results for this experiment will be compared to an ad hoc schedule as well. It is expected that driving ad hoc becomes more viable with larger battery capacities.

The experiment was run using similar parameters as previous experiments: all runways available, all airlines, all 73W aircraft, 2 taxibots with 1 charger, fully deterministic and no fairness taken into account. Six different battery capacities were used. Using the same taxibot parameters as previously, it was determined that a trip to the Polderbaan and back would require approximately 90kWh. Assuming a battery range between 20% and 80%, this would require a minimum battery capacity of 150kWh. The remaining battery capacities are multiples of 150kWh, up until 750kWh, as that was the largest battery capacity deemed realistic (the largest electric truck battery in Europe has a capacity of 680kWh [7]). One fictitious battery capacity of 30,000 kWh was added, which is considered equivalent to a diesel taxibot (diesel fuel and emissions are not taken into account in the optimization).

1. Results and Analysis

Figure 11 shows the results of the fuel savings for different battery capacities. The scheduled, actual and ad hoc results are shown. By definition, the actual performance is worse than the scheduled performance, but interestingly the ad hoc performance is better than the 'actual' performance in both the small battery (150kWh) case and the diesel case. In the larger battery capacity cases, ad hoc performs worse than the actual executed schedule.

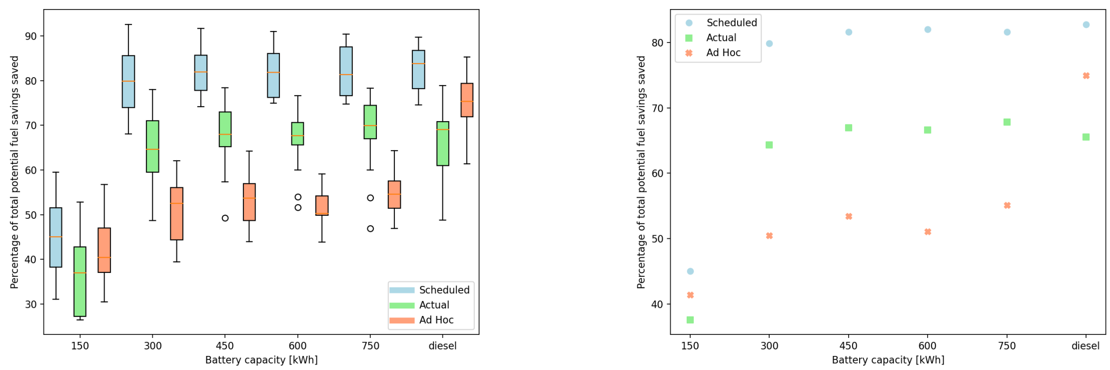

(a) Boxplot of the percentage of potential fuel savings saved **(b)** Average percentage of potential fuel savings saved

Figure 11. Fuel saving results for the different battery capacity

It was expected that ad hoc would perform better in the 'diesel' case, as the taxibots are free to go to the next compatible flight without charging limitations, so taxibot productivity would be high. However, why does ad hoc perform better with small battery capacities as well? Looking at the boxplot in Figure 11a, it can be seen that the scheduled plan for 150kWh still performs better than the ad hoc version, as it should. The ad hoc algorithm requires charging until 80% when it does charge. However, as mentioned before, 150kWh is the minimum battery capacity to go to the Polderbaan and back when keeping the battery state of charge between 20% and 80%. This means that, if a taxibot is desirable for a trip to the Polderbaan, charging until 80% would always need to be done, regardless of using the ad hoc or optimized schedule.

To verify this, the scheduled battery percentages for the two taxibots over time is plotted in Figure 12. It can be seen that the taxibots do not always charge until the maximum allowed (80%), but generally they do. It can also be seen that the red line, taxibot b0, is attempting to charge between 800 and 1000 minutes past midnight, however, it is hampered by the other taxibot, which also wants to charge. This means that, because so much charging is necessary with this low battery capacity, planning the charging still requires the two taxibots to wait for one another. Then, when this optimized schedule is applied to the actual flight schedule, some flights are bound to drop out due to delays. Therefore, for this small battery capacity, ad hoc performs better than the actual schedule.

Figures 11a and 11b also show that increasing the battery capacity beyond 300kWh does not increase the percentage of total potential fuel savings saved significantly. Even with a diesel taxibot, not much more can be saved, despite not needing to charge. It appears that the two taxibots have reached the limit of taxibot trips possible with the given flight schedule. Roughly 80% of the potential fuel savings is saved with a battery capacity of 300kWh, and it is likely that the remaining flights coincide with the taxibot trips that are being performed by the two taxibots. This would explain the marginal increase in performance for higher battery capacities. To verify this, the exact same experiment is performed with only one taxibot. It must be noted that with one taxibot, waiting for a charger to become available is no longer necessary, so other conclusions drawn from Figure 11 can not be compared directly.

Figure 13 shows the results of the experiment using only one taxibot. Although there is a slight upward trend with higher battery capacities, it is still not as defined as one might expect. When looking at the results for the actual fuel savings, with a battery capacity of 450kWh, on average only 0.6% less fuel could be saved

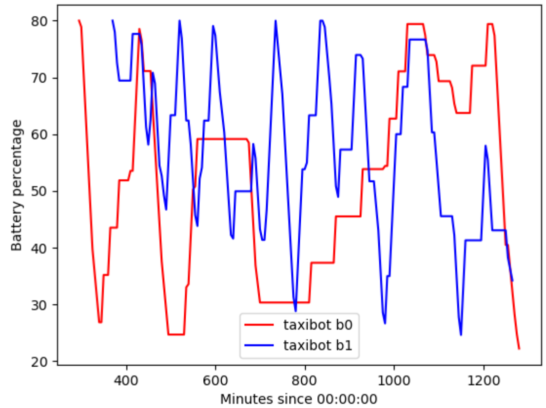

Figure 12. Battery percentages over time on March 3rd 2022, with a battery capacity of 150kWh.

than when using a diesel taxibot. It should be noted that the charger available did not change along with the taxibot battery capacity. This means that the charging speed is still the same, so regardless of battery capacity, 5 minutes of charging yields the same amount of energy. Nonetheless, the initial stint with a larger battery can be longer than with a smaller battery.

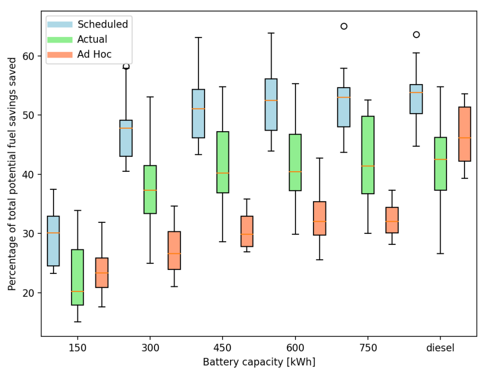

(a) Boxplot of the percentage of potential fuel savings saved

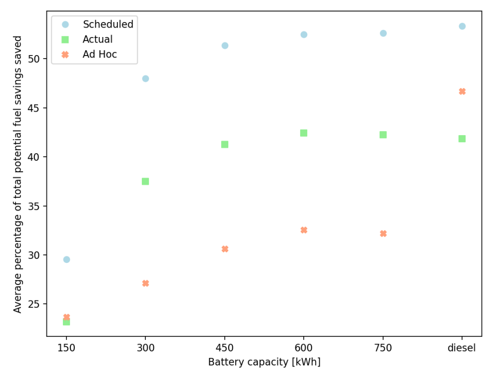

(b) Average percentage of potential fuel savings saved

Figure 13. Fuel saving results for the different battery capacities with only one taxibot

From theses results it can be concluded that a significantly larger battery does not yield proportional additional fuel savings. A battery larger than 450kWh (in combination with the charger for these experiments) would most likely not be worth it. However, more investigation should be done in terms of the charger. Different combinations of battery capacities and charging speeds should be tested. When choosing a battery capacity, battery degradation should be taken into account, which this model does not do.

B. Charger Locations

Currently, Schiphol does not yet own any electric taxibots. However, when they do, there must be a place where the taxibots can charge. This requires some infrastructure to be constructed. It can be imagined that making this infrastructure somewhere further away in the field is cheaper and easier than making this infrastructure close to, or even at the gates. However, it is expected that charging at the gates is more efficient, as fewer deadhead trips are needed for charging; a charger would always be located at the start or end of a trip. Therefore another six experiment types are devised. In each experiment, there will be two identical chargers. Where the chargers are located is denoted in Table 2. As a reminder, Figure 14 shows where each node is located.

As can be seen from Table 2 and Figure 14, it was chosen to have a configuration with both chargers in bays, both chargers in the field by the remote deicing platform, both chargers in the field by the cargo platform, one charger by the deicing platform and one by the cargo platform, one by the bay and one by the cargo platform, and last, a configuration with one charger in the bay and one by the deicing platform. These locations are theoretical.

Table 2. Locations of the chargers in different configurations

	config 1	config 2	config 3	config 4	config 5	config 6
Location 1	35	40	12	40	35	19
Location 2	19	40	12	12	12	40

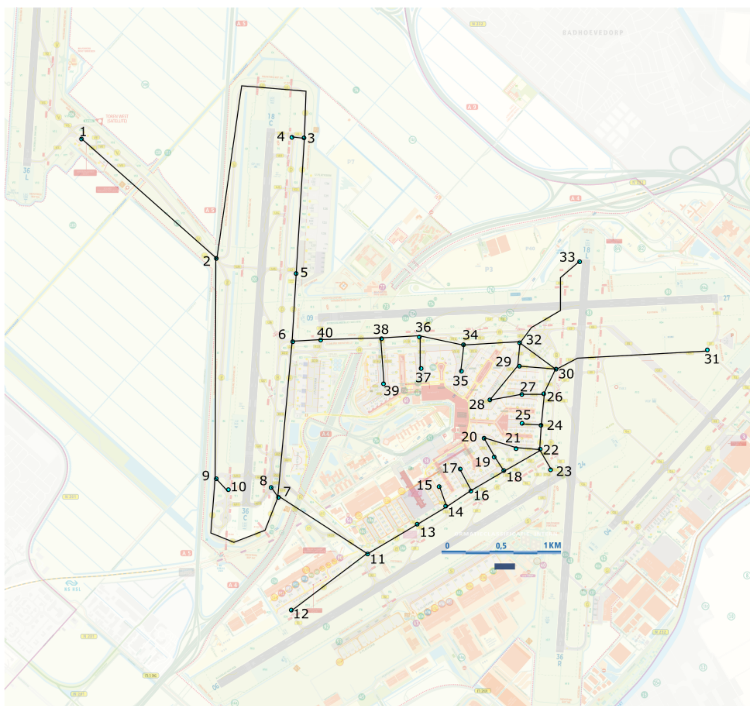

Figure 14. The graph made of AAS projected onto the map of AAS.

Other than the charger locations, the experiment set up is the same as the initial experiment for the battery capacities, with a 300kWh battery capacity for the two taxibots. However, as it is known that additional chargers increases the complexity of the model significantly, each instance is given a time limit of 40 minutes (instead of 20). It will stop short of the 40 minutes if a MIP gap of 5% is reached.

1. Results and Analysis

Figure 15 shows the results of the experiments performed. Figure 15a shows the boxplots for the different charger locations. From there, it can be seen that there is very little difference between the different configurations. One could argue that configuration 3 performs worst, as the interquartile range lies the lowest. However, configuration 1 has the lowest outlier. When looking at the averages for the two weeks in Figure 15b, it seems as though configuration 6 performs much better than the other configurations, and that configuration 3 performs worst. However, the difference between the two averages is only 3%, which could be considered negligible considering the termination criterion of a 5% MIP gap. This experiment therefore shows that charger locations do not really matter as much as initially expected. Nonetheless, charging at the gates would be convenient for drivers. However, it can be imagined that placing chargers at the gates is a large and complex operation. With these results, placing chargers further away from the gates, in places with more room for large infrastructural changes can be justified. There is little to no difference.

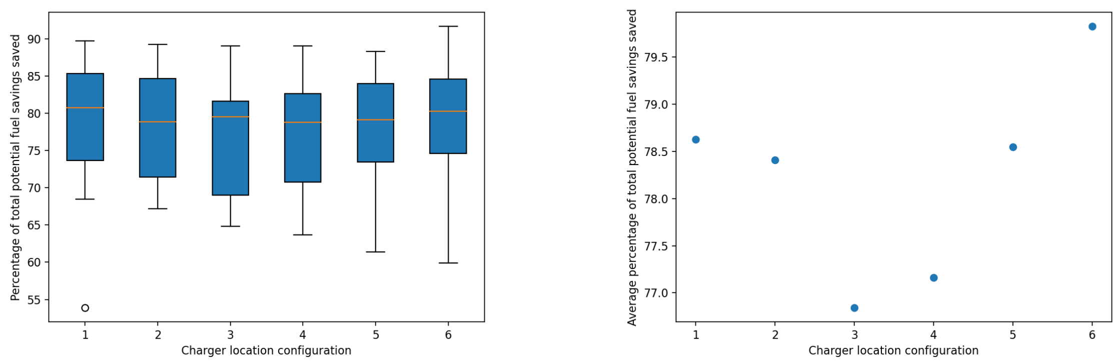

(a) Boxplot of the percentage of potential fuel savings saved **(b)** Average percentage of potential fuel savings saved

Figure 15. Fuel saving results for the different charger configurations

To understand where the different percentages come from, the percentage potential fuel savings saved is also plotted for each configuration for each day in Figure 16. It can directly be seen that each day performs differently; March 2nd and 10th perform rather well in all six configurations, while March 8th and 12th perform worse. There are many factors that affect this, such as flight schedule and runway configuration. In an attempt to explain some results, the latter was investigated further. However, as runway configurations change many times throughout the day, no conclusion could be drawn from those. It is advised to investigate the performance from day to day for more days, and a more controlled input in terms of runway configurations to draw more complete conclusions.

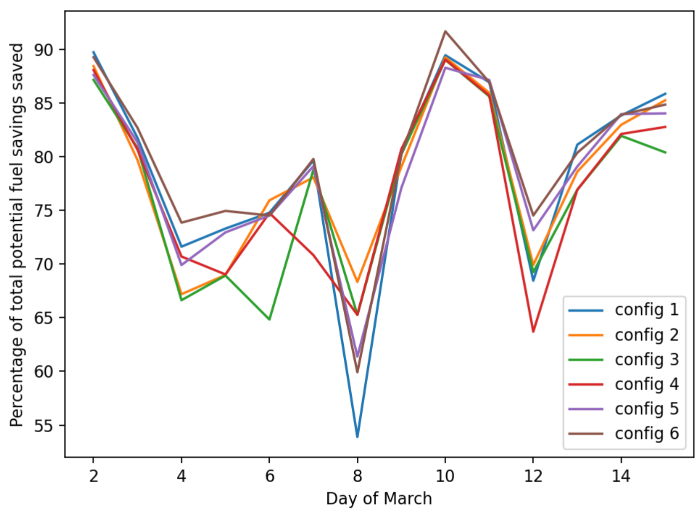

Figure 16. Percentage fuel savings for the different days in March.