MSc Transport, Infrastructure & Logistics

The Impact of TaxiBot Operations on Ground Traffic Flow at Amsterdam Airport Schiphol

R. M. N. (Ruben) Beumer

Master Thesis

TU Delft

Faculty of Civil Engineering and Geosciences

MSc Transport, Infrastructure & Logistics

The Impact of TaxiBot Operations on Ground Traffic Flow at Amsterdam Airport Schiphol

by

R. M. N. (Ruben) Beumer
4529804

June 19, 2024

In partial fulfillment of the requirements
for the degree of Master of Science
at the Delft University of Technology

The Impact of TaxiBot Operations on Ground Traffic Flow at Amsterdam Airport Schiphol

June 19, 2024

R. M. N. (Ruben) Beumer
Student nr: 4529804

Thesis committee:

Committee chair	Dr. J.A. (Jan Anne) Annema	Technology, Policy and Management
Daily supervisors	Dr.ir. A.J. (Adam) Pel	Civil Engineering and Geosciences
	Ir. P.C. (Paul) Roling	Aerospace Engineering
Additional supervisor	Dr. S. (Srinath) Mahesh	Civil Engineering and Geosciences
Royal Schiphol Group	P. (Philip) Kerbusch, MSc	Data and Analytics department

Cover image of TaxiBot and Boeing 737-800 by KLM. Retrieved from https://news.klm.com/klm-trials-sustainable-taxiing/

Preface

Looking back to when I started my graduation internship at Schiphol airport in October of 2023, I would not have expected to learn so much, and have so much fun, in such a short time. The experiences of driving around the apron, and even visiting the two TaxiBots in real life (and driving in them!) are memories that I will cherish for a long time. Countless people have contributed to this very formative part of my life, that has culminated in this master thesis lying before you.

First and foremost, I would like to thank the supervisors at Schiphol that guided me on a daily basis. Philip, Frank, and Sarah: thank you. Without your support on scoping the research, providing me with access to information, and helping me reach key people within the company, this thesis would be of far lower quality than what has been achieved now. It also brings me a lot of satisfaction to see the value that the insights from this work have brought to the sustainable taxiing team.

Of course my thanks also goes to the extended Schiphol community: all the people in the Innovation Hub and Data and Analytics departments, for making my time at the office so enjoyable. Thanks for the support and the fun Mario Kart sessions! Special thanks also goes to my fellow graduation interns, for being able to share experiences and challenges. I hope I was able to help you as much as you all helped me!

I also want to thank my supervisors at the university, for guiding me through the academic processes and keeping me on track. Adam, thanks for your questions and feedback on my work which really made me focus on the critical aspects of the simulations. It helped me a lot. Paul, thanks for our progress meetings and all the valuable background information you were able to provide on airport processes and aircraft information. We sometimes spent more time nerding out about aircraft than discussing my research, but I enjoyed it greatly. Srinath, thanks for letting me be your first ever graduation student! Your questions and feedback regarding emissions were helpful, and I hope I was also able to help you in gaining supervision experience. And of course Jan Anne, thank you for chairing the supervision committee. It feels quite special to me to receive both my Bachelors (Technische Bestuurskunde, 2018-2021) and soon my Masters degree from you.

Of course, I also want to thank my fellow students of the thesis room back at the university. Those coffee breaks that sometimes took hours, the delirious singing at 20:00 after we had been working way too long, helping each other through difficult moments, celebrating successes, playing little games, enjoying movie nights in that very same thesis room, partying (quite) hard whenever the opportunities presented themselves, and recently even starting the cycling group rides. There are some wonderful posters and quotes on the pin-up board, but the most memorable has to be a frog sitting on a rock, with a large cup of coffee, accompanied by the text "It's okay if all you did today was survive". Quite striking, and very true. Eva, Mathijs, Roxana, Blandine, Thaddäus, Madeline, Simon, Wouter, and the others: thanks!

Last, but definitely not least, the rest of my friends and family, for the countless little moments of guidance and support and for allowing me to take my mind off things. Especially for those times where I found my work-life balance shifting a little too far into the work-work direction, which probably happened more than I was aware of. And Anna, you probably remember that I sometimes joked that I could now draw the Schiphol taxiway layout by heart, from memory. Feel free to hold me to that!

Ruben Beumer
Delft, June 19, 2024

Summary

Schiphol airport, one of Europe's busiest, serves as a crucial hub with significant environmental and operational challenges due to its high traffic volume. A recent ruling by the Netherlands Labour Authority (NLA) has determined that emission levels – in particular those of ultrafine particles (UFP) – are too high in the apron bays between the piers of Schiphol. The NLA has established a 'green zone' around the apron bays where UFP emissions need to be significantly reduced. TaxiBots, which are powerful electric tow trucks capable of towing aircraft at their normal taxiing speeds, offer a potential solution by towing aircraft out of the bays or to the runways, eliminating the need for aircraft to use their engines.

Earlier research on TaxiBot operations at smaller airports has been conducted by Khammash et al. (2017) and Salihu (2020), but it is still unknown what the impact of TaxiBot operations will be on congested, complex airports. Furthermore, these studies primarily focused on emission reduction and cost-effectiveness, leaving a gap in understanding the potential congestion and impact on taxiing times. This study aims to fill that scientific gap.

The main research objective of this thesis is to investigate the impact of TaxiBot operations on the flow and emissions of taxiing aircraft at Schiphol airport. In particular, the research focuses on current patterns and bottlenecks in the traffic flow, the comparison of TaxiBot taxiing process times compared to those of regular operations, the flow inside apron bays, congestion around the unloading locations near runways, and the impact of TaxiBot operations on the emissions of CO2 and UFP at the airport.

To perform this research, a discrete event simulation model was developed using the Simio software package. The main input data for this model consists of real-life operational radar data from Schiphol for two selected days, each characterized by high activity on the Polderbaan runway and the airport overall. Nine experiments were designed, ranging from simple outbound operations to the Polderbaan runway, to operations including inbound movements, more runways, towing aircraft only to the edges of this defined 'green zone', and finally combining all policies into one maximum-complexity experiment.

The simulation results indicate that TaxiBots can reduce fuel consumption and associated emissions by up to 76% per towing mission, equating to a 32.9% reduction in total airport emissions. Flights with the longest taxiing times, particularly those to the furthest runways, experience the greatest reduction in CO2 emissions. Experiments permitting aircraft to start their engines during towing, after exiting the green zone to save time at unloading stations, result in outbound taxiing process times that are on average 1 to 1.5 minutes faster than the reference scenario. Moreover, the reduced time at the unloading stations, due to engine start-up during towing, significantly impacts waiting times, nearly completely eliminating the need for queuing at these stations. Inbound towing also shows promising results for emission reduction, but leads to additional taxi-in times of 2.5 to 3.25 minutes. Since TaxiBot operations allow aircraft to skip the engine start-up process in the bay, an average of three minutes per flight is saved inside apron bays. This reduced time can create additional gate capacity for the airport. Due to a lack of data on how much UFP is emitted by an aircraft engine, no quantitative analysis regarding UFP emissions could be performed. To mitigate this lack of data, heat maps displaying emission types (engine status) and their locations at the airport were generated.. Analysis of these heat maps shows that full-scale TaxiBot operations can have a significant contribution to the reduction of UFP emissions within the apron bays.

The findings suggest that implementing TaxiBot operations at Schiphol airport is both feasible and beneficial. Key benefits include reduced outbound taxiing time (1-1.5 minutes) and significant CO2 and fuel reductions (60-70%) per flight, with negligible impact on airport ground traffic flow given the TaxiBot's top speed of 22 knots. It is recommended to start aircraft engines during towing for all operational policies due to its positive impacts on outbound taxiing times, unloading station congestion, and CO2 emissions (reducing total airport emissions by an additional 1% compared to stationary start-up scenarios).

The findings are generalizable to other airports in a number of ways. The two relationships of reduction of outbound taxiing times to taxiing distance, and CO2 reduction to towing durations can be used to estimate the impact of TaxiBot operations at other airports. The findings regarding congestion at the unloading stations near the runway can also be generalized, because runways around the world have capacity limits of around 40 aircraft per hour (Airports Council International, 2023), and as long as an airport can make four unloading spots available near a runway and can allow engines to be started during towing, TaxiBot operations will not cause signification congestion near the runways.

Future work should focus on quantifying the additional gate capacity resulting from reduced time in bays, measuring UFP emissions once relevant data becomes available, and investigating the impact of TaxiBots on service road traffic at the airport from congestion and infrastructural perspectives.

This thesis demonstrates that TaxiBots have a promising potential to improve environmental sustainability and operational efficiency at large and complex airports. By effectively reducing taxiing emissions and improving the flow of aircraft on the ground, TaxiBots could be a key component in the future of green airport operations.

Contents

List of figures

List of tables

List of abbreviations

Abbreviation	**Definition**
AIBT	Actual in-block time, the time that an aircraft arrives at the stand. See also: AOBT.
AOBT	Actual off-block time, the time that an aircraft left the stand. Other types of off-block times (and in-block times) are EOBT (expected OBT), SOBT (scheduled OBT), and TOBT (target OBT). Similar abbreviations exist for in-block times (AIBT, EIBT, etc.), for when aircraft arrive at the stand.
APU	Auxiliary power unit, a smaller engine on an aircraft that can provide electrical power, for functions other than propulsion.
ATC	Air Traffic Control, the people in charge of aircraft movements on the ground and in the air.
CISS	Central Information System Schiphol, the central database storing a lot of flight information like aircraft types and departure times.
DES	Discrete Event Simulation.
EASA	European Union Aviation Safety Agency
IATA	International Air Transport Association, the international organisation representing airlines and related parties.
kts	Knots, or nautical miles per hour. This is a unit of speed commonly used in aviation. 1 kts = 0.514 m/s = 1.852 km/h
NLA	Netherlands Labour Authority (Dutch: *Nederlandse Arbeidsinspectie*). The Netherlands Labour Authority ensures healthy and safe working conditions in The Netherlands.
OTP	On-time performance, the delay (or lack thereof) of an aircraft. Negative values indicate an aircraft being ahead of time, positive values indicate delay.
SAF	Sustainable aviation fuel, a type of advanced bio-fuel usable by aircraft engines that reduces CO2 impact by 75% (IATA, 2023).
UFP	Ultrafine particulate matter ($PM_{0.1}$).

1 Introduction

This chapter presents the introduction to this thesis. First, a background section is dedicated to providing context around current airport operations, what sustainable taxiing entails, and what a TaxiBot is. Then, the objective of the research is defined in section 2, along with the research questions. Section 3 explains the scope of the research, and section 4 will detail the contribution of this work to the academic world and the aviation sector. Section 5 presents the structure of the rest of the chapters in this thesis.

1.1 Background: Schiphol airport, sustainable taxiing, and TaxiBots

Schiphol airport

Schiphol airport is situated to the south west of Amsterdam. It is the largest airport in the Netherlands, and among the busiest in Europe. The airport saw 442,000 flight movements in 2023 (Royal Schiphol Group, 2024c), welcoming over 61 million passengers. The airport forms an important hub for AirFrance-KLM, as well as serving as an important operating base for EasyJet, Transavia, Delta Air Lines, and many others. These airlines and their networks make Schiphol the best connected airport in Europe, and one of the best connected in the world (Royal Schiphol Group, 2024a).

The airport has six runways, of which usually three are in use simultaneously (in either a two arriving one departing, or one departing two arriving configuration), depending on wind conditions. The airport has a single-terminal concept, with eight piers (A to H, named counterclockwise) containing over 200 stands. Servicing six runways from eight piers, the airport has a complicated network of taxiways to make sure every aircraft can get to where it needs to go. Figure 1.1 below shows the overview of Schiphol's runway, taxiway and apron layout.

The process of an aircraft at an airport

The usual process of an aircraft performing a passenger flight, once it touches down on one of the six runways, is as follows. After exiting the runway, the aircraft will get instructions from ATC (air traffic control) on how to reach its destination at the airport. Driving on the ground is called 'taxiing', which is where 'taxiways' (the roads aircraft drive on) get their name from. The aircraft taxies over the instructed taxiways to reach its parking location, called a 'stand'. Usually, this is where passengers can also board or leave the aircraft (for the passengers, this is then called a 'gate'). The area where all the stands are located is occasionally called an 'apron'. For Schiphol specifically, these areas are also referred to as 'gate bays' or 'apron bays', due to them being surrounded by two piers.

Figure 1.1: The layout of Schiphol (RZjets, n.d.)

Once stationary, the aircraft will shut down its engines. Sometimes the aircraft still needs electrical power, for which a special engine called the 'auxiliary power unit' (APU) is used. This engine –located in the back of the aircraft – is a versatile part of the aircraft: it is also used to provide air-conditioning and is necessary to start the main engines.

At the stand, the aircraft gets replenished and refuelled, and its baggage offloaded, in a process called the 'turnaround'. Once the passengers and baggage have left the aircraft, and new passengers, baggage, fuel, supplies etc. have been taken onboard, the aircraft is ready to depart again. Because aircraft cannot drive backwards, they get 'pushed back' by a pushback truck. This truck positions the aircraft on the closest taxiway, after which they can start their engines. This usually takes anywhere from three to seven minutes, depending on the type of aircraft. Once the engines are running, the aircraft is ready to taxi over the various taxiways (again, following ATC instructions) to a runway, and once it has received clearance to take off, can depart to its next destination.

Aircraft sizes

These aircraft come in many shapes and sizes. The main (passenger) aircraft manufacturers are Boeing and Airbus, but the most important distinction between aircraft types for airports is their size. This size difference can be substantial. Aircraft types are generally classified in two categories: narrowbodies and widebodies, referring to the widths of their fuselages. A comparison between the two can be seen in the figure below. Usually, airports have a more complicated aircraft sizing system with more categories – Schiphol uses nine (Royal Schiphol Group, n.d.-b).

Figure 1.2: A Boeing 737 (narrowbody, bottom) and Boeing 777 (widebody, top) next to each other (Watts, A (Airlines.net), 2013)

Sustainable taxiing and the green zone

The consequence of all this aircraft activity is that it causes a lot of air pollution and even leads to traffic congestion. The air pollution is so severe, that the Dutch labour authority NLA (in Dutch: arbeidsinspectie) has ruled that the emission of exhaust gases has to be reduced significantly, within an area demarcated as the 'green zone' (see figure 1.3). This zone was established because this is where the majority of Schiphol's workers are active, and where the exposure to emissions, especially ultrafine particulate matter (UFP, $PM_{0.1}$) is greatest.

The majority of emissions within this zone are caused by the pushback procedure explained earlier in this section. Sustainable taxiing provides a way to move these aircraft out of the green zone (and beyond) without them needing to start their engines. As an additional benefit, the three to seven minutes that an aircraft is usually stationary within the apron bays is no longer needed either, meaning that there is potential for a lot of freed up congestion within these highly congested areas. This results in new aircraft being able to reach and leave their stands faster, and ultimately more aircraft can make use of the airport.

There are various forms of sustainable taxiing, like electrified nose landing gears, adjusted procedures to taxi with fewer engines running, and a process called dispatch towing. Dispatch towing is when aircraft are towed by a pushback truck (also called a tow truck), allowing them to move without using their engines for longer distances than the normal pushback procedure. For example, the truck can tow the aircraft all the way to its runway, to the edge of the bay, or to another location somewhere in between. These forms of sustainable taxiing will be explored in more detail in the literature review, in chapter 2.

Figure 1.3: The green zone of Schiphol (Netherlands Labour Authority, 2023a)

The TaxiBot

A TaxiBot is a special type of tow truck (Royal Schiphol Group, 2024b). It is more powerful than a regular one, meaning that aircraft can be towed at the speeds they would normally taxi, up to 45 kilometers per hour, or 23 knots (nautical miles per hour). Regular tow trucks are usually much slower, travelling about 20-30 kilometers per hour maximum. Another special feature of the TaxiBot is that it can be controlled by the pilot. That way, ATC can still give instructions to the pilot, and the pilot can taxi around the airport as they would normally do, just without using their engines. Contrary to what the name might suggest, the TaxiBot is not a robot; there is always a driver in the cabin that is able to control the TaxiBot and drive it around whenever it is not connected to an aircraft. The 'Bot' portion of the name comes from the semi-robotic mechanisms that let the vehicle imitate a nosewheel's movements, allowing pilots to control the TaxiBot like they would their own aircraft. Figure 1.4 shows a TaxiBot connected to an aircraft.

Figure 1.4: A TaxiBot connected to a Transavia Boeing 737 at Schiphol (Ground Handling International, 2022)

Currently, only narrowbody versions of the TaxiBot exist, and operations are certified for most but (crucially) not all narrowbody aircraft types – most notably, at the time of writing, the Boeing 737 Max series (Royal Schiphol Group, n.d.-d). A widebody TaxiBot is in development, and the manufacturer is also hard at work to certify the TaxiBot for more aircraft types (Israel Aerospace Industries, n.d.).

TaxiBots in operation

Schiphol currently owns two of these TaxiBots, and has used them to perform some trials and dry-runs (Royal Schiphol Group, 2021). Initial findings confirmed fuel savings of 50%-65% compared to standard taxiing procedures. The airport is currently working on the first 'TaxiBotting' missions where the aircraft will actually take off from the runway, after it has been towed. Schiphol is not the only airport where TaxiBot operations are being tested. Airports in Paris (Groupe ADP, 2024), Frankfurt (Airport Technology, 2015) and Brussels (Brussels Airport, 2024) are also performing their own trials, and two airports in India (Delhi, Bengaluru) (The Hindu Bureau, 2023) are actually already using TaxiBots in their daily operations. The first commercial passenger flight, an Air India Airbus A320, was towed by a TaxiBot in 2019. In May 2021, Delhi airport reported that they had carried out the 1,000th TaxiBotted movement (Airports Council International, 2021). As of May 2024, over 3800 missions have been performed in India, saving 710 tons of fuel, and preventing an estimated 2,245 tons of CO2 emissions (TaxiBot-India, 2024), for an average of just under 190 kg of fuel per flight.

Initial plans for TaxiBotting at Schiphol are focused on towing aircraft outbound from their stands to the Polderbaan (runway 18R-36L), in the north-west corner of Schiphol. This mission setup has been selected due to the long taxi distances involved (anywhere from four to ten kilometers, depending on stand location and route). This way, the highest fuel savings can be achieved. After the implementation of towing to the Polderbaan, additional runways (the Zwanenburgbaan, 18C-36C) and towing policies (to the edge of the green zone) would be added to the operations.

The aim for this thesis is to simulate TaxiBot operations that align with these plans as much as possible, in order to investigate what their impact would be on ground traffic and emissions. This is done because it is currently unclear if the airport's infrastructure will be able to accommodate the added complexity (and number of vehicles) that TaxiBot operations would introduce. Additionally, the research will assess what the environmental contributions would be in terms of CO2 impact, as well as UFP emissions to comply with the NLA's green zone ruling.

1.2 Objective and research questions

TaxiBots are able to tow aircraft at their normal taxiing speeds, without the need for them to use their engines. Some research has already been performed on the impact of TaxiBot operations at airports (section 2.2.2), but operations at large complex airports have not been investigated. This is relevant to know, because sustainable taxiing at large airports can lead to major sustainability gains, but the complexity of these airports also creates uncertainty regarding the impact that sustainable taxiing will have on the efficiency and on-time performance of daily operations.

This thesis therefore aims to investigate the impact of sustainable taxiing operations using TaxiBots on ground traffic flow a large and complex airport like Schiphol, by means of developing a Discrete Event Simulation (DES) model to simulate the airside ground traffic flow based on real-life radar movements. Some earlier research on this topic at smaller airports has been done by Khammash et al. (2017) and Salihu (2020), but it is still unknown what the impact of TaxiBot operations will be on congested, complex airports.

In the simulation, special attention is dedicated to process times at the stand and traffic flow within apron bays, as well as any congestion that might occur around the TaxiBot unloading locations and their capacity to meet (un)loading demand. Additionally, the travel times, distances, and engine data recorded by the model are used calculate the impact of sustainable taxiing operations on CO2 and ultrafine particle (UFP; $PM_{0.1}$) emissions at the airport.

Various levels of implementation of TaxiBot operations will be explored in a number of different scenarios, in order to analyse the impact that different levels of implementation will have ground traffic flow and emissions.

The main research question of this thesis is:

What is the potential impact of TaxiBot operations on the flow and emissions of taxiing aircraft at Schiphol airport?

To be able to comprehensively answer the main research question, several sub-questions are defined. Hypotheses for the questions – developed in collaboration with Schiphol airport – are listed below the research questions.

1. *What are current patterns and bottlenecks in the traffic flow at Schiphol airport?*
    - H1a: No bottlenecks are expected to be found, because air traffic control has already dealt with routing issues that could have arisen.
    - H1b: Traffic flow and taxiing speeds in and around the apron bays are expected to be lower than further out on the airfield.
2. *How do TaxiBot process times compare to those of regular operations?*
    - H2a: Outbound times are roughly equal.
    - H2b: Taxiing time is a bit longer.
    - H2c: Time is saved due to engine startup happening during towing operations.
    - H2d: Inbound times are slightly longer, due to TaxiBot loading.
3. *How do the TaxiBot operations affect apron bay flow?*
    - H3a: Sustainable taxiing operations allow aircraft to leave the apron bays quicker thus mitigating bay blockages.
    - H3b: Better flow in the apron bays will lead to higher on-time performance.
4. *How are the unloading locations near the Polderbaan affected by the envisioned operations?*
    - H4a: (Un)loading will lead to no/negligible queues.
    - H4b: The flow of in/outbound aircraft is not hindered when (un)loading inbound/outbound flights.
5. *What are the impacts of TaxiBot operations on emissions at the airport?*
    - H5a: TaxiBot operations lead to an overall reduction in CO2 emissions at Schiphol.
    - H5b: TaxiBot operations lead to an overall reduction in UFP emissions at Schiphol.
    - H5c: The relocation of UFP emissions due to engine start-up outside the apron bays, leads to fewer UFP in the apron bays, and thereby better air quality for the people working there.

All the above research questions will be answered via simulations using the DES model. The workings of this model will be further detailed in chapter 3, building upon work performed by earlier studies which are reviewed in chapter 2.

1.3 Scope

Airports are complex infrastructure hubs. Therefore it is important to carefully define the scope of the research. This thesis is focused on the taxiing process at Schiphol airport, based on real-life movements observed via the Schiphol ground radar. These movements serve as a foundation upon which TaxiBot operations are built, to evaluate their impact on traffic flow as realistically as possible. Each simulated experiment will encompass one day of traffic, with traffic levels that can be labelled as 'busy for 2023 operations'. No disruptions or adverse (weather) condition scenarios are sought out, as disruptions fall outside the scope of the research.

Within the scope, special attention is paid to the pushback process, the flow of aircraft within apron bays, as well as the various (un)loading destinations and any congestion around them. Additional analysis of emissions based on aircraft and TaxiBot movements is included, but detailed modelling (like dispersion, airstream velocity, or engine gas chemistry) is out of scope.

Several topics are explicitly out of scope. Firstly, no TaxiBot assignment modelling will take place. The goal of this study is not to determine how much can be achieved with a set number of TaxiBots (or how many would be required for a determined level of service), but rather to evaluate the impact, if large-scale sustainable taxiing operations were to take place. For detailed research on this topic at Schiphol airport, please see van Winkel (2023). Related to this, TaxiBot (re)scheduling and delay recovery is out of scope. It is assumed that every flight that requires a TaxiBot will have one available where and when it is needed. This means that charging, refuelling and any mechanical issues or faults are also out of scope. It is assumed that every vehicle works as intended and is available when needed. However, TaxiBot mission start and end times will be recorded, allowing for a calculation of the maximum number of active TaxiBots at one time. While this is not a fully accurate approach, it will provide an indication of the TaxiBot fleet required for the analysed operations.

Secondly, service road traffic (e.g. refueling, catering, baggage vehicles) will not be modelled. The simple yet unfortunate reason for this is that there is no (sufficiently detailed) radar data available for these vehicles. The lack of traffic interaction for TaxiBots on service roads is taken into account during simulations by lowering the TaxiBot's non-towing speed. There currently are some compatibility issues with TaxiBot vehicles encountering opposing traffic on some of the service roads. These infrastructural limitations are also placed out of scope, as Schiphol airport is already actively pursuing solutions to mitigate these compatibility issues.

Finally, aircraft that are maintenance-towed (repositioning aircraft from one stand to another by a tow truck, without passengers) on the airport are not included in the simulation, because these movements did not appear in the radar data. The transponder aboard the tow truck supplies this location data for the towing movement, but that data is not something that was available for this study. The main consequence for the model for these maintenance tows being out of scope, is that a few moving aircraft will be missing in the simulation. This is not expected to have a negligible impact, because maintenance-towed aircraft always give way to other aircraft, according to Schiphol policy.

1.4 Relevance and academic contribution

The relevance of this research is twofold. From an academic perspective, new insights are gained into how smaller and larger airports respond differently to substantial changes in daily operations, and the research adds to the existing body of literature on the impact of TaxiBot operations on emissions at the airport (chapter 4.4). For the aviation sector, this thesis serves a valuable purpose of estimating the benefits and drawbacks of TaxiBot operations before having to make financial and operational commitments to deploying them in real-world scenarios. An additional academic contribution is the successful implementation of real-world data in a simulation model (chapter 3.1.2), where this 'simulated reality' then forms the basis on which new behaviour is modelled (chapter 3.2).

Earlier research by Khammash et al. (2017) and Salihu (2020) has focused on TaxiBot operations at the airports of Lisbon and Montréal respectively, which both have much reduced traffic volumes when compared to Schiphol airport. Their focuses have been on emissions and the business case (cost) of TaxiBot operations, and resulted in an initial body of knowledge that approximates the impact of TaxiBot operations. The aim of this thesis is to augment this existing knowledge with insights into the flow and congestion of taxiing traffic, and increased realism through the use of real-life radar movements. This contributes to better insights by shedding light on a new topic of research (traffic flow, chapter 4.1), as well as providing increased trust in model findings through the use of real-life runway use, gate assignment, and aircraft routing and speeds that are accurate for each unique aircraft.

While the model is based on Schiphol airport, both the methodological approach as well as the model findings are generalizable to other airports. If the radar data and taxiing network is available, similar simulations can be performed at other airports using the same method. Similarly, the outcomes of the simulation (taxiing times, TaxiBot towing distances, and corresponding environmental impacts) are all based on travel distances, travel times, and engine-on durations. If this information is available for other airports, these statistics can be compared to the input data of the simulation in this thesis, to draw conclusions on the potential impact of TaxiBot operations at other airports. As an example, let's stipulate that a TaxiBot towing operation is five kilometers long, saving 200 kg of fuel for the outbound movement. Another airport, where the envisioned towing operation would be only two kilometers, can then estimate their own fuel savings based on those data points (40 kg fuel per towing kilometer, in the fictional example).

1.5 Structure

The structure of this thesis is as follows. Chapter 2 presents the literature related to TaxiBots and sustainable taxiing, and earlier work regarding the simulation thereof. The chapter sections detail the literature search process, then the analysis, and finally the conclusions and discussion. The data sources, constructed experiments, and model output of the research are all detailed in chapter 3. The simulation results along with their analyses are then presented in chapter 4. The findings are discussed in chapter 5, after which conclusions are drawn in chapter 6. A scientific article written about this thesis can be found in appendix C.

2 Literature study

This chapter contains a literature review of the current state of scientific knowledge on the various types of sustainable taxiing and associated topics. Section 1 will detail the search process and present an overview of the collected literature. Section 2 will discuss and analyse the literature, based on various themes. Conclusions from these findings will be presented and discussed in section 3.

The literature review leads to the research gap and objective of this thesis, namely that there exists a knowledge gap of how the increased complexity on taxiway networks of large airports affects the airports' ground traffic patterns if sustainable taxiing would be introduced.

2.1 Search strategy and overview of collected works

2.1.1 Purpose and scope

The purpose of this literature review is to collect and analyse the existing body of literature regarding the topics of sustainable taxiing using dispatch towing, and airport ground traffic congestion. Given the ruling of the Netherlands Labour Authority (2023b) that Schiphol airport needs to reduce its ultrafine particle (UFP) emissions, and the airport's own wishes to reduce these as well as CO_2(-equivalent) emissions, these topics are of great importance as they will allow aircraft to reduce the time that their engines are running.

The scope of the literature search encompasses the two previously mentioned topics, as well as other innovative taxiing systems. These encompass sustainability enhancements in taxiing through scheduling and routing, and sustainable taxiing by using electric motors in the landing gear. Various taxiing analysis techniques like discrete event simulation (DES) and mixed-integer linear programming (MILP) are also included in the search scope, to gain a richer understanding of the methodologies used in this field of research.

2.1.2 Literature search methodology

Due to the novelty of sustainable taxiing as a concept, the literature search is not limited to a specific time frame. However, outdated software and methods in simulation techniques are excluded from review. Multiple types of literature, including journals, conference papers, PhD dissertations, but also master theses are included in the search to ensure a comprehensive collection of literature on the topics.

The literature has been collected from multiple search engines, using multiple detailed search queries. Scopus and Google Scholar were used to collect the main body of academic literature on the topic. The TU Delft repository (consisting of theses and research publications) was also consulted to collect research that has already been done at TU Delft on sustainable taxiing at Schiphol airport. Finally, some further papers were found through both forwards and backwards snowballing through the literature. The collected papers and their search methodology can be seen in table 2.1 below. Naturally, many papers appeared in multiple search queries. They are displayed below in only one of the queries.

Table 2.1: Literature search queries and their yields

Search engine	**Query**	**Papers yielded**
Scopus	"Discrete event simulation" AND "airport"	Jen et al., 2022
		Zbakh et al., 2023
	"Dispatch towing"	Cao et al., 2023
		Di Mascio et al., 2022
		Khammash et al., 2017
	"Electric taxiing"	Groot & Roling, 2022
		Zoutendijk et al., 2023
Google Scholar	Airport electric taxi discrete event simulation	Corey & Clymer, 1991
		Lai et al., 2021
		Ouerghi, 2008
		Salihu, 2020
		Wollenheit & Mühlhausen, 2013
	"Dispatch towing"	Gualandi, 2014
		Hein & Baumann, 2016
	Electric taxiing system	Hospodka, 2014
		Lukic et al., 2018
		Lukic et al., 2019
	Gate congestion "airport"	Khadilkar & Balakrishnan, 2014
		Salihu et al., 2021
	Sustainable taxiing	Zhang et al., 2019a
		Zhang et al., 2019b
	Taxiing simulation	Kariya et al., 2011
		Kariya et al., 2013
		Mori, 2012
	"Process mining" and "Discrete event simulation"	Antunes et al., 2019
		Jadrić et al., 2020
		Liu, 2015
		Rashid & Louis, 2022
TUD Repository	TaxiBot	Benda, 2020
		Tindemans, 2021
		van Winkel, 2023
	Dispatch towing	Clemens, 2023
		Soepnel, 2015
Snowballing	(Backward and forward)	Ahmadi, 2019
		Postorino et al., 2016
		Roling et al., 2015
		Sirigu et al., 2018
		Soltani et al., 2020
		van Baaren & Roling, 2019
		Vaishnav, 2014
		van Oosterom et al., 2023

2.2 Analysis

The literature can be divided in multiple themes. These themes are: earlier work regarding sustainable taxiing at Schiphol, dispatch towing, previous discrete event simulation of TaxiBot operations, congestion management and other forms of sustainable taxiing, and methodology-related literature. Table A.1 in appendix A shows the papers, their main topic, method, and themes they apply to.

2.2.1 Earlier sustainable taxiing research at Schiphol

Five master students at TU Delft faculties have performed earlier work on sustainable taxiing projects at Schiphol. Soepnel (2015) looked at electric taxiing systems in the landing gear of aircraft, and found that gate capacity would increase with this innovation because conventional push-back and engine-start processes would no longer be necessary. Electrified landing gears are a different process to dispatch towing with TaxiBots, but the effects are expected to be similar, especially with regard to apron traffic impact. Benda (2020) and Tindemans (2021) in their theses have explored other aspects of TaxiBot operations at Schiphol airport, related to route planning, vehicle assignment and schedule disruptions. Benda (2020) built an agent-based model to find that traffic on Schiphol's taxiways will increase with the use of TaxiBots, but that this can be accommodated without affecting safety. Tindemans (2021) considered the impact of schedule disruptions on TaxiBot operations using a mixed-integer linear programming (MILP) model. Short-term tactical planning can allocate TaxiBots to 48% of flights, which is only a minor loss compared to 48.5% of flights in the strategic planning model that was built. Van Winkel (2023) then expanded on this work by building a tactical planning MILP model that could optimize TaxiBot allocation in terms of fuel savings, schedule robustness or fairness among participating airlines. Finally, Clemens (2023) performed a societal cost-benefit analysis on various business cases to determine that the implementation of dispatch towing at Schiphol is a positive development for society. His findings suggest that electric tows to the edge of the Schiphol apron yield a bigger benefit than towing the aircraft all the way to the Zwanenburgbaan (Runway 18C-36C) for departure.

There is also some other academic literature published by TU Delft on the topic of sustainable taxiing at Schiphol. Roling et al. (2015) found that electric taxiing at Schiphol should reach a speed of at least 17.5 kts (9 m/s) in order to avoid unacceptable departure delays to other departing traffic. This research was focused on electrified landing gear rather than tow tractors but the findings should extend to both methods. The fuel savings of implementing these forms of sustainable taxiing can save 82% of the fuel used for taxiing at Schiphol (van Baaren & Roling, 2019). Medium category aircraft (e.g. the Boeing 737 and Airbus A320 families) would see the biggest benefits. Other research has shown that sustainable taxiing is also most suitable for shorter haul flights to major airports, since these flights carry relatively more fuel dedicated to taxiing across the large airports in Amsterdam and the flight's destination (Groot & Roling, 2022). Some recent research by van Oosterom et al. (2023) based on traffic volumes in the winter of 2019 has shown that Schiphol would need around 38 TaxiBots to service 913 flights: 26 small TaxiBots for 750 narrowbodies, and 12 large TaxiBots for the 163 widebody flights that day. Their research assumed electric TaxiBots so these figures include the need for charging, and the researchers also found that these relationships scale linearly as the number of flights increase.

2.2.2 Optimization of dispatch towing

A few papers have looked at the benefits and drawbacks of TaxiBots when the innovation first came to market. The main advantages, as already mentioned, are significant fuel savings for the airlines, as well as a resulting reduction in CO_2 and UFP emissions for the airport and its surroundings (Postorino et al., 2016; Di Mascio et al., 2022). Drawbacks of the system are the vehicle price, increased complexity for scheduling and assignment, and a more complicated ground traffic picture (Gualandi, 2014; Hospodka, 2014). The fuel saving effects will be biggest for major airports, where taxi times are longer. Furthermore, because the engines are not running, taxiing operations should become quieter (Hein & Baumann, 2016). However, noise is not expected to be an important aspect for the TaxiBot's success because the majority of an aircraft's noise emissions are during take-off and climb-out.

Zoutendijk et al. (2023) conclude from a literature survey that there are two main ways of solving routing problems: mixed-integer linear programming and simulation. Sirigu et al. (2018) developed algorithms based on neural networks and graph theory. They found that all algorithms perform adequately, but that the graph theory-based algorithms are much more efficient. One of the biggest challenges of this type of problem is conflict avoidance due to the increased number of vehicles in play. Soltani et al. (2020) therefore developed a MILP-model and applied it successfully to the case of Montreal airport.

2.2.3 Previous discrete event simulation of TaxiBot operations

Another common way to analyse dispatch towing systems is via discrete event simulation. Khammash et al. (2017) used this technique to simulate ground movements at Lisbon airport. Lisbon airport is a medium-sized airport in Portugal, with one runway and 220,000 yearly aircraft movements in 2023 (VINCI Airports, 2024, p.7). The researchers found fuel savings of 6% per deployed TaxiBot along with CO2 emission reductions of 5% per TaxiBot, with no observed increases in total taxiing times. Salihu (2020) also used discrete event simulation to assess the impact of dispatch towing operations at Montreal airport, another medium sized airport with two main runways and 200,000 yearly aircraft movements in 2023 (Aéroports de Montréal, 2024, p.5). They found that 22 TaxiBots would be needed to service all flights in the study period. Due to the minimization of fuel costs, this scenario would also lead to the lowest operational cost, compared to partial service scenarios. Further research by the same researchers introduces uncertainty to the problem, resulting in a less definitive conclusion about the cost-effectiveness of TaxiBot operations in the Montreal case study (Salihu et al., 2021). Their results show that 'TaxiBotting' all flights would add 4.2 minutes to the taxi-in process that currently costs 10.9 minutes (+38.5%), and would add 2 minutes to the current 15.1 minute taxi-out process (+13.2%).

While the results of the previous studies are very useful as a first indication, the studies did make some unrealistic assumptions and simplifications. The main areas where this may impact findings are in non-context aware taxiing speeds, (simple) delay estimation, and runway and gate assignment. Their modelling choices, along with the intentions for this research thesis, are displayed in the table below.

Table 2.2: Different modelling choices of TaxiBot discrete-event simulations

Aspect	**Khammash et al., 2017**	**Salihu et al., 2021**
TaxiBot level of implementation	Scenarios for 2 narrowbody (NB) TaxiBots, 4 NB's, and 4 NB + 1 widebody (WB).	16-30 TaxiBots modelled. Scenarios for towing partial/all aircraft.
Taxiing speed aircraft	15 kts = 7.7 m/s	7 m/s = 13.6 kts
Taxiing speed TaxiBot	Unclear, likely same as aircraft	7 m/s loaded, 4 m/s (7.8 kts) unloaded
Runway assignment	Only Runway 03 in consideration	Random proportional for the available runways
Gate assignment	Based on carrier, aircraft type, turnaround time	Assigned to gates in ascending order
Route finding	Shortest path	Shortest path
Conflict detection & avoidance	Not mentioned	None on taxiways, aircraft can overlap (pass through) on intersections
TaxiBot's role in ground traffic:	Not mentioned (TaxiBot disappears after arriving at runway)	Not considered, they can overlap with other model objects
Separation between aircraft	60 meters	30 meters
Arrivals & delay	Not modelled, focus on departing flights	Schedule from flightradar24.com, delay simulated random triangular(-15,0,30) minutes

Turnaround time	According to flight schedule	NB: random uniform(26,51) minutesWB: random uniform(25,130) minutes
Pushback & engine start	Random uniform, 3-5 minutes	4 minutes constant
Modelling demarcation choices	Taxi-in not considered, because their base scenario showed a bottleneck on taxi-out; no speed difference between corners and straight line taxiing	Tow trucks consume no fuel, even when towing; no speed difference between corners and straight line taxiing

This research thesis aims to improve on those assumptions, and thus provide more accurate findings. Additionally, this thesis is focused on a much larger and more complex airport (six runways, 442,000 flight movements in 2023 (Royal Schiphol Group, 2024c), complex gate apron structure, with dynamically changing active runways) which will by itself contribute to pushing the academic boundary forwards, and gaining further insights on how complex airports react to a shift in traffic operations. Thirdly, this thesis has a main focus on traffic flow of TaxiBot operations with a secondary focus on emissions (CO2 and UFP), whereas the earlier two studies have mainly focused on potential emission reduction and cost-effectiveness.

2.2.4 Congestion management and other forms of sustainable taxiing

Lukic et al. (2018) and Lukic et al. (2019) performed literature reviews to analyse the various sustainable taxiing options in more detail. The main conclusion to be drawn is that there are three major forms of sustainable taxiing: dispatch towing, electrified landing gear taxiing, and single-engine taxiing. All three are significantly better than the status quo and the best solution largely depends on the problem owner. Airports can use a fleet of TaxiBots to introduce dispatch towing, airlines can modify their fleets to enable sustainable taxiing around the world, and single-engine taxiing is a general policy advisory that can be implemented everywhere with little to no adaptation (Vaishnav, 2014). Electrified landing gear taxiing has the best emission reductions on the ground, albeit at the cost of extra weight and thus reduced in-flight fuel performance (Cao et al., 2023), and in terms of operational efficiency it also is the quickest due to the aircraft now being able to execute its own pushback (Wollenheit & Mühlhausen, 2013; Zhang et al., 2019a).

Another form of saving fuel while taxiing is eliminating waiting time and congestion as much as possible. Kariya et al. (2011) created a discrete event simulation model based on actual operational data to analyse taxiing patterns at Tokyo Haneda airport. They suggest adaptations to the schedule and an aircraft's actual off-block time (AOBT) to ease congestion and improve the airport's sustainability. A paper published two years later fine-tuned their strategies and provided further evaluation of results, with the main conclusions staying the same (Kariya et al., 2013). Various models have been created by Mori (2012), Zhang et al. (2019b), and Zbakh et al. (2023) that all quantify and minimize the taxiing congestion (and maximize capacity) by adjusting the schedules of airports around the world. Lai et al. (2021) investigated a specific case at JFK airport in New York where they used discrete event simulation to locate a known bottleneck, and identify ways to ease the congestion. They identified the taxi-out, and runway queueing before departure as the biggest factors causing the delays. A similar case for Boston Logan airport was investigated, where an algorithm was developed to control taxiing congestion. In that case, the control of pushback times was identified as the most critical short-term strategy to improve the departure efficiency (Khadilkar & Balakrishnan, 2014). These last two findings are significant for TaxiBot operations at Schiphol, since one of the policies that are being considered is to start TaxiBotting at pushback, and unload from the TaxiBot very close to the runway, shortly before take-off. TaxiBot punctuality could therefore play an important role in the on-time performance (OTP) and congestion at Schiphol airport.

2.2.5 Simulation and the importance of data quality

As already discussed in sections 2.2.1 to 2.2.4, discrete event simulation has been applied to (sustainable) taxiing problems by many different researchers, for many different cases (Kariya et al., 2011, 2013; Mori, 2012; Wollenheit & Mühlhausen, 2013; Khammash et al., 2017; Salihu, 2020; Salihu et al., 2021; Lai et al., 2021; Zbakh et al., 2023). The method has also been used for de-icing strategy analysis (Jen et al., 2022). This shows that discrete event simulation is an effective tool to conduct research in airport operations, as is also affirmed by more theoretically-based works by Corey and Clymer (1991) and Ouerghi (2008).

Modern data collection methods can be very valuable to enrich these simulations with accurate, real-life data. These big collections of data can be used to extract information about patterns, durations, variations and abnormalities in existing airport processes via a method called process mining (Van Der Aalst et al., 2012). A couple of new innovative frameworks have been developed in order to connect the output of process mining to the input of simulation models. Liu (2015) developed a plugin for process mining software that would format its outputs in easy-to-read Excel workbooks that can be imported into some DES software packages. Antunes et al. (2019) successfully demonstrated the concept by applying this two-method process to hospital waiting rooms, using simulation to compare the original process with an optimized solution. A similar approach was taken by Jadrić et al. (2020), who used two different cases (a data-poor and data-rich example) to emphasize the importance of "event logs that are clearly defined and refer to a case (i.e. process instance) and an activity (i.e. step in the process)". The method was then also applied to a real-life case in construction monitoring where the findings demonstrate that process mining methods effectively identify the process model using event data gathered from the field, and that consistently adjusting the input parameters of the DES model leads to precise and reliable productivity estimates (Rashid & Louis, 2022).

While this thesis will not use process mining explicitly, the approach is very similar. 'Event log' data (in the form of timestamped positional radar data) is collected and transformed to patterns and durations (segment routes and times), and then imported into capable simulation packages (Simio), to enable the simulation of high-quality realistic experiments.

2.3 Conclusion and discussion

The existing literature has shown that there are four main forms of sustainable taxiing: single-engine taxiing, sustainable routing and scheduling, electrified on-board taxi systems, and dispatch towing. All four have demonstrated significant improvements over conventional taxiing procedures, with dispatch towing and electrified on-board systems performing the best (Lukic et al., 2019; Zhang et al., 2019a). Airports with high taxiing times, and short haul flights where a relatively large portion of fuel is dedicated to taxiing would receive the most benefit of sustainable taxiing (van Baaren & Roling, 2019; Groot & Roling, 2022). That makes Schiphol airport very suitable for the implementation of a dispatch towing TaxiBot system.

A lot of research has been done in the fields of scheduling, routing, and assignment of dispatch towing systems. The routing problem can be solved, but particular attention needs to be dedicated to conflict avoidance (Soltani et al., 2020). Achieving optimized solutions for dispatch towing assignments has posed a greater challenge, but disruptions and the associated short-term adaptations have only a small effect on the efficiency of operations (Tindemans, 2021; van Winkel, 2023).

Simulations regarding bottlenecks and congestion have shown that push-back and take-off waiting time are the biggest causes of taxiway congestion and delays (Khadilkar & Balakrishnan, 2014; Lai et al., 2021). These two moments are also important steps in the TaxiBot process, where the aircraft is loaded and unloaded.

An innovative way of providing input for these simulations is to obtain real-life operational data and translate that into detailed information on the airport's processes through a method called process mining (Antunes et al., 2019). This thesis will not use that method explicitly, but the use of radar data to act as a foundation for simulations shows many similarities to it.

These results naturally lead to the research gap and objective of this thesis, that there exists a knowledge gap of how the increased complexity on taxiway networks of large airports affects the airports' ground traffic patterns if sustainable taxiing would be introduced. In the case of Schiphol specifically, its gate layout is such that TaxiBot operations are expected to also have a positive impact on the flow of aircraft in gate bays.

3 Data and methods

This chapter will introduce the data and methodology used in this thesis, as well as detailing which experiments will be simulated. First, the data sources and processing approach will be presented, in which Schiphol ground radar data is converted into taxiway segments and speed information. Secondly, the way this data is used in the simulation model is explained. Section 3 first presents how the selected simulation days were determined, and which experiments will be performed. The final section is dedicated to the types of results that will be generated by the simulation model, and how those results will be used to answer the research questions.

3.1 Model input

3.1.1 Data sources

To make the simulation model as realistic as possible, a broad selection of data sources provided by Schiphol have been used in this thesis.

Table 3.1: Data sources

Data source	**Information in source**	**Purpose**
Schiphol airport map (Basiskaart) (Royal Schiphol Group, n.d.-a)	A correctly scaled map of Schiphol's runways, taxiways, stands and service roads.	To construct the network of nodes and links so that it reflects the real-life airport.
Ground Radar at Schiphol *(not publicly accessible)*	Timestamped information (every second) of an aircraft's latitude, longitude, altitude, aircraft registration, flight number, and arrival/departure indicator.	To reconstruct the radar routes and speeds that every aircraft has taken.
Geofence polygons *(not publicly accessible)*	A geojson file containing segments of runways, taxiways, stands and other elements, with names and positional information, to provide accurate and unique locations.	To map each aircraft's latitude and longitude to a segment (runway, taxiway, stand).
Central Information System Schiphol (CISS) *(not publicly accessible)*	The central Schiphol database, containing stand/gate assignment, aircraft types, departure and arrival times, gate changes, etc.	To correctly assign aircraft information in the simulation, and to ensure aircraft in the simulation arrive and leave at the same time that they did in real life.
TaxiBot mission portal *(not publicly accessible)*	Loading and unloading durations of all TaxiBot operations, at airports around the world.	To accurately model the loading and unloading times in the simulation.
Fuel flow information and aircraft counts *(not publicly accessible)*	Per aircraft type: visit count for 2023, fuel flows for all engines at idle thrust (used for taxiing) and for the APU.	To calculate engine emissions.

These data sources are combined to be able to reconstruct a flight's journey through Schiphol airport, all the way from landing at the runway and displaying the correct aircraft type in the simulation, to taxiing at the correct speeds and the correct route, to the right stand, leaving the stand at the correct time, and again departing using the correct routes, speeds, taxiways and runways.

Figure 3.1 below shows a snapshot of the raw Casper Ground Radar data upon import. The Schiphol airport map, figure 3.2a, can be accessed via the link in the references, as it is publicly available (Royal Schiphol

Group, n.d.-a). An overview of the geofence polygons used to map the coordinates to airport locations is shown in figure 3.2b.

Index	ts	acreg	fltnr	arrdep	lon	lat	alt	spd	casperid
0	2023-11-06 00:00:00+01:00	PHEZM	KL1368	A	4.76216	52.3034	0	nan	7409122
1	2023-11-06 00:00:00+01:00	PHHBK	HV6118	A	4.76799	52.3055	0	nan	7409257
2	2023-11-06 00:00:00+01:00	PHHXJ	HV6886	A	4.77118	52.31	nan	nan	7409259
3	2023-11-06 00:00:00+01:00	PHXRX	HV5354	A	4.76959	52.3117	nan	20.85	7409261
4	2023-11-06 00:00:01+01:00	PHEZM	KL1368	A	4.76216	52.3034	0	nan	7409122
5	2023-11-06 00:00:01+01:00	PHHBK	HV6118	A	4.76799	52.3055	0	nan	7409257
6	2023-11-06 00:00:01+01:00	PHHXJ	HV6886	A	4.77118	52.31	nan	nan	7409259
7	2023-11-06 00:00:01+01:00	PHXRX	HV5354	A	4.76952	52.3116	nan	20.18	7409261

Figure 3.1: Radar data before any processing

(a) Schiphol Basiskaart (version Sep 2023)

(b) Geofence polygons

Figure 3.2: Schiphol map and corresponding polygons

3.1.2 Ground radar data processing

The Casper Ground Radar data is not immediately ready to be used in the simulation. Data processing was necessary for two reasons: to turn the stream of GPS coordinates into a route of segments that could be imported into the simulation for the aircraft to follow, and to supplement the radar data with the necessary flight information. Some information related to the aircraft had to be added from CISS, and the radar information was sometimes incomplete (when aircraft disappeared from the radar due to transponders being shut down), or contained some positional errors (due to GPS-drift at stands, for example).

Firstly, the GPS data had to be linked to information about the aircraft. Flight numbers could not be used because they are different for arriving and departing flights of the same aircraft, and the aircraft registration could not be used either, because some aircraft visit the airport multiple times in a day. This called for the creation of a new identifier, called a TrackID, to give every radar track its own label. This TrackID was formed by combining the inbound flight number, with the aircraft registration, and outbound flight number at the end. To account for delayed flights, the schedule date of the flight was inserted before each flight number. As an example, aircraft PH-BCG performing flights KL1123 inbound and KL1124 outbound on the 29th of May, would result in a radar TrackID of *'29xKL1234_PHBCG_29xKL1235'*. Flight movements that did not have an inbound or outbound leg (because they stayed the night at Schiphol, for example) had that portion of the TrackID replaced with 'none' (*'29xKL1234_PHBCG_none'*).

The second step was to then assign the GPS coordinates to airside segments (figure 3.2b), while also taking an aircraft's altitude value into account to filter out overflying segments. This step of reducing GPS coordinates to segments is necessary, in order to later enable entities in the simulation to follow a route from segment to segment. The resolution of the segment is such that every junction, taxiway lane, aircraft stand and runway has its own segment. After these steps, all data rows that were not in a segment, were at too high of an altitude (the ground radar also captures some movements before landing and after take-off) could be filtered out of the dataset.

What followed was lot of error checking and handling of exceptions, like GPS drift at a stand, causing aircraft to appear to have visited multiple stands during its turnaround process (see figure 3.3). Some other examples of difficulties were transponders being turned on too late or shut down too early (thus missing segment information), and repeating segments when an aircraft stood stationary on top of a border between two segments for a while (causing the radar drift to jump back and forth between the two segments, resulting in a 'Segment1, Segment2, Segment1, Segment2' instead of just Segment1 followed by Segment2).

Figure 3.3: GPS drift at a stand causing radar labelling issues

As final steps, the list of radar segments per TrackID was checked for validity ("are these two segments in the route actually connected to each other in real life?"), and the time spent in each segment was split in 'seconds driving' and 'seconds stopped' groups, based on the GPS speed of the aircraft (with a threshold of 1.5 kts for determining 'driving' or 'stopped'). This last step is necessary to correctly mimic the aircraft's driving and waiting behaviour in the simulation.

The output of these radar processing steps is a dataframe of timestamps, radar positions (taxiway segment locations), driving seconds, and stopped seconds. This data would then be connected to the simulation network to also calculate path distance until the next node, and the resulting taxiing speed. A snapshot of this dataframe can be seen in figure 3.4 below.

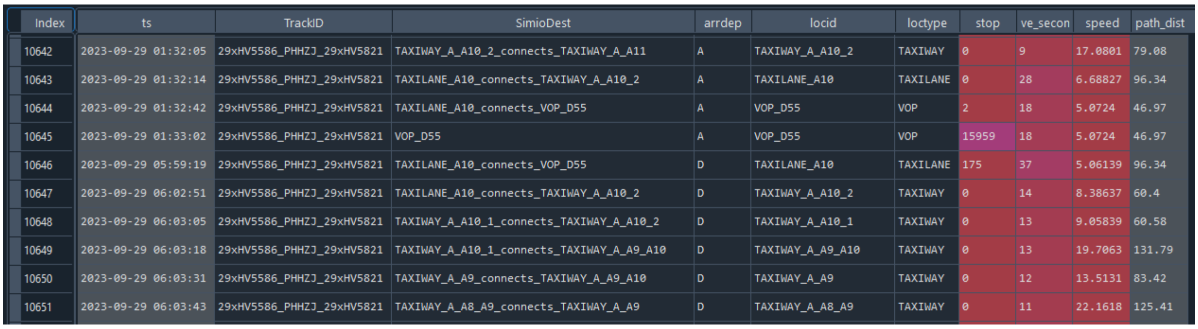

Index	ts	TrackID	SimioDest	arrdep	locid	loctype	stop	ve_secon	speed	path_dist
10642	2023-09-29 01:32:05	29xHV5586_PHHZJ_29xHV5821	TAXIWAY_A_A10_2_connects_TAXIWAY_A_A11	A	TAXIWAY_A_A10_2	TAXIWAY	0	9	17.0801	79.08
10643	2023-09-29 01:32:14	29xHV5586_PHHZJ_29xHV5821	TAXILANE_A10_connects_TAXIWAY_A_A10_2	A	TAXILANE_A10	TAXILANE	0	28	6.68827	96.34
10644	2023-09-29 01:32:42	29xHV5586_PHHZJ_29xHV5821	TAXILANE_A10_connects_VOP_D55	A	VOP_D55	VOP	2	18	5.0724	46.97
10645	2023-09-29 01:33:02	29xHV5586_PHHZJ_29xHV5821	VOP_D55	A	VOP_D55	VOP	15959	18	5.0724	46.97
10646	2023-09-29 05:59:19	29xHV5586_PHHZJ_29xHV5821	TAXILANE_A10_connects_VOP_D55	D	TAXILANE_A10	TAXILANE	175	37	5.06139	96.34
10647	2023-09-29 06:02:51	29xHV5586_PHHZJ_29xHV5821	TAXILANE_A10_connects_TAXIWAY_A_A10_2	D	TAXIWAY_A_A10_2	TAXIWAY	0	14	8.38637	60.4
10648	2023-09-29 06:03:05	29xHV5586_PHHZJ_29xHV5821	TAXIWAY_A_A10_1_connects_TAXIWAY_A_A10_2	D	TAXIWAY_A_A10_1	TAXIWAY	0	13	9.05839	60.58
10649	2023-09-29 06:03:18	29xHV5586_PHHZJ_29xHV5821	TAXIWAY_A_A10_1_connects_TAXIWAY_A_A9_A10	D	TAXIWAY_A_A9_A10	TAXIWAY	0	13	19.7063	131.79
10650	2023-09-29 06:03:31	29xHV5586_PHHZJ_29xHV5821	TAXIWAY_A_A9_connects_TAXIWAY_A_A9_A10	D	TAXIWAY_A_A9	TAXIWAY	0	12	13.5131	83.42
10651	2023-09-29 06:03:43	29xHV5586_PHHZJ_29xHV5821	TAXIWAY_A_A8_A9_connects_TAXIWAY_A_A9	D	TAXIWAY_A_A8_A9	TAXIWAY	0	11	22.1618	125.41

Figure 3.4: Radar data after processing

Next to this routing dataframe, an input dataframe for the model was also created. This dataframe tells the simulation when the aircraft should be created, and also contained supplementary information like their aircraft category, arrival and departure runways, arrival and departure stands (sometimes those differ), departure time from the stand, and also if any TaxiBot action should take place for the outbound and inbound movements. More about that later, in section 3.3. That data was partially gathered from CISS, and partially inferred from the radar routes.

Additionally, aircraft were assigned a size category. The model would get too complex if every aircraft type (e.g. Airbus A321, Boeing 737-800, Embraer 175, etc.) would be included, and so every aircraft type was assigned to one of three categories: regional, narrowbody, and widebody. This split was made according to the categories in the Aircraft Stand table (Royal Schiphol Group, n.d.-b). Regional aircraft are defined as categories 1 to 3, narrowbodies as category 4, and widebodies as categories 5 to 9.

3.2 Simulation model

Simio Academic Edition (Simio, n.d.) is used as the simulation software. This is a Discrete Event Simulation (DES) package. DES is a type of simulation where real life processes are simulated, by breaking them down into specific events, that happen at discrete times. Everything that happens is treated as a series of events, and variables can be assigned and updated at every step. This makes it possible to both model processes in fine detail, yet still have a fast runtime, because the model is able to skip time periods in which nothing happens.

Simio in particular was selected for this research due to its ability to import and export large amounts of data, create detailed processes to accurately model processes like the TaxiBot (un)loading steps and turnaround at the stand, and its ability to have both a 'fast-forward' calculation mode as well as a 3D animation ability for visualisation, allowing for images (and videos) like can be seen in figure 3.5.

Figure 3.5: The simulation model in Simio

3.2.1 Simulation mechanisms

The simulation's foundation consists of a set of nodes and links that has been drawn over the Schiphol airport map, which was scaled up so that one meter on the airport map would correspond to one meter distance in the simulation environment. The nodes are placed on the intersections of radar polygons (figure 3.2b), and are connected by links whenever an aircraft in real life would be able to drive from one polygon to another, see figures 3.6a and 3.6b. The green rectangles in figure 3.6a were created as a development step, to indicate where model nodes should be placed in the simulation model, and for use in the radar processing to see if two segments had a valid connection.

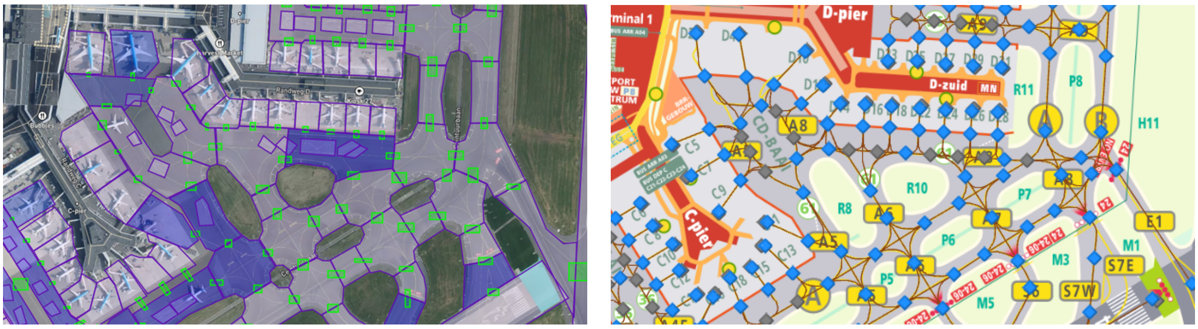

(a) Polygons between the C and D piers

(b) The same network, implemented in Simio

Figure 3.6: Creating the simulation network based on polygons

Figure 3.6b shows both blue and grey nodes. Blue nodes are called TransferNodes, and contain routing logic as well as the possibility to run custom processes once entities enter or exit them. The grey nodes are called BasicNodes, and they simply allow for multiple links to intersect each other.

The created dataframes as described in section 3.1.2 deliver all the data that is required for the model to run. They are imported into Simio's "data tables" section, and the model mechanisms are set up in such a way that all aircraft in the simulation (called entities) are assigned their corresponding statistics from those data tables. Every time an entity enters or leaves a TransferNode, several processes would be triggered. Some examples of processes are assigning the next destination in the entity's route sequence, holding the entity stationary for its determined 'stop seconds', assigning it its new speed, updating the total distance travelled metrics, etc.

Some special nodes have additional processes that are triggered. Nodes representing a stand, for example, would also hold the entity stationary until the actual off-block time (AOBT; the time the aircraft left the stand according to the radar data and CISS database). Additionally, some processes related to measuring the inbound and outbound process times would be triggered. Other examples of special nodes are nodes at the edge of a runway (where inbound and outbound process times are started and ended), or nodes at the edge of the apron bays (again, to measure results). Some other special nodes in the network are related to TaxiBot routing sequences, where entities that meet certain conditions are told to deviate from their radar routes, to instead follow a route that leads them to a TaxiBot (un)loading point, or that lets them follow the TaxiBot policy route. More on that in section 3.3.3

All the while, the aircraft would still be following their radar speeds, to keep the simulation as realistic as possible. However, some TaxiBot routes require the aircraft to visit locations on the airport where it has never been, and thus no radar speeds are available. To account for that, a 'dry run' was performed in the simulation, where every path in the model recorded the speeds of all aircraft that travelled on it, to create a 'speed map' of average speeds per taxi path. Per path, separate speeds were determined for the regional, narrowbody and widebody aircraft categories in the simulation. In the real simulation runs, that speed map could then be used to still provide realistic taxiing speeds for an aircraft, even if the aircraft itself did not have speed data available.

3.2.2 Simulated TaxiBot properties

Model properties relating to TaxiBots could not be derived from radar data or the CISS database. Certain assumptions and choices had to be made for things like the maximum towing speed, driving speed when the TaxiBot is not towing an aircraft, and what the loading and unloading times were.

The TaxiBot's internal speed limit is configured for 23 knots of taxiing speed. However, experience from Schiphol's earlier trials has shown that it is more optimal to operate the TaxiBot at a maximum speed of 22 knots. Therefore, this property was reflected in the simulation. When an aircraft is being towed by a TaxiBot, it will drive at its radar speed (or the speed from the speed map, as explained earlier), with a maximum limit of 22 knots. The non-towing speed is also listed as a maximum of 23 knots. However, a speed limit exists on the service roads around the Schiphol apron of 30 kilometers per hour (16 kts), and 60 kilometers per hour (32 kts) on the service roads further out on the airfield. Therefore, this speed limit of 23 kts would not accurately reflect the speeds a TaxiBot would drive in real life, for example when driving slowly around corners, and maneuvering amongst other service traffic like refueling trucks and baggage vehicles. This other service road traffic is out of scope for the simulation, and therefore the decision has been made to model the non-towing driving speed of the TaxiBot as 15 kts, to account for this traffic interaction, speed limit, and lower cornering speed.

Fortunately, real-life data does exist for the loading and unloading times of the TaxiBot. Because of Schiphol's collaboration with TaxiBot manufacturer SAS, they have access to a shared data portal where data of all TaxiBot missions worldwide is uploaded. The distributions of TaxiBot loading and unloading times were taken from this portal, and a function was modelled that best approximated this real-life distribution. This function takes the form of a Random Triangular distribution, indicated as Random.Triang(a,b,c). Value a is the minimum value that occurs, value b is the most frequently occurring value, and value c is the maximum value. The triangle is then formed so that the area under the curve is equal to one.

Every time a TaxiBot is loading or unloading from an aircraft in the simulation, a random draw is made from these distributions to realistically model the (un)loading time. Because this TaxiBot data itself is not publicly available, it can not be included in the plots in this thesis. However, the approximated random distributions are displayed in the figures below.

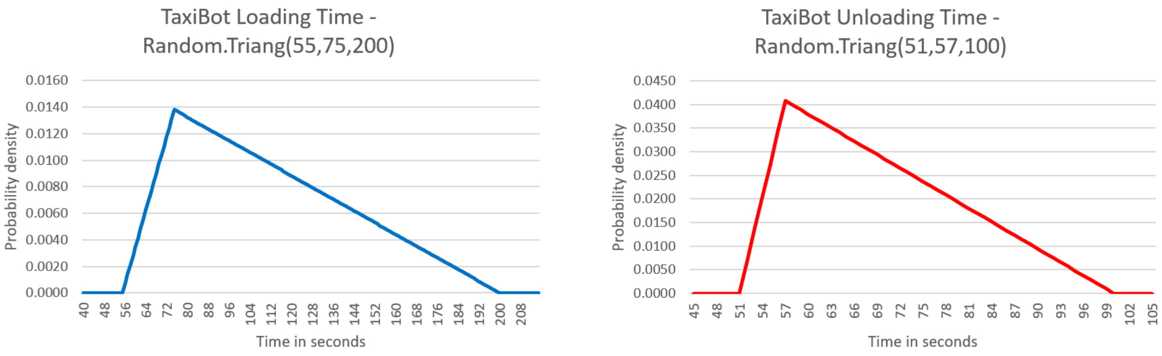

(a) Loading

(b) Unloading

Figure 3.7: TaxiBot loading and unloading times, based on operational data

3.3 Experiments

To determine which experiments would be run, it first needs to be determined which days are most suitable for analysis, after which an initial data exploration will take place. Following that, various TaxiBot towing policies will be defined, and finally, the experiments to simulate along with their motivations will be presented.

3.3.1 Selecting days for analysis

As mentioned in the introduction, the main current focus of TaxiBot concept development at Schiphol airport is towing flights to the Polderbaan, the furthest away runway (see figure 1.1). Simultaneously, the aim is to investigate if these operations are feasible, even under peak loads.

The objective for selecting days is therefore twofold; the day must be a busy one (meaning a high flight count), and there must be lots of Polderbaan activity. This activity can be outbound, a combination of outbound and inbound traffic, and solely inbound. However, since inbound towing is seen as a future scenario, and will never happen without also including outbound towing operations, this third option can be discarded.

Flight count data was combined with radar information, to gather the number of flight movements, Polderbaan arrivals, and Polderbaan departures for every day of 2023. Initially the goal was to simulate days in 2019 (due to higher traffic counts from before the COVID-19 pandemic), but this was not possible due to the unavailability of radar data for 2019.

To assess the most optimal days for analysis, two 'Polderbaan scores' were developed, where the flight count of that day was multiplied by the Polderbaan departure count (for the optimal outbound day), and where the flight count of that day was multiplied by the Polderbaan departure count plus the Polderbaan arrival count (for the optimal combination day). The maximum scores of these two metrics in 2023 would determine the days to be simulated. They are shown in the table below.

Table 3.2: Selected days for simulation

Day of year	Flight movements	Polderbaan departures	Polderbaan arrivals	Departure score	Combination score	Selection
14 Jun 2023	1,411	456	0	643,416	0	Departures
29 Sep 2023	1,427	233	227	332,491	107,703,477,139	Combined usage

A single day per experiment – rather than a range of days to test each policy on – is deemed sufficient for the following reasons. Most importantly, every flight undergoing TaxiBot towing within the analysed day is essentially a mini-experiment of the towing policy. All aircraft undergoing that policy then form the variation within the dataset based on aircraft type, route, taxiing speeds, emission reduction, waiting time at the unloading stations, etc.

Additionally, the selected day will already be representative of usual traffic volumes and split of aircraft types. The station congestion levels are determined by the aircraft mix (how many spaces they take up), unloading time at the station, and runway usage, with a maximum capacity of 40 movements per hour (Airports Council International, 2023). Of these three factors, only the runway usage would change on different days, but with the selected days already hitting the maximum capacities for the 2023, no additional days are required for the analysis.

Finally, the addition of extra days would not introduce any new circumstances, except maybe different runway use configurations. However, the fact that the selected days are selected precisely for their desired runway configurations makes this a moot point.

3.3.2 Data exploration of selected days

The radar data processing script has been run on the radar data of the two selected days. Not all flight movements were able to be included in the final radar dataset, due to some irrecoverable errors related to transponders being shut down for significant portions of the taxiing journey, private jets and general aviation traffic showing erratic transponder behaviour at the stands, and other similar issues where a runway or stand could not be determined. The statistics of how many movements had to be excluded from the simulation dataframes is shown in the table below. The high retention rates demonstrate that the effort to reduce errors and handle exceptions has been worthwhile.

Table 3.3: Flight movements in the simulation

Simulation day	Original movements	Movements dropped	Retention rate	Dropped from Polderbaan
14 Jun 2023	1,411	53	96.2%	11 departures
29 Sep 2023	1,427	39	97.3%	4 departures, 5 arrivals

Radar speed data for these days has been collected, and plotted in the two figures 3.8a and 3.8b below. The color split has been selected so that there is a clear distinction between slow and medium-speed traffic, and the upper limit of 23 kts has been selected to allow for initial analysis of any areas where the TaxiBots maximum speed limit might be a traffic bottleneck. Due to the plot being a heat map, the figures also provide an indication of which taxiways and runways were used, and how often.

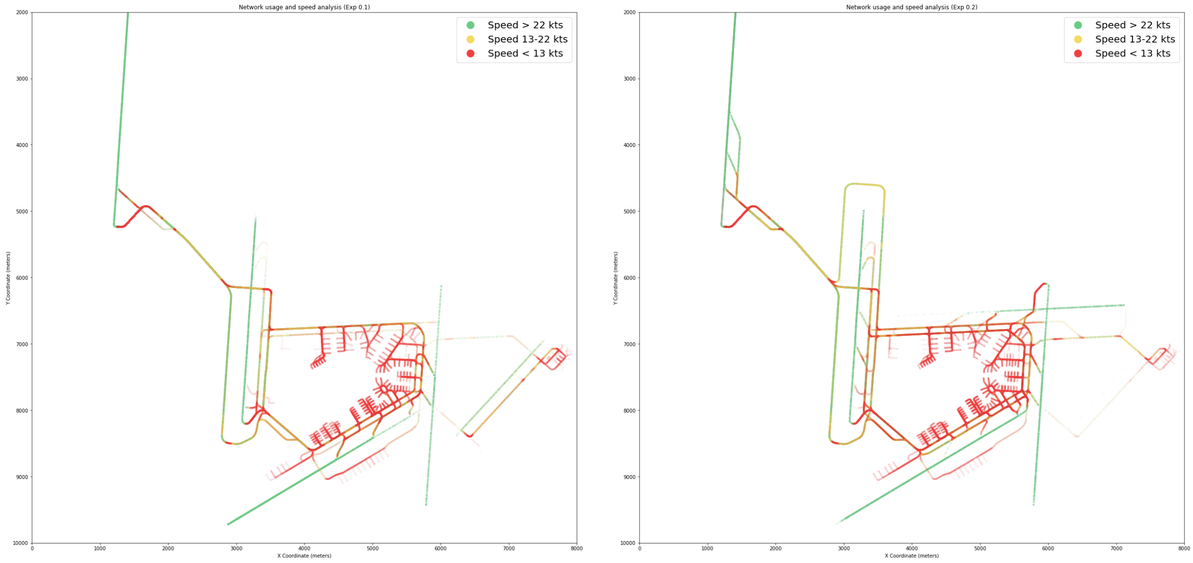

(a) 14 Jun 2023

(b) 29 Sep 2023

Figure 3.8: Radar speeds for the selected simulation days

The two plots logically show very similar results. Speeds around the apron bays are low, further out on the airfield they become a bit higher, sometimes even surpassing the speed limit of the TaxiBot. And as can be seen, cornering speeds or runway entry points also show much reduced speeds.

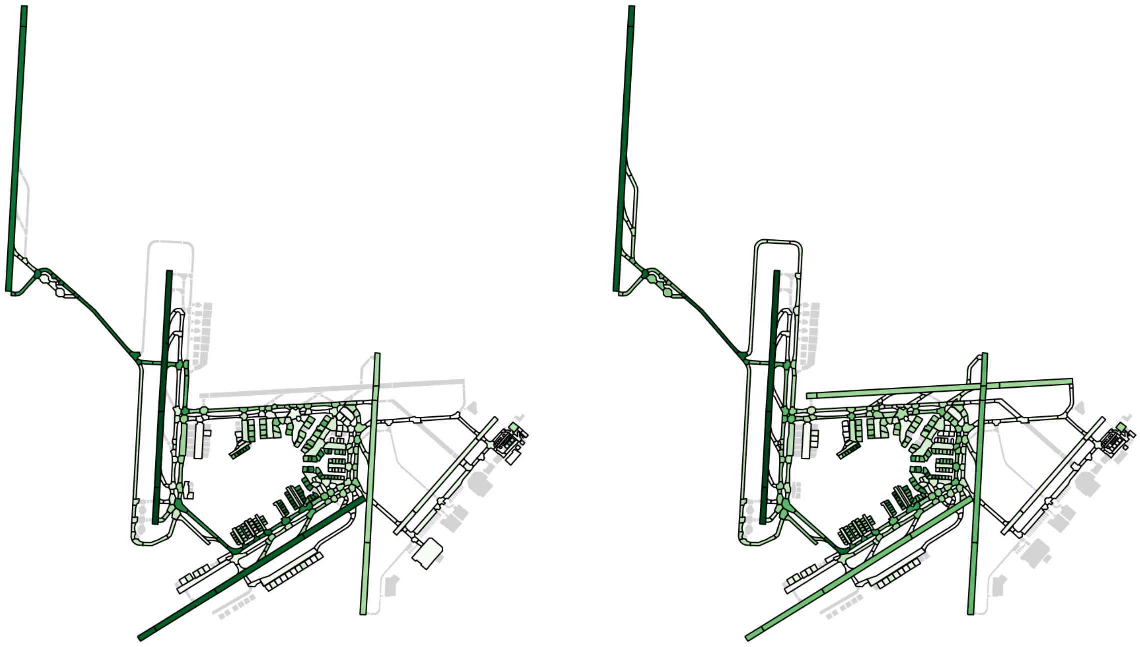

(a) Reference scenario 14 June

(b) Reference scenario 29 September

Figure 3.9: The segment frequencies of reference scenarios (darker is more frequent)

The two plots in figure 3.9 show how frequently each segment was visited in each of the radar days. A darker segment is visited more frequently than a lighter one, and the grey segments were not visited at all. It can be seen that aircraft stands on the south side of the green zone (platform A, piers B and C) see much more use than stands on the north-east side. This is explained by narrowbodies having a shorter turnaround (and thus more aircraft making use of the same stand), but the traffic implication of this is that the taxiways around this area will also see more traffic, especially taxiway Q, which is the single taxiway located just south-west of the A-platform.

Finally, exploring the runway usage of these two studied days will provide some context for the simulated experiments. It can be seen in figure 3.10 that the Polderbaan (36L-18R, indicated with an orange column label) is in use for departures for nearly the whole day on 14 June, and that there is a much more varied runway usage on 29 September. Interestingly, the runway usage almost reverses entirely at 1pm, indicating a dramatic change in wind conditions that likely occurred.

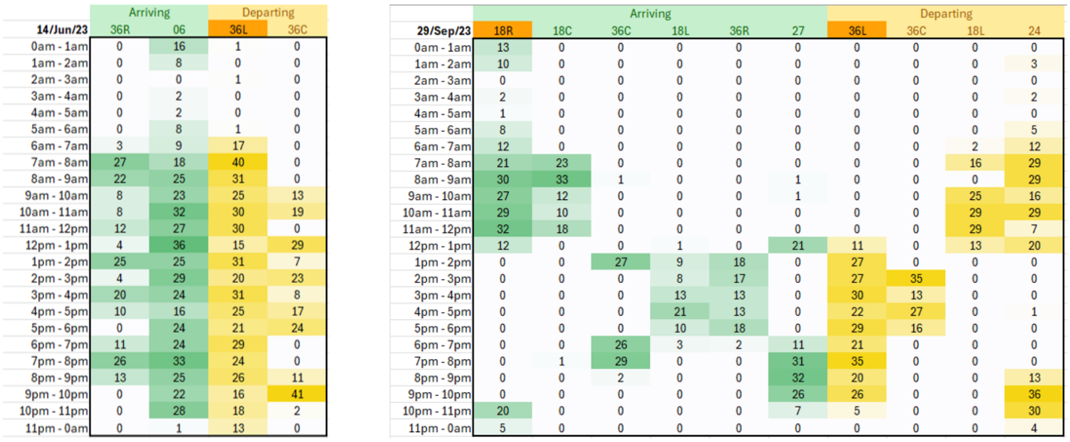

	Arriving		Departing	
14/Jun/23	36R	06	36L	36C
0am - 1am	0	16	1	0
1am - 2am	0	8	0	0
2am - 3am	0	0	1	0
3am - 4am	0	2	0	0
4am - 5am	0	2	0	0
5am - 6am	0	8	1	0
6am - 7am	3	9	17	0
7am - 8am	27	18	40	0
8am - 9am	22	25	31	0
9am - 10am	8	23	25	13
10am - 11am	8	32	30	19
11am - 12pm	12	27	30	0
12pm - 1pm	4	36	15	29
1pm - 2pm	25	25	31	7
2pm - 3pm	4	29	20	23
3pm - 4pm	20	24	31	8
4pm - 5pm	10	16	25	17
5pm - 6pm	0	24	21	24
6pm - 7pm	11	24	29	0
7pm - 8pm	26	33	24	0
8pm - 9pm	13	25	26	11
9pm - 10pm	0	22	16	41
10pm - 11pm	0	28	18	2
11pm - 0am	0	1	13	0

	Arriving						Departing			
29/Sep/23	18R	18C	36C	18L	36R	27	36L	36C	18L	24
0am - 1am	13	0	0	0	0	0	0	0	0	0
1am - 2am	10	0	0	0	0	0	0	0	0	3
2am - 3am	0	0	0	0	0	0	0	0	0	0
3am - 4am	2	0	0	0	0	0	0	0	0	2
4am - 5am	1	0	0	0	0	0	0	0	0	0
5am - 6am	8	0	0	0	0	0	0	0	0	5
6am - 7am	12	0	0	0	0	0	0	0	2	12
7am - 8am	21	23	0	0	0	0	0	0	16	29
8am - 9am	30	33	1	0	0	1	0	0	0	29
9am - 10am	27	12	0	0	0	1	0	0	25	16
10am - 11am	29	10	0	0	0	0	0	0	29	29
11am - 12pm	32	18	0	0	0	0	0	0	29	7
12pm - 1pm	12	0	0	1	0	21	11	0	13	20
1pm - 2pm	0	0	27	9	18	0	27	0	0	0
2pm - 3pm	0	0	0	8	17	0	27	35	0	0
3pm - 4pm	0	0	0	13	13	0	30	13	0	0
4pm - 5pm	0	0	0	21	13	0	22	27	0	1
5pm - 6pm	0	0	0	10	18	0	29	16	0	0
6pm - 7pm	0	0	26	3	2	11	21	0	0	0
7pm - 8pm	0	1	29	0	0	31	35	0	0	0
8pm - 9pm	0	0	2	0	0	32	20	0	0	13
9pm - 10pm	0	0	0	0	0	26	26	0	0	36
10pm - 11pm	20	0	0	0	0	7	5	0	0	30
11pm - 0am	5	0	0	0	0	0	0	0	0	4

Figure 3.10: Runway usage on the selected days

3.3.3 Towing policies

The TaxiBot towing policies can be divided in two categories: towing to a runway, and towing out of the apron bays. They both have different purposes. Towing to a runway saves the most fuel, and thus CO2 and UFP emissions. Towing out of the apron bays has the purpose of complying with the NLA's green zone ruling (Netherlands Labour Authority, 2023a) in the most operationally efficient way: shorter distances mean more towing activities, given a fixed number of TaxiBots. Here, aircraft are only brought to the edge of the green zone, where they can start their engines and continue on their way.

Since the focus of TaxiBot operations lies on green zone emission reduction, and this thesis investigates the traffic impact thereof, only aircraft that depart from stands inside this green zone are considered for TaxiBot operations. Aircraft outside this zone, like general aviation and the various cargo platforms, continue operating as usual.

There are nine possible towing options that are considered as policies: six for runways (figure 3.11a), and three for bay edge towing (figure 3.11b).

The six runway policies are inbound and outbound actions for three runway ends; the Polderbaan runway, the north side of the Zwanenburgbaan runway (at the U-platform), and the south side of the Zwanenburgbaan runway (at the P4 and P5 holding areas). Please note that the runway assignment for flights in the simulation does not change; only flights that originally made use of these runways will be considered for these towing policies.

The three locations selected for bay edge towing are the J-platform located to the north-west of the piers, the P-platform located to the north-east of the piers, and the R-platform located to the south-west. Which of the three bay edge towing policies is selected, depends on the departure stand of the aircraft. The aircraft is always towed to the closest towing destination according to the cyan dotted lines in figure 3.11b, with one exception: traffic departing from runway 18L (right next to the P-platform) will always be towed to the P-platform. A similar case could be made for towing all traffic bound for runway 09 to the J-platform, but since this runway was not in use for departures in the simulation days, this scenario does not occur.

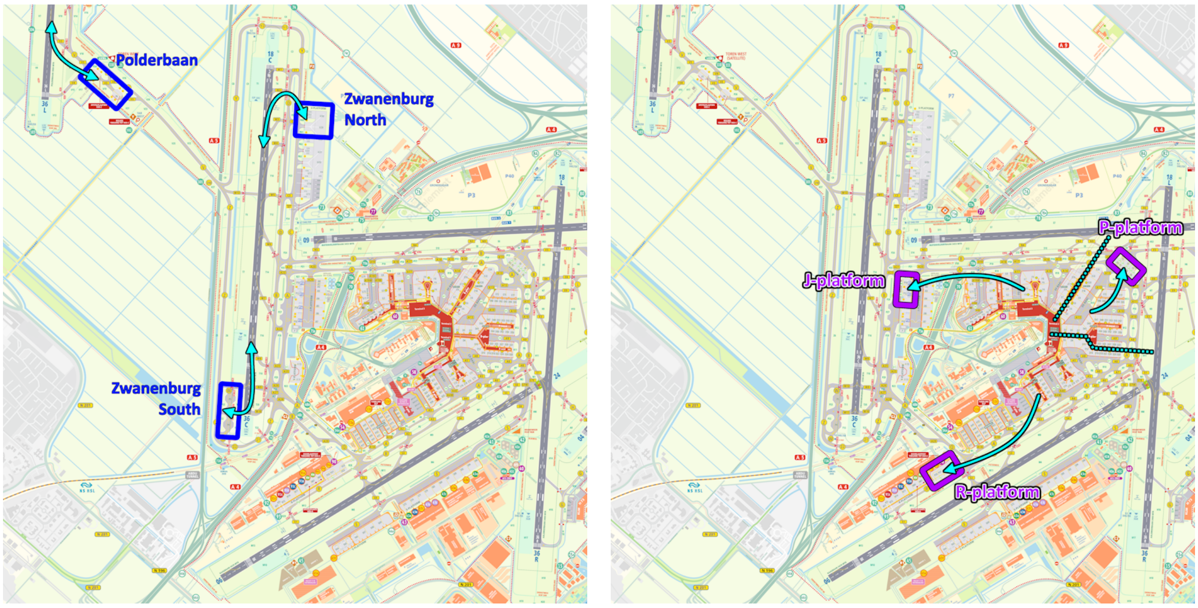

(a) Towing towards the runways

(b) Towing out of the bays

Figure 3.11: TaxiBot towing destinations

The bay edge towing policies are only considered for outbound traffic. Towing inbound traffic from the bay edge to the stand is deemed unfeasible from a conceptual perspective, for the following reason: according to Schiphol operational experts, aircraft engines need to cool down for four minutes before they are able to be shut down. Inbound taxiing durations from the central runways (where the runway towing policies from the previous paragraphs would not apply) are usually shorter than – or close to – that. In the meantime, these engines would still be producing emissions. Waiting for the engines to cool down and shut down would therefore nearly double the operational time. As a result, this option is a non-starter from both operational as well as emission reduction points of view.

Besides selecting a towing location, there is also the policy option of whether or not to allow the aircraft to start their engines during the towing movement. Since this process can take anywhere from three to seven minutes, a lot of time could be saved by starting this process while the aircraft is still being towed, rather than while occupying precious unloading station capacity. The assumption is made for this research that engine starting during towing is possible for all aircraft types and TaxiBots. The engine start moment would be timed in such a way that the engines finished their startup at the moment that the TaxiBot has finished unloading the aircraft. However, to still respect the aim of reducing emissions in the apron bays, these engine start-up moments cannot happen before an aircraft has left the apron bay, even if that results in a longer waiting time at the unloading station.

For the simulation, each of the six towing destinations is assumed to have the same capacity: two holdings, with two stations per holding. Each holding is able to accommodate two narrowbody aircraft (one per station), or one widebody aircraft (occupies both stations, due to size), see figure 3.12. This results in three options: four narrowbodies can be (un)loading at the same time, or two widebodies, or one widebody and two narrowbodies. The simulated station assignment logic for narrowbody aircraft ensures that a narrowbody aircraft will always try to select the spot next to another narrowbody aircraft, with the purpose of keeping the second holding available for a widebody aircraft to prevent unnecessary waiting.

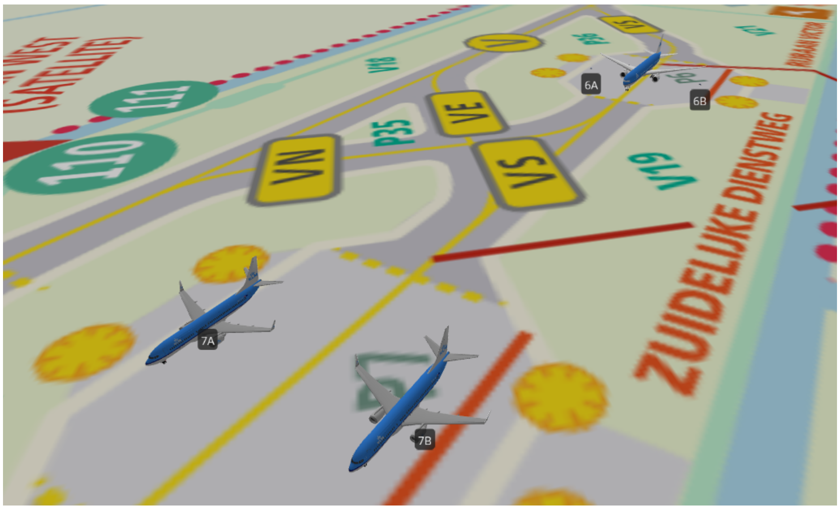

Figure 3.12: Each station can accommodate two narrowbodies, or one widebody aircraft

As a final note, it is important to emphasize that these policies do not represent any official intentions or planned procedures by Royal Schiphol Group, despite being developed in consultation with Schiphol operational experts. Except for outbound towing to the Polderbaan, all operational plans outlined here are intended solely for the purpose of this thesis.

3.3.4 Experiments to simulate

Combining the above towing policies and the selected days for simulation, nine experiments and two reference scenarios are drawn up. The reference scenarios simulate the airport as it was that day; purely according to the radar, with no TaxiBot operations. The scenarios are presented in the below table. Each change relative to the previous row is highlighted in bold.

Table 3.4: Experiments

#	Towing destinations	Direction	Radar day	Aircraft	Startup
0.1	Reference scenario	-	14 Jun	-	-
0.2	Reference scenario	-	29 Sep	-	-
1	Polderbaan only	Outbound	14 Jun	Narrowbody only	Stationary
2	Polderbaan only	Outbound	14 Jun	Narrowbody only	**In motion**
3	Polderbaan only	Outbound	14 Jun	**All types**	In motion
4	Polderbaan only	**In and out**	**29 Sep**	All types	In motion
5	**Polderbaan & Zwanenburgbaan**	In and out	29 Sep	All types	In motion
6	**Bay edges (platforms J, P, R)**	**Outbound**	29 Sep	All types	**Stationary**
7	Bay edges (platforms J, P, R)	Outbound	29 Sep	All types	**In motion**
8	**All**	Outbound	29 Sep	All types	In motion
9	All	**In and out**	29 Sep	All types	In motion

The nine TaxiBot experiments start with conditions that are currently being trialled in real life at Schiphol airport. Gradually, the experiments are scaled up in complexity. All results per experiment will be compared to the appropriate reference scenario, to determine the impact of the TaxiBot operations in the experiment. These impacts can then be compared to the impact of other experiments in order to draw conclusions on which TaxiBot operations perform best.

The first experiment is about current Schiphol trials: outbound towing of narrowbody aircraft from the stand to the Polderbaan, with a stationary startup at the unloading station. The second experiment is identical, except for allowing startup in motion. This allows for a direct comparison and thus a direct evaluation of this policy. Experiment 3 then adds widebody aircraft to the TaxiBot operations, to investigate their impact on emissions and station occupation. For these first three experiments, the 14th of June (maximum polderbaan outbound movements) is selected as the radar day.

The other six experiments use the radar day of 29 September. Experiment 4 also investigates the Polderbaan, this time with a mix of inbound and outbound towing. The impact of additional runway operations (Zwanenburgbaan, both directions) are then investigated in experiment 5.

Experiments 6 to 9 are about towing aircraft to the edge of the green zone. Experiments 6 and 7 once again investigate this impact for both options of starting engines while stationary, and while in motion. Experiment 8 then also includes outbound runway operations, to introduce a mix of operations, and also explore the effect of alleviating pressure on the unloading stations, by spreading capacity over more stations around the airport. Experiment 9, finally, adds inbound runway operations, for a 'maximum operations' scenario.

3.4 Model output

The aim of this thesis – see the research questions in section 1.2 – is to investigate multiple types of impacts of TaxiBot operations: flow and on time performance of both TaxiBotted and non-TaxiBotted aircraft movements, congestion around the (un)loading stations, and emissions in the form of CO2 and UFP. Additionally, reporting on the operations of the TaxiBots themselves will bring insight into the scale and feasibility of the operations.

This results in four categories of results: flow and on-time performance (OTP), TaxiBot operations, station congestion, and emissions. Before defining results for the emissions category, it is necessary to first understand the fuel usage at the airport, as well as the calculations for converting fuel into CO2 and UFP emissions. This will be done in section 3.4.1, after which the model outputs will be discussed in section 3.4.2.

3.4.1 Calculating engine durations and resulting emissions

Engine durations and assumptions

First, it is necessary to determine when an aircraft's engines and APU are running, and when they are not. As the radar data or CISS database contain no engine usage information, this has to be based on assumptions. For these assumptions, some experts within Schiphol and even a KLM pilot have been consulted, though it should be stressed that – like the towing policies – the assumptions here are purely for use in this thesis, and do not reflect the positions of those companies. Additionally, slight inaccuracies in absolute emission calculations are not the end of the world, because they are being compared to a reference scenario. For example, if both the reference scenario and the TaxiBot experiment overestimate APU fuel use by 5%, the reduction factor will remain the same.

Assumptions related to APU usage:

- For inbound movements, the APU is turned on one minute before arriving at the stand (AIBT).
- For inbound movements, the APU is turned off three minutes after arriving at the stand. The assumption is made that the aircraft can then get its power from a ground power unit.
- For inbound movements with a TaxiBot, the APU is turned on one minute before the start of the TaxiBot loading process, and turned off as usual (after three minutes) at the stand.
- For outbound movements, the APU is turned on five minutes before pushing back from the stand (AOBT). The assumption is made that the aircraft can be connected to ground power until then. In reality at Schiphol, the rule is not five minutes before AOBT but five minutes before target (engine) start-up approval time (TSAT). This timestamp is not included in the model, so AOBT is used as a good approximation.
- For outbound movements, the APU is turned off upon the completion of the pushback (and thus engine start) process.
- For outbound movements with a TaxiBot, the APU is turned off after the engines are started.

Assumptions related to engine usage:

- For inbound movements, engines are turned 'on' upon exiting the runway (runway engine use is out of scope).
- For inbound movements, engines are turned off upon arriving at the stand (AIBT). If the total engine time from runway exit until stand is less than four minutes, this runtime is increased to four minutes (due to the required cool down time before shutdown of four minutes, as mentioned earlier in section 3.3.3).
- For inbound movements with a TaxiBot, the engines are shut down upon the completion of the TaxiBot loading process. The same four minute minimum applies.
- For outbound movements, the engines are turned on at the start of the pushback process.
- For outbound movements, the engines are turned 'off' upon entering the runway (runway engine use is out of scope).

- For outbound movements with a TaxiBot, the time engines are turned on is dependent on the 'engine start during towing' policy. If engines are started when at the unloading station, then the engine start moment is the start of the TaxiBot unloading process. If engines can be started during the tow, then the engine start time is defined in such a way that engines have finished starting up when TaxiBot unloading completes ($T_{EngineStart} = T_{TaxiBotUnloaded} - T_{EngineStartDuration}$). The one exception to this calculation, is that engines may not be started inside the green zone. So if $T_{EngineStart}$ is earlier in time than $T_{ExitedBay}$, then instead $T_{ExitedBay}$ is taken as the start time, and the aircraft will wait at the unloading station until the engines have finished starting up.

Fuel usage

Next, these engine durations have to be translated to amounts of fuel. This data is available at Schiphol per aircraft type for both the APU and engines, averaged per engine variation (*background: many aircraft types have multiple engine supplier options, with slightly varying statistics regarding power, fuel flow, etc.*). Unfortunately, the simulation model works with aircraft categories rather than types, so the types had to be assigned to the categories in the same way as described in the last paragraph of section 3.1.2 (Royal Schiphol Group, n.d.-b).

Along with the fuel flow per aircraft type, a count for each aircraft type of how often it has visited Schiphol in 2023 was used. Per aircraft type, the following data points are now combined: category, visit count for 2023, fuel flow for all engines at idle thrust (used for taxiing), and fuel flow for the APU.

Then, weighted averages for each aircraft size category can be established, for both APU fuel flow and engine fuel flow. The visit counts are used as the weights. The result is an overview of fuel flow in kilograms per minute for each aircraft category, see table 3.5 below.

The assumption is made that all aircraft taxiing activity at the airport happens at idle thrust, using all available engines (so no single-engine taxiing procedures, as explored in section 2.2.4). Any engine activity on the runways is out of scope for the emission calculations, as this is a taxiing and ground traffic simulation model.

Table 3.5: Fuel flow per aircraft category

Category	Engine fuel flow [kg/min]	APU fuel flow [kg/min]
Regional	8.924	1.800
Narrowbody	12.863	1.936
Widebody	36.171	4.792

Converting fuel usage to CO2

Converting fuel usage to CO2 is relatively straightforward, because the APU and the engines make use of the same type of fuel, called Jet-A1. The standard conversion factor for a kilogram of Jet-A1 fuel to a kilogram of CO2 emission by an aircraft engine is defined by IATA, the International Air Transport Association, as 3.16 kg CO2 per kg Jet-A1 (IATA, 2022, p.8).

That being said, the European Union Aviation Safety Agency EASA, as part of the European Union, has defined several targets to reduce aviation emissions, by way of introducing a type of advanced bio-fuel called 'sustainable aviation fuel', or SAF. For 2025, EASA mandates that 2% of all fuel used in air transport in Europe should be SAF (EASA, n.d.). IATA mentions that SAF can reduce CO2 emissions by up to 80% compared to regular Jet-A1 (IATA, 2023).

Combining these policies and percentages, the modelling decision is made to assume a 2% SAF mix, having a 75% reduction of CO2. This results in the fuel mix having a CO2 impact of 3.113 kg of CO2 per kg of fuel.

Introducing this factor into the fuel flows determined earlier, CO2 emissions per minute of an engine running can be calculated, see the table below. As previous, the engine statistics are for all engines running.

Table 3.6: CO2 emissions per minute, per aircraft category

Category	Engine CO2 emissions [kg/min]	APU CO2 emissions [kg/min]
Regional	27.778	5.602
Narrowbody	40.038	6.027
Widebody	112.586	14.916

Ultrafine-particle data

Despite multiple UFP-measuring experiments happening at Schiphol currently, it has not been possible to get rough estimates, let alone quantitative data, on how much UFP is emitted by an aircraft engine. This is because the UFP measurements currently happening at the airport (Royal Schiphol Group, n.d.-c) are focused on measuring the extent to which workers are exposed to ultrafine particles, rather than precisely how many particles come from one specific engine. This data is not usable because the measurement locations are not positioned directly behind (and close enough to) one engine, and do not control for factors like APU usage, winds dispersing ultrafine particles, and ultrafine particles emitted by other passing aircraft hitting the sensors.

The UFP measuring devices found levels to be too high for safe working conditions, leading to the NLA's ruling that engine emissions in the apron bays need to be reduced. Schiphol is working on multiple solutions for that, with TaxiBot operations being one of them. That makes it important to somehow attempt to quantify the impact that various TaxiBot operational policies would have on the UFP-situation in the apron bays.

The solution to this dilemma is as follows. While it is not known how much UFP comes out of an engine, it is known which aircraft have their engines turned on at any given time, and what their position is. To show the impact of UFP emissions around the airport – particularly within apron bays – heat maps can be generated, distinguishing each engine mode differently.

The assumptions underlying this heat map are that the engine start-up is the most harmful period regarding UFP emissions, both in absolute number of particles emitted, as well as the fact that these occur within the apron bays. A second assumption is that the APU emits much less UFP than an engine. This assumption is based on the fact that the APU consumes 5.0 to 7.5 times less fuel (table 3.5). A third assumption is that the UFP that is emitted has a local impact, and that it is not blown around the airfield by winds and engine exhausts. The final assumption regarding these heat maps is that the UFP emission per second is a constant value, and not one that changes as time, engine temperature, or other conditions change – with the exception of the engine startup process.

The aim is to shed some initial insight into how UFP emissions would change for each experiment of different TaxiBot operation policies, rather than to definitively conclude on the effectiveness of TaxiBot operations on ultrafine particles within apron bays. A complicated emission dispersion model is too complex for this research due to the variations in wind speed, gusts and direction, engine exhaust streams of passing aircraft, and UFP levels emerging from those aircraft, and is therefore out of scope.

CO2 emissions of a TaxiBot

The TaxiBots that Schiphol owns have diesel-electric generators in them, powered by two V10 engines (per TaxiBot). However, future versions of the TaxiBot will be electric vehicles powered by batteries, so for this research the same will be assumed. As mentioned in the previous section, these electric TaxiBots are assumed to have no UFP emissions, but for the sake of completeness, its CO2 contributions will be included in the model. This is because electric vehicles are zero-emission locally, but the power generated for their batteries still contributes to global CO2 emissions. The Dutch Bureau of Statistics (CBS) calculated that in 2022, 0.27 kilograms of CO2 were emitted for every kWh of electricity generated in the Netherlands (Centraal Bureau voor de Statistiek, 2023).

Schiphol has received a brochure for the narrowbody version of this envisioned electric TaxiBot, and while not all details can be made public, the specifications listed in the brochure form the basis of the CO2 calculation for the TaxiBot. Here, the battery size, expected mission count and assumed distance per mission are first calculated to a certain value of energy use per distance (kWh/km) for the narrowbody TaxiBot, before continuing the calculation using the following assumptions. The result of this calculation for the CO2 contribution of the TaxiBot can be seen in table 3.7 below.

- Widebody pushback trucks (non-TaxiBots) use 8 liters of diesel fuel per pushback, and narrowbody trucks use 6 liters. Therefore, the assumption is made that widebody TaxiBots will use 1.33 times more electricity and thus need a 1.33 times larger battery than is advertised for the narrowbody TaxiBot.
- The energy spent while towing also depends on the weight of the towed aircraft. For this, 85% of the maximum take-off weight (MTOW) of a Boeing 737-800 is considered for narrowbody TaxiBots, and 85% of the MTOW of a Boeing 787-9 Dreamliner for widebody TaxiBots.
- The weight of the TaxiBot itself is assumed to be 30,000 kg and 42,960 kg for narrowbodies and widebodies respectively, based on the prevalent AM350 and AM500 aircraft pushback trucks in operation at Schiphol, and a 120% upscale factor to account for the battery pack.

Table 3.7: CO2 emissions per TaxiBot

Category	CO2 during non-tow [kg/km]	CO2 during tow [kg/km]
Narrowbody TaxiBot	0.355	1.063
Widebody TaxiBot	0.276	1.614

3.4.2 Performance indicators per analysis category

Flow and on-time performance

The flow and on-time performance is quantified using three statistics: inbound process time, outbound process time, and time out of bay. Traffic counters or other conventional traffic flow analysis methods are not required for the purpose of this thesis, because any impact on flow would translate to higher taxiing process times, which is the ultimate concern for Schiphol airport.

The inbound and outbound process times measure the time difference between the moment an aircraft exits the runway onto the taxiway and the moment it arrives at its stand, and vice versa. Naturally, runway crossings as part of the taxiing route are ignored for these calculations.

The 'time out of bay' measurement is calculated from the time difference between the departure time at the stand, and when an aircraft has exited one of the 29 TransferNodes in the network that connect a bay entry to the A/B taxiway ring-road around the Schiphol central apron.

These three statistics together will enable insights into inbound taxiing time, outbound taxiing time, and flow within the apron bays. A 'time into bay' statistic is not necessary because this process is the same for TaxiBot and non-TaxiBot operations.

TaxiBot operations

Per TaxiBot mission, multiple statistics are recorded. Distance travelled and travel time are recorded for both the towing and non-towing parts of the mission. Non-towing movement consists of travelling from the depot – fictionally positioned at the current location of the KLM pushback truck depot – to the aircraft, and after the mission, back from the aircraft to the depot. Because TaxiBot assignment/optimization is not within scope of this thesis, any complicated routing (e.g. depot-mission-mission-depot) scenarios are not considered. When not towing an aircraft, the TaxiBot makes use of service roads.

The towing and non-towing times are combined into a total mission time statistic, with their start and end times also recorded. An additional five minutes is added to this duration, owing to the assumption that a TaxiBot should be present at the mission start location five minutes before the aircraft is ready to be connected to the TaxiBot. This 'total mission time' allows for the calculation of the number of active TaxiBots at any one time in the simulation, and therefore also the maximum number of active TaxiBots. This statistic is used to approximate the number of TaxiBots required for these operations. This remains an approximation, as mission optimization is not applied and no charging or refuelling is considered. Finally, also a count of total performed missions for the whole airport is recorded.

Station congestion

To investigate the congestion around stations, two main measurements are considered: the waiting time for a station to become available, and the time at station (including waiting time) from requesting a station until release. This data is recorded per aircraft in each simulation experiment, and will be mainly reported as an average. To gain additional insight in the distribution of waiting times for a station and the time at a station, the 25th percentile, 75th percentile, and maximum time are also reported.

The congestion at the stations can also be considered from the perspective of the unloading destination (four stations) itself. For this, an additional two measurements are recorded, first of which is the average occupation utilization of each station over the day. This occupation rate is adjusted for activity at Schiphol: very little happens between 1 am and 5 am, and therefore this 24-hour rate is recalculated to a 20-hour rate.

The last measurement is the total time that all stations are occupied, and this can be thought of as a clock that starts running once all stations at a destination are occupied. This would mean that a next aircraft to arrive would not immediately have a station available, so that it would have to wait.

Emissions

The complexity of the emission calculations mainly lies in the engine durations and produced CO2, as presented in section 3.4.1. The final step in these calculations is to measure the time that the engines and APU are on, and then use the determined conversion factors to calculate kilograms of emitted CO2 per aircraft. A similar calculation is made for TaxiBots, this time using distances driven in towing and non-towing modes. The total CO2 emitted for an aircraft at Schiphol can then be calculated by adding up the CO2 contributions of the APU, the main engines, and the TaxiBot (if the aircraft had TaxiBot activity).

The UFP heat map is generated using positions and engine mode statuses, with each engine mode being a different color. The X and Y position in the plot represent the aircraft's location, while the color and size of the dot display the engine mode as follows: a large red dot is used for an engine start-up happening, a medium-sized yellow dot for aircraft engines running, and a small green dot is used for the APU running. Observations where all engines are turned off will not be plotted. TaxiBot vehicles themselves, being electric vehicles, are assumed to have zero (or negligible) UFP emissions.

These heat maps will not result in any quantitative reductions of UFP emissions, but they will show how the distribution of engine emissions at the airport is transformed by the various TaxiBot policy experiments.

An overview of the collected results per analysis category is displayed in the table below:

Table 3.8: Overview of results per category

Category	**Result**
Flow and OTP	Inbound process time Outbound process time Time out of bay
TaxiBot operations	Distance travelled towing Distance travelled not towing Time spent towing Time spent not towing Total mission time Maximum number of active TaxiBots Time of day of maximum active TaxiBots Total mission count
Station congestion	Waiting time at station Total time at station Station occupancy rate Time all stations occupied
Emissions	CO2 of aircraft engines in kg CO2 of TaxiBots in kg Total CO2 emitted for an aircraft's journey through Schiphol Engine mode timestamps (UFP) X,Y coordinates of aircraft every 5 seconds (UFP)

3.4.3 Approach for answering research questions

The collected results presented in the previous section allow the research questions to be answered. An overview of the research questions of this thesis are listed below.

What is the potential impact of TaxiBot operations on the flow and emissions of taxiing aircraft at Schiphol airport?

1. *What are current patterns and bottlenecks in the traffic flow at Schiphol airport?*
2. *How do TaxiBot process times compare to those of regular operations?*
3. *How do the TaxiBot operations affect apron bay flow?*
4. *How are the unloading locations near the Polderbaan affected by the envisioned operations?*
5. *What are the impacts of TaxiBot operations on emissions at the airport?*

Research question 1 is an exploratory one. Radar location-speed plots (heat maps) of traffic movements for the reference scenario experiments will be able to answer this question, no TaxiBot operation simulations are needed.

Research question 2 is about inbound and outbound process times. The flow and OTP-related results compared to reference scenarios will largely answer this question, supplemented by results related to engine start-up while stationary or in motion. This last policy alternative will be analysed by comparing the model results of experiments 1 and 2, and experiments 6 and 7. The experimental setups for these models are consistent across all model settings, except for the engine start policy. Experiments 1 and 2 focus on towing narrowbody aircraft to a runway, while experiments 6 and 7 investigate towing all aircraft types to the edge of the green zone. This dual approach enables the analysis for both long-distance towing (long time for engines to warm up) as well as stress-tested bay edge towing (due to all departures being towed to a few select locations).

Research question 3 relates to apron flow. The 'time out of bay' measurement for aircraft in the various experiments covers this question. Since the process of departing the stand and towing out of the bay is the same for aircraft across all experiments, similar results are expected across the experiments.

Research question 4 is about unloading locations and queueing. The station congestion results are used to answer this question. The OTP of non-TaxiBotted flights will show whether the unloading at the stations has hindered the flow. Comparing experiments 1 to 2 and 6 to 7 regarding the engine start-up policy (as already mentioned earlier in this section) will again be a valuable direct comparison. Furthermore, experiment 6 is seen as the ultimate stress-test for station congestion, given that all departing aircraft within the green zone will be brought to only 3 unloading locations, where they will also spend a long time (starting their engines). Experiments 7 (engine start during tow) and 8 (also towing to runways) alleviate the 'pressure' on the unloading stations, and the results of these three experiments can be used to draw conclusions about acceptable levels of station congestion.

Finally, research question 5 is about emissions. The CO2 calculations in each experiment will provide quantitative insight, and the proposed UFP heat maps are used to provide (provisional) insight, even in the absence of quantitative UFP data. These results are generated for each experiment, and especially the CO2 reduction will allow for direct comparisons between each experiment to assess environmental performance. Experiments 6 to 9 are the ones where every aircraft in the green zone – regardless of departure runway – is towed, and so these are expected to have the most significant UFP results.

4 Results

This chapter will introduce the results of the experiments as described in chapter 3. Per category of results, various model results and plots will be presented for all experiments at once, in order to draw conclusions on the differences between them. The results are generated by taking the results from an experiment, comparing them to the results from the reference scenario simulation run, and reporting on the differences between the two.

For easy reference when analysing the results, table 3.4 with experiment descriptions is repeated here:

#	Towing destinations	Direction	Radar day	Aircraft	Startup
0.1	Reference scenario	-	14 Jun	-	-
0.2	Reference scenario	-	29 Sep	-	-
1	Polderbaan only	Outbound	14 Jun	Narrowbody only	Stationary
2	Polderbaan only	Outbound	14 Jun	Narrowbody only	**In motion**
3	Polderbaan only	Outbound	14 Jun	**All types**	In motion
4	Polderbaan only	**In and out**	**29 Sep**	All types	In motion
5	**Polderbaan & Zwanenburgbaan**	In and out	29 Sep	All types	In motion
6	**Bay edges (platforms J, P, R)**	**Outbound**	29 Sep	All types	**Stationary**
7	Bay edges (platforms J, P, R)	Outbound	29 Sep	All types	**In motion**
8	**All**	Outbound	29 Sep	All types	In motion
9	All	**In and out**	29 Sep	All types	In motion

4.1 On-time performance and flow

Table 4.1 shows the average process times for inbound taxiing from runway to stand, outbound taxiing from stand to runway, and the time out of bay, for all aircraft, only TaxiBotted aircraft, and non-TaxiBotted aircraft. All on-time performance (OTP) times are relative to the reference scenarios in which no TaxiBot operations occurred.

Table 4.1: Process times for on-time performance, relative to reference scenarios

	All aircraft			Non-TaxiBotted aircraft			TaxiBotted aircraft		
	Inbound	**Outbound**	**Out of bay**	**Inbound**	**Outbound**	**Out of bay**	**Inbound**	**Outbound**	**Out of bay**
Exp 1	-00:01	+00:21	-01:22	-00:01	-00:05	00:00	–	+00:52	-03:04
Exp 2	-00:01	-00:34	-01:22	-00:01	-00:05	00:00	–	-01:12	-03:04
Exp 3	-00:03	-00:48	-01:47	-00:03	00:00	00:00	–	-01:30	-03:18
Exp 4	+00:38	-00:28	-00:58	-00:01	-00:00	00:00	+02:27	-01:31	-03:13
Exp 5	+01:32	-00:46	-01:18	-00:01	-00:01	00:00	+03:12	-01:53	-03:13
Exp 6	-00:02	+01:24	-02:44	-00:02	-00:26	–	–	+01:33	-02:57
Exp 7	+00:01	-00:29	-02:44	+00:01	-00:26	–	–	-00:30	-02:57
Exp 8	+00:01	-00:37	-02:44	+00:01	-00:23	–	–	-00:38	-02:57
Exp 9	+01:31	-00:34	-02:44	-00:01	-00:21	–	+03:10	-00:35	-02:57

Several patterns become apparent from the data in the table. Firstly, all TaxiBotting inbound operations add a significant amount of time to the taxiing duration. Next, outbound TaxiBot operations always result in a reduction of outbound time, except for when engines are started up while standing still at the unloading station (experiments 1 and 6). Third, the time gains for leaving the bay are fairly constant at a reduction of three minutes.

For non-TaxiBotted aircraft, the changes for all three measurements are negligible. Experiments 6 to 9 see a reduction in outbound times, this is caused by a few outliers of flights on the General Aviation apron (K-platform). The Q1 (25th percentile), median and Q3 (75th percentile) data points of the distributions of differences in outbound times of the experiments compared to the reference scenario, are all 0. It is likely that this pattern is always in the data (also for other experiments), but that this becomes apparent now that all outbound movements in the green zone get TaxiBotted, and thus disappear from this result category.

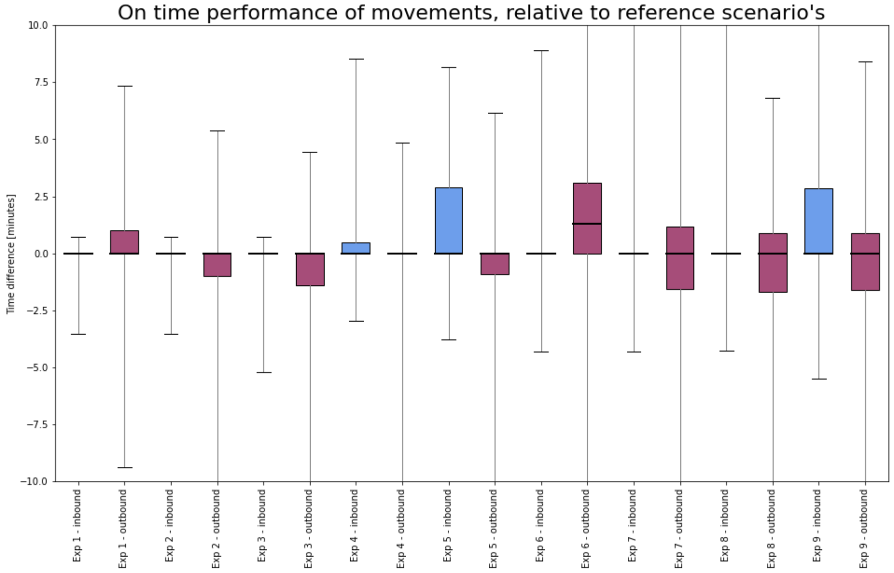

Figure 4.1: Inbound and outbound on-time performance distributions per experiment

Figures 4.1 shows the data from the first two columns of table 4.1 (and their distributions) as a box plot. Some whiskers have been cut-off to maintain detail in the colored boxes. Blue boxes represent inbound movements, and purple ones show outbound. The thicker black stripe represents the median. This is often 0, except for Experiment 6, where it has a value of 1.31 minutes (+1:19). This indicates that towing aircraft to the edge of the bays and starting engines stationary results in significant outbound delays; this will be explored more in section 4.3

The distribution of outliers for the on-time performance measurements can be seen in figure 4.2. The box plots from the previous figure are marked in a black outline.

It is apparent that the outliers are very large, both positively and negatively. This means that some flights took a lot longer for their inbound (blue) or outbound (pink) taxiing times, and that other flights took a lot less time. As can be seen in the figure, outbound taxiing results tend to have a greater spread. The causes for these outliers can be divided in two categories: TaxiBot operations and traffic interactions. Outliers are not caused by major routing changes; as flights in the experiment and reference simulations follow the same route around the airfield to the same runway or gate. Additionally, no strong patterns or correlations were found between certain routes and bigger or smaller outlier values.

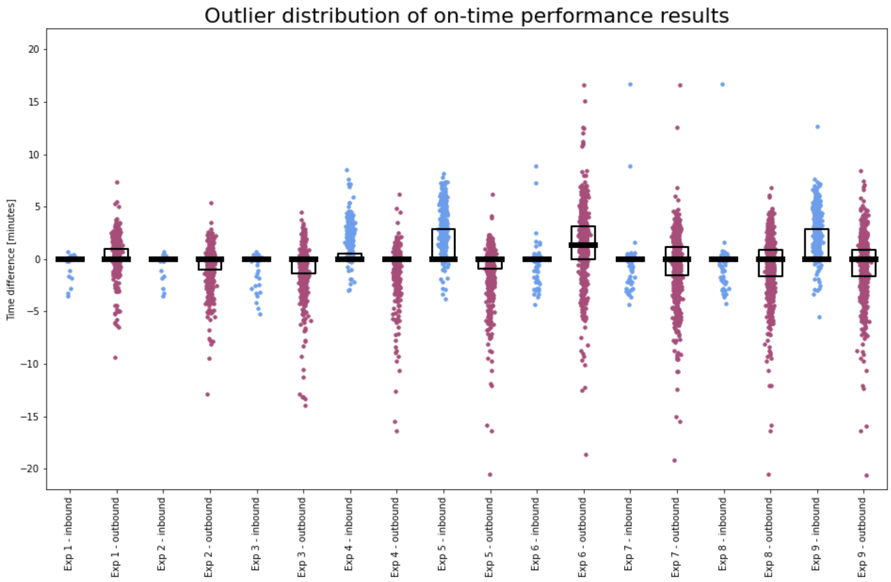

Figure 4.2: Distributions of inbound and outbound on-time performance results per experiment

Some flights see large OTP gains (or losses) as a result of 'extreme' TaxiBot operation impacts, like long queueing times or being able to forego a long pushback procedure by starting engines during the tow (N.B. the differences between experiments 1 and 2, and 6 and 7).

Traffic interactions can also cause large differences between the reference simulation and the TaxiBot experiment operations. Some flights have had to give way to a lot of other traffic in an experiment simulation, where they did not have to do so in the reference simulation (and vice versa). Runway sequencing (waiting to be granted take-off clearance) or being placed on a remote holding for multiple minutes in the reference simulation can also lead to operational differences between the reference and experiment simulations.

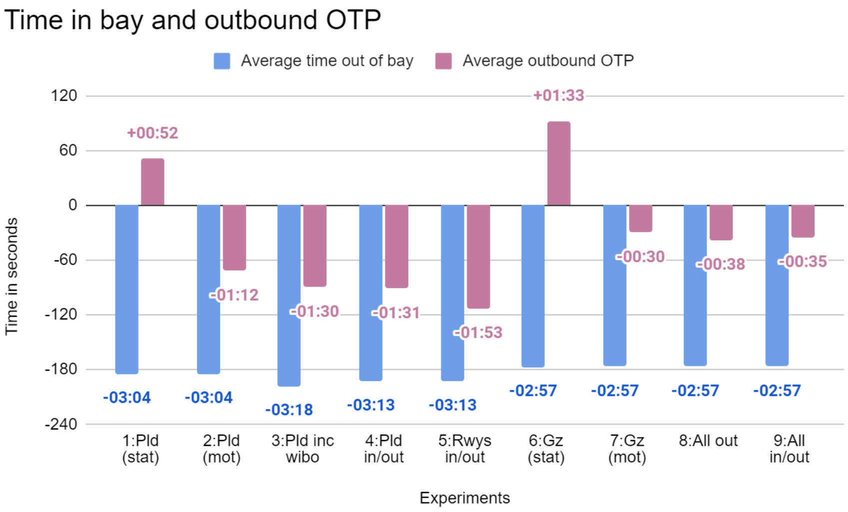

Figure 4.3: The impact of time in bay reduction on outbound on-time performance

Figure 4.3 shows the impact of stationary engine startup on outbound OTP. This plot clearly shows that the time out of bay is the main cause of the time gains. Some of this time is lost again in the unloading process, but if the TaxiBot is able to start the engines while taxiing (all experiments except 1 and 6), the aircraft will be faster than in the reference (no-TaxiBot) scenario.

Figure 4.4 displays the sum of 'time out of bay' reductions for all aircraft per experiment. The results show that up to 36 hours of bay capacity can be reclaimed by TaxiBot operations, without major impact on on-time performance. The variation in experiments 1-5 is explained by the number of TaxiBotted flights. The more flights get towed out of the bay, the higher the bay capacity savings will be. This simulation model does not include a gate planning (or re-planning) step, so no conclusions can be drawn on actual created capacity at the airport, but this is a promising area for future work.

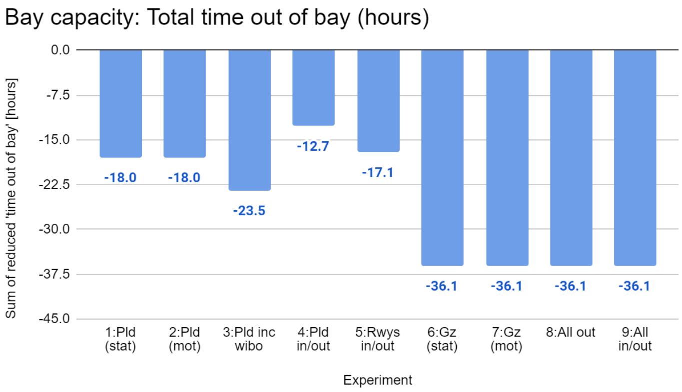

Figure 4.4: Reduction in time spent inside apron bays per experiment

The impact of TaxiBots on the flow on the taxiways is not explicitly quantified, but can still be assessed by an analysis of radar speeds in the reference scenarios.

The two plots in figure 4.5 display the radar speed heat maps of the reference scenarios, as earlier displayed in section 3.3.2. Data points in green indicate speeds above 22 kts, the upper limit of the TaxiBots. Red data points are speeds lower than 13 kts, and can be considered as aircraft travelling slowly.

The speeds around the apron bays, runway entry locations, and tight corners are a lot lower than those further out on the straight taxiways. There are also very few locations on the map where the prevailing taxiing speeds are faster (green) than the 22 kts speed limit of the TaxiBot. As a result, the vast majority of any on-time performance differences will have been caused by process steps like (un)loading the TaxiBot, starting engines, queueing for a station, and skipping the usual pushback process, rather than any type of congestion or traffic jam caused by aircraft being stuck behind a TaxiBot towing an aircraft.

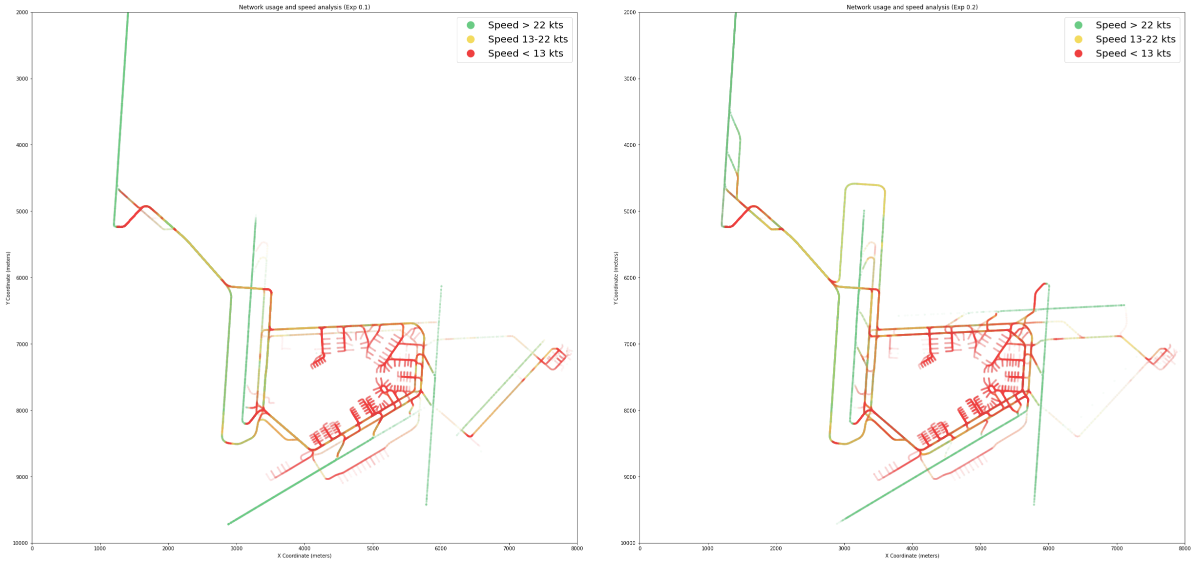

(a) Reference scenario 14 June

(b) Reference scenario 29 September

Figure 4.5: The radar speeds of reference scenarios (current day operations)

4.2 TaxiBot operations

Table 4.2 shows the statistics for the TaxiBot operations, from the perspective of the TaxiBots. The results for distance and time are recorded per TaxiBot mission (an inbound or outbound towing mission), and then averaged to be displayed in the table.

The distances are split in distance travelled towing an aircraft while driving on the taxiways, and distance travelled over service roads, travelling from the depot to the aircraft (and vice versa). As mentioned in section 3.4.2, assignment optimization is out of scope and so no complicated mission-to-mission routing exists. As a result, the non-towing distances will be overestimated compared to what an optimized model could achieve.

Table 4.2: TaxiBot simulation results (distances and times averaged for all missions)

	Distance [m]		**Time [mm:ss]**			**Operational**		
	Dist towing	**Dist no tow**	**Time towing**	**Time no tow**	**Time busy**	**Mission count**	**Peak TaxiBots**	**Peak time**
Exp 1	6076.7	10314.3	16:24	27:18	43:36	351	30	07:17 am
Exp 2	6111.2	10343.2	14:24	27:18	41:42	351	29	07:17 am
Exp 3	6066.8	10218.2	15:00	27:06	42:06	428	29	07:17 am
Exp 4	6639.6	10251.6	15:36	27:06	42:42	446	26	11:48 am
Exp 5	5632.2	8496.1	13:42	23:24	37:00	705	43	18:48 pm
Exp 6	1463.5	4854.6	09:30	15:30	25:00	649	28	07:07 am
Exp 7	1452.7	4847.3	07:18	15:30	22:48	649	27	07:07 am
Exp 8	3822.2	6901.7	11:18	19:54	31:12	649	35	21:26 pm
Exp 9	4338.3	7328.7	11:48	20:48	32:36	1033	51	08:04 am

Towing times are split in time towing and time not towing, combining into a 'time busy' statistic that adds them together, to represent the full time that the TaxiBot was occupied with this aircraft towing mission. The non-towing time includes a five minute waiting time, to account for the fact that a TaxiBot needs to be present at the mission pick-up location five minutes before the aircraft is ready to depart. The start and end times of the 'time busy' time window are used to calculate how many TaxiBots were active at one point in time. The maximum value and the associated time are presented in the final two columns of the table. Logically, most of these times are during the morning peak, which is the busiest time at Schiphol.

An interesting result is that for the majority of towing policies, the peak number of TaxiBots lies around 30. This is very close to the 38 TaxiBots that van Oosterom et al. (2023) calculated with their model, especially when considering that the simulation model does not account for charging the electric TaxiBots. The research by van Oosterom et al. (2023) found that 38 TaxiBots were needed for 913 flights; 26 small TaxiBots for 750 narrowbodies, and 12 large TaxiBots for 163 widebodies. Their research setup was closest to Experiment 5 of this simulation model, where all flights are towed to the runways (the main difference being that they considered all runways, versus this simulation only considering three runway destinations, see figure 3.11a). In the simulation, experiment 5 requires 41 narrowbody and 12 widebody TaxiBots to service 587 and 118 flights respectively. Direct comparison is difficult to due to the aforementioned factors of optimization, charging and slightly different model setups and the fact that different days were analysed, but the figures are in the same ballpark.

The split between narrowbody and widebody TaxiBots for all experiments in the simulations are shown in figure 4.6 and table 4.3.

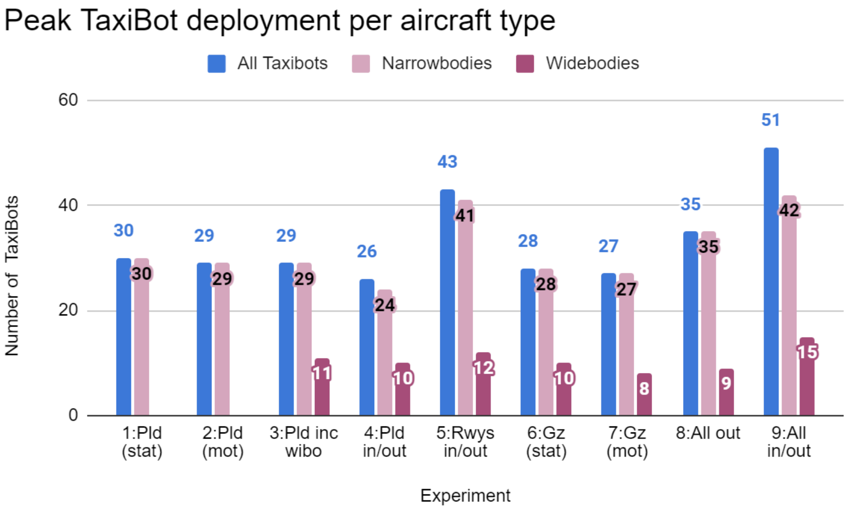

Figure 4.6: Peak TaxiBots active per size type

Table 4.3: TaxiBot deployment results

	All TaxiBots			Narrowbodies			Widebodies		
	Missions	**Peak active**	**Peak time**	**Missions**	**Peak active**	**Peak time**	**Missions**	**Peak active**	**Peak time**
Exp 1	351	30	07:17 am	351	30	07:17 am	–	–	–
Exp 2	351	29	07:17 am	351	29	07:17 am	–	–	–
Exp 3	428	29	07:17 am	351	29	07:17 am	77	11	10:38 am
Exp 4	446	26	11:48 am	357	24	19:01 pm	89	10	10:04 am
Exp 5	705	43	18:48 pm	587	41	18:48 pm	118	12	07:45 am
Exp 6	649	28	07:07 am	553	28	07:07 am	96	10	10:58 am
Exp 7	649	27	07:07 am	553	27	07:06 am	96	8	10:54 am
Exp 8	649	35	21:26 pm	553	35	21:26 pm	96	9	14:05 pm
Exp 9	1033	51	08:04 am	865	42	08:05 am	168	15	10:02 am

Figure 4.7 below presents a visual overview of some selected mission statistics. The shorter average towing distances for experiments 6 and 7 are reflected in shorter mission times and less TaxiBots are required, when comparing the results to experiments 8 and 9 where runways are also included as mission destinations.

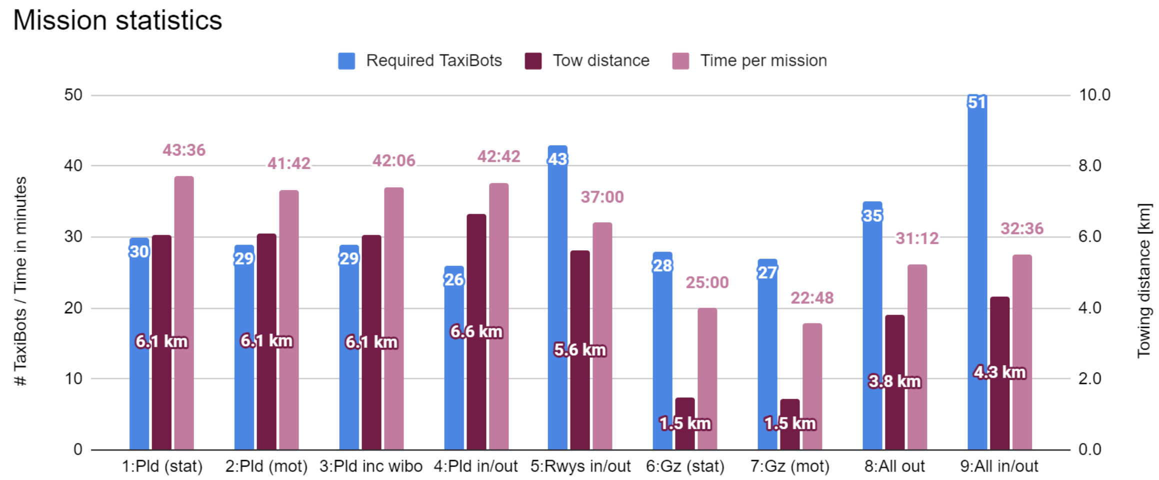

Figure 4.7: TaxiBot mission statistics, per experiment

4.3 Station congestion

This section will analyse the congestion around the (un)loading stations. The simulation results are shown in table 4.4, for both the queueing process before being assigned a station, and the total time spent queueing and at the station. For better insight in the distributions of these times, the Q1 (25th percentile), Q3 (75th percentile), and maximum values are also reported.

The queueing times, especially the Q1-Q3 ranges all being 0, show that aircraft rarely have to wait to get an (un)loading station assigned. The variation in average times therefore is caused by outlier values. A similar picture exists for total times at stations, where the average is driven up by maximum values as outliers.

Table 4.4: Station congestion results per experiment

	Queueing time			**Total time at station**			**Station occupancy**	
	Mean	**[Q1-Q3]**	**Max**	**Mean**	**[Q1-Q3]**	**Max**	**Occupancy rate**	**Time all occupied**
Exp 1	00:03	[00:00 - 00:00]	01:36	04:08	[03:15 - 04:49]	08:11	29.9%	70:30
Exp 2	00:00	[00:00 - 00:00]	00:00	02:10	[01:39 - 02:30]	05:35	15.9%	03:18
Exp 3	00:04	[00:00 - 00:00]	02:13	02:21	[01:46 - 02:45]	05:35	24.8%	59:36
Exp 4	00:02	[00:00 - 00:00]	03:08	02:20	[01:46 - 02:45]	05:38	27.1%	39:06
Exp 5	00:01	[00:00 - 00:00]	03:08	02:16	[01:44 - 02:39]	05:38	13.4%	14:24
Exp 6	00:12	[00:00 - 00:00]	14:32	03:54	[02:49 - 04:31]	19:22	20.1%	48:42
Exp 7	00:02	[00:00 - 00:00]	03:45	01:45	[01:03 - 02:12]	06:31	9.7%	12:18
Exp 8	00:03	[00:00 - 00:00]	03:45	02:03	[01:25 - 02:30]	06:31	6.6%	09:48
Exp 9	00:02	[00:00 - 00:00]	03:45	02:10	[01:35 - 02:38]	06:31	9.5%	11:54

The occupancy rates represent how much of the time a station was occupied. This value is recorded per station, then averaged for all stations that were available in the simulation run, and reported in the table.

The 'Time all occupied' column represents a total time of how often an (un)loading location (consisting of four stations) was unavailable for arriving aircraft, resulting in them having to queue. Separate data per station for each experiment can be found in appendix B, section B.1.

These results paint an interesting picture, in the fact that all occupancy rates are below 30%. In other words: the stations were unoccupied for the vast majority of the time. However, the times when all stations were occupied can get quite high, in experiment 1 even more than 1 hour and 10 minutes. This indicates that the regular levels of congestion around the stations are quite low, but that some peak moments can still lead to queueing.

Comparing experiments 7 to 8 (towing everything to the edge of the bay, versus to the runways where possible) shows the impact of distributing the unloading demand over more locations at the airport. Occupancy rates drop by a third, and the times that all stations are occupied also drop by around 20%. Interestingly, the maximum durations for queue times and times at stations are not impacted that much.

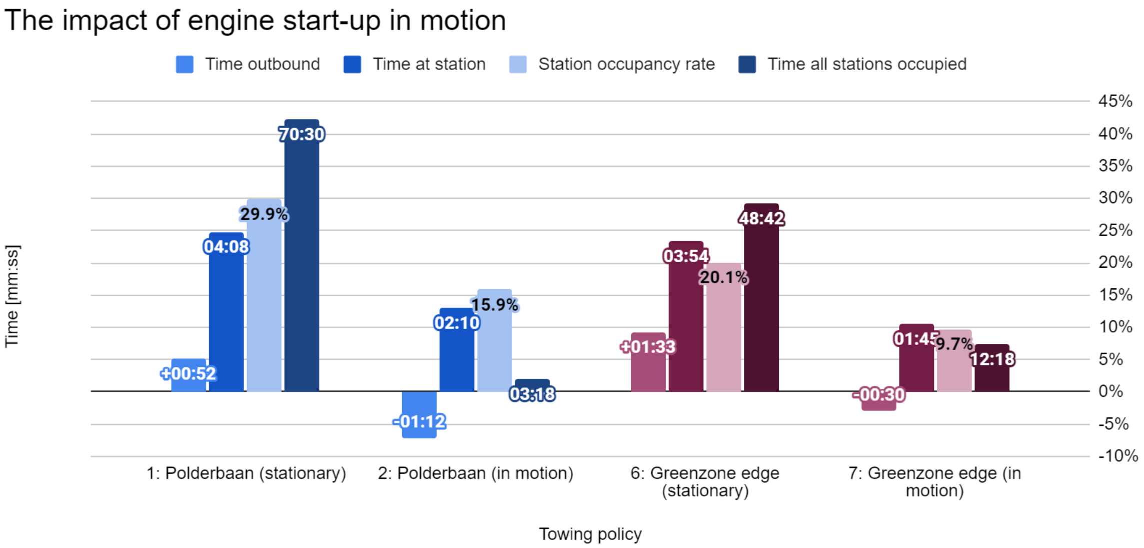

Figure 4.8: The impact of engine start-up in motion on station congestion

The impact of starting engines while stationary, or starting them during the tow already also becomes quite apparent, as presented in figure 4.8.

The plot compares two sets of experiments to each other; experiments 1 to 2, and 6 to 7. These experiments – narrowbody outbound to the Polderbaan (in blue), and towing all flights to the bay edges (in pink) – are identical except for the engine start-up policy.

The time at station drops by around 2 minutes in both scenarios, also resulting in 2 minutes difference in the outbound times. This is a major finding of the simulations: allowing engines to start up during the tow, can reduce outbound times by 30 seconds to 1 minute, instead of delaying them by 1 to 1.5 minutes! Additionally, the time that all stations are occupied also gets reduced by a significant amount; 95% and 75% for the sets of experiments respectively.

4.4 Emissions

The emission analysis of the simulations consists of quantitative CO2 calculations, and more qualitative UFP emission location heat maps.

CO2 emissions

Table 4.5 shows the impact of the TaxiBot operation experiments on CO2 emitted at the airport. The values are reported for all aircraft, as well as for aircraft which had an inbound or outbound movement (or both) TaxiBotted.

Table 4.5: CO2 results per experiment

	All aircraft			**TaxiBotted aircraft**				
	Engines on	**APU on**	**CO2 impact**	**Engines on**	**APU on**	**CO2 impact**	**CO2 [Q1-Q3]**	**Max CO2**
Exp 1	-04:34	+04:55	-13.6%	-10:11	+11:03	-37.0%	[-27.0% ; -48.9%]	-72.4%
Exp 2	-04:37	+04:02	-14.3%	-10:16	+09:05	-38.9%	[-29.5% ; -50.2%]	-73.8%
Exp 3	-05:53	+05:03	-27.2%	-10:49	+09:19	-41.3%	[-31.0% ; -50.4%]	-74.9%
Exp 4	-06:07	+07:03	-25.1%	-12:15	+14:10	-40.4%	[-29.2% ; -44.6%]	-74.4%
Exp 5	-07:42	+09:48	-30.3%	-10:43	+13:41	-37.9%	[-22.0% ; -45.0%]	-75.7%
Exp 6	-01:43	+03:06	-6.5%	-01:49	+03:20	-7.4%	[+5.1% ; -13.6%]	-54.4%
Exp 7	-01:45	+01:17	-7.8%	-01:51	+01:23	-8.9%	[+3.1% ; -15.7%]	-54.8%
Exp 8	-05:07	+04:32	-20.2%	-05:30	+04:53	-23.3%	[-3.2% ; -36.9%]	-74.4%
Exp 9	-08:15	+10:32	-32.9%	-08:52	+11:22	-36.5%	[-15.6% ; -43.7%]	-75.8%

Logically, CO2 impact increases throughout the experiments as TaxiBot activity increases. When comparing experiment 2 and 3, the only difference in policy (towing outbound to the Polderbaan, startup in motion) is that experiment 3 also includes towing widebodies. While there are many more narrowbody missions than widebody missions for experiment 3 (351 vs 77, see table 4.3), the CO2 impact nearly doubles. This means that including widebody aircraft in TaxiBot operations is a very worthwhile decision, when aiming to reduce CO2 emissions.

To determine how much CO2 can be saved by TaxiBotting an inbound or outbound flight movement, the max CO2 savings can be considered. This is because the 'TaxiBotted aircraft' category also contains aircraft that TaxiBot either inbound or outbound but use conventional taxiing for the other leg, making the average value of impact harder to interpret. Some flights only have one observed movement (inbound or outbound) because the aircraft stayed overnight at Schiphol. For these flights, if the observed movement is TaxiBotted, the overall CO2 savings for that aircraft are much higher compared to flights also using conventional taxiing for the other leg, leading to the max CO2 impact being the best indicator.

From this column of max CO2 impact values, it can be concluded that CO2 (and fuel usage) reductions of 72% to 75% are achievable when towing to the Polderbaan – this will always be the max CO2 saving, because the towing distance is the highest – and that savings of up to 53% are possible for towing to the edge of the bays – likely for flights departing runway 18L, which is situated very close to the P-platform used for unloading. This second type of CO2 impact is lower because the APU that was running at the stand starts playing a bigger role proportionally speaking, the shorter the taxiing distance is.

Speaking of engine run times, the values in the table and in figure 4.9 below show that the added APU run times are not always equal to the saved engine times. A few of the reasons causing these discrepancies, are extra engine times related to cooling down after landing before shutdown, overlap between engines and

APU during the engine start-up process, reduced APU usage by starting engines during towing and therefore spending less time stationary at the unloading location (compare experiment 1 to 2, and 6 to 7), and changes to APU run times courtesy of the reduced (or additional) outbound taxiing times as described in section 4.1 about On-time performance and flow. Allowing the engines to start during the tow also has a slight positive effect on CO2, reducing total airport emissions by 0.7 and 1.3% respectively.

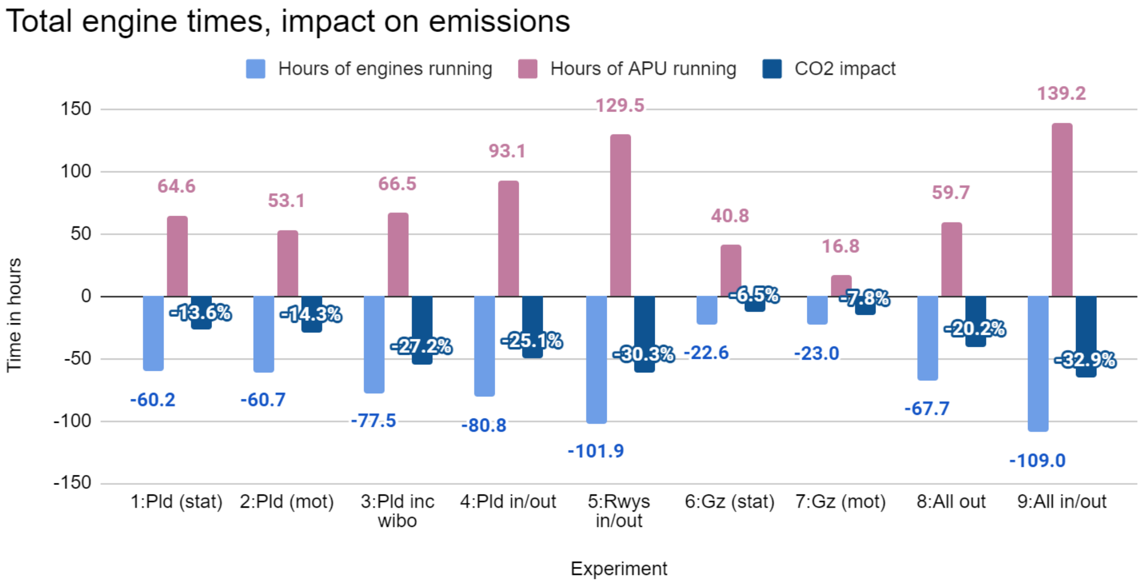

Figure 4.9: Engine times, and the impact on CO2 emissions

The figure above also shows the CO2 savings that the various TaxiBot operation policies would have on the emissions for the whole airport. The maximum impact is had in experiment 9 (where as much gets TaxiBotted as possible), where 32.9% of emissions are reduced, but it is interesting that even the runways-only scenario in experiment 5 already comes very close to this figure, with much less operational complexity. The impact of towing out of the bays on CO2 is much smaller comparatively (experiments 6 and 7), even though the mission count (705 vs. 649, table 4.2) is not that much less. This leads to the – perhaps self-evident – conclusion that the main contributor to CO2 savings is simply the engine running time saved, most often caused by the distance that aircraft are towed: an average of 5.6 kilometers in experiment 5, and 1.4 kilometers in experiment 6.

The impact of sustainable aviation fuel on CO2 emission reduction

In chapter 3, section 3.4.1, the calculated fuel usage figures were converted to CO2 emissions by assuming a mix of 98% Jet-A1 aircraft fuel, and 2% sustainable aviation fuel (SAF), resulting in an overall CO2 emission factor of 3.113 kilograms of emitted CO2 for every 1 kilogram of fuel burn. The results from simulations showed that CO2 emissions due to TaxiBot operations ranged from -6.5% to -32.9%, see table 4.5.

To investigate the impact of sustainable aviation fuel on CO2 emission reduction, the operations in experiment 9 (32.9% CO2 reduction) and the reference scenario of 29 September were also performed for fuel mixes of 0%, 15% and 50% SAF. The resulting CO2 emission factors are then 3.160, 2.805 and 1.975 kilograms of emitted CO2 for every 1 kilogram of fuel burn, respectively. TaxiBots run on electricity, so their CO2 emissions remain unaffected. The results of these runs are shown in the table below.

Table 4.6: The impact of sustainable aviation fuel on CO2 emission reduction

Fuel mix	CO2 emission factor [kg/kg]	Resulting total CO2 reductions
0% SAF	3.160	-32.868%
2% SAF	3.113	-32.854%
15% SAF	2.805	-32.751%
50% SAF	1.975	-32.315%

As can be seen in table 4.6, the changes in CO2 reduction are minimal, even for major changes in the amount of sustainable aviation fuel used at the airport.

The reason for this is that flights are compared to a reference scenario where the same fuel mix percentage is used. The only reason why the CO2 emission reductions slightly decrease for increasing amounts of SAF, is because of the impact of the TaxiBot's emissions: as the fuel becomes more sustainable, the total CO2 emissions of aircraft get smaller, and so the TaxiBot emissions start contributing more, proportionally speaking. The conclusion from these findings is that the CO2 reductions really are a result of TaxiBot operations, and that the SAF fuel mix used at airports does not influence these results.

UFP heat maps

The heat maps of ultrafine particle emissions give an idea of where emissions occur, but they can not say much regarding quantity of emissions. The intention for providing these heat maps is partially to provide insight for this thesis, but also for future researchers to use the data from these simulations to quantify the impacts of UFP emissions when more is known about how much UFP is actually emitted by an aircraft engine. The most interesting contrasts between experiments will be highlighted in this section, and the heat maps of all experiments are available in appendix B, section B.2.

As mentioned earlier in chapter 3, section 3.4.2, the UFP heat maps are generated using aircraft positions and their engine statuses, with each engine mode being a different color. The X and Y position in the plot represent the aircraft's location, while the color and size of the dot display the engine mode as follows: a large red dot is used for an engine start-up happening, a medium-sized yellow dot for aircraft engines running, and a small green dot is used for the APU running. Observations where all engines are turned off are not plotted. Ultrafine particle emissions from TaxiBot vehicles themselves, being electric vehicles, are assumed to be zero (or negligible) and are therefore also not plotted.

The reference scenarios of experiments 0.1 and 0.2 are shown in the two figures below. It becomes apparent how much of the harmful engine startup emissions currently occur inside the NLA green zone, within the apron bays.

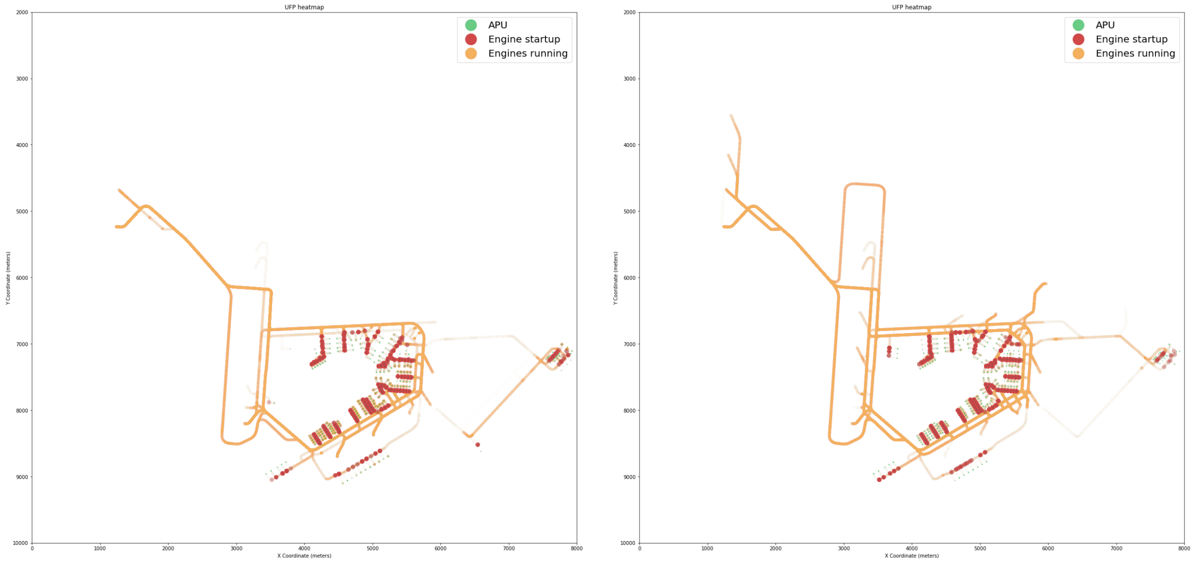

(a) Reference scenario 14 June

(b) Reference scenario 29 September

Figure 4.10: The UFP impact of reference scenarios (current day operations)

The two plots in figure 4.11 show the first two experiments. When comparing these to figure 4.10a, it can be seen that the impact on UFP emissions in the bays is not that noticeable. Every TaxiBotted flight helps, but the vast majority of engine start-ups still occur in the bays.

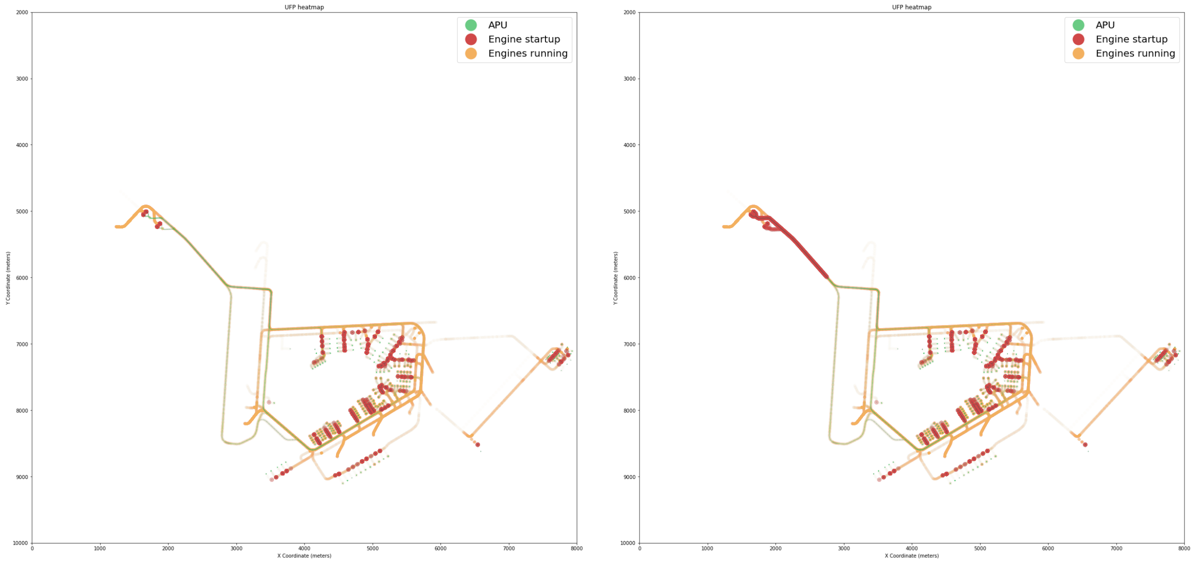

(a) Startup stationary at holding (Experiment 1)

(b) Startup in motion (Experiment 2)

Figure 4.11: The UFP impact of startup in motion

The main difference between the two plots in figure 4.11 is the location of the engine start-up starting point. The advantages from an OTP and queueing perspective have already been highlighted, and in terms of UFP emissions, there is no negative impact since the 'earlier' emission of UFP from the engine is still far enough from the green zone to cause harm. Therefore, it can be concluded that allowing aircraft to start their engines during the tow to the Polderbaan has many positives, and no drawbacks that become apparent from these simulations.

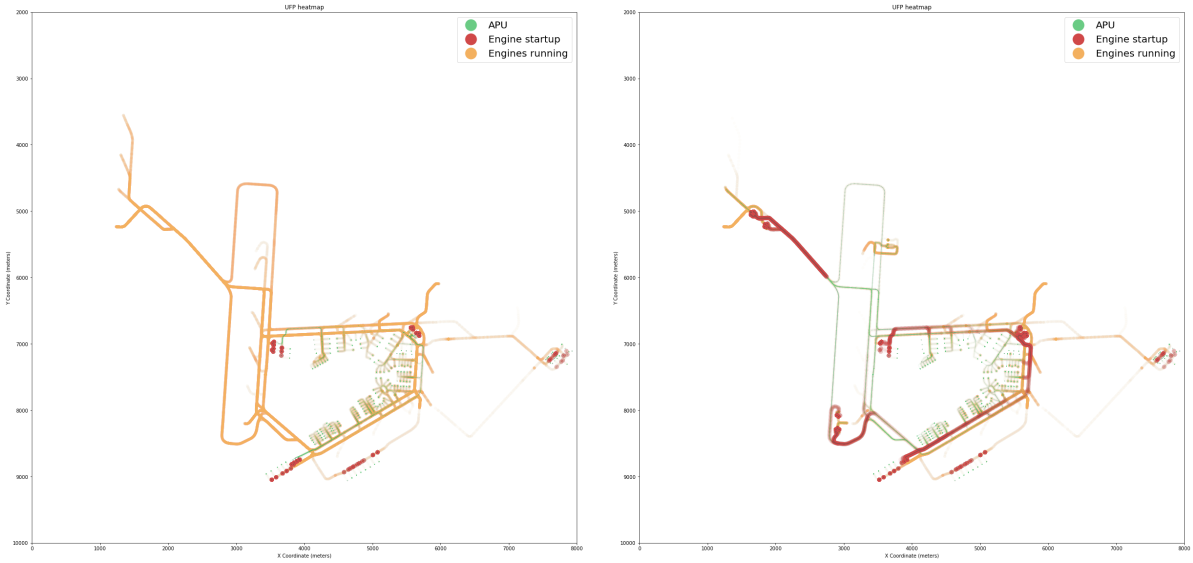

(a) Maximum UFP impact (Experiment 6)

(b) Maximum CO2 impact (Experiment 9)

Figure 4.12: The UFP impact of two promising experiments

The two heat maps in figure 4.12 show the emissions of experiment 6 and 9 (respectively: towing all outbound flights to the edges of the green zone with stationary startup, and towing to runways inbound and outbound with the rest being towed to the green zone edges, while starting all engines during the towing processes).

The results of experiment 6 in figure 4.12a are drastic, when compared to 4.10b. A lot of UFP is removed from the bays, and the stationary startup also means that the engine start-up emissions remain limited to the unloading locations. Operational downsides of this experiment have been detailed in earlier sections of this chapter, however.

The UFP heat map of experiment 9 in figure 4.12b shows the emissions for the scenario that reduces the maximum amount of CO2. A similar picture to that of experiment 2 (figure 4.11b) is seen regarding the Polderbaan unloading, and while the bays are also cleared of engine start-up UFP like in experiment 6, the engine start-up while driving results in a lot of UFP emissions happening on the ring-road around the apron bays. There currently is not enough data about the spread of ultrafine particles to know if the emissions there can reach the workers inside the apron bays, and therefore it cannot be said if this policy is acceptable or if it is too harmful.

5 Discussion

This chapter will discuss the results presented in the last chapter. Firstly, a section is dedicated to reflecting on the model validity. Secondly, key findings are highlighted, and then their implications on airport operations are considered. Section 4 details some limitations of the research, and section 5 proposes avenues for future work.

5.1 Model validation

There are multiple ways to validate a model, the most common of which being to compare the model reference scenario to reality, and to have the model peer reviewed. In this research, both methods are applied, and the model is found to be valid.

First of all, the model is inherently valid by replicating historic radar data and other real-life sources, like TaxiBot (un)loading times. This automatically ensures that realistic runway usage, routing, gate assignment, and taxiing speeds and (un)loading times are represented in the model. The modelled situations therefore naturally correspond to real life, because when no TaxiBot operations are simulated, a simulation of real life movements remains.

Secondly, the model shows high validity by reproducing findings earlier determined by other calculations, and real-life TaxiBot trials at Schiphol. Taxiing times, distances, and 'time busy' (total time a TaxiBot is occupied with towing one flight to the Polderbaan) that are used as assumptions for the TaxiBot trials, all corroborate the distances and times found in these simulations. Additionally, after the trials at Schiphol in the COVID-19 pandemic, fuel savings were calculated to be up to 50-65% for the outbound movement (stand to runway) of the towed flight (Royal Schiphol Group, 2021). These findings correspond to the results presented in table 4.5, experiment 1, where the 75% percentile of CO2 reductions is 49% and the maximum is 72% (NB: the CO2 values in that table are for the aircraft's whole journey through Schiphol, so including inbound movements that are not TaxiBotted).

Thirdly, the model is also validated by peer review. The findings of the simulation have been presented to the Sustainable Taxiing Core Team; a collaboration effort between Schiphol airport, participating airlines, and other relevant parties all working on the TaxiBot operations. These operational experts have expertises ranging from ground handling, to Schiphol operations, and even one pilot certified for both Boeing 737 narrowbody and Boeing 787 widebody aircraft. After learning how the model works, and being told about the results of the first experiment (that reproduced the earlier trials), they engaged with the other results in a way that showed real trust and belief in the findings.

As a second aspect of the peer review, the TaxiBot towing policies and unloading locations have also been assessed by consulting with various pushback, towing, and operational experts at Schiphol that were simultaneously working on similar projects for real-life operations. The policies in this thesis were not necessarily provided by them, but they agreed that these towing routes could feasibly be used in reality in the future.

While the model does not replicate any runs to determine uncertainty in any findings, this is not expected to be a concern. As mentioned in chapter 3, section 3.3.1, each simulated day consists of roughly 1400 flight movements. Each TaxiBot towing mission that occurs within that simulated day is essentially a 'mini-experiment' of the same towing operation under slightly different conditions (e.g. departure gate, traffic situation, unloading station congestion, radar speeds). All these observations together form the dataset that has lead to the results presented in this research, and therefore the findings are – in a way – already based on many replications of the same types of movements.

5.2 Key findings

The key findings of this research relate to taxiing times of inbound and outbound movements, engine start-up policy, traffic flow, emission reduction, and station congestion.

Inbound taxiing times see a significant amount of time added to the taxiing process due to the TaxiBot loading process, with not much time being saved elsewhere in the process to recover this time loss. Outbound taxiing times see either a delay or a time savings depending on engine start-up policy. Some scenarios see aircraft saving an average of 1 to 1.5 minutes (instead of adding 0.5 to 1 minutes) of outbound taxiing times, resulting in better on-time performance while also saving on fuel burn and emissions.

Allowing aircraft engines to start-up during the outbound towing operation saves significant time at the unloading station. This both reduces congestion at the unloading station, and improves the on-time performance of the aircraft. Additionally, this time saving also leads to a small additional reduction of CO2 emissions, while (likely) not causing additional UFP exposure.

The impact of TaxiBots on traffic flow is minimal to negligible, due to the TaxiBots ability to achieve towing speeds up to 22 kts. Taxiing speed analysis of aircraft has shown that this speed is rarely exceeded. This allows the TaxiBot to fit in with other aircraft in most cases, instead of being a hindrance to them.

Significant reductions in CO2 emissions were achieved across various experiments, with reductions ranging from 6.5% to 32.9% for different towing policies. The highest reductions occur for the longest towing movements in terms of time and distance. Ultrafine particle heat maps indicate a noticeable shift in UFP emissions away from the apron bays, potentially reducing exposure of ground crew. More data on quantified UFP emissions by aircraft is needed in order to confirm these initial results.

Station congestion is manageable for all performed experiments. The policy of engine start-up during the tow sees significant impacts on queueing and station occupancy times, and is therefore strongly recommended. Additionally, including more towing destinations in the policy helps spread the unloading demand over the airport, further reducing queueing times.

5.3 Implication of results

Simulation and modelling choices

Discrete-event simulation has shown to be a suitable technique for the analysis of large-scale operational changes at complex airports. While other optimization techniques like mixed-integer linear programming (MILP), machine learning algorithms, capacity analysis and data analysis of current processes all have their own strengths and weaknesses, they are not as effective as simulation in analysing the interactions that emerge when multiple processes are changed simultaneously, and in observing the resulting outcomes.

The set of results from the simulations create an alluring perspective of future operations. It should be noted however that simulations remain just that: simulations. They are based on assumptions and mechanisms for which attempts have been made to make them as realistic as possible, but comparisons with real-life operations, like early trials, will have to be made to finally validate the findings and confirm the results.

Feasibility and best policies

Taking the above into account, the results from the simulations have shown that full-scale TaxiBot operations at Schiphol are feasible, given that the made assumptions are valid, and that any aspects placed out of scope will not lead to operational problems. As highlighted in the previous section, the engine start-up during the towing operation has a significant positive impact on capacity at the (un)loading stations, and also leads to positive results for on-time performance and CO2 reduction. Therefore, this seems to be a critical aspect to consider for full-scale operations.

Traffic congestion should not be a concern for TaxiBot operations. If the advertised maximum speed of 22 kts can be achieved reliably in daily operations, then the TaxiBot will have limited to no negative impact on air traffic flow around the airport. Emission analysis has also shown that it is worthwhile to deploy TaxiBot operations, with fuel and CO2 reductions up to 60-70% per taxiing movement.

It is difficult to select a 'best practice' towing policy for the experiments performed in this research, because this is dependent on the objective, and the number of TaxiBots available. Towing out of bays has shown to have a large impact on UFP exposure within the apron bays, while having a limited impacts on CO2 reduction and resulting in negative on-time performance changes. On the other hand, towing aircraft for large distances to the runways leads to the most CO2 reduction, but also results in a long mission time, meaning that fewer aircraft can be towed out of bays with the same number of TaxiBots. It is up to each airport to define their own objective, and design their TaxiBot policy accordingly. One strong recommendation that remains valid for every situation, is that starting engines during the tow is a great idea.

Generalizability to other airports

Most of these findings are generalizable to other airports. Relationships between taxiing time gains to taxiing distance and CO2 reductions to towing time can easily be extrapolated to other airports. Similarly, taxiing speed analysis at other airports (and critically, identifying how often aircraft exceed the TaxiBot's speed limit) can also confirm whether the TaxiBot will lead to operational issues regarding traffic flow, or if this will not be an area of concern.

The conclusions regarding congestion at the (un)loading stations are also valid for other airports, because the maximum runway capacity in terms of departures is around 40 aircraft per hour, or one every 90 seconds (Airports Council International, 2023). This is achieved by the outbound operations at the Polderbaan on the first radar day (see figure 3.10) and did not lead to problems for experiments 1-3 (see tables 4.4 and B.2). As long as an airport can make four unloading spots available at a runway and can allow engine start-up while in motion, this will not cause operational problems.

5.4 Limitations of study

Despite the efforts to make the simulation model as all-encompassing and realistic as possible, this study has several limitations.

The first limitation relates to the simulation model's modelling of ground traffic, TaxiBots and the pushback process. No ground service vehicles like refuelling and catering trucks are simulated, and no TaxiBot assignment or charging need is incorporated, leading to reduced accuracy of service road traffic and thus less accurate findings regarding the number of TaxiBots required. Finally, a lack of data of pushback trucks has also led to an incomplete data picture of the pushback route and duration.

The second limitation is that many assumptions are made about the functioning of TaxiBots. The model assumes optimal reliability and availability, and expects that all service road infrastructure is usable for the TaxiBot. Another limitation is that TaxiBot towing speeds are derived from the aircraft's original radar speeds, and limited at 22 kts. This assumes similar acceleration and deceleration profiles, something that might be different in real-life.

Third, assumptions about emissions might have impacted the results. The model assumes that when aircraft engines are running, this goes for all engines, at idle thrust. Secondly, assumptions about the engine run times as presented in section 3.4.1 might (unintentionally) differ from reality. However, since emissions are all compared to reference scenarios – in which the same assumptions were made – the comparative results might only have been impacted in a much more limited way. Finally, the assumption that UFP emissions from the taxiway A/B ring-road around the piers will not impact UFP levels inside the bays has had an impact on the conclusions regarding UFP emissions. Should this assumption be incorrect, than the viability of the 'engine start-up during tow' policy might need to be re-evaluated.

Next, the study lacks explicit sensitivity analysis. There is implicit testing of variation in the model by assigning the same TaxiBot policies to many aircraft per experiment, but no stochastic testing has been performed.

The simulation has stayed within the boundaries of the historical data, in the sense that no weather impact or unforeseen delays have been modelled, beyond what was naturally included in the radar data. Therefore, the impact that TaxiBot operations might have on delay generation and delay recovery remains unknown. This also means that the gate planning, routing and runway assignment stays as it was, even though the changed operations might have led air traffic controllers to assign different stands or routes based on congestion patterns at the airport.

Finally, the study's results are based on data from Schiphol airport, for two days. Although the findings can apply to other airports, the unique features of each airport might cause different results when using similar TaxiBot operations.

5.5 Recommendations for future work

The research and its findings have exposed several avenues for further research to gain a more complete insight into the impacts of dispatch towing using TaxiBots. The four avenues are gate planning and capacity, emissions of ultrafine particles, service road infrastructure analysis, and simulation detail.

Gate planning and capacity

The main avenue for improvement lies in gate planning and capacity analysis. As demonstrated in the previous chapter, TaxiBot operations have a serious impact on time spent in the apron bays. This reduction in time leads to added gate capacity, but it is not known to what extent this leads to gate availability. Aircraft at Schiphol sometimes have to wait until their stand becomes available (a so-called 'wachter in het veld', or 'waiter in the (air)field'), or are placed on a remote buffer and have their passengers unloaded via buses –

not an ideal experience. TaxiBot operations could potentially free up enough time at gates so that these waiting or buffered aircraft can be assigned to a conventional aircraft stand at a pier.

Additionally, aircraft sometimes have to wait at the gate for a departure slot while they are already able to depart, occupying additional gate space. Waiting further out in the airfield does is not preferred, because it would waste fuel. With TaxiBot towing this could become an option, which has the potential to create even more gate capacity at the airport.

A third aspect of possible gate planning improvements is that TaxiBot operations might reduce pushback conflicts and bay blockages. When an aircraft performs its pushback in current operations, other aircraft in the bay often have to wait, because their path is now blocked. If the engine start-up can be skipped, these conflicts and blockages might reduce or disappear, leading to even more capacity gain. A first-stage of this type of research would be exploratory; first quantifying how often this occurs and what the impact is on aircraft that have to wait at the stand, even though they might be ready for departure.

UFP emissions

For emissions, the main gap where future work could add additional insight is in quantifying ultrafine particle exposure. Until data becomes available that details how much UFP is emitted by various types of aircraft engines and APUs, it will remain difficult to assess the impacts of sustainable taxiing policies on UFP using simulation.

Once that data is known, the UFP impact can be optimized during real-life sustainable taxiing operations by prioritising flight movements that will expose ground crew working in the green zone to the most UFP.

Service road infrastructure analysis

The third aspect of future work relates to the service road network and the traffic using it, which was placed out of scope for this research and instead it was assumed that service roads would not cause any operational issues. Nevertheless, there are some limitations regarding the widths of service roads around the piers that might make TaxiBot operations difficult, and from a traffic perspective things can also get more complicated if an additional 30 to 50 TaxiBots are expected to make use of the semi-congested service road network around the airport.

Another aspect of this service road infrastructure analysis, is to evaluate if the infrastructure can be redesigned in simple but effective ways, and what the impact of that would be on capacity and on-time performance. Purpose-designed unloading locations, or wider service road lanes could provide a real improvement, for relatively small investment.

Simulation detail

Finally, some aspects of the simulation detail could be improved even further. The main area of interest here lies again in the apron bays. The pushback route and duration of each aircraft was hard to determine, because no radar data of pushback trucks was available. That made it difficult to precisely determine when each aircraft had their pushback, and where they stood stationary to start their engines. Improving this aspect of the simulation will lead to increased realism of processes and traffic interaction in the bays, resulting in more accurate taxiing times.

Secondly, model intelligence can be improved, for routing and (as mentioned earlier) gate planning. Aircraft now follow their radar tracks to their gates that were assigned on the day, but TaxiBot operations change the situation at the airport in such a way, that air traffic control might have given different – more optimal – routing and gate instructions, for example to taxi around congestion, or to instead use a gate that leads to fewer conflicts.

A final way to increase simulation realism is to build an assignment model for TaxiBot mission assignment that also includes charging needs of the vehicles. That way, a more accurate picture can be constructed of the number of required TaxiBots per policy experiment.

6 Conclusion

This thesis has investigated the potential of TaxiBot operations to enhance sustainability without adversely affecting ground traffic flow at Schiphol airport, employing a discrete-event simulation model derived from radar data for the busiest days of 2023 on the Polderbaan runway. A range of TaxiBot policies were examined across nine experiments, yielding valuable insights into the effectiveness of TaxiBot operations and identifying potential challenges.

The main research question and sub-questions are repeated below.

What is the potential impact of TaxiBot operations on the flow and emissions of taxiing aircraft at Schiphol airport?

1. *What are current patterns and bottlenecks in the traffic flow at Schiphol airport?*
2. *How do TaxiBot process times compare to those of regular operations?*
3. *How do the TaxiBot operations affect apron bay flow?*
4. *How are the unloading locations near the Polderbaan affected by the envisioned operations?*
5. *What are the impacts of TaxiBot operations on emissions at the airport?*

Sub-question 1 relating to current patterns and bottlenecks was investigated using radar speed plots of the airport. Speeds on and around the apron were significantly lower than those farther out on the field, with the only bottlenecks in flow identified near runway entry points, where aircraft awaited take-off clearance.

Analysis of sub-question 2 about taxiing times found that outbound times can be higher or lower than the taxiing times in the reference scenario, depending on engine start-up policy: start-up during the tow can save up to 1.5 minutes of taxiing time. Inbound times were found to be 2.50 to 3.25 minutes longer than in reference scenarios.

Apron bay flow was investigated for sub-question 3, and highlighted significant time gains which have the potential to create additional gate capacity at the airport. More research is needed to quantify these gains as a result of the time saved.

The unloading locations are the subject of sub-question 4, and once again the engine start-up policy turns out to be critical. Manageable queueing occurs if engines must start up while stationary, with this congestion nearly vanishing if engines can be started during the tow. Additionally, spreading the unloading demand over multiple locations at the airport (multiple runways combined with towing to edges of the green zone) further reduces station load. Analysis of taxiing times for aircraft not involved in TaxiBot operations revealed no significant deviations from the reference scenario.

Finally, emissions were analysed in sub-question 5. CO2 emissions decreased by up to 76% per flight and up to 32.9% for the airport overall. The results indicate that the primary driver of CO2 reduction is the decrease in taxiing time, recommending the servicing of the furthest runways as optimal for CO2 reduction. UFP analysis is limited by the lack of quantitative emission data, yet emission heat maps indicate that significant relocation of ultrafine particle emissions is possible.

Addressing the main research question, it can be concluded that by implementing TaxiBot operations, Schiphol airport can achieve significant emission reduction with minimal negative, or even with positive impacts on flow and traffic performance at the airport.

The most impactful finding that supports this conclusion is that the policy of allowing aircraft to start their engines during towing contributes greatly to the flow. This results in a two-minute improvement in outbound taxiing time, significantly reduces queuing and congestion around the unloading stations, and also has a minor additional impact on CO2 reduction. Another key finding is that TaxiBot operations have a negligible impact on the flow of conventional traffic around the towed aircraft, attributed to the TaxiBot's ability to sustain a maximum towing speed of 22 knots.

Analysis regarding the policy choice between towing aircraft to the edge of the green zone or towing them all the way to the runway has shown that runway towing is much more impactful for on-time performance and CO2 reduction, and that a combination of the two policies – towing to runways where possible, and otherwise to the edge of the green zone – performs better than just green zone towing, because unloading demand is now spread over multiple unloading locations at the airport which reduces congestion around the stations. Thus, should green zone towing be selected for UFP priority reasons, a scaled-up operation incorporating runway towing is more likely to succeed than towing solely to the edge of the green zone.

A high degree of generalizability towards other (complex) airports is expected for these findings. Correlations between taxiing time savings and taxiing distances, and between CO2 reductions and towing times, can be used to predict outcomes for other airports. Additionally, the runway capacity for inbound and outbound movements is independent of specific airport characteristics; thus, if an airport can provide four unloading stations next to a runway and incorporate engine start-up during towing into their TaxiBot operations, similar queuing and unloading capacity characteristics can be anticipated.

Future research is particularly recommended in the areas of gate planning and airport capacity. TaxiBot-towed aircraft spend three minutes less in the apron bays compared to normal operations; however, the simulation model is not designed to translate these gains into optimized gate planning or to determine the extent of the increase in airport capacity.

Although there are some limitations to the research setup, primarily concerning the sub-optimal analysis of service road traffic and TaxiBot assignment, as well as the lack of adaptability in gate planning and taxi routing under the new operations, the implications of the findings are noteworthy. The simulation has demonstrated that operationally, full-scale TaxiBot deployment at large and complex airports is feasible. The success of implementation is ensured by allowing engines to be started during towing, and by confirming that the TaxiBots top speed of 22 kts is both achievable in daily operations and that average taxiing speeds at the airport do not frequently exceed these 22 kts.

This indicates that TaxiBot operations have the potential to be a significant driver for airport sustainability. By decreasing fuel consumption and CO2 emissions by up to 76%, and reducing taxiing times by up to 1.5 minutes per flight, airports equipped with TaxiBots can substantially impact the future of the air transport sector.

References

Ahmadi, S. (2019). *Green airport operations: Conflict and collision free taxiing using electric powered towing alternatives* (Unpublished doctoral dissertation). Concordia University.

Airport Technology. (2015, February 19). *Frankfurt airport deploys fuel-saving taxibot for taxiing aircraft.* Retrieved from `https://www.airport-technology.com/news/newsfrankfurt-airport-deploys-fuel-saving-TaxiBot-for-taxiing-aircraft-4517031/`

Airports Council International. (2021, May 7). *Delhi airport registers 1,000 taxibot movements.* Retrieved from `https://www.aci-asiapac.aero/media-centre/news/delhi-airport-registers-1-000-TaxiBot-movements`

Airports Council International. (2023). *Guidance on airport capacity declarations.* Retrieved from `https://www.aci-europe.org/downloads/publications/ACI%20Guidance%20on%20Airport%20Capacity%20Declarations.pdf`

Antunes, B. B., Manresa, A., Bastos, L. S., Marchesi, J. F., & Hamacher, S. (2019). A solution framework based on process mining, optimization, and discrete-event simulation to improve queue performance in an emergency department. In *Business process management workshops: Bpm 2019 international workshops, vienna, austria, september 1–6, 2019, revised selected papers 17* (pp. 583–594).

Aéroports de Montréal. (2024, February 3). *2023 aéroports de montréal passenger statistics.* Retrieved from `https://www.admtl.com/sites/default/files/2024/ADM_Statsdet_2023EN.pdf`

Benda, B. (2020). Agent-based modelling and analysis of non-autonomous airport ground surface operations. *Master of Science Thesis*.

Brussels Airport. (2024, April 3). *Brussels airport and tui fly are testing the taxibot for sustainable taxiing.* Retrieved from `https://www.brusselsairport.be/en/pressroom/news/brussels-airport-and-tui-test-the-TaxiBot`

Cao, F., Tang, T.-Q., Gao, Y., You, F., & Zhang, J. (2023). Calculation and analysis of new taxiing methods on aircraft fuel consumption and pollutant emissions. *Energy*, *277*, 127618.

Centraal Bureau voor de Statistiek. (2023, December 20). *Rendementen, co2-emissie elektriciteitsproductie, 2022.* Retrieved from `https://www.cbs.nl/nl-nl/maatwerk/2023/51/rendementen-co2-emissie-elektriciteitsproductie-2022`

Clemens, P. (2023). To tow or not to tow: A social cost-benefit analysis of engine-off dispatch aircraft towing at amsterdam airport schiphol. *Master of Science Thesis*.

Corey, P. D., & Clymer, J. R. (1991). Discrete event simulation of object movement and interactions. *Simulation*, *56*(3), 167–174.

Di Mascio, P., Corazza, M. V., Rosa, N. R., & Moretti, L. (2022). Optimization of aircraft taxiing strategies to reduce the impacts of landing and take-off cycle at airports. *Sustainability*, *14*(15), 9692.

EASA. (n.d.). *Fit for 55 and ReFuelEU Aviation.* Retrieved 2024-05-24, from `https://www.easa.europa.eu/en/light/topics/fit-55-and-refueleu-aviation`

Groot, M., & Roling, P. C. (2022). The potential impact of electric aircraft taxiing: A probabilistic analysis and fleet assignment optimization. In *Aiaa aviation 2022 forum* (p. 3919).

Ground Handling International. (2022, February 18). *Schiphol becomes first european airport to invest in taxibots.* Retrieved from `https://www.groundhandlinginternational.com/content/news/schiphol-becomes-first-european-airport-to-invest-in-TaxiBots`

Groupe ADP. (2024, May). *Groupe adp innovation hub.* Retrieved from `https://presse.groupeadp.fr/wp-content/uploads/2024/05/cf3e0d24fde11b588004c9aaf37885ca.pdf`

Gualandi, E. (2014). Taxibot: Analisi dei benefici derivanti dall'introduzione di un veicolo semi-robotico per dispatch towing (taxibot: Analysis of the benefits of introducing a semi-robotic vehicle for dispatch towing). *Master of Science Thesis*.

Hein, K., & Baumann, S. (2016). Acoustical comparison of conventional taxiing and dispatch towing-taxibot's contribution to ground noise abatement. In *30th congress of the international council of the aeronautical sciences (icas)* (pp. 1–7).

Hospodka, J. (2014). Electric taxiing–taxibot system. *MAD-Magazine of Aviation Development*, *2*(10), 17–20.

IATA. (2022, April 19). *Iata carbon offset program: Frequently asked questions.* Retrieved from `https://www.iata.org/contentassets/922ebc4cbcd24c4d9fd55933e7070947/icop_faq_general-for-airline-participants.pdf`

IATA. (2023, May). *Net zero 2050: sustainable aviation fuels.* Retrieved from `https://www.iata.org/en/iata-repository/pressroom/fact-sheets/fact-sheet---alternative-fuels/`

Israel Aerospace Industries. (n.d.). *Taxibot certifications.* Retrieved from `https://TaxiBot-international.com/certifications/`

Jadrić, M., Pašalić, I. N., & Ćukušić, M. (2020). Process mining contributions to discreteevent simulation modelling. *Business Systems Research: International journal of the Society for Advancing Innovation and Research in Economy*, *11*(2), 51–72.

Jen, H.-C., Huff, B. L., LeBoulluec, A. K., Nasirian, B., Bum Kim, S., Rosenberger, J. M., & Chen, V. C. (2022). A discrete-event simulation tool for airport deicing activities: Dallas-fort worth international airport. *Simulation*, *98*(12), 1097–1114.

Kariya, Y., Mase, T., Yoshihara, S., & Ota, J. (2011). Analysis of congestion of taxiing aircraft at a large airport. In *2011 ieee international conference on robotics and biomimetics* (pp. 180–185).

Kariya, Y., Yahagi, H., Takehisa, M., Yoshihara, S., Ogata, T., Hara, T., & Ota, J. (2013). Modeling and designing aircraft taxiing patterns for a large airport. *Advanced Robotics*, *27*(14), 1059–1072.

Khadilkar, H., & Balakrishnan, H. (2014). Network congestion control of airport surface operations. *Journal of Guidance, Control, and Dynamics*, *37*(3), 933–940.

Khammash, L., Mantecchini, L., & Reis, V. (2017). Micro-simulation of airport taxiing procedures to improve operation sustainability: Application of semi-robotic towing tractor. In *2017 5th ieee international conference on models and technologies for intelligent transportation systems (mt-its)* (pp. 616–621).

Lai, J., Che, L., & Kashef, R. (2021). Bottleneck analysis in jfk using discrete event simulation: An airport queuing model. In *2021 ieee international smart cities conference (isc2)* (pp. 1–7).

Liu, S. T. (2015). *Integrating process mining with discrete-event simulation modeling*. Brigham Young University.

Lukic, M., Giangrande, P., Hebala, A., Nuzzo, S., & Galea, M. (2019). Review, challenges, and future developments of electric taxiing systems. *IEEE Transactions on Transportation Electrification*, *5*(4), 1441–1457.

Lukic, M., Hebala, A., Giangrande, P., Klumpner, C., Nuzzo, S., Chen, G., ... Galea, M. (2018). State of the art of electric taxiing systems. In *2018 ieee international conference on electrical systems for aircraft, railway, ship propulsion and road vehicles & international transportation electrification conference (esars-itec)* (pp. 1–6).

Mori, R. (2012). Aircraft ground-taxiing model for congested airport using cellular automata. *IEEE Transactions on Intelligent Transportation Systems*, *14*(1), 180–188.

Netherlands Labour Authority. (2023a, March 14). *Op schiphol moeten diesel aangedreven arbeidsmiddelen worden vervangen.* Retrieved from `https://www.nlarbeidsinspectie.nl/`

`actueel/nieuws/2023/03/14/op-schiphol-moeten-diesel-aangedreven-arbeidsmiddelen-worden-vervangen`

Netherlands Labour Authority. (2023b, May). *Ultrafine particles in the vicinity of schiphol airport affect health.* Retrieved from `https://zoek.officielebekendmakingen.nl/stcrt-2023-14547.html`

Ouerghi, N. (2008). *Modeling airport ground operations using discrete event simulation (des) and x3d visualization* (Unpublished doctoral dissertation). Monterey, California. Naval Postgraduate School.

Postorino, M. N., Mantecchini, L., & Gualandi, E. (2016). Integration between aircraft and handling vehicles during taxiing procedures to improve airport sustainability. *International Journal of Transport Development and Integration*, *1*(1), 28–42.

Rashid, K. M., & Louis, J. (2022). Integrating process mining with discrete-event simulation for dynamic productivity estimation in heavy civil construction operations. *Algorithms*, *15*(5), 173.

Roling, P. C., Sillekens, P., Curran, R., & Wilder, W. D. (2015). The effects of electric taxi systems on airport surface congestion. In *15th aiaa aviation technology, integration, and operations conference* (p. 2592).

Royal Schiphol Group. (n.d.-a). *Aircraft process maps.* Retrieved from `https://www.schiphol.nl/en/operations/page/maps/`

Royal Schiphol Group. (n.d.-b). *Aircraft stand allocation and properties [download: Overview aircraft types].* Retrieved 2024-05-29, from `https://www.schiphol.nl/en/operations/page/allocation-aircraft-stands/`

Royal Schiphol Group. (n.d.-c). *New research lab for ultrafine particles on apron.* Retrieved 2024-05-24, from `https://www.schiphol.nl/en/innovation/blog/new-research-lab-for-ultrafine-particles-on-the-apron/`

Royal Schiphol Group. (n.d.-d). *What is sustainable taxiing? (part 1).* Retrieved 2024-05-24, from `https://www.schiphol.nl/en/innovation/blog/what-is-sustainable-taxiing-part-1`

Royal Schiphol Group. (2021). *Sustainable taxiing: Taxibot trial.* Retrieved from `https://www.schiphol.nl/en/innovation/page/sustainable-taxiing-TaxiBot-trial/`

Royal Schiphol Group. (2024a, May 3). *Europe's best connected airport.* Retrieved from `https://www.schiphol.nl/en/blog/europes-best-connected-airport/`

Royal Schiphol Group. (2024b, February 29). *It takes two taxibots to tango!* Retrieved from `https://www.schiphol.nl/en/innovation/blog/sustainable-taxiing-it-takes-two-TaxiBots-to-tango/`

Royal Schiphol Group. (2024c, January 5). *Schiphol welcomed 61.7 million travellers in 2023.* Retrieved from `https://news.schiphol.com/schiphol-welcomed-617-million-travellers-in-2023/`

RZjets. (n.d.). *Schiphol layout.* Retrieved from `https://rzjets.net/airports/?code=AMS#_1`

Salihu, A. L. (2020). *Impact of on-ground taxiing with electric powered tow-trucks on congestion, cost, and carbon emissions at montreal-trudeau international airport* (Unpublished doctoral dissertation). Concordia University.

Salihu, A. L., Lloyd, S. M., & Akgunduz, A. (2021). Electrification of airport taxiway operations: A simulation framework for analyzing congestion and cost. *Transportation Research Part D: Transport and Environment*, *97*, 102962.

Simio. (n.d.). *Simio simulation software.* Retrieved from `https://www.simio.com/software/simulation-software.php`

Sirigu, G., Cassaro, M., Battipede, M., & Gili, P. (2018). Autonomous taxi operations: algorithms for the solution of the routing problem. In *2018 aiaa information systems-aiaa infotech@ aerospace* (p. 2143).

Soepnel, S. (2015). Impact of electric taxi systems on airport apron operations and gate congestion at aas. *Master of Science Thesis*.

Soltani, M., Ahmadi, S., Akgunduz, A., & Bhuiyan, N. (2020). An eco-friendly aircraft taxiing approach with collision and conflict avoidance. *Transportation Research Part C: Emerging Technologies*, *121*, 102872.

TaxiBot-India. (2024). *Taxibot india update.* Retrieved 2024-05-24, from `http://TaxiBot-india.com/services/`

The Hindu Bureau. (2023, April 13). *https://www.thehindu.com/news/cities/bangalore/air-india-to-launch-taxibot-operations-at-bengaluru-and-delhi-airports-for-a320s/article66732807.ece.* Retrieved from `https://www.thehindu.com/news/cities/bangalore/air-india-to-launch-TaxiBot-operations-at-bengaluru-and-delhi-airports-for-a320s/article66732807.ece`

Tindemans, B. (2021). A greedy approach to the minimisation of deviations of the dynamic vehicle routing problem with electric taxiing systems. *Master of Science Thesis*.

Vaishnav, P. (2014). Costs and benefits of reducing fuel burn and emissions from taxiing aircraft: Low-hanging fruit? *Transportation Research Record*, *2400*(1), 65–77.

van Baaren, E., & Roling, P. C. (2019). Design of a zero emission aircraft towing system. In *Aiaa aviation 2019 forum* (p. 2932).

Van Der Aalst, W., Adriansyah, A., De Medeiros, A. K. A., Arcieri, F., Baier, T., Blickle, T., . . . others (2012). Process mining manifesto. In *Business process management workshops: Bpm 2011 international workshops, clermont-ferrand, france, august 29, 2011, revised selected papers, part i 9* (pp. 169–194).

van Oosterom, S., Mitici, M., & Hoekstra, J. (2023). Dispatching a fleet of electric towing vehicles for aircraft taxiing with conflict avoidance and efficient battery charging. *Transportation Research Part C: Emerging Technologies*, *147*, 103995.

van Winkel, C. (2023). Tactical taxibot planning at amsterdam airport schiphol under uncertainty. *Master of Science Thesis*.

VINCI Airports. (2024, January 16). *Vinci airports – traffic at 31 december 2023.* Retrieved from `https://www.vinci.com/commun/communiques.nsf/6D529D6EECC7BC2AC1258AA6003CE67E/$file/vinci-airports--traffic-31-december-2023.pdf`

Watts, A (Airlines.net). (2013, December 11). *Boeing 737-823 - american airlines.* Retrieved from `https://www.airliners.net/photo/American-Airlines/Boeing-737-823/2359182/L`

Wollenheit, R., & Mühlhausen, T. (2013). Operational and environmental assessment of electric taxi based on fast-time simulation. *Transportation research record*, *2336*(1), 36–42.

Zbakh, D., Benhadou, M., Benkacem, A., & Lyhyaoui, A. (2023, September). Airport traffic optimisation and airdrome analysis using mathematical modelling. *Indonesian Journal of Electrical Engineering and Computer Science*, *31*(3), 1744.

Zhang, M., Huang, Q., Liu, S., & Li, H. (2019a). Assessment method of fuel consumption and emissions of aircraft during taxiing on airport surface under given meteorological conditions. *Sustainability*, *11*(21), 6110.

Zhang, M., Huang, Q., Liu, S., & Li, H. (2019b). Multi-objective optimization of aircraft taxiing on the airport surface with consideration to taxiing conflicts and the airport environment. *Sustainability*, *11*(23), 6728.

Zoutendijk, M., Mitici, M., & Hoekstra, J. (2023). An investigation of operational management solutions and challenges for electric taxiing of aircraft. *Research in Transportation Business & Management*, *49*, 101019.

A Literature themes

Table A.1: Literature topics, methods and associated themes

Paper	Topic	Method	Schiphol	Dispatch tow.	Cong. & ST	Methods
Ahmadi, 2019	Taxi control optimization	MILP		×		
Antunes et al., 2019	Process mining & DES theory	PM & DES				×
Benda, 2020	TaxiBot towing	Agent-based	×	×		×
Cao et al., 2023	Emissions of new taxi methods	Emission calc.		×	×	
Clemens, 2023	Dispatch towing at AAS	CBA	×	×		
Corey & Clymer, 1991	DES theory	DES				×
Di Mascio et al., 2022	Reducing emissions	Emission calc.		×		
Groot & Roling, 2022	On-board motors for taxiing	MILP	×		×	
Gualandi, 2014	Emissions of TaxiBot operations	Emission calc.		×		
Hein & Baumann, 2016	Noise of TaxiBot operations	Noise analysis		×		
Hospodka, 2014	TaxiBot as a new innovation	Concept		×		
Jadrić et al., 2020	PM & DES theory	PM & DES				×
Jen et al., 2022	De-icing DES	DES				×
Kariya et al., 2011	Simulating taxiing	DES			×	×
Kariya et al., 2013	Congestion of taxiing aircraft	DES			×	×
Khadilkar & Balakrishnan, 2014	Congestion during taxiing	Numerical sim			×	
Khammash et al., 2017	Simulation of ground movements	DES		×		×
Lai et al., 2021	Taxiing simulation	DES			×	×
Liu, 2015	PM & DES theory	PM & DES				×
Lukic et al., 2018	Electric taxiing systems	Literature rev.		×	×	
Lukic et al., 2019	Electric taxiing systems	Literature rev.		×	×	
Mori, 2012	Ground traffic modelling	DES			×	×

Paper	**Topic**	**Method**	**Schiphol**	**Dispatch tow.**	**Cong. & ST**	**Methods**
Ouerghi, 2008	Simulating ground operations	DES				×
Postorino et al., 2016	Emissions modelling	Emission calc.		×		
Rashid & Louis, 2022	PM & DES theory	PM & DES				×
Roling et al., 2015	Electric taxiing at Schiphol	Numerical sim	×		×	
Salihu, 2020	Congestion during dispatch towing	DES		×		×
Salihu et al., 2021	Taxiway congestion	DES		×		×
Sirigu et al., 2018	Taxiing route selection	Algorithm		×		
Soepnel, 2015	Gate planning at Schiphol	VBA & excel	×		×	
Soltani et al., 2020	Taxiing route conflicts	MILP		×		
Tindemans, 2021	TaxiBot traffic at Schiphol	MILP GVRSP	×	×		
Vaishnav, 2014	Emissions modelling	Emission calc.			×	
van Baaren & Roling, 2019	Feasibility of dispatch towing	MILP	×	×		
van Oosterom et al., 2023	Dispatchment of towing trucks	MILP	×	×		
van Winkel, 2023	TaxiBot allocation planning	MILP EVSP	×	×		
Wollenheit & Mühlhausen, 2013	Electric taxiing simulation	DES			×	×
Zbakh et al., 2023	Ground traffic modelling	DES			×	×
Zhang et al., 2019a	Emissions of taxi methods	Emission calc.			×	
Zhang et al., 2019b	Taxiing optimization	MILP			×	×
Zoutendijk et al., 2023	Sustainable taxiing methods	Literature rev.		×		×

B Supplementary results

B.1 Statistics per (un)loading location

Table B.1: (Un)loading station statistics for experiment 1

Unloading location	**TaxiBot station**	**Occupation rate**	**Seize count**	**Avg mins per seize**	**Mins all occupied**
Polderbaan	6A	9.01%	37	3.51	70.48
Polderbaan	6B	16.54%	69	3.45	70.48
Polderbaan	7A	31.10%	105	4.27	70.48
Polderbaan	7B	42.88%	140	4.41	70.48

Table B.2: (Un)loading station statistics for experiment 2

Unloading location	**TaxiBot station**	**Occupation rate**	**Seize count**	**Avg mins per seize**	**Mins all occupied**
Polderbaan	6A	1.52%	11	1.99	3.26
Polderbaan	6B	5.86%	50	1.69	3.26
Polderbaan	7A	15.57%	100	2.24	3.26
Polderbaan	7B	29.96%	190	2.27	3.26

Table B.3: (Un)loading station statistics for experiment 3

Unloading location	**TaxiBot station**	**Occupation rate**	**Seize count**	**Avg mins per seize**	**Mins all occupied**
Polderbaan	6A	14.16%	79	2.58	59.56
Polderbaan	6B	16.99%	105	2.33	59.56
Polderbaan	7A	18.81%	119	2.28	59.56
Polderbaan	7B	32.55%	202	2.32	59.56

Table B.4: (Un)loading station statistics for experiment 4

Unloading location	**TaxiBot station**	**Occupation rate**	**Seize count**	**Avg mins per seize**	**Mins all occupied**
Polderbaan	6A	16.30%	77	3.05	39.07
Polderbaan	6B	19.80%	102	2.80	39.07
Polderbaan	7A	17.98%	120	2.16	39.07
Polderbaan	7B	36.33%	236	2.22	39.07

Table B.5: (Un)loading station statistics for experiment 5

Unloading location	**TaxiBot station**	**Occupation rate**	**Seize count**	**Avg mins per seize**	**Mins all occupied**
Polderbaan	6A	16.65%	78	3.07	37.92
Polderbaan	6B	19.76%	100	2.85	37.92
Polderbaan	7A	17.71%	118	2.16	37.92
Polderbaan	7B	36.20%	239	2.18	37.92
Zwanenburgbaan North	29A	4.66%	30	2.24	0.36
Zwanenburgbaan North	29B	7.43%	49	2.18	0.36

Zwanenburgbaan North	30A	0.12%	1	1.69	0.36
Zwanenburgbaan North	30B	0.97%	6	2.34	0.36
Zwanenburgbaan South	4A	14.26%	104	1.97	4.86
Zwanenburgbaan South	4B	7.77%	54	2.07	4.86
Zwanenburgbaan South	5A	5.39%	28	2.77	4.86
Zwanenburgbaan South	5B	3.14%	16	2.82	4.86

Table B.6: (Un)loading station statistics for experiment 6

Unloading location	**TaxiBot station**	**Occupation rate**	**Seize count**	**Avg mins per seize**	**Mins all occupied**
J-Platform	10A	22.62%	73	4.46	20.98
J-Platform	10B	13.76%	41	4.83	20.98
J-Platform	12A	7.41%	19	5.61	20.98
J-Platform	12B	4.53%	10	6.52	20.98
P-Platform	1A	19.74%	68	4.18	105.10
P-Platform	1B	31.84%	118	3.89	105.10
P-Platform	3A	16.41%	56	4.22	105.10
P-Platform	3B	12.86%	41	4.52	105.10
R-Platform	20A	4.87%	20	3.51	20.06
R-Platform	20B	10.50%	42	3.60	20.06
R-Platform	21A	21.25%	96	3.19	20.06
R-Platform	21B	34.78%	161	3.11	20.06

Table B.7: (Un)loading station statistics for experiment 7

Unloading location	**TaxiBot station**	**Occupation rate**	**Seize count**	**Avg mins per seize**	**Mins all occupied**
J-Platform	10A	11.30%	82	1.98	3.61
J-Platform	10B	6.00%	38	2.27	3.61
J-Platform	12A	2.61%	15	2.51	3.61
J-Platform	12B	1.59%	8	2.86	3.61
P-Platform	1A	13.55%	75	2.60	32.76
P-Platform	1B	21.04%	141	2.15	32.76
P-Platform	3A	7.20%	43	2.41	32.76
P-Platform	3B	5.18%	24	3.11	32.76
R-Platform	20A	0.26%	2	1.90	0.43
R-Platform	20B	1.64%	18	1.31	0.43
R-Platform	21A	6.60%	73	1.30	0.43
R-Platform	21B	20.15%	226	1.28	0.43

Table B.8: (Un)loading station statistics for experiment 8

Unloading location	**TaxiBot station**	**Occupation rate**	**Seize count**	**Avg mins per seize**	**Mins all occupied**
Polderbaan	6A	6.07%	33	2.65	20.58
Polderbaan	6B	8.06%	51	2.28	20.58
Polderbaan	7A	9.48%	65	2.10	20.58
Polderbaan	7B	19.02%	123	2.23	20.58
Zwanenburgbaan South	4A	5.95%	48	1.79	2.35
Zwanenburgbaan South	4B	2.68%	21	1.84	2.35
Zwanenburgbaan South	5A	3.45%	17	2.92	2.35
Zwanenburgbaan South	5B	1.87%	9	3.00	2.35
J-Platform	10A	5.26%	34	2.23	1.47
J-Platform	10B	3.20%	15	3.07	1.47
J-Platform	12A	0.67%	4	2.42	1.47
J-Platform	12B	0.42%	2	3.03	1.47
P-Platform	1A	8.77%	51	2.48	24.74
P-Platform	1B	13.92%	94	2.13	24.74
P-Platform	3A	4.77%	31	2.22	24.74
P-Platform	3B	3.48%	17	2.95	24.74
R-Platform	20A	0.00%	0	0.00	0.00
R-Platform	20B	0.34%	3	1.65	0.00
R-Platform	21A	2.92%	28	1.50	0.00
R-Platform	21B	9.59%	99	1.40	0.00

Table B.9: (Un)loading station statistics for experiment 9

Unloading location	**TaxiBot station**	**Occupation rate**	**Seize count**	**Avg mins per seize**	**Mins all occupied**
Polderbaan	6B	19.82%	99	2.88	40.68
Polderbaan	6A	17.01%	78	3.14	40.68
Polderbaan	7A	18.17%	120	2.18	40.68
Polderbaan	7B	36.41%	238	2.20	40.68
Zwanenburgbaan North	30B	1.04%	6	2.51	0.00
Zwanenburgbaan North	30A	0.00%	0	0.00	0.00
Zwanenburgbaan North	29B	7.50%	50	2.16	0.00
Zwanenburgbaan North	29A	4.57%	30	2.19	0.00
Zwanenburgbaan South	5B	2.34%	12	2.81	4.75
Zwanenburgbaan South	5A	5.22%	30	2.50	4.75
Zwanenburgbaan South	4B	8.64%	56	2.22	4.75
Zwanenburgbaan South	4A	14.69%	104	2.03	4.75
J-Platform	10A	5.30%	34	2.25	1.47
J-Platform	10B	3.20%	15	3.07	1.47
J-Platform	12A	0.65%	4	2.35	1.47
J-Platform	12B	0.43%	2	3.09	1.47
P-Platform	3B	3.61%	17	3.05	24.59
P-Platform	3A	4.96%	28	2.55	24.59
P-Platform	1A	8.60%	52	2.38	24.59
P-Platform	1B	13.82%	96	2.07	24.59
R-Platform	20A	0.00%	0	0.00	0.00
R-Platform	20B	0.45%	4	1.63	0.00
R-Platform	21A	2.98%	29	1.48	0.00
R-Platform	21B	9.60%	97	1.42	0.00

B.2 UFP heat maps

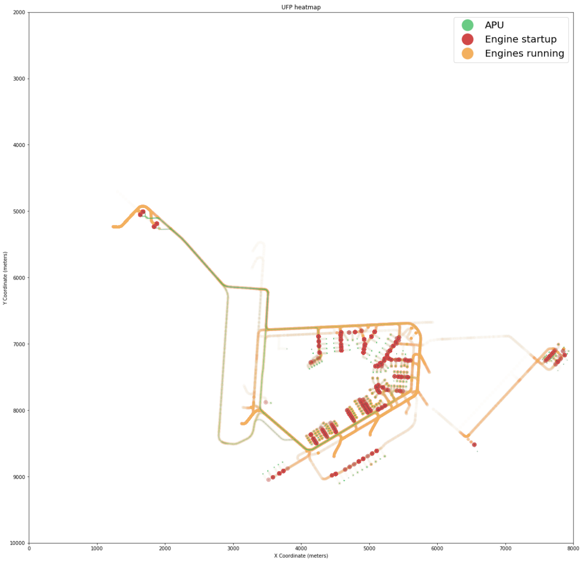

Figure B.1: UFP Heat map - Experiment 1

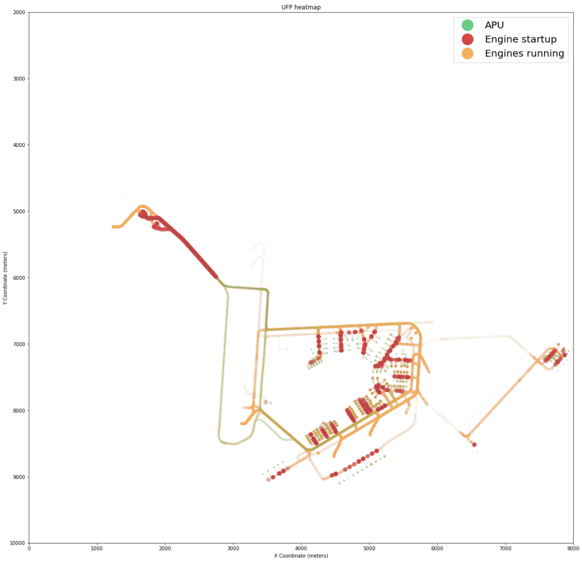

Figure B.2: UFP Heat map - Experiment 2

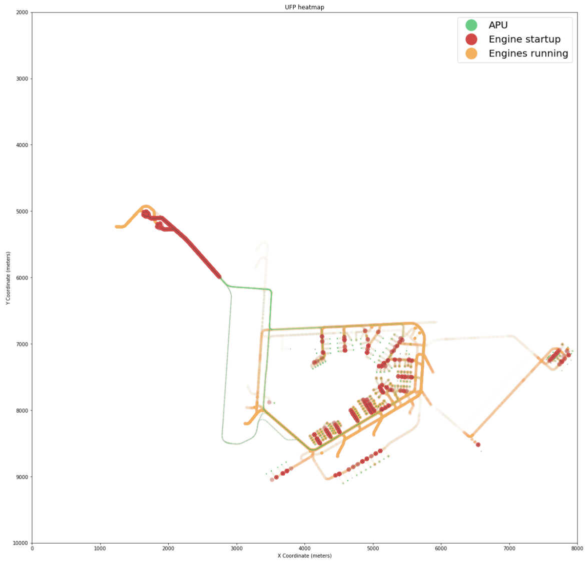

Figure B.3: UFP Heat map - Experiment 3

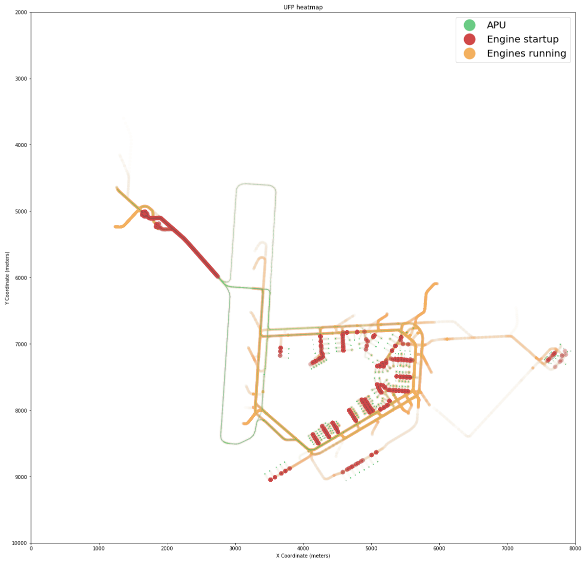

Figure B.4: UFP Heat map - Experiment 4

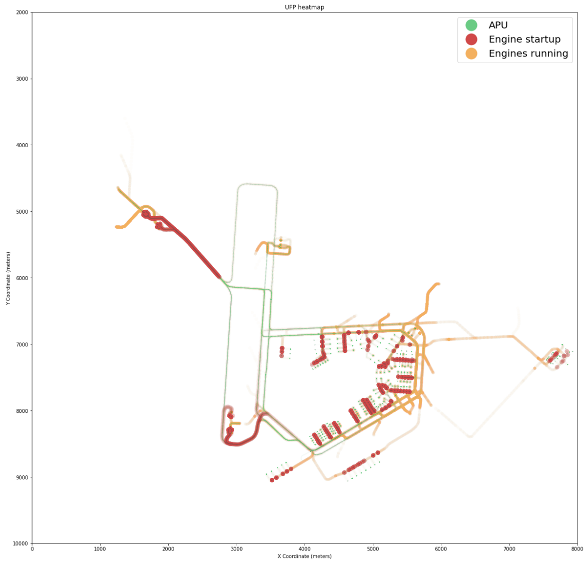

Figure B.5: UFP Heat map - Experiment 5

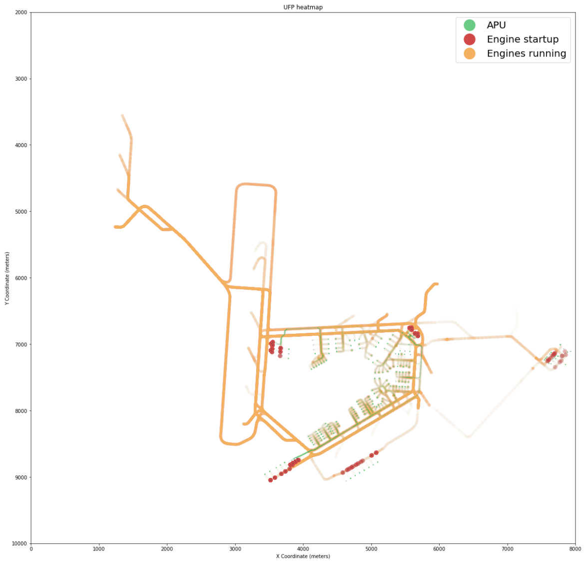

Figure B.6: UFP Heat map - Experiment 6

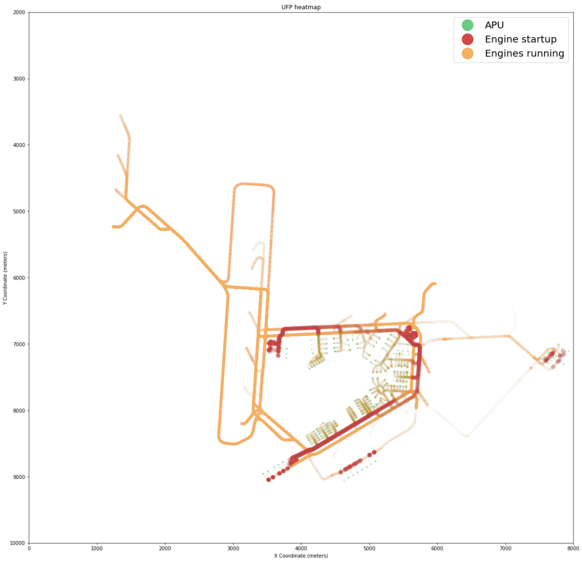

Figure B.7: UFP Heat map - Experiment 7

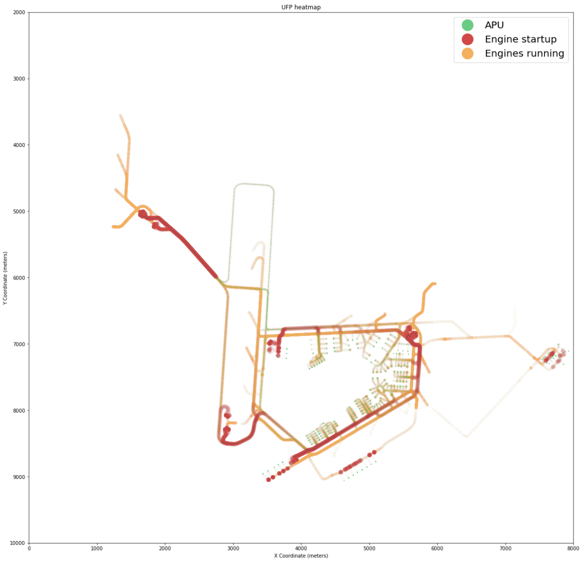

Figure B.8: UFP Heat map - Experiment 8

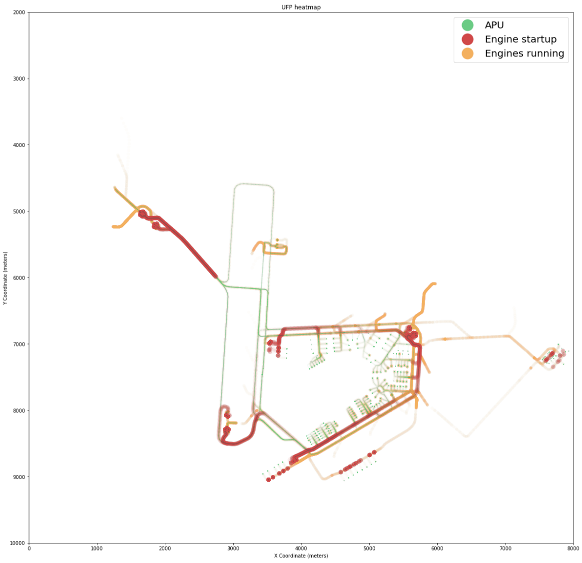

Figure B.9: UFP Heat map - Experiment 9

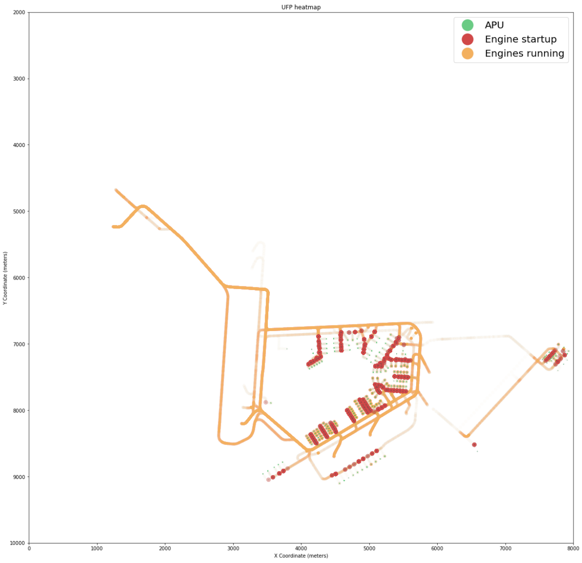

Figure B.10: Reference scenario 14 June

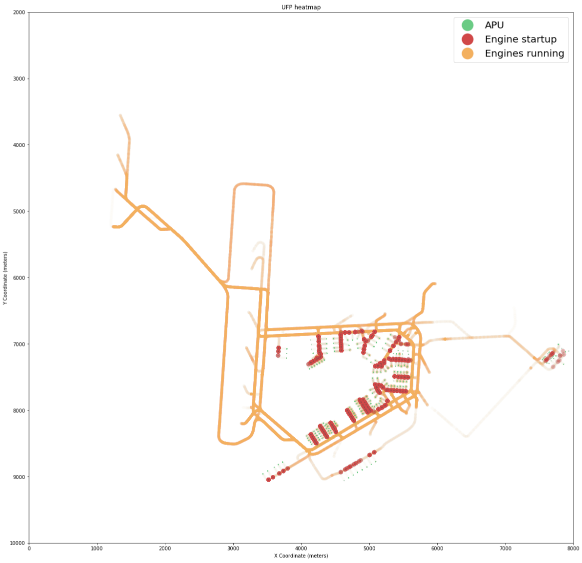

Figure B.11: Reference scenario 29 September

C Scientific paper

The Impact of TaxiBot Operations on Ground Traffic Flow at Amsterdam Airport Schiphol

R. M. N. (Ruben) Beumer (4529804)

June 19, 2024
Delft University of Technology *(Jaffalaan 5, Delft, The Netherlands)*

Keywords: Sustainable taxiing, TaxiBots, discrete event simulation, radar data, Schiphol airport

Abstract

Sustainable taxiing using TaxiBot vehicles is expected to contribute greatly to ambitions to reduce emissions and improve air quality in apron bays for airports around the world. However, the impact of large-scale dispatch towing operations at complex airports has never been investigated. This study creates a realistic simulation of ground traffic at Schiphol airport based on real-life radar data for different towing policies, to assess the impact on flow and on-time performance, congestion, and emissions.

Results show that TaxiBot towing can decrease fuel consumption and associated CO2 emissions at Schiphol airport by up to 76% per towing movement, and up to 32.9% of total airport emissions. Experiments in which aircraft are allowed to start their engines during the towing movement to save time at the unloading stations, result in outbound taxiing times that are 1 to 1.5 minutes faster than the reference scenario. Inbound times were found to become 2.50 to 3.25 minutes longer, but still contribute to CO2 savings. Apron bay flow was investigated and highlighted significant time gains of three minutes per outbound aircraft. This has the potential to create additional gate capacity at the airport, but further research is needed to quantify these gains. Congestion at unloading stations is found to be minimal, especially in experiments where engines can be started during the tow. Ultrafine-particle heat maps of towing policies show that significant relocation of emissions is possible, improving air quality in bays.

1. Introduction

This paper explores sustainable taxiing at Schiphol airport, particularly the use of TaxiBots and their impact on traffic flow and emissions. Schiphol, located southwest of Amsterdam, is the largest airport in the Netherlands and one of Europe's busiest. In 2023, it managed 442,000 flight movements and served over 61 million passengers (Royal Schiphol Group, 2024c), making it a critical hub for airlines such as AirFrance-KLM and Delta Air Lines. The airport's infrastructure, featuring six runways and a single terminal with eight piers, results in complex taxiing traffic flows.

A recent ruling by the Netherlands Labour Authority (2023a) has said that the emission of exhaust gases has to be reduced significantly, within an area around the piers demarcated as the 'green zone'. This zone was established because this is where the majority of Schiphol's workers are active, and where the exposure to emissions, especially ultrafine particulate matter

(UFP, $PM_{0.1}$) is greatest. Sustainable taxiing using TaxiBots – powerful tow trucks that can tow aircraft at their normal taxiing speeds – provides a way to move these aircraft out of the green zone without them needing to start their engines.

The aim of this study is to investigate the impact of sustainable taxiing operations using TaxiBots on ground traffic flow a large and complex airport like Schiphol, by means of developing a Discrete Event Simulation (DES) model to simulate the airside ground traffic flow based on real-life radar movements. Some earlier research on this topic at smaller airports has been done by Khammash et al. (2017) and Salihu (2020), but it is still unknown what the impact of TaxiBot operations will be on congested, complex airports. This is relevant to know, because sustainable taxiing at large airports can lead to major sustainability gains, but the complexity of these airports also creates uncertainty regarding the impact that sustainable taxiing will have on the efficiency and on-time performance of daily operations.

The main research question is to investigate the potential impact of TaxiBot operations on the flow and emissions of taxiing aircraft at Schiphol airport. Sub-questions are dedicated to current patterns and bottlenecks in the traffic flow at Schiphol airport, how TaxiBot taxiing process times compare to those of regular operations, how TaxiBot operations affect flow in apron bays, how the unloading locations near runways are affected, and what the impact on emissions of both CO2 and UFP is.

This thesis is focused on the taxiing process at Schiphol airport, based on real-life movements observed via the Schiphol ground radar. These movements serve as a foundation of the simulation upon which TaxiBot operations are modelled, to evaluate their impact on traffic flow as realistically as possible. Each simulated experiment will encompass one day of traffic, based on the busiest days for the Polderbaan runway (18R-36L) of 2023.

Within the scope, special attention is paid to the pushback process, the flow of aircraft within apron bays, and the various (un)loading destinations and any congestion around them. Additional analysis of emissions based on aircraft and TaxiBot movements is included, but detailed dispersion modelling is out of scope. The topics of disruptions, TaxiBot assignment, and service road traffic modelling are also out of scope.

2. Literature

Various themes of literature are relevant for this research. These themes are: earlier work regarding sustainable taxiing at Schiphol, dispatch towing, previous discrete event simulation of TaxiBot operations, and congestion management and other forms of sustainable taxiing. They will be analysed in the next sections.

2.1 Earlier sustainable taxiing research at Schiphol

Benda (2020) and Tindemans (2021) in their theses have explored aspects of TaxiBot operations at Schiphol airport related to route planning, vehicle assignment and schedule disruptions. Benda (2020) built an agent-based model to find that traffic on Schiphol's taxiways will increase with the use of TaxiBots, but that this can be accommodated without affecting safety. Tindemans (2021) considered the impact of schedule disruptions on TaxiBot operations using a mixed-integer linear programming (MILP) model. Short-term tactical planning can allocate TaxiBots to 48% of flights, which is only a minor loss compared to 48.5% of flights in the strategic planning model that was built. Van Winkel (2023) then expanded on this work by building a tactical planning MILP model that could optimize TaxiBot allocation in terms of fuel savings, schedule robustness or fairness among participating airlines. Finally, Clemens (2023) performed a societal cost-benefit analysis to determine the implementation of dispatch towing at Schiphol. His findings suggest that electric tows to the edge of the Schiphol apron yield a bigger (financial) benefit than towing the aircraft all the way to the Zwanenburgbaan (Runway 18C-36C) for departure.

There is also some other academic literature published by TU Delft on the topic of sustainable taxiing at Schiphol. Roling et al. (2015) found that electric taxiing at Schiphol should reach a speed of at least 17.5 kts (9 m/s) in order to avoid unacceptable departure delays to other departing traffic. This research was focused on electrified nose wheels rather than tow tractors but the findings should extend to both methods. The fuel savings of implementing these forms of sustainable taxiing can save 82% of the fuel used for taxiing at Schiphol (van Baaren & Roling, 2019). Other research by Groot and Roling (2022) has shown that sustainable taxiing is most suitable for shorter haul flights to major airports, since these flights carry relatively more fuel dedicated to taxiing across the large airports in Amsterdam and the flight's destination. Some recent research by van Oosterom et al. (2023) based on traffic volumes in the winter of 2019 has shown that Schiphol would need around 38 TaxiBots to service 913 flights: 26 small TaxiBots for 750 narrowbodies, and 12 large TaxiBots for the 163 widebody flights that day. Their research assumed electric TaxiBots so these figures include the need for charging, and the researchers also found that these relationships scale linearly as the number of flights increase.

2.2 Optimization of dispatch towing

A few papers have looked at the benefits and drawbacks of TaxiBots when the innovation first came to market. The main advantages, as already mentioned, are significant fuel savings for the airlines, as well as a resulting reduction in CO_2 and UFP emissions for the airport and its surroundings (Postorino et al., 2016; Di Mascio et al., 2022). Drawbacks of the system are the vehicle price, increased complexity for scheduling and assignment, and a more complicated ground traffic picture (Gualandi, 2014; Hospodka, 2014). The fuel saving effects will be biggest for major airports, where taxi times are longer. Furthermore, because the engines are not running, taxiing operations should become quieter (Hein & Baumann, 2016).

Zoutendijk et al. (2023) conclude from a literature survey that there are two main ways of solving routing problems: mixed-integer linear programming and simulation. Sirigu et al. (2018) developed algorithms based on neural networks and graph theory. They found that all algorithms perform adequately, but that the graph theory-based algorithms are much more efficient. One of the biggest challenges of this type of problem is conflict avoidance due to the increased number of vehicles in play. Soltani et al. (2020) therefore developed a MILP-model and applied it successfully to the case of Montreal airport.

2.3 Discrete event simulation of TaxiBot operations

Another common way to analyse dispatch towing systems is via discrete event simulation. Khammash et al. (2017) used this technique to simulate ground movements at Lisbon airport. Lisbon airport is a medium-sized airport in Portugal, with one runway and 220,000 yearly aircraft movements in 2023 (VINCI Airports, 2024, p.7). The researchers found fuel savings of 6% per deployed TaxiBot along with CO2 emission reductions of 5% per TaxiBot, with no observed increases in total taxiing times. Salihu (2020) also used discrete event simulation to assess the impact of dispatch towing operations at Montreal airport, another medium sized airport with two main runways and 200,000 yearly aircraft movements in 2023 (Aéroports de Montréal, 2024, p.5). They found that 22 TaxiBots would be needed to service all flights in the study period. Due to the minimization of fuel costs, this scenario would also lead to the lowest operational cost, compared to partial service scenarios. Further research by the same researchers introduces uncertainty to the problem, resulting in a less definitive conclusion about the cost-effectiveness of TaxiBot operations in the Montreal case study (Salihu et al., 2021). Their results show that 'TaxiBotting' all flights would add 4.2 minutes to the taxi-in process that currently costs 10.9 minutes (+38.5%), and would add two minutes to the current 15.1 minute taxi-out process (+13.2%).

While the results of the previous studies are very useful as a first indication, the studies did make some unrealistic assumptions and simplifications. The main areas where this may impact findings are in constant, non-context aware taxiing speeds (13.6 and 15 kts) and routes (shortest-path), (simple) delay estimation, and unrealistic runway and gate assignments.

This research aims to improve on those assumptions, and thus provide more accurate findings. Additionally, the focus lies on a much larger and more complex airport (Amsterdam Schiphol) which will by itself contribute to pushing the academic boundary forwards, and gaining further insights on how complex airports react to a shift in traffic operations. Thirdly, this paper has a main focus on traffic flow of TaxiBot operations with a secondary focus on emissions (CO2 and UFP), whereas the earlier two studies have mainly focused on potential emission reduction and cost-effectiveness of operations.

2.4 Congestion management and other forms of sustainable taxiing

Lukic et al. (2018) and Lukic et al. (2019) performed literature reviews to analyse the various sustainable taxiing options in more detail. The main conclusion to be drawn is that there are three major forms of sustainable taxiing: dispatch towing, electrified nose landing gear (NLG) taxiing, and single-engine taxiing. Electrified NLG taxiing has the best emission reductions on the ground, albeit at the cost of extra weight and thus reduced in-flight fuel performance (Cao et al., 2023), and in terms of operational efficiency it also is the quickest due to the aircraft now being able to execute its own pushback (Wollenheit & Mühlhausen, 2013; Zhang et al., 2019a).

Taxiing can also be made more sustainable by optimising processes. Kariya et al. (2011) created a discrete event simulation model to analyse taxiing patterns at Tokyo Haneda airport. They suggest adaptations to the schedule and an aircraft's actual off-block time (AOBT) to ease congestion and improve the airport's sustainability. Various models have been created by Mori (2012), Zhang et al. (2019b), and Zbakh et al. (2023) that all quantify and minimize the taxiing congestion (and maximize capacity) by adjusting the schedules of airports around the world. The control of pushback times was identified by Khadilkar and Balakrishnan (2014) as the most critical short-term strategy to improve the departure efficiency.

3. Data and Methods

The simulation model uses many different sources of data to build up realistic experiments. This chapter will introduce these sources, detail their usage, present how the simulation model is constructed and finally present the experiments that are performed and explain how the results from the simulations are generated.

3.1 Model input

The following data sources are used for the simulation:

- Schiphol airport map / *Basiskaart*, for constructing the network on a 1:1 scale.
- Groud radar data of GPS coordinates, timestamps and aircraft identifiers, to reconstruct the real-life radar routes and speeds.
- Geofence polygons of airport locations, to assign GPS coordinates to taxiways. An example is shown in figure C.1.
- Central information system Schiphol (CISS), to supplement the radar data with more information about the aircraft, like it's actual departure time at the game.
- TaxiBot mission portal, to simulate accurate loading and unloading durations.
- Fuel flow information and aircraft counts at Schiphol, to realistically calculate fuel usage and CO2 emissions.

The GPS tracks were processed and cleaned up to remove errors and mislabeled locations (e.g. when overflying polygons after take-off), resulting in a continuous, valid list of runway and taxiway segments from the runway to the gate for the arriving flight, and then from the gate back to the runway for the departing flight of that same aircraft.

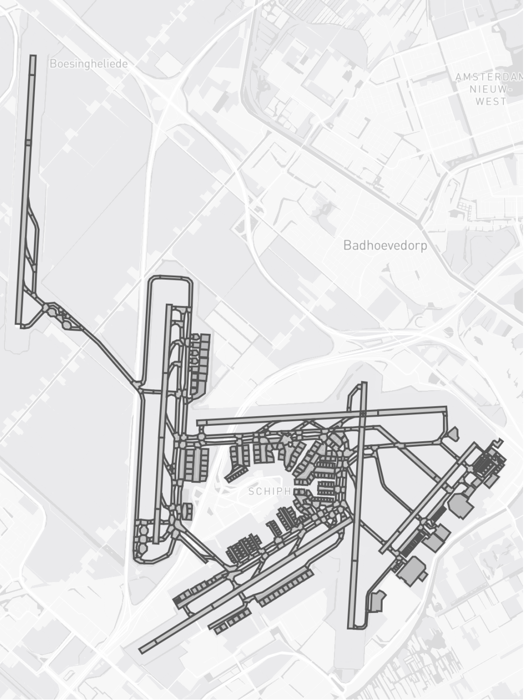

Figure C.1: Geofence polygons

Along with the radar route, some supplementary information was included in the simulation, like gate departure information (actual off-block time, AOBT) and whether any TaxiBot towing should take place for the inbound and outbound movements.

Aircraft were also assigned a size category, because the model would get too complex if every aircraft type (e.g. Airbus A321, Boeing 737-800, Embraer 175, etc.) would be included. Every aircraft type was assigned to one of three categories: regional, narrowbody, and widebody.

This split was made according to the categories in the Aircraft Stand table (Royal Schiphol Group, n.d.-b). Regional aircraft are defined as categories 1 to 3, narrowbodies as category 4, and widebodies as categories 5 to 9.

3.2 Simulation model

The discrete-event simulation is created in Simio (Simio, n.d.). The simulation's foundation consists of a set of nodes and links that has been drawn over the Schiphol airport map, which was scaled up so that one meter on the airport map would correspond to one meter distance in the simulation environment. The nodes are placed on the intersections of radar polygons (figure C.1), and are connected by links whenever an aircraft in real life would be able to drive from one polygon to another, see figures C.2a and C.2b.

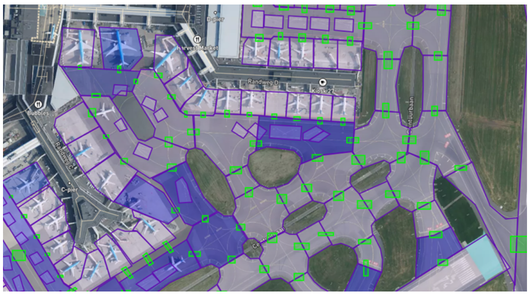

(a) Polygons between the C and D piers

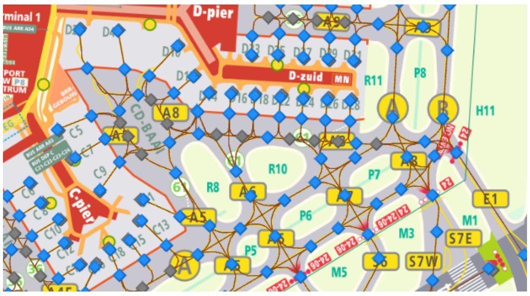

(b) The same network, implemented in Simio

Figure C.2: Creating the simulation network based on polygons

Model properties relating to TaxiBots could not be derived from radar data or the CISS database. Certain assumptions and choices had to be made for aspects like the maximum towing speed, driving speed when the TaxiBot is not towing an aircraft, and what the loading and unloading times were.

The TaxiBot's internal speed limit is configured for 22 knots of taxiing speed. When an aircraft is being towed by a TaxiBot it will drive at its radar speed, with a maximum limit of 22 knots. The non-towing speed of the TaxiBot is set to a constant 15 knots rather than its 22 knots limit, to account for driving slowly around corners, traffic interactions, and speed limits of 16 kts on the service roads around the Schiphol apron.

The values of TaxiBot loading and unloading times were taken from the TaxiBot mission portal that contains durations from TaxiBot missions around the world, and a function was modelled that best approximates this real-life distribution. This probability density function takes the form of a Random Triangular distribution, indicated as Random.Triang(a,b,c). Value a is the minimum value that occurs, value b is the most frequently occurring value, and value c is the maximum value. The triangle is then formed so that the area under the curve is equal to 1. Every time a TaxiBot is loading or unloading from an aircraft in the simulation, a random draw is made from these distributions to realistically model the loading time.

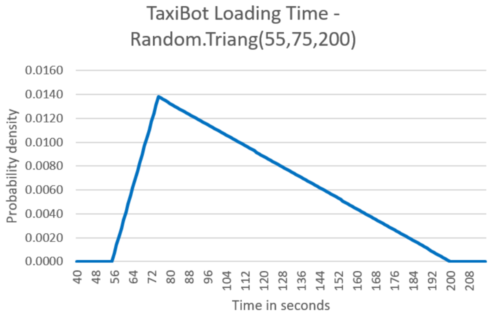

(a) Loading

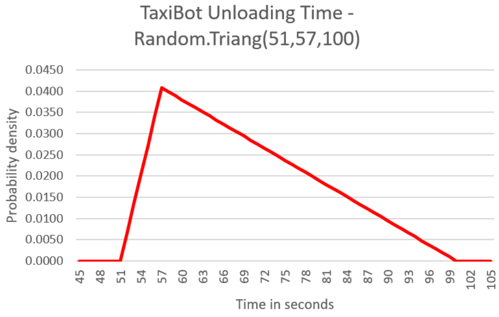

(b) Unloading

Figure C.3: TaxiBot loading and unloading times, based on real-world operational data

3.3 Experiments

To determine which experiments would be run, it first needs to be determined which days are most suitable for analysis. The current focus of concept TaxiBot operations at Schiphol airport is towing flights to the Polderbaan (18R-36L, the furthest away runway. The aim is to investigate if these operations are feasible, even on the busiest days.

3.3.1 Selecting days for analysis

The objective for selecting days is therefore twofold; the day must be a busy one (meaning a high flight count), and there must be lots of Polderbaan activity. This activity can be outbound, or a combination of outbound and inbound traffic.

Flight count data was combined with radar information, to gather the number of flight movements, Polderbaan arrivals, and Polderbaan departures for every day of 2023. Two 'scores' were developed, where the flight count of that day was multiplied by the Polderbaan departure count (for the optimal outbound day), and where the flight count of that day was multiplied by the Polderbaan departure count plus the Polderbaan arrival count (for the optimal combination day). The maximum scores of these two metrics in 2023 would determine the days to be simulated, and these turned out to be 14 June for outbound day, and 29 September for the combination day. They are shown in table C.1.

Table C.1: Selected days for simulation

Date	Total flights	Polderbaan Departures	Arrivals	Selection
14 June	1411	456	0	Departures
29 Sep	1427	233	227	Combined use

A single day – rather than a range of days to test each policy on – is deemed sufficient for the following reasons. Most importantly, every flight undergoing TaxiBot towing within the analysed day is essentially a mini-experiment of the towing policy. All aircraft undergoing that policy then form the variation within the dataset based on aircraft type, route, taxiing speeds, emission reduction, waiting time at the unloading stations, etc. Additionally, the selected day will already be representative of usual traffic volumes and split of aircraft types. The station congestion levels are determined by the aircraft mix (how many spaces they take up), unloading time at the station, and runway usage, with a maximum capacity of 40 movements per hour (Airports Council International, 2023). Of these three factors, only the runway usage would change on different days, but with the selected days already hitting the maximum capacities for the 2023, no additional days are required for the analysis. Finally, the addition of extra days would not introduce any new circumstances, except maybe different runway use configurations. However, the fact that the selected days are selected precisely for their desired runway configurations makes this a moot point.

Not all flight movements were able to be included in the final radar dataset, due to some irrecoverable errors related to transponders being shut down for significant portions of the taxiing journey, private jets and general aviation traffic showing erratic transponder behaviour at the stands, and other similar issues where a runway or stand could not be determined. The statistics of how many movements had to be excluded from the simulation dataframes is shown in table C.2. The high retention rates demonstrate that the effort to reduce errors and handle exceptions has been worthwhile.

Table C.2: Flight movements in the simulation

Date	Total flights	Flights dropped	Retention rate
14 June	1411	53	96.2%
29 Sep	1427	39	97.3%

From the 53 dropped flights on 14 June, there were eleven departing flights using the Polderbaan, and 0 arrivals. From the 39 dropped flights on 29 September, there were four departing flights using the Polderbaan, and five arrivals.

3.3.2 TaxiBot towing policies

The TaxiBot towing policies can be divided in two categories: towing to a runway, and towing out of the apron bays. They both have different purposes. Towing to a runway saves the most fuel, and thus CO2 and UFP emissions. Towing out of the apron bays has the purpose of complying with the NLA's green zone ruling (fig. C.5) in the most operationally efficient way: shorter distances mean more towing activities, given a fixed number of TaxiBots. Here, aircraft are only brought to the edge of the green zone, where they can start their engines and continue on their way. Since the focus of TaxiBot operations lies on green zone emission reduction, only aircraft that depart from stands inside this green zone are considered for TaxiBot operations.

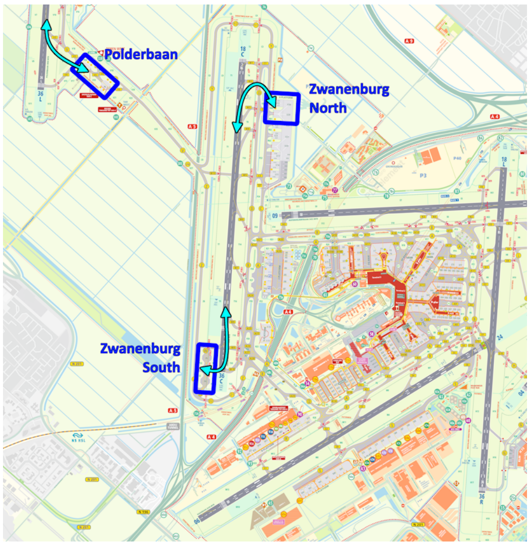

(a) Towing towards the runways

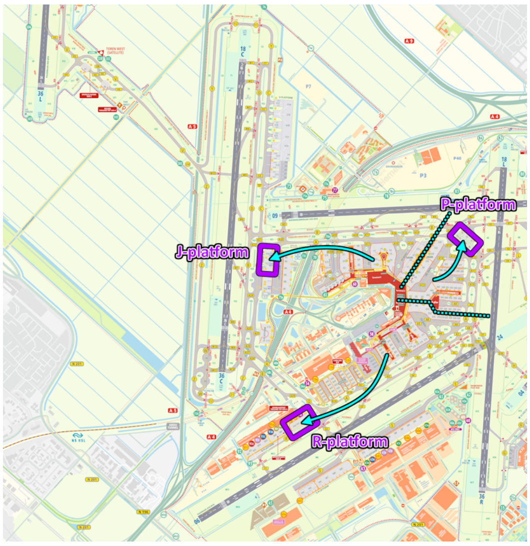

(b) Towing out of the bays

Figure C.4: TaxiBot towing destinations

There are nine possible towing options that are considered as policies: six for runways (figure C.4a), and three for bay edge towing (figure C.4b), which are only considered for outbound movements.

Besides selecting a towing location, there is also the policy option to allow the aircraft to start their engines during the towing movement. This process could save anywhere from three to seven minutes. The engine start moment would be timed in such a way that the engines finished their startup at the moment that the TaxiBot has finished unloading the aircraft. However, these engine start-up moments cannot happen before an aircraft has left the apron bay.

For the simulation, each of the six towing destinations is assumed to have the same capacity: two holdings, with two stations per holding. Each holding is able to accommodate two narrowbody aircraft (one per station), or one widebody aircraft (occupies both stations, due to size). The simulated station assignment logic for narrowbody aircraft ensures that a narrowbody aircraft will always try to select the spot next to another narrowbody aircraft, with the purpose of keeping the second holding available for a widebody aircraft to prevent unnecessary waiting.

Figure C.5: The green zone of Schiphol (Netherlands Labour Authority, 2023a)

3.3.3 Experiments to simulate

Combining the above towing policies and the selected days for simulation, nine experiments and two reference scenarios are drawn up. The reference scenarios simulate the airport as it was that day; purely according to the radar, with no TaxiBot operations. The scenarios are presented in table C.3. Each change relative to the previous row is highlighted in bold.

The nine TaxiBot experiments start with conditions that are currently being trialled in real life at Schiphol airport. Gradually, the experiments are scaled up in complexity. All results per experiment will be compared to the appropriate reference scenario, to determine the impact of the TaxiBot operations in the experiment. These impacts can then be compared to the impact of other experiments in order to draw conclusions on which TaxiBot operations perform best.

Table C.3: Experiments

#	Towing destinations	Direction	Radar day	Aircraft	Startup
0.1	Reference scenario	-	14 Jun	-	-
0.2	Reference scenario	-	29 Sep	-	-
1	Polderbaan only	Outbound	14 Jun	Narrowbody only	Stationary
2	Polderbaan only	Outbound	14 Jun	Narrowbody only	**In motion**
3	Polderbaan only	Outbound	14 Jun	**All types**	In motion
4	Polderbaan only	**In and out**	**29 Sep**	All types	In motion
5	**Polderbaan & Zwanenburgbaan**	In and out	29 Sep	All types	In motion
6	**Bay edges (platforms J, P, R)**	**Outbound**	29 Sep	All types	**Stationary**
7	Bay edges (platforms J, P, R)	Outbound	29 Sep	All types	**In motion**
8	**All**	Outbound	29 Sep	All types	In motion
9	All	**In and out**	29 Sep	All types	In motion

3.4 Model output

The results from the simulation can be split in four categories: flow and on-time performance (OTP), TaxiBot operations, station congestion, and emissions. Before defining results for the emissions category, it is necessary to first understand the fuel usage at the airport, as well as the calculations for converting fuel into CO2 emissions.

3.4.1 Calculating engine durations and resulting emissions

Assumptions are made for the moments when engines and APUs are turned on and off. These durations are then multiplied by fuel flows per aircraft category to determine used fuel, and finally converted into CO2 emissions.

For inbound movements, engines are turned 'on' upon exiting the runway (runway engine use is out of scope because this is a taxiing simulation), and turned off upon arriving at the stand. If the total engine time from runway exit until stand is less than four minutes, this runtime is increased to four minutes, due to the required cool down time before shutdown of four minutes. The APU is turned on one minute before arriving at the stand, and turned off three minutes after arriving. The assumption is made that the aircraft can then get its power from a ground power unit.

For inbound movements with a TaxiBot, the engines are shut down upon the completion of the TaxiBot loading process. The same four minute minimum applies. The APU is turned on one minute before the start of the TaxiBot loading process, and turned off as usual (after three minutes) at the stand.

For outbound movements, the engines are turned on at the start of the pushback process and turned 'off' upon entering the runway, since runway engine use is out of scope. The APU is turned on five minutes before pushing back from the stand (AOBT). The assumption is made that the aircraft can be connected to ground power until then. The APU is turned off upon the completion of the pushback (and thus engine start) process.

For outbound movements with a TaxiBot, the time engines are turned on is dependent on the 'engine start during towing' policy. If engines are started when at the unloading station, then the engine start moment is the start of the TaxiBot unloading process. If engines can be started during the tow, then the engine start time is defined in such a way that engines have finished starting up when TaxiBot unloading completes, with the exception that engines may not be started inside the green zone. The APU is turned off after the engines are started.

Fuel flows per aircraft category are presented in table C.4. These values are derived from Schiphol data about fuel flows per aircraft type, and taken as a weighted average per category with weights being the number of times those aircraft visit Schiphol.

Table C.4: Fuel flow per aircraft category

Category	Engines fuel flow [kg/min]	APU fuel flow [kg/min]
Regional	8.924	1.800
Narrowbody	12.863	1.936
Widebody	36.171	4.792

Converting fuel usage to CO2 is relatively straightforward, because the APU and the engines make use of the same type of fuel, called Jet-A1. This fuel contains 3.16 kg of CO2 per kg of fuel (IATA, 2022, p.8). The assumption is made that 2% (EASA, n.d.) of sustainable aviation fuel (SAF) is mixed with Jet-A1, and that SAF reduces CO2 by 75% (IATA, 2023). This results in the fuel mix having a CO2 impact of 3.113 kg of CO2 per kg of fuel. The resulting CO2 emissions are shown in the table below.

Table C.5: CO2 emissions per minute, per aircraft category

Category	Engines CO2 [kg/min]	APU CO2 [kg/min]
Regional	27.778	5.602
Narrowbody	40.038	6.027
Widebody	112.586	14.916

Emissions of TaxiBots are also calculated, based on the battery capacity, towing range, and estimated energy expenditure per tow with data from the TaxiBot's manufacturer. The CO2 emissions of a TaxiBot are listed in table C.6.

Table C.6: CO2 emissions per TaxiBot

TaxiBot type	No-tow CO2 [kg/km]	Towing CO2 [kg/km]
Narrowbody	0.355	1.063
Widebody	0.276	1.614

No data is available for how many ultrafine particles are emitted by an aircraft engine, this is currently a topic of ongoing research. To still provide insights into how UFP emissions are affected by the towing policies, heat maps will be made of an aircraft's location and engine status (APU running, engines running, engine start-up in progress).

3.4.3 Performance indicators per analysis category

Alongside CO2 and UFP emissions, the simulation records multiple taxiing statistics for each result category. An overview is provided in table C.7.

For the TaxiBot operations category, the towing and non-towing times are combined into a total mission time statistic, with their start and end times also recorded. An additional five minutes is added to this duration, owing to the assumption that a TaxiBot should be present at the mission start location five minutes before the aircraft is ready to be connected to the TaxiBot. This 'total mission time' – in particular its start and end times – allows for the calculation of the number of active TaxiBots at any one time in the simulation, and therefore also the maximum number of active TaxiBots. This statistic is used to approximate the number of TaxiBots required for these operations. This remains an approximation, as mission optimization is not applied and no charging or refuelling is considered.

Table C.7: Overview of results per category

Category	Result
Flow and OTP	Inbound taxiing time Outbound taxiing time Time out of bay
TaxiBot operations	Distance travelled towing Distance travelled not towing Time spent towing Time spent not towing Total mission time Maximum number of active TaxiBots Time of day of maximum active TaxiBots Total mission count
Station congestion	Waiting time at station Total time at station Station occupancy rate Time all stations occupied
Emissions	CO2 of aircraft engines and APU in kg CO2 of TaxiBots in kg Total CO2 emitted per aircraft Engine mode timestamps (UFP) X,Y coordinates of aircraft every 5 seconds (UFP)

4. Results

Results of the simulations will be presented for each of the four categories that were determined in the previous section.

4.1 On-time performance and flow

Several patterns become apparent from the data in table C.8. Firstly, all TaxiBotting inbound operations add a significant amount of time to the taxiing duration. Next, outbound TaxiBot operations always result in a reduction of outbound time, except for when engines are started up while standing still at the unloading station (experiments 1 and 6). Third, the time gains for leaving the bay are fairly constant at a reduction of three minutes. For non-TaxiBotted aircraft, the changes for all three measurements are negligible.

Table C.8: Process times for on-time performance, relative to reference scenarios

	All aircraft			Non-TaxiBotted aircraft			TaxiBotted aircraft		
	Inbound	**Outbound**	**Out of bay**	**Inbound**	**Outbound**	**Out of bay**	**Inbound**	**Outbound**	**Out of bay**
Exp 1	-00:01	+00:21	-01:22	-00:01	-00:05	00:00	–	+00:52	-03:04
Exp 2	-00:01	-00:34	-01:22	-00:01	-00:05	00:00	–	-01:12	-03:04
Exp 3	-00:03	-00:48	-01:47	-00:03	00:00	00:00	–	-01:30	-03:18
Exp 4	+00:38	-00:28	-00:58	-00:01	-00:00	00:00	+02:27	-01:31	-03:13
Exp 5	+01:32	-00:46	-01:18	-00:01	-00:01	00:00	+03:12	-01:53	-03:13
Exp 6	-00:02	+01:24	-02:44	-00:02	-00:26	–	–	+01:33	-02:57
Exp 7	+00:01	-00:29	-02:44	+00:01	-00:26	–	–	-00:30	-02:57
Exp 8	+00:01	-00:37	-02:44	+00:01	-00:23	–	–	-00:38	-02:57
Exp 9	+01:31	-00:34	-02:44	-00:01	-00:21	–	+03:10	-00:35	-02:57

Figures C.6 shows the data from the first two columns of table C.8 (and their distributions) as a box plot. Blue boxes represent inbound movements, and purple ones show outbound. The thicker black stripe represents the median. For Experiment 6, it has a value of 1.31 minutes (+1:19). This indicates that towing aircraft to the edge of the bays and starting engines stationary results in significant outbound delays; this will be explored more in section 4.3.

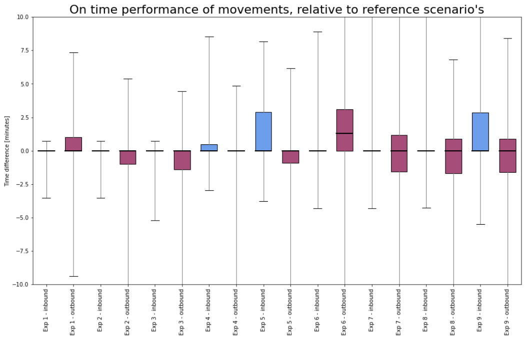

Figure C.6: Inbound and outbound on-time performance distributions per experiment

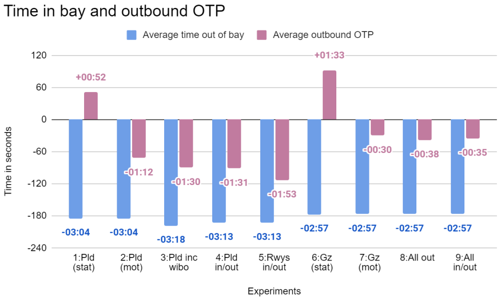

Figure C.7: The impact of time in bay reduction on outbound on-time performance

Figure C.7 shows the impact of stationary engine startup on outbound OTP. This plot clearly shows that the time out of bay is the main cause of the time gains. Some of this time is lost again in the unloading process, but if the TaxiBot is able to start the engines while taxiing (all experiments except 1 and 6), the aircraft will be faster than in the reference (no-TaxiBot) scenario.

The impact of TaxiBots on the flow on the taxiways is not explicitly quantified, but can still be assessed by an analysis of radar speeds in the reference scenarios, like in figure C.8.

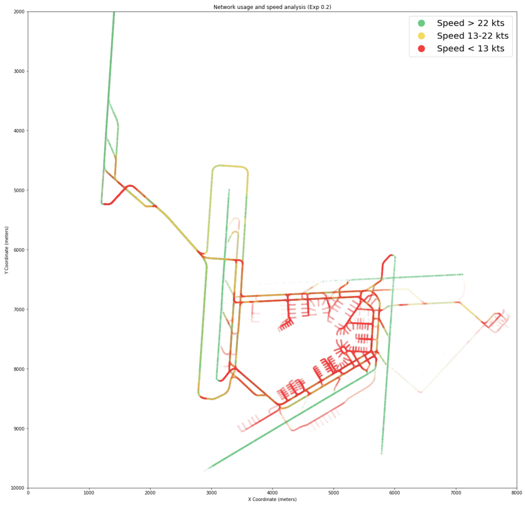

Figure C.8: The radar speeds of the reference scenario 0.2 (29 September 2023)

The speeds around the apron bays, runway entry locations, and tight corners are lower than those further out on the straight taxiways. There are also very few locations on the map where the prevailing taxiing speeds are faster (green) than the 22 kts speed limit of the TaxiBot. As a result, the vast majority of any on-time performance differences will have been caused by process steps like (un)loading the TaxiBot, and skipping the usual pushback process, rather than any type of congestion caused by TaxiBot towing.

4.2 TaxiBot operations

Figure C.9 presents a visual overview of TaxiBot mission statistics. The shorter average towing distances for experiments 6 and 7 are reflected in shorter mission times and fewer TaxiBots are required, when comparing the results to experiments 8 and 9 where runways are also included as mission destinations. Time per mission consists of time towing and time not towing (the TaxiBot driving to and from the mission), to represent the full time that the TaxiBot was occupied with this aircraft towing mission. The mission time includes a five minute waiting time, to account for the fact that a TaxiBot needs to be present at the mission pick-up location five minutes before the aircraft is ready to depart.

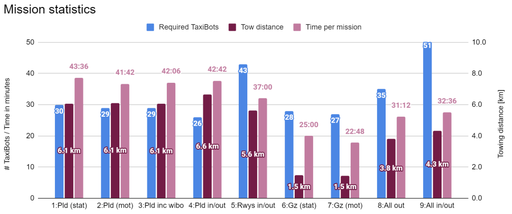

Figure C.9: TaxiBot mission statistics, per experiment

The split between narrowbody and widebody TaxiBots for all experiments in the simulations are shown in table C.10. For the majority of towing policies, the peak number of TaxiBots lies around 30. This is very close to the 38 TaxiBots that van Oosterom et al. (2023) calculated with their model, especially when considering that the simulation model does not account for charging the electric TaxiBots. They found that 38 TaxiBots were needed for 913 flights; 26 small TaxiBots for 750 narrowbodies, and 12 large TaxiBots for 163 widebodies. Their research setup was closest to Experiment 5 of this simulation model, where all flights are towed to the runways (the main difference being that they considered all runways, versus this simulation only considering three runway destinations). In the simulation, experiment 5 requires 41 narrowbody and 12 widebody TaxiBots to service 587 and 118 flights respectively.

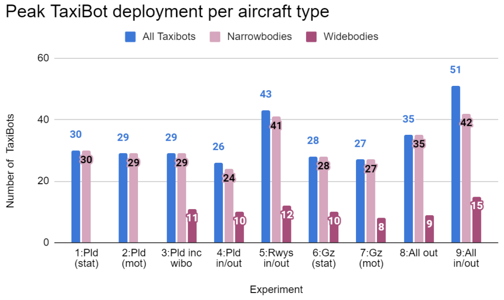

Figure C.10: Peak TaxiBots active per size type

4.3 Station congestion

The simulation results are shown in table C.9, for both the queueing process before being assigned a station, and the total time spent queueing and at the station.

Table C.9: Station congestion results per experiment

	Queueing time			Total time at station			Station occupancy	
	Mean	[Q1-Q3]	Max	Mean	[Q1-Q3]	Max	Occupancy rate	Time all occupied
Exp 1	00:03	[00:00 - 00:00]	01:36	04:08	[03:15 - 04:49]	08:11	29.9%	70:30
Exp 2	00:00	[00:00 - 00:00]	00:00	02:10	[01:39 - 02:30]	05:35	15.9%	03:18
Exp 3	00:04	[00:00 - 00:00]	02:13	02:21	[01:46 - 02:45]	05:35	24.8%	59:36
Exp 4	00:02	[00:00 - 00:00]	03:08	02:20	[01:46 - 02:45]	05:38	27.1%	39:06
Exp 5	00:01	[00:00 - 00:00]	03:08	02:16	[01:44 - 02:39]	05:38	13.4%	14:24
Exp 6	00:12	[00:00 - 00:00]	14:32	03:54	[02:49 - 04:31]	19:22	20.1%	48:42
Exp 7	00:02	[00:00 - 00:00]	03:45	01:45	[01:03 - 02:12]	06:31	9.7%	12:18
Exp 8	00:03	[00:00 - 00:00]	03:45	02:03	[01:25 - 02:30]	06:31	6.6%	09:48
Exp 9	00:02	[00:00 - 00:00]	03:45	02:10	[01:35 - 02:38]	06:31	9.5%	11:54

For better insight in the distributions of these times, the Q1 (25th percentile), Q3 (75th percentile), and maximum values are also reported. The queueing times, especially the Q1-Q3 ranges all being 0, show that aircraft rarely have to wait to get an (un)loading station assigned. The variation in average times therefore is caused by outlier values.

The occupancy rates represent how much of the time a station was occupied. This value is recorded per station, then averaged for all stations that were available in the simulation run, and reported in the table. The 'Time all occupied' column represents a total time of how often an (un)loading location (consisting of four stations) was unavailable for arriving aircraft, resulting in them having to queue. All occupancy rates are below 30%, meaning that the stations were unoccupied for the majority of the time. However, the times when all stations were occupied can get quite high, in experiment 1 even more than 1 hour and 10 minutes. This indicates that the regular levels of congestion around the stations are quite low, but that some peak moments can still lead to queueing.

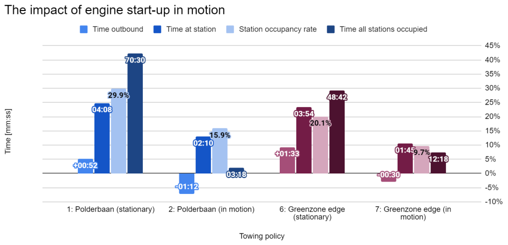

Figure C.11: The impact of engine start-up in motion on station congestion

The impact of starting engines while stationary, or starting them during the tow already also becomes quite apparent, as presented in figure C.11. The plot compares two sets of experiments to each other; experiments 1 to 2 (blue), and 6 to 7 (pink). These experiment sets are identical except for the engine start-up policy. The time at station drops by around two minutes in both scenarios, also resulting in two minutes difference in the outbound times. Additionally, the time that all stations are occupied also gets reduced by a significant amount; 95% and 75% for the sets of experiments respectively.

4.4 Emissions

4.4.1 CO2 emissions

Table C.10 shows the impact of the TaxiBot operation experiments on CO2 emitted at the airport. When comparing experiment 2 and 3, the only difference in policy (towing outbound to the Polderbaan, startup in motion) is that experiment 3 also includes towing widebodies. While there are many more narrowbody missions than widebody missions for experiment 3 (351 vs 77), the CO2 impact nearly doubles.

From the column of max CO2 impact values, it can be concluded that CO2 (and fuel usage) reductions of 72% to 75% are achievable when towing to the Polderbaan, and that savings of up to 53% are possible for towing to the edge of the bays. Allowing the engines to start during the tow also has a slight positive effect on CO2, reducing total airport emissions by 0.7 and 1.3% respectively.

Table C.10: CO2 results per experiment

	All aircraft			TaxiBotted aircraft				
	Engines on	**APU on**	**CO2 impact**	**Engines on**	**APU on**	**CO2 impact**	**CO2 [Q1-Q3]**	**Max CO2**
Exp 1	-04:34	+04:55	-13.6%	-10:11	+11:03	-37.0%	[-27.0% ; -48.9%]	-72.4%
Exp 2	-04:37	+04:02	-14.3%	-10:16	+09:05	-38.9%	[-29.5% ; -50.2%]	-73.8%
Exp 3	-05:53	+05:03	-27.2%	-10:49	+09:19	-41.3%	[-31.0% ; -50.4%]	-74.9%
Exp 4	-06:07	+07:03	-25.1%	-12:15	+14:10	-40.4%	[-29.2% ; -44.6%]	-74.4%
Exp 5	-07:42	+09:48	-30.3%	-10:43	+13:41	-37.9%	[-22.0% ; -45.0%]	-75.7%
Exp 6	-01:43	+03:06	-6.5%	-01:49	+03:20	-7.4%	[+5.1% ; -13.6%]	-54.4%
Exp 7	-01:45	+01:17	-7.8%	-01:51	+01:23	-8.9%	[+3.1% ; -15.7%]	-54.8%
Exp 8	-05:07	+04:32	-20.2%	-05:30	+04:53	-23.3%	[-3.2% ; -36.9%]	-74.4%
Exp 9	-08:15	+10:32	-32.9%	-08:52	+11:22	-36.5%	[-15.6% ; -43.7%]	-75.8%

Figure C.12 shows the CO2 savings that the various TaxiBot operation policies would have on the emissions for the whole airport. The maximum impact is had in experiment 9 (where as much gets TaxiBotted as possible), where 32.9% of emissions are reduced, but it is interesting that even the runways-only scenario in experiment 5 already comes very close to this figure, with much less operational complexity. The CO2 savings for experiments 6 and 7 are much lower, leading to the conclusion that the main contributor to CO2 savings is simply the engine running time saved, most often caused by the distance that aircraft are towed.

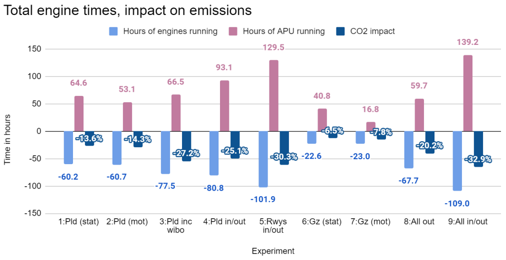

Figure C.12: Engine times, and the impact on CO2 emissions

4.4.2 UFP heat maps

The heat maps of ultrafine particle emissions give an idea of where emissions occur, but they can not say much regarding quantity of emissions. The intention for providing these heat maps is partially to provide insight for this thesis, but also for future researchers to use the data from these simulations to quantify the impacts of UFP emissions when more is known about how much UFP is actually emitted by an aircraft engine. The most interesting contrasts between experiments will be highlighted and discussed.

The reference scenario of 29 September (Experiment 0.2) is shown in figure C.13. It becomes apparent how much of the harmful engine startup emissions currently occur inside the NLA green zone, within the apron bays.

The two heat maps in figures C.14 and C.15 show the emissions of experiment 6 and 9 (respectively: towing all outbound flights to the edges of the green zone with stationary startup, and towing to runways inbound and outbound with the rest being towed to the green zone edges, while starting all engines during the towing processes).

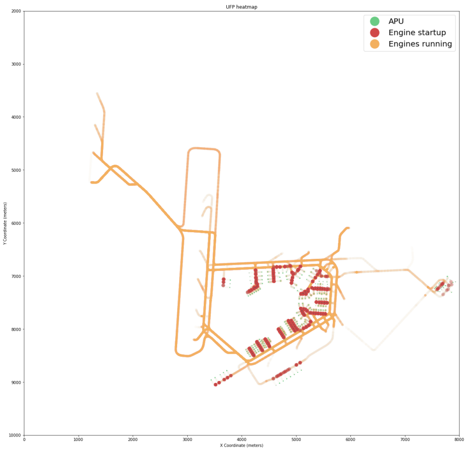

Figure C.13: Reference scenario 29 September

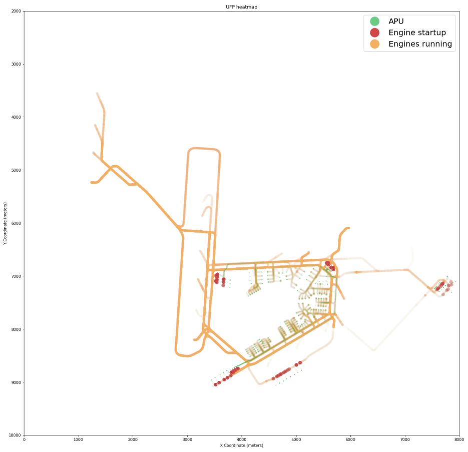

Figure C.14: Maximum UFP impact (Experiment 6)

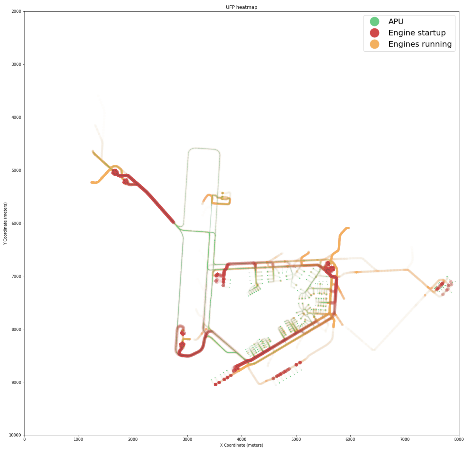

Figure C.15: Maximum CO2 impact (Experiment 9)

The results of experiment 6 in figure C.14 are drastic, when compared to C.13. A lot of UFP is removed from the bays, and the stationary startup also means that the engine start-up emissions remain limited to the unloading locations. Operational downsides of this experiment have been detailed in earlier sections of this chapter, however.

The UFP heat map of experiment 9 in figure C.15 shows the emissions for the scenario that reduces the maximum amount of CO2. Engine startup during towing is clearly visible at the Polderbaan, and while the bays are also cleared of engine start-up UFP like in experiment 6, the engine start-up while driving results in a lot of UFP emissions on the ring-road around the apron bays.

5. Discussion

5.1 Model validity

Model validity is ensured by the way that the simulation model is constructed – using a lot of real world data – and observed by the results it produces in reference scenarios. When no TaxiBot operations are modelled, a simulation or real-life movements as they were remains. After the trials at Schiphol in the COVID-19 pandemic, fuel savings were calculated to be up to 50-65% for the outbound movement (stand to runway) of the towed flight (Royal Schiphol Group, 2021). These findings correspond to the results presented in table C.10, experiment 1, where the 75% percentile of CO2 reductions is 49% and the maximum is 72%. The model is also validated by peer review. Schiphol operational experts and even a KLM pilot have been shown the model, and agreed with its workings.

5.2 Key findings and implications

The key findings of the research are as follows. Flights using the polderbaan can reduce their ground-based CO2 emissions by over 70%, while also reducing outbound taxiing times by an average of 1 minute and 30 seconds per flight. Queueing at the unloading stations does not cause operational problems, and is often almost entirely prevented if engine startup can take place during the towing operation. Additionally, start-up during the tow saves an average of two minutes of outbound process time compared to stationary start-up operations (causing the 1:30 time saving, rather than a 0:30 delay). Inbound towing saves on CO2 emissions but also leads to an inbound taxiing time delay of over 2.5 minutes per aircraft. Its ultrafine particle emission gains are also minimal, due to a required four minute engine cool down before engines are able to be shut down.

The main implication of these results is that TaxiBot operations on large scales, at a complex airport like Amsterdam Schiphol are feasible, making allowance for the assumptions that have been made and topics that were placed out of scope like service road traffic compatibility. Taxiway traffic congestion, nor congestion at the unloading stations, should not be a concern for operations.

5.3 Generalizability, limitations, and future work

A high degree of generalizability towards other (complex) airports is expected for these findings. Correlations of taxiing time savings to taxiing distances, and of CO2 reductions to towing times can be extended to predict the results for other airports. Additionally, runway capacity of inbound and outbound movements is independent of airport characteristics, so as long as an airport can make four unloading stations available next to a runway, and can incorporate engine start-up during towing into their TaxiBot operations, similar queueing and unloading capacity characteristics can be expected.

Limitations of the study are that the model does not simulate any service road traffic, leaves TaxiBot mission assignment and charging needs of the towing vehicles out of scope, and that the model assumes perfect reliability for both aircraft and TaxiBots. Additionally, assumptions about fuel usage and TaxiBot acceleration might have impacted the results. The model assumes taxiing using all engines at idle thrust, which is accurate but not completely representative of reality.

Future work can build on the findings from this research in three ways. Firstly, the impact of the freed up time in apron bays on gate planning can be investigated, to see if TaxiBot operations would also make a meaningful impact on airport capacity. Secondly, service road traffic and related infrastructure is left out of scope, and therefore not considered as a possible operational limitation or impediment. The widths of the service roads and possible mitigations of problems can be investigated, as well as the likelihood of other service road traffic leading to unexpected TaxiBot delays. The third topic for further research regards UFP emissions. Once data becomes available that details how much UFP is emitted by various types of aircraft engines and APUs, the UFP impact of real-life TaxiBot operations can be optimized by prioritising flight movements that will expose ground crew working in the green zone to the most ultrafine particle emissions, to have the maximum positive impact on health of workers at the airport.

6. Conclusions

Analysis of taxiing times found that outbound times can be less or more than the taxiing times in the reference scenario, depending on engine start-up policy: start-up during the tow can save up to 1.5 minutes of taxiing time. Inbound times were found to be 2.50 to 3.25 minutes longer than in reference scenarios.

Apron bay flow was investigated and highlighted significant time gains of three minutes per outbound aircraft on average. This has the potential to create additional gate capacity at the airport, but more research is needed to quantify these gains.

The engine start-up policy turns out to be critical for the (un)loading locations. Some (manageable) queueing forms if engines are required to start up while stationary, with this congestion almost completely disappearing if engines are able to be started during the tow. Additionally, spreading the unloading demand over multiple locations at the airport (multiple runways combined with towing to edges of the green zone) further reduces station load. Taxiing time analysis of aircraft not involved with TaxiBot operations saw no significantly different results from the reference scenario.

Emissions were the final focus area of the research. CO2 emissions reduced for all experiments by up to 76% per flight and 32.9% for the airport as a whole. Results show that reduction of taxiing time is the main driver of CO2 reduction, leading to the recommendation that servicing the furthest runways is optimal from a CO2 reduction perspective. UFP analysis is hindered by the absence of quantitative emission data, but emission heat maps have shown that significant relocation of emissions is possible.

The potential impact of TaxiBot operations on the flow and emissions of taxiing aircraft at Schiphol airport is that significant emission reduction can be achieved with minimal negative, or even positive impacts on flow and on-time performance at the airport.

References

Airports Council International. (2023). *Guidance on airport capacity declarations.* Retrieved from `https://www.aci-europe.org/downloads/publications/ACI%20Guidance%20on%20Airport%20Capacity%20Declarations.pdf`

Aéroports de Montréal. (2024, February 3). *2023 aéroports de montréal passenger statistics.* Retrieved from `https://www.admtl.com/sites/default/files/2024/ADM_Statsdet_2023EN.pdf`

Benda, B. (2020). Agent-based modelling and analysis of non-autonomous airport ground surface operations. *Master of Science Thesis.*

Cao, F., Tang, T.-Q., Gao, Y., You, F., & Zhang, J. (2023). Calculation and analysis of new taxiing methods on aircraft fuel consumption and pollutant emissions. *Energy, 277*, 127618.

Clemens, P. (2023). To tow or not to tow: A social cost-benefit analysis of engine-off dispatch aircraft towing at Amsterdam airport Schiphol. *Master of Science Thesis.*

Di Mascio, P., Corazza, M. V., Rosa, N. R., & Moretti, L. (2022). Optimization of aircraft taxiing strategies to reduce the impacts of landing and take-off cycle at airports. *Sustainability, 14*(15), 9692.

EASA. (n.d.). *Fit for 55 and ReFuelEU aviation*. Retrieved 2024-05-24, from `https://www.easa.europa.eu/en/light/topics/fit-55-and-refueleu-aviation`

Groot, M., & Roling, P. C. (2022). The potential impact of electric aircraft taxiing: A probabilistic analysis and fleet assignment optimization. In *Aiaa aviation 2022 forum* (p. 3919).

Gualandi, E. (2014). TaxiBot: Analisi dei benefici derivanti dall'introduzione di un veicolo semi-robotico per dispatch towing (TaxiBot: Analysis of the benefits of introducing a semi-robotic vehicle for dispatch towing). *Master of Science Thesis.*

Hein, K., & Baumann, S. (2016). Acoustical comparison of conventional taxiing and dispatch towing - TaxiBot's contribution to ground noise abatement. In *30th congress of the international council of the aeronautical sciences (icas)* (pp. 1–7).

Hospodka, J. (2014). Electric taxiing–TaxiBot system.

MAD-Magazine of Aviation Development, 2(10), 17–20.

IATA. (2022, April 19). *Iata carbon offset program: Frequently asked questions.* Retrieved from `https://www.iata.org/contentassets/922ebc4cbcd24c4d9fd55933e7070947/icop_faq_general-for-airline-participants.pdf`

IATA. (2023, May). *Net zero 2050: sustainable aviation fuels.* Retrieved from `https://www.iata.org/en/iata-repository/pressroom/fact-sheets/fact-sheet---alternative-fuels/`

Kariya, Y., Mase, T., Yoshihara, S., & Ota, J. (2011). Analysis of congestion of taxiing aircraft at a large airport. In *2011 ieee international conference on robotics and biomimetics* (pp. 180–185).

Khadilkar, H., & Balakrishnan, H. (2014). Network congestion control of airport surface operations. *Journal of Guidance, Control, and Dynamics, 37*(3), 933–940.

Khammash, L., Mantecchini, L., & Reis, V. (2017). Micro-simulation of airport taxiing procedures to improve operation sustainability: Application of semi-robotic towing tractor. In *2017 5th ieee international conference on models and technologies for intelligent transportation systems (mt-its)* (pp. 616–621).

Lukic, M., Giangrande, P., Hebala, A., Nuzzo, S., & Galea, M. (2019). Review, challenges, and future developments of electric taxiing systems. *IEEE Transactions on Transportation Electrification, 5*(4), 1441–1457.

Lukic, M., Hebala, A., Giangrande, P., Klumpner, C., Nuzzo, S., Chen, G., ... Galea, M. (2018). State of the art of electric taxiing systems. *In 2018 ieee international conference on electrical systems for aircraft, railway, ship propulsion and road vehicles & international transportation electrification conference (esars-itec)* (pp. 1–6).

Mori, R. (2012). Aircraft ground-taxiing model for congested airport using cellular automata. *IEEE Transactions on Intelligent Transportation Systems, 14*(1), 180–188.

Netherlands Labour Authority. (2023a, March 14). *Op schiphol moeten diesel aangedreven arbeidsmiddelen worden vervangen.* Retrieved from `https://www.nlarbeidsinspectie.nl/actueel/nieuws/2023/03/14/op-schiphol-moeten-diesel-aangedreven-arbeidsmiddelen-worden-vervangen`

Postorino, M. N., Mantecchini, L., & Gualandi, E. (2016). Integration between aircraft and handling vehicles during taxiing procedures to improve airport sustainability. *International Journal of Transport Development and Integration, 1*(1), 28–42.

Roling, P. C., Sillekens, P., Curran, R., & Wilder, W. D. (2015). The effects of electric taxi systems on airport surface congestion. In *15th aiaa aviation technology, integration, and operations conference* (p. 2592).

Royal Schiphol Group. (n.d.-b). *Aircraft stand allocation and properties [download: Overview aircraft types].* Retrieved 2024-05-29, from `https://www.schiphol.nl/en/operations/page/allocation-aircraft-stands/`

Royal Schiphol Group. (2021). *Sustainable taxiing: TaxiBot trial.* Retrieved from `https://www.schiphol.nl/en/innovation/page/sustainable-taxiing-TaxiBot-trial/`

Royal Schiphol Group. (2024c, January 5). *Schiphol welcomed 61.7 million travellers in 2023.* Retrieved from `https://news.schiphol.com/schiphol-welcomed-617-million-travellers-in-2023/`

Salihu, A. L. (2020). *Impact of on-ground taxiing with electric powered tow-trucks on congestion, cost, and carbon emissions at montreal-trudeau international airport* (Unpublished doctoral dissertation). Concordia University.

Salihu, A. L., Lloyd, S. M., & Akgunduz, A. (2021). Electrification of airport taxiway operations: A simulation framework for analyzing congestion and cost. *Transportation Research Part D: Transport and Environment, 97*, 102962.

Simio. (n.d.). *Simio simulation software.* Retrieved from `https://www.simio.com/software/simulation-software.php`

Sirigu, G., Cassaro, M., Battipede, M., & Gili, P. (2018). Autonomous taxi operations: algorithms for the solution of the routing problem. In *2018 aiaa information systems-aiaa infotech@ aerospace* (p. 2143).

Soltani, M., Ahmadi, S., Akgunduz, A., & Bhuiyan, N. (2020). An eco-friendly aircraft taxiing approach with collision and conflict avoidance. *Transportation Research Part C: Emerging Technologies, 121*, 102872.

Tindemans, B. (2021). A greedy approach to the minimisation of deviations of the dynamic vehicle routing problem with electric taxiing systems. *Master of Science Thesis.*

van Baaren, E., & Roling, P. C. (2019). Design of a zero emission aircraft towing system. In *Aiaa aviation 2019 forum* (p. 2932).

van Oosterom, S., Mitici, M., & Hoekstra, J. (2023). Dispatching a fleet of electric towing vehicles for aircraft taxiing with conflict avoidance and efficient battery charging. *Transportation Research Part C: Emerging Technologies, 147*, 103995.

van Winkel, C. (2023). Tactical TaxiBot planning at amsterdam airport schiphol under uncertainty. *Master of Science Thesis.*

VINCI Airports. (2024, January 16). *Vinci airports – traffic at 31 december 2023*. Retrieved from `https://www.vinci.com/commun/communiques.nsf/6D529D6EECC7BC2AC1258AA6003CE67E/\protect\T1\textdollarfile/vinci-airports--traffic-31-december-2023.pdf`

Wollenheit, R., & Mühlhausen, T. (2013). Operational and environmental assessment of electric taxi based on fast-time simulation. *Transportation research record, 2336*(1), 36–42.

Zbakh, D., Benhadou, M., Benkacem, A., & Lyhyaoui, A. (2023, September). Airport traffic optimisation and airdrome analysis using mathematical modelling. *Indonesian Journal of Electrical Engineering and Computer Science, 31*(3), 1744.

Zhang, M., Huang, Q., Liu, S., & Li, H. (2019a). Assessment method of fuel consumption and emissions of aircraft during taxiing on airport surface under given meteorological conditions. *Sustainability, 11*(21), 6110.

Zhang, M., Huang, Q., Liu, S., & Li, H. (2019b). Multi-objective optimization of aircraft taxiing on the airport surface with consideration to taxiing conflicts and the airport environment. *Sustainability, 11*(23), 6728.

Zoutendijk, M., Mitici, M., & Hoekstra, J. (2023). An investigation of operational management solutions and challenges for electric taxiing of aircraft. *Research in Transportation Business & Management, 49*, 101019.

Schiphol
Group

TU Delft
Delft
University of
Technology