

Maximising Taxibot Throughput under Mixed Ground Operations: A Case Study of Departing Traffic at Brussels Airport

Emile Schoemans

Student number: 02006755

Supervisor: Prof. Dries Goossens

Master's dissertation submitted in order to obtain the degree of Master of Science in Business Engineering - Data Analytics

Academic year 2025-2026

Confidentiality Statement

This master's dissertation contains confidential information and/or confidential research results proprietary to Ghent University or third parties. It is strictly forbidden to publish, cite or make public in any way this master's dissertation or any part thereof without the express written permission of Ghent University. Under no circumstance may this master's dissertation be communicated to or put at the disposal of third parties. Photocopying or duplicating it in any other way is strictly prohibited.

I declare that the research was conducted in accordance with the rules governing scientific and academic integrity. I have read, and acted in accordance with, the Code of Ethics of the Faculty, the ethical code research with participants and personal data (if relevant) and the guidelines of the faculty about the use of GenAI in the master's dissertation.

Abstract

Aircraft taxiing with the main engines running accounts for a substantial share of an airport's ground emissions, because jet engines are inefficient at the low thrust settings used on the surface. Taxibots, external towing vehicles that pull an aircraft from the gate to the runway with its engines off, can cut taxi-out fuel burn by up to eighty per cent, and operational trials at Frankfurt, Amsterdam Schiphol and Brussels have confirmed their feasibility. Existing scheduling studies, however, ask how large a fleet would be needed to tow every aircraft. This thesis asks the inverse and, for an airport introducing the technology gradually, more practical question: given a small fixed fleet, how many departing aircraft can it serve in a single operational day, and what fuel and CO₂ savings follow? The problem is formulated as a scheduling model that assigns a fixed number of taxibots to departing narrow-body aircraft under release-time and deadline constraints, accounting for the full service cycle including the unproductive return drive, and for the capacity of the start-up zones where towed aircraft restart their engines. It is solved both exactly, as an integer linear programme, and with a greedy construction heuristic in two ordering variants. This model is built on ground movement data for July 2025 at Brussels Airport, provided by Skeyes, and its parameters and operational assumptions were validated with the same partner. A small fleet captures most of the achievable effect. Three taxibots serve about half of the eligible departures on a busy day and five serve two thirds. At three taxibots this avoids roughly a third of the taxi-out CO₂ of the eligible fleet, or twenty-four tonnes. Beyond five vehicles the fleet ceases to be the binding constraint and returns diminish, a saturation point that coincides with the operational maximum identified independently by Skeyes. The heuristic matches the exact model to within one aircraft across this range while computing in milliseconds, making it the appropriate tool for operational use.

Acknowledgements

I would like to express my sincere gratitude to Professor Dries Goossens for his supervision, guidance and valuable feedback throughout the development of this thesis. I would also like to thank Lisa Garcia Tercero for her support, constructive feedback and assistance during the research process.

I am especially grateful to François Claude and Luc Decock from Skeyes for their time, operational expertise and valuable insights into ground operations and the practical implementation of taxibot operations at Brussels Airport. Their input was essential for understanding the operational context of this research and for validating several of the assumptions used in the model.

Finally, I would like to thank everyone who supported me throughout the completion of this thesis.

Contents

Abstract	3
Acknowledgements	4
Declaration on the Use of Artificial Intelligence	8
1 Introduction	13
2 Literature Study	16
2.1 The airport ground movement problem	16
2.2 Engine-off taxiing technologies and their measured benefits	16
2.3 Scheduling and dispatching of towing vehicle fleets	18
2.4 Assignment and scheduling methods	20
2.5 Synthesis and research gap	22
3 Problem Description	24
3.1 Problem Overview	24
3.2 Operational Setting	25
3.3 Airport Surface Network	26
3.4 Taxibot Service Cycle	27
3.5 Start-up Zones at the Detachment Points	28
3.6 Key Problem Characteristics	30
3.7 Assumptions and Scope Delimitation	31
3.8 Scenario Design	32
4 Data Description	34

4.1	Data Source and Context	34
4.2	Ground Movement Dataset	34
4.3	Aircraft Types and Taxibot Eligibility	36
4.4	Airport Surface Network and Traversal Times	37
4.5	Data Limitations	39
4.6	Derived Variables	40
5	Methodology	42
5.1	Solution Approach Overview	42
5.2	Mathematical Formulation	43
5.2.1	Sets, parameters and decision variables	43
5.2.2	Objective function	45
5.2.3	Constraints	46
5.2.4	Implemented model and solution protocol	49
5.3	Scheduling Heuristics	49
5.3.1	Algorithm	50
5.3.2	Ordering variants	50
5.3.3	Constraints enforced and comparison protocol	51
5.4	Parameter Setting	51
5.5	Scenario Design	53
5.6	Fuel and Emissions Comparison	54
6	Results	56
6.1	Experimental setup	56
6.2	Throughput versus fleet size	56

6.3	Heuristic versus exact model	58
6.4	The start-up pool	59
6.5	Sensitivity analysis	60
6.6	Fuel and CO ₂ effect	63
6.7	Summary of findings	64
7	Discussion	65
7.1	Interpretation of the main findings	65
7.2	Position relative to the literature	66
7.3	Operational feasibility at Brussels Airport	66
7.4	Limitations	67
8	Conclusion and future work	69
8.1	Conclusion	69
8.2	Future work	70
	References	71
	Appendix	77

Declaration on the Use of Artificial Intelligence

Generative artificial intelligence tools, in particular ChatGPT and Claude, were used during the preparation of this thesis as supporting tools. They were used to assist with language editing and clarity, LaTeX formatting, and the conversion of existing thesis content into LaTeX. They were also used as supporting tools during programming, including assistance with developing and debugging functions and with the creation and formatting of visualisations.

All research decisions, model design, data analysis, calculations, interpretation of results, and conclusions remain the work of the author. AI-generated suggestions and code were reviewed, tested and, where necessary, modified or rejected by the author. The author takes full responsibility for the accuracy, originality, and final content of this thesis.

List of Figures

1	Stakeholder-objective matrix for airport ground operations (based on personal communication with Skeyes).	25
2	The airport surface network used in my model, overlaid on the aerodrome chart of Brussels Airport, with the start-up zones marked. . . .	27
3	Phases of the taxibot service cycle for taxi-out (adapted from Postorino et al., 2017).	28
4	Location of the two start-up zones at the runway 25R threshold: the B1 zone and the W42 platform.	29
5	Simplified illustration of an airport surface graph, with gate nodes, taxiway nodes and runway entry. Created with Canva.	39
6	Aircraft served against fleet size, pool model, 9 July 2025. Open markers indicate ILP solutions that reached the time limit.	56
7	Taxibot service cycles over the operational day, greedy pool model, $K = 3$, 9 July 2025.	59
8	Sensitivity of aircraft served to the towing speed factor, $K = 3$, 9 July 2025.	61
9	Sensitivity of aircraft served to the drive-back time, pool model, greedy EDF, 4 July 2025.	61
10	Sensitivity of aircraft served to the engine warm-up time, $K = 8$, 9 July 2025.	62
11	Net CO ₂ saved against fleet size, in kilograms per day (left) and as a percentage of the eligible baseline (right), 9 July 2025.	63
12	A.7: Total departures, taxibot-eligible departures and aircraft served by the greedy EDF heuristic at $K = 3$, per day, July 2025 (top), and the served share of eligible departures (bottom).	76
13	A.8: Net CO ₂ saved per day at $K = 3$, first week of July 2025, for the pool and no-pool configurations with a diesel and an electric taxibot fleet.	77
14	A.9: Per-aircraft ILP schedule on 04/07/2025 (part 1).	78

15	A.9: Per-aircraft ILP schedule on 04/07/2025 (part 2).	79
----	--	----

List of Tables

1	Fuel consumption and CO ₂ emissions by taxiing method for a narrow-body aircraft over a 15-minute taxi-out. Adapted from Camilleri & Batra (2021).	17
2	Columns of the ground movement dataset used in my model.	35
3	Taxi/idle fuel flow per engine at the 7% thrust setting (European Union Aviation Safety Agency, 2026).	36
4	Symbols used in the network representation.	38
5	Variables derived from the ground movement dataset. All times are in seconds.	41
6	Sets, parameters and decision variables of the optimisation model. . . .	43
7	Base parameter values used in the fleet-size sweep.	52
8	Share of pool-eligible departures served, per day and fleet size, first week of July 2025. N is the number of pool-eligible departures that day, an absolute count; all other columns are percentages of N	57
9	Greedy heuristic against the exact model, pool configuration, 9 July 2025, ILP time limit 180 s per instance.	58
10	A.1: Greedy heuristic against the exact model, 4 July 2025, configuration without start-up zones, ILP time limit 3600 s per instance.	74
11	A.2: Aircraft served by the greedy EDF heuristic per day and fleet size, first week of July 2025, configuration without start-up zones. $K = 0$ is the conventional baseline in which no aircraft is towed.	74
12	A.3: Aircraft served by the greedy heuristic in release-time order per day and fleet size, first week of July 2025, configuration without start-up zones. N is the number of taxibot-eligible departures that day. . . .	75
13	A.4: Sensitivity of aircraft served to the B1 start-up zone capacity, pool model, $K = 3$, W42 capacity 2, 9 July 2025, of 205 pool-eligible departures. . . .	75
14	A.5: Sensitivity of aircraft served to the engine warm-up time, $K = 3$, 9 July 2025.	75

15	A.6: Sensitivity of aircraft served to the drive-back time and to the W42 capacity, pool model, greedy EDF, 4 July 2025, of 211 pool-eligible departures.	76
----	---	----

1 Introduction

Aviation is one of the harder sectors to decarbonise (International Council on Clean Transportation, 2022). Most public and research attention is directed at the airborne phase of flight, such as the cruise, the climb, and the search for sustainable fuels, yet a substantial and often overlooked share of an aircraft’s emissions is produced while it is still on the ground. The reason is a basic mismatch: jet engines are optimised for the high-thrust conditions of cruise, which makes them disproportionately inefficient at the very low thrust settings used while taxiing (Nikoleris et al., 2011). An engine idling its way from the gate to the runway therefore burns fuel out of proportion to the small amount of movement it produces. As a result, ground movements account for roughly 5–20% of an airport’s total carbon footprint depending on its size and traffic mix (Camilleri & Batra, 2021), and for a narrow-body aircraft such as the Airbus A320 the taxi phase alone can represent 5–10% of total fuel burn (Postorino et al., 2017). Because this fuel is consumed at low speed on a confined surface rather than in the air, ground operations are also one of the more tractable places to intervene: meaningful gains can be made here without redesigning the aircraft, the engine, or the fuel.

To see where these savings come from, it helps to look first at how an aircraft currently reaches the runway. In conventional operations a pushback tractor pushes the aircraft back from its stand, after which it taxis to the runway under its own power with the main engines running at roughly 7% idle thrust (Khadilkar & Balakrishnan, 2012). For the narrow-body twin-engine aircraft that dominate European traffic this means both engines run for the entire taxi-out, a practice referred to throughout this thesis as dual-engine taxiing. A partial improvement is single-engine taxiing, where one engine is shut down for most of the taxi, but this adds crew workload and reduces manoeuvrability and is not always practical (Camilleri & Batra, 2021).

A more far-reaching alternative is to take the engines out of the taxi process altogether. A taxibot is an external towing vehicle that, unlike an ordinary pushback tractor, tows the aircraft from pushback at the gate along its entire route to a release point near the runway with the main engines switched off, while the pilot keeps directional control from the cockpit (Lukic et al., 2019). The engines are started and warmed up only once the aircraft has been released near the threshold. The most established example, the TaxiBot developed by Israel Aerospace Industries and TLD, has been certified by EASA for Boeing 737 and Airbus A320 operations and has been in commercial service at Frankfurt Airport since 2014 (Lukic et al., 2019). A fully electric variant, the e-taxibot, removes the towing vehicle’s own ground-level emissions, leaving only the upstream emissions from generating the electricity used to charge it. Replacing

dual-engine taxiing with taxibot towing can cut taxi-out fuel consumption by up to roughly 80% (Camilleri & Batra, 2021), and it is against dual-engine taxiing as the baseline that every alternative in this thesis is measured.

My thesis studies the deployment of such a fleet at Brussels Airport, in collaboration with Skeyes, the Belgian air navigation service provider that manages the airport’s air traffic and provided the operational ground movement data used here. Brussels Airport, located in Zaventem about 12 km northeast of the city centre and identified by the ICAO (International Civil Aviation Organization) code EBBR (IATA BRU), is Belgium’s principal international airport and handles the large majority of the country’s scheduled flights, with more than 112,000 departing flights recorded in 2017 (Bous-sauw & Vanoutrive, 2019). It is a relevant case for three reasons. First, its traffic is dominated by exactly the A320- and 737-family aircraft for which taxibot certification exists, so a large share of its movements is in principle eligible for towing. Second, the Skeyes collaboration gives access to real operational data, including gate assignments, pushback and taxi timestamps, taxiway routes and aircraft types, allowing this model to be calibrated against genuine movements rather than synthetic assumptions. Third, Brussels is a mature airport whose layout, like that of most existing airports, predates any notion of towing operations; whether and where the taxibot concept fits infrastructure that was never designed for it is a question with relevance far beyond this one airport, and it returns in the conclusion of this thesis. Brussels Airport has itself already tested a single TaxiBot operationally, but no study has examined how a fleet of them should be scheduled there.

The central question is therefore: given a fixed fleet of taxibots, how many departing aircraft can be served in a single operational day, how does this change with the size of the fleet, and how does it translate into fuel and emissions savings compared with the current dual-engine baseline? The number of taxibots is treated as a fixed scenario parameter and varied across scenarios, so that the analysis reveals not only what a given fleet delivers but also where adding vehicles stops paying off. The objective deliberately maximises the number of aircraft served rather than minimising fuel directly: the two are strongly aligned, since every towed aircraft saves fuel, and the throughput objective matches the operational priorities of the air navigation service provider, an argument developed in Section 3. Fuel and CO₂ savings then follow as a direct consequence of the schedule.

Answering this is more involved than counting how many aircraft a fleet could in theory tow. Each taxibot serves only one aircraft at a time. After releasing it near the runway, it must drive back across the airport before picking up the next one, and this drive-back is unproductive time that still limits how often a taxibot can be reused. Every aircraft

has a limited time window: it cannot be served before it is ready to push back, and it must be released early enough to make its departure slot. The space near the runway where aircraft are released and start their engines is limited, so deliveries cannot all be bunched together. And because towed and conventionally taxiing aircraft share the same taxiway network, the schedule must respect the realities of that shared surface. The result is a combinatorial optimisation problem at the meeting point of vehicle scheduling under limited resources and airport surface operations. The scope of this study is bounded to departures from runway 25R, the airport’s dominant departure runway, during the summer season, and to the narrow-body aircraft eligible for towing; the full delimitation and its motivation are given in Section 3.

The thesis makes two contributions. Methodologically, it formulates and solves the fixed-fleet, throughput-maximising taxibot scheduling problem in a mixed-operation setting, a framing that inverts the existing literature, which requires all aircraft to be towed and minimises the fleet needed. Practically, it provides Skeyes with a quantified and operationally validated assessment of what taxibot deployment could deliver at Brussels Airport: the results show that a small fleet already captures most of the achievable effect, serving about half of the eligible departures with three vehicles and avoiding roughly a third of their taxi-out CO₂, with diminishing returns beyond approximately five vehicles.

The remainder of this thesis is structured as follows. Section 2 reviews the literature on engine-off taxiing and towing-fleet scheduling. Section 3 describes the problem and its scope, Section 4 the data, and Section 5 the mathematical formulation and the scheduling models. Section 6 presents the results, Section 7 discusses their operational meaning and limitations, and Section 8 concludes.

2 Literature Study

2.1 The airport ground movement problem

The movement of aircraft between gates and runways is studied in the operations research literature as the airport ground movement problem: the task of routing and scheduling aircraft across a shared taxiway network such that each reaches its runway or gate on time, without conflicts, and at minimum cost in delay or fuel. Atkin et al. (2010) survey this field and show that it has been addressed with a wide range of methods, from mixed-integer programming formulations that route aircraft through a network under separation constraints, to genetic algorithms and other heuristics that trade optimality for computation speed. Their survey also documents a shift in the objectives considered: where early work concentrated on minimising taxi times and delays, later studies increasingly include fuel burn and emissions, reflecting the environmental pressure on the sector. A related line of work improves the inputs to these models; Ravizza et al. (2013), for example, combine a statistical approach with a ground movement model to estimate taxi times more accurately, showing that speeds vary substantially between airports and even between routes at the same airport.

Within this broader problem, this thesis concerns one specific and relatively new resource: the taxibot. Deploying a fleet of towing vehicles adds a scheduling layer on top of the ground movement problem, because the taxibots themselves are scarce, must be assigned to individual departures, and occupy the same taxiway network as conventionally taxiing aircraft. The setting for this study is Brussels Airport, in collaboration with Skeyes, whose operational ground movement data underpins this model developed in later sections; the operational context and the dataset are described in Sections 3 and 4. The remainder of this section reviews what previous studies have found on engine-off taxiing technologies and their benefits (Section 2.2), how fleets of towing vehicles have been scheduled and dispatched in earlier work (Section 2.3), and which general routing and scheduling methods those studies build on (Section 2.4), before identifying the gap this thesis addresses (Section 2.5).

2.2 Engine-off taxiing technologies and their measured benefits

Several technologies have been developed to reduce or eliminate main-engine use during taxiing, and their benefits have been quantified in a series of studies. The simplest option requiring no equipment is single-engine taxiing. Single-engine taxiing is the

practice of shutting down one of the aircraft’s engines and taxiing on the remaining one, so only half the engines burn fuel during the taxi phase. Stettler et al. (2018) analysed its impact using real flight data and found fuel savings of roughly 20–30% on the taxi-out, but also showed why the technique is only situationally usable: the shut-down engine needs several minutes of warm-up before take-off, and the running engine requires higher thrust to break away and turn, which constrains its use at high weight and in confined apron areas. Single-engine taxiing therefore delivers a partial saving at best, which is why research attention has shifted to removing the engines from the taxi process altogether.

Lukic et al. (2019) review the resulting landscape of electric taxiing systems and distinguish two families: on-board systems, in which electric motors are integrated into the landing gear and powered by the APU (such as the WheelTug system and the EGTS demonstrated by Safran and Honeywell), and external systems, in which a towing vehicle moves the aircraft. On-board systems add permanent weight that is carried on every flight, which penalises fuel burn airborne; external towing avoids this, at the cost of introducing a fleet of vehicles that must be managed on the airport surface. The most mature external system is the TaxiBot, certified by EASA for Boeing 737 and Airbus A320 operations and in commercial service at Frankfurt Airport since 2014 (Lukic et al., 2019).

The savings from external towing have been quantified from several angles. Camilleri & Batra (2021) compare the main taxi-out methods for a narrow-body aircraft over a standard 15-minute route; their figures, summarised in Table 1, show that a diesel taxibot cuts fuel burn by roughly 80% against the dual-engine baseline, and that a fully electric taxibot removes ground-level consumption entirely, leaving around 11 kg of CO₂ in upstream grid emissions per operation.

Table 1: Fuel consumption and CO₂ emissions by taxiing method for a narrow-body aircraft over a 15-minute taxi-out. Adapted from Camilleri & Batra (2021).

Taxiing method	Fuel consumed (kg)	CO₂ emitted (kg)
Full dual-engine	102.2	322.4
Single-engine	79.9	252.1
Diesel taxibot	19.7	62.3
Electric taxibot (EU grid)	0	~11.3

Postorino et al. (2017) go beyond per-operation figures by simulating taxibot deployment at Bologna Airport for a full traffic sample. Two of their findings shape this thesis directly. First, only 61% of movements were considered eligible for taxibot service, because of departure-delay risk and airline participation, which means any realistic model must handle mixed operations in which towed and conventionally taxiing aircraft share

the network. Second, a taxibot is tied up for considerably longer than the tow itself. It also contains coupling, decoupling and the return drive. Savings estimates built on towing time alone therefore overstate what a fleet can actually deliver.

Finally, the technology has moved from simulation studies into operational trials, including in the exact setting of this thesis. Similar trials have been conducted at Amsterdam Schiphol, where the TaxiBot has been tested as part of the airport’s ambition to make sustainable taxiing a standard procedure (Schiphol, 2021). In March 2024, as part of the European Stargate programme, Brussels Airport and TUI fly ran a three-month operational trial of a biodiesel TaxiBot, the first in Belgium (Brussels airport, 2024). The airport reported reductions of up to 85% in emissions (CO₂, NO_x and ultrafine particles) and up to 60% in noise during towed taxiing, with a towed trajectory consuming around 3 litres of biodiesel in the vehicle against roughly 40 litres of jet fuel that the aircraft’s own engines would have burned. Both trials demonstrated operational feasibility, but each involved a single vehicle and neither addressed how a fleet should be scheduled across a day of departures, the question taken up in the next section and in this thesis.

2.3 Scheduling and dispatching of towing vehicle fleets

A smaller but rapidly growing body of work addresses the question that follows once towing technology is proven: how a fleet of towing vehicles should be operated. Zoutendijk, Mitici & Hoekstra (2023) survey this literature and classify the existing models into four groups: vehicle routing models for the tugs themselves, models assigning vehicles to flights, fleet sizing models, and battery charging optimisation models. Their review also identifies the field’s main open challenge as robustness: towing schedules are sensitive to flight delays, and most models to date assume a deterministic schedule known in advance.

The earliest operational studies focused on feasibility and system design. Sirigu et al. (2018) develop routing algorithms for a fleet of autonomous tugs, in which aircraft are towed to the runway closest to the apron, the paved area around the terminal where aircraft are parked, loaded and serviced, and cover the final segment on their own engines, a structure that anticipates the detachment-point logic used in later work. van Baaren & Roling (2019) design a zero-emission towing system for Schiphol and estimate the fleet needed to serve its traffic, and Salihu et al. (2021) build a simulation framework for electrified taxiing that quantifies the congestion and cost effects of adding towing vehicles to the taxiway system, showing that the vehicles themselves become a source of surface traffic that must be managed.

The most complete optimisation treatment to date is by van Oosterom et al. (2023), who propose an end-to-end framework for dispatching electric towing vehicles at a large airport. Their mixed-integer formulation integrates three layers that earlier work treated separately: routing of the vehicles across the taxiway network with minimum separation maintained at all times, assignment of vehicles to towing tasks, and scheduling of battery recharging. Applied to a full day at Amsterdam Schiphol with 913 flights, their model finds that a fleet of 38 vehicles is required to tow all aircraft, with each vehicle towing for around four hours per day and charging demand concentrated just before the evening peak. In related work, van Oosterom & Mitici (2022) show how strongly the required fleet depends on battery technology: 41, 29 or 24 vehicles are needed under battery capacities of 100, 320 and 500 kWh respectively. Zoutendijk et al. (2023b) extend the framework to disruption management, proposing a model that reassigns vehicles when flights are delayed while keeping the schedule stable for operators; on a busy Schiphol day, 55% of vehicle assignments could be kept unchanged between time steps without reducing the number of aircraft towed.

A further operational insight from these studies concerns the interface between towing and the runway itself. Delivering aircraft to the runway faster than it can launch them saves nothing: it merely relocates the queue to the runway entry area, where space is limited, and the aircraft burn fuel waiting with engines running after detachment. Farhadi et al. (2014) make the general point that the efficiency of runway operations is largely determined upstream, by how aircraft are assigned and fed toward the runway. For towing operations specifically, Postorino et al. (2017) address this through the position of the detachment point, the location where the vehicle releases the aircraft and the engines are started for the final segment. Its placement is a design trade-off: closer to the runway maximises engine-off time per aircraft but concentrates traffic in the constrained area near the threshold, while further away eases congestion at the cost of savings. This trade-off directly motivates the holding-pool construct near the runway in the model of this thesis, described in Section 3.

Two characteristics of this literature stand out for the present thesis. First, nearly all of it is built on Amsterdam Schiphol and centres on electric vehicles, with battery charging as a core modelling concern. Second, and more fundamentally, these studies pose the fleet-sizing question in one direction: they fix the requirement that all aircraft are towed and minimise the fleet needed to achieve it. This thesis poses the inverse and, for an airport introducing taxibots gradually, arguably more practical question: given a small fixed fleet, how many departing aircraft can be served? That framing makes mixed operations, in which most aircraft still taxi conventionally, the central feature of the problem rather than an edge case, and it connects the towing literature above to the assignment and scheduling methods reviewed next.

2.4 Assignment and scheduling methods

The towing studies of the previous section build on a longer tradition of vehicle routing and scheduling research, from which this thesis draws its solution methods. The structural parent of the taxibot problem is the Vehicle Routing Problem with Time Windows (VRPTW), in which a limited fleet serves customers who each require service within a given time interval. Solomon (1987) introduced the standard construction heuristics for this problem, including insertion methods in which a customer is added to a route only if it can be served within its own window without pushing any later customer past its deadline, and found that heuristics weighing both distance and time outperform purely distance-based ones. The problem is computationally hard in a strong sense: even finding a feasible solution with a fixed fleet size is NP-complete (Savelsbergh, 1985). This hardness result, together with the extensive evidence collected in later overviews (Bräysy & Gendreau, 2005; Cordeau et al., 2002) that construction heuristics reach high-quality solutions at a fraction of the computational cost, is the direct justification for the heuristic solution approach in this thesis. The taxibot problem maps onto the VRPTW with two twists: each “customer” (a departing aircraft) requires a full attach-tow-detach-return cycle rather than a stop, and unserved customers are permitted, since aircraft without a taxibot simply taxi conventionally.

For exact approaches, the closest analogues come from airport scheduling. Roling & Visser (2008) formulate airport surface traffic planning as a mixed-integer linear programme that routes and times aircraft over the taxiway network to minimise taxi delay, and demonstrate it on a generic airport layout; their work established MILP as workable for surface planning but also showed how quickly instance size becomes limiting. Gupta et al. (2009) apply a MILP to departure scheduling, sequencing aircraft near the runway under separation constraints. The most transferable finding is from Farhadi et al. (2014), who study runway capacity management at Doha and show that optimising the assignment of aircraft to resources while keeping the sequencing rule simple, such as first-come-first-served, already yields near-optimal solutions. This supports the design of the model in this thesis, where assignment is the optimised decision and sequencing follows chronologically, since optimising both jointly for a full day is computationally prohibitive.

A recurring modelling challenge across these studies is preventing conflicts between vehicles that share the taxiway network. The literature handles this in a small number of standard ways. The most common is to track occupation intervals per network segment and impose that no two vehicles occupy the same segment simultaneously, formulated in exact models as disjunctive either-or constraints linearised with big-M

terms and ordering variables (Postorino et al., 2017; Roling & Visser, 2008). A second approach enforces conflict-freedom through the routing itself: van Oosterom et al. (2023) integrate minimum separation distances directly into the path planning of their towing vehicles, so that generated routes are conflict-free by construction rather than checked afterwards. Sequencing-based methods form a third family, in which conflicts at a shared resource are resolved by ordering the users and imposing minimum time gaps between successive ones, as Gupta et al. (2009) do for departures at the runway. These approaches trade off differently: interval-based constraints are general but add many binary variables, which contributes to the poor scaling of exact models, while separation-aware routing and sequencing rules scale better but constrain the solution space. The model in this thesis follows the interval-based logic, enforced procedurally within a heuristic rather than through binary variables, precisely to avoid that scaling cost.

Two further strands inform my model’s structure. The fleet-sizing literature examines the marginal value of additional vehicles; Farhadi et al. (2014) document the typical pattern of strong initial gains followed by diminishing returns as capacity grows, after which performance depends more on scheduling strategy than on further capacity. The same pattern is expected across taxibot fleet sizes, and identifying the point where it sets in at Brussels is one of the aims of the scenario analysis. Resource-constrained scheduling theory (Brucker et al., 1999) explains why: when a resource is scarce, the bottleneck is availability itself, and once it is abundant the binding constraints shift elsewhere, here to the aircraft time windows. Postorino et al. (2017) confirm the practical relevance of this for towing operations specifically, showing that the unproductive return leg makes the true capacity of a towing vehicle considerably smaller than towing-time arithmetic suggests, which is why the full service cycle, including the drive-back, is modelled explicitly in this thesis.

The order in which a construction heuristic considers its jobs is itself a design choice, and two priority rules are standard in the scheduling literature. The first is to process jobs in the order they become available, which in an airport context corresponds to first-come-first-served. This is the *de facto* discipline in air traffic control and the reference against which alternative sequences are measured, because deviation from the FCFS order is the usual operational measure of equity between operators (Bennell et al., 2011). Farhadi et al. (2014) show, moreover, that keeping such a simple sequencing rule while optimising the assignment already yields near-optimal solutions. The second rule orders jobs by their deadline rather than their availability, and it carries stronger theoretical support. Alshaikhli & Alshaikhli (2024) showed that sequencing by earliest due date minimises maximum lateness on a single machine, and Horn (1974) extended the result to jobs with individual release dates: when preemption is allowed, always

serving the available job with the earliest deadline solves the problem optimally and provides a feasibility test for whether all deadlines can be met at all. The same rule is the classical optimality result in real-time systems, where earliest-deadline-first is guaranteed to meet every deadline whenever any schedule can (Liu & Layland, 1973). These guarantees do not carry over directly to the problem of this thesis, which has several parallel servers, no preemption, and a service duration that includes an unproductive return leg; the deadline rule therefore becomes a heuristic rather than an exact method. Its motivation nonetheless remains the one the theory identifies: aircraft whose windows close soonest are committed first, so that a taxibot is not occupied by a less urgent aircraft at the moment a tightly constrained one becomes ready. Comparing the two rules on the same instances quantifies what the operationally natural order costs relative to the deadline-driven one, and this comparison is reported in Section 6.

2.5 Synthesis and research gap

Three conclusions follow from the literature reviewed above. First, engine-off taxiing works and its benefits are well quantified: external towing cuts taxi-out fuel burn by roughly 80% compared with dual-engine taxiing (Camilleri & Batra, 2021), and operational trials at Frankfurt, Schiphol and Brussels have confirmed feasibility and savings in practice (Brussels airport, 2024). Second, the operation of towing fleets has begun to receive dedicated modelling attention, with mixed-integer formulations that integrate vehicle routing, assignment and battery charging, demonstrated almost exclusively at Amsterdam Schiphol (van Oosterom et al., 2023); Zoutendijk et al. (2023a). Third, the underlying methods are mature: the VRPTW literature supplies construction heuristics whose quality is well documented (Bräysy & Gendreau, 2005; Solomon, 1987), airport MILP studies show that exact optimisation is workable on constrained instances (Gupta et al., 2009; Roling & Visser, 2008), and the finding of Farhadi et al. (2014) that near-optimal solutions are reached by optimising assignment while keeping sequencing simple provides direct design guidance.

The gap sits at the junction of these strands. The Schiphol studies pose the fleet question in one direction: tow every aircraft and minimise the fleet required, with battery management at the core. No published study answers the inverse question that an airport introducing taxibots actually faces: given a small, fixed fleet, how many departing aircraft can be served in a day, when the majority of traffic continues to taxi conventionally on the same network? This mixed-operation setting is supported by the eligibility evidence of Postorino et al. (2017), who found that only 61% of movements at Bologna qualified for towing, yet it appears in existing models only as a boundary

case. Moreover, no scheduling study of any kind has been applied to Brussels Airport, where the TaxiBot has been trialled operationally (Section 2.2) but never studied as a fleet scheduling problem (F. Claude, personal communication, December 17, 2025). This thesis addresses both gaps: it formulates the fixed-fleet, throughput-maximising scheduling problem with mixed operations, and applies it to real ground movement data from Brussels Airport.

3 Problem Description

3.1 Problem Overview

This section describes the optimisation problem addressed in this thesis. The goal is to determine how a fixed fleet of taxibots can be assigned and scheduled to serve departing aircraft at Brussels Airport, maximising the number of aircraft towed by taxibots rather than taxiing under their own engine power.

The problem is set in a mixed-operation context: not all departing aircraft are served by taxibots. Aircraft that are not assigned to a taxibot continue to taxi conventionally using engine power. The objective is to maximise the total number of aircraft served by the taxibot fleet in a single operational day. The environmental benefit, reduced fuel consumption and CO₂ emissions, is not a separate objective but a direct effect of this choice: every aircraft that is towed instead of taxiing on its own engines saves fuel, so serving more aircraft saves more fuel. The size of that effect is quantified in the fuel comparison in Section 5.

The choice to maximise the number of aircraft served, rather than to minimise fuel consumption directly, deserves motivation, since both are defensible formulations. Three arguments support the throughput objective. First, the two objectives are strongly aligned: fuel savings scale with the number of towed aircraft, so a schedule that serves more aircraft saves more fuel, and the cases where they diverge (assigning a taxibot to one long route instead of two short ones) are marginal at the fleet sizes considered here. Second, the operational stakeholders weigh objectives differently, and for the air navigation service provider the dominant priorities are safety and a smooth, orderly flow of traffic on the ground rather than fuel as such (L. Decock, personal communication, April 21, 2026). Figure 1 shows a stakeholder-objective matrix provided by Skeyes, setting the main airport users against the objectives of smooth flow, safety, fuel saving, emissions and noise; it illustrates that fuel minimisation is primarily an airline objective, while the party that would operate and dispatch the taxibots is oriented toward flow. A throughput objective matches the perspective of the intended user of this research. Third, these objectives interact: more aircraft means more passengers, baggage and surface traffic, which raises congestion and with it the risk of delay and reduced safety. A formulation that maximises served aircraft while respecting hard time-window and separation constraints keeps that interaction explicit, whereas a pure fuel objective would leave flow entirely unmodelled. A formulation that minimises taxi fuel directly remains a valid alternative and is noted as future work.

	Smooth flow (minimal delays)	Safety	Fuel saving	Emissions: CO ₂ reduction	Emissions: particulate matter	Noise
Airport operator						
Airlines						
Ground handlers						
Air traffic control (tower)						

Figure 1: Stakeholder–objective matrix for airport ground operations (based on personal communication with Skeyes).

3.2 Operational Setting

This model is applied to Brussels Airport during its summer season, the busiest operational period of the year. Brussels Airport is Belgium’s principal international airport and handles the large majority of the country’s scheduled flights (see Section 1). All departures in this model use runway 25R, the dominant departure runway in the standard westerly configuration (see Section 1). Towed aircraft are delivered to a holding pool near the runway where they start and warm up their engines before lining up for departure independently; this pool is described in Section 3.5.

The operational data is provided by Skeyes under a non-disclosure agreement. The dataset contains departure records including pushback clearance times, actual taxiway paths, gate positions, aircraft types and take-off timestamps; a full description follows in Section 4. My model considers a single operational day for which the set of departing aircraft is known in advance, including each aircraft’s gate area, departure window and taxiway route. The number of taxibots available is a fixed input that is varied across scenarios ($K = 1, 2, \dots, 8$) to analyse how fleet size affects the number of aircraft served.

3.3 Airport Surface Network

The physical layout of the airport’s taxiway system is represented as a network of nodes and connections: nodes correspond to relevant locations on the airport surface (apron areas, taxiway intersections, holding points, runway entry points and the detachment point near runway 25R), and connections correspond to the taxiway segments between them. Figure 2 shows this network overlaid on the airport layout, with the departure runway and the holding pool location marked. Each connection has a traversal time based on its physical length and the speed of the vehicle using it: the standard taxi speed for engine-powered aircraft, a slower towing speed for taxibot-towed aircraft, and the taxibot’s own speed for the drive-back. Some nodes in the network do not correspond to operational locations but are placed along curved taxiway segments purely to make the computed distances follow the actual geometry of the taxiways; they play no role beyond accurate distance calculation.

An important modelling choice is that each aircraft’s route through this network is taken as given: This model uses the taxiway path each aircraft actually followed, as recorded in the data, rather than choosing routes itself. This is deliberate. Actual routes reflect operational practice at Brussels Airport, including standard taxi routings and controller preferences that a shortest-path calculation would not reproduce, so re-optimising routes would make the modelled traffic less realistic, not more. It also keeps the problem focused: the decision studied here is which aircraft each taxibot serves and when, not which way they drive. Interactions that the given routes do not capture, such as crossings of an active runway and interference from arriving traffic, are represented through penalty times added to the affected connections, described in Section 4.



3.4 Taxibot Service Cycle




Figure 3: Phases of the taxibot service cycle for taxi-out (adapted from Postorino et al., 2017).

From a scheduling perspective, the essential feature of this cycle is that it fully occupies the taxibot: attachment, towing, detachment and the drive-back must all be completed before the next service can begin. The drive-back deserves emphasis. During the return trip the taxibot serves no aircraft, but the time is unavoidable, and ignoring it would produce schedules in which a taxibot is assigned to a gate before it could physically have returned from its previous job. The total occupied time per service is therefore considerably longer than the towing alone, which is what makes the fleet a genuinely limited resource: even when many aircraft are ready to depart, the number that can be served in a given period is bounded by how quickly the taxibots complete their full cycles and become available again.

3.5 Start-up Zones at the Detachment Points

Aircraft towed by a taxibot are delivered to a start-up zone near the runway 25R threshold, where they stand clear of the active taxiway while starting and warming up their engines before joining the departure sequence. Such zones are not an invention of this model; remote holding and engine run-up positions near runway thresholds exist at major airports, and the taxibot literature places the detachment point at exactly

such a position to keep engine start-up from blocking through traffic (Postorino et al., 2017). The recorded runway entry points support this placement: over 95% of the 25R departures in the data enter the runway via the B1 or W41 holding points, so my model uses two zones, one adjacent to the B1 flow, located on runway 19 next to 25R, and one on the W42 platform next to W41. Figure 4 shows the location of both zones on the airport surface.




Figure 4: Location of the two start-up zones at the runway 25R threshold: the B1 zone and the W42 platform.

The capacities of the two zones were validated with Skeyes (F. Claude, personal communication, 2026). The W41/W42 platform can physically accommodate up to four aircraft, and offers the taxibot a return path between the aircraft lines; it is represented in this model as a single start-up zone of capacity four. The B1 zone is assigned a capacity of two aircraft as a physical upper bound. Skeyes notes, however, that a taxibot cannot return from the B1 taxiway without either using the runway, which is operationally unacceptable, or infrastructure works, so the operational feasibility of the B1 zone is more doubtful than that of W42, a point taken up in the discussion.

An aircraft leaves its zone toward the runway as soon as its engines are started and warmed up and the departure sequence allows it; a zone does not need to be full before aircraft are released. The capacity works in the other direction: when a zone is

occupied to capacity, no additional towed aircraft can be delivered to it until a position frees up, which can delay taxibot cycles and, through them, the whole schedule. Both capacities are varied in the sensitivity analysis.

The purpose of the start-up zones follows from the mixed operation. A towed aircraft arrives at the runway with cold engines and needs approximately three minutes of engine start and warm-up before it can line up, whereas a conventionally taxiing aircraft arrives with its engines already running and ready for departure. If the towed aircraft warmed up at the runway holding point itself, it would stand in the path of exactly that ready traffic: an aircraft that began taxiing later would have to wait behind an aircraft that cannot yet move. The start-up zone removes this conflict by taking the warming-up aircraft out of the flow, so that ready traffic can pass and enter the runway while the towed aircraft completes its warm-up beside the path rather than on it. The zones are nonetheless a design option rather than a requirement of the concept, and my model is therefore evaluated in two configurations: one in which the taxibot tows each aircraft along its full route to the holding point, where detachment and warm-up take place, and one in which towing ends at the start-up zone, the aircraft warms up there and covers the final segment on its own power. Comparing the two quantifies what the zones contribute, both to the traffic flow and to the taxibot's cycle time, a trade-off presented in the results.

3.6 Key Problem Characteristics

The taxibot scheduling problem described above has the following key characteristics that shape this modelling approach:

- **Assignment under limited resources.** A fixed number of taxibots must be assigned to a subset of departing aircraft. Each taxibot serves one aircraft at a time, and the sequential service cycle bounds the fleet's capacity.
- **Time windows.** Each aircraft has an earliest possible pickup time, derived from its pushback clearance, and a latest time by which its towing must be completed so that engine start, warm-up and line-up still fit before its departure time. Because each aircraft's route is known in advance, its towing duration is known before scheduling, and this latest-completion requirement translates directly into a latest possible start of service. An aircraft that cannot be served within its window taxis conventionally.
- **No overlap on a taxibot.** A taxibot assigned to aircraft i cannot begin serving aircraft j until the full cycle for i , including the drive-back, is complete.

- **Conflicts on shared infrastructure.** Two aircraft cannot occupy the same taxiway segment at the same time, and a minimum time separation applies between successive users of a segment. Nodes that lie physically close together form mutual-exclusion zones that cannot be occupied simultaneously. Limited queuing is allowed on the final segments approaching the runway, reflecting the real holding pattern before departure.
- **Pool capacity.** The holding pool blocks new deliveries when occupied to capacity.
- **Mixed operations.** Aircraft not assigned to taxibots taxi on their own engines over the same network, creating interactions between the two groups.

A mathematical formulation of this problem, including the objective, decision variables and constraints, is presented in Section 5.

3.7 Assumptions and Scope Delimitation

My model operates under the following assumptions, each with its motivation:

- **Single runway configuration.** All departures use runway 25R. The wind at Brussels Airport blows predominantly from the southwest and the airport’s Preferential Runway System is built around this, so under normal conditions the westerly runways are the most frequently used, with 07L/25R as the primary runway; alternative configurations are activated mainly under strong easterly or southerly winds (Brussels Airport Traffic Control, n.d.). Modelling 25R therefore captures the configuration in which the airport operates most of the time, while one runway keeps the problem manageable.
- **Departures only.** Arriving aircraft are not modelled individually. Including them would require arrival data that is not part of the dataset and would roughly double the traffic to be conflict-checked; their effect on departing traffic is instead approximated through penalty times at runway crossing points and periodic delays along routes, which captures the delay they cause without modelling their movements.
- **Routes as given.** Each aircraft follows its recorded taxiway path, for the realism reasons set out in Section 3.3.

- **Fixed service durations.** Attachment, detachment and drive-back are constants, following the values used in the taxibot scheduling literature (see Section 5.4); their route-to-route variation is small compared with the towing time itself.
- **Towing speed from literature.** Taxibots tow more slowly than aircraft taxi on their engines, to limit stress on the nose landing gear (TLD Group, n.d.). The towing speed is set via a speed factor taken from published taxibot studies (Salihu et al., 2021), detailed in Section 5.4, and varied in sensitivity analysis.
- **Narrow-body aircraft only.** Taxibot service is restricted to A320- and B737-family aircraft, the categories for which TaxiBot certification exists; wide-body aircraft appear only as conventionally taxiing traffic.
- **Known schedule.** The departure schedule is deterministic and known in advance. This model answers a planning question, how many aircraft a fleet can serve on a given day, for which the realised schedule is the natural input; handling delays and disruptions in real time is a separate problem noted as future work.
- **Emissions as an effect.** Fuel and CO₂ savings follow from the schedule and are quantified in the fuel comparison in Section 5; they are not part of the objective, for the reasons given in Section 3.1.
- **Weather not considered.** Weather-driven operational changes are excluded, consistent with the single-configuration assumption.
- **Single-airport scope.** This model treats Brussels Airport in isolation and does not capture how delays or aircraft rotations would propagate to connecting flights or other airports in the network, effects that are known to be substantial (Pyrriotis et al., 2013). This keeps the analysis focused on the ground operation itself.

3.8 Scenario Design

My model is run for fleet sizes $K = 1$ to 8, and each scenario is compared against the conventional baseline in which all aircraft taxi on dual engines via the same routes. Results are reported as the number of aircraft served, the proportion of departures covered, and the resulting fuel and CO₂ savings. The distinction between diesel and electric taxibots does not affect the schedule, since both operate identically in this model; it enters only in the fuel comparison, where the taxibot’s own consumption

differs by propulsion type. Because no fuel consumption data is available for the individual movements in the dataset, this comparison is based on consumption values from the literature and is therefore an approximation; the estimation approach is described in Section 5.

4 Data Description

4.1 Data Source and Context

The data used in this thesis is provided by Skeyes, which plays a dual role in this research: as the operational partner that provided the ground movement data used to calibrate and validate this model, and as the intended end user of the research findings. The data was made available under a non-disclosure agreement following a collaboration established in December 2025 (F. Claude, personal communication, December 17, 2025).

Three datasets are used: a ground movement dataset with one record per departing aircraft (Section 4.2), a reference dataset of the most commonly used aircraft types at Brussels Airport (Section 4.3), and a manually constructed representation of the taxiway network (Section 4.4). The section closes with the data limitations (Section 4.5) and the variables derived from the raw data for this model (Section 4.6).

4.2 Ground Movement Dataset

The primary dataset is a tabular Excel file containing one record per departing aircraft movement at Brussels Airport. Each record captures the aircraft’s identity, type, gate position, taxiway route, and a sequence of timestamps marking key moments in the departure process. The full dataset contains 27 columns; those used in my model are listed in Table 2, grouped into identification, aircraft, route and time fields. Columns not used in this model, such as weight and wake turbulence categories, redundant time fields and parking-related fields, are omitted.

Table 2: Columns of the ground movement dataset used in my model.

Column	Description	Role in model
ACT. DATE	Date of the departure	Selection of operational days
A/D	Arrival/departure indicator	Filter on departures
CALLSIGN	Flight identifier (e.g. AEF393B)	Identification, unique per date
ACFT	ICAO aircraft type code (e.g. A320)	Taxibot eligibility and fuel comparison
STAND	Parking position (e.g. 234R)	Mapped to the starting node of the route
APRON	Apron area identifier (e.g. 02NE)	Groups stands to apron-level nodes
RWY	Departure runway	Filter on the 25R configuration
INT	Runway entry intersection (e.g. B1)	Final node before the runway
ACTUAL TAXIWAY PATH	Sequence of taxiway segments (e.g. “S O7 B1E B1 25R”)	Defines the route through the network
PUSH CLR	Pushback clearance time	Release time r_i
BLK TIME	Time the aircraft leaves the stand	Fallback for missing PUSH clr
TAXI CLR	Taxi clearance time	Start of rolling, speed calibration
END TAXI	End of the taxi phase	Speed calibration
T/O TIME	Actual take-off time	Deadline derivation

Three fields play a central role. The **pushback clearance** (PUSH clr) is used as the release time r_i of each departing aircraft i : the earliest moment at which service can begin. It marks the point at which the aircraft is operationally ready to leave, which in the taxibot scenario is the earliest moment a taxibot could begin attachment. Where PUSH clr is missing, the block time is used as a fallback.

The **actual taxiway path** is a string with the ordered sequence of taxiway segments the aircraft followed from gate to runway; these paths are mapped onto the network of Section 4.4 to define the route of each aircraft, from which its towing time is computed.

The **stand** identifies where the aircraft was parked. This model does not distinguish individual stands: stands are grouped by their apron area, and each route starts at the corresponding apron-level node, the first taxiway node of the recorded path. This

simplification is acceptable because the distance from an individual stand to its apron exit is short and similar across stands, and it avoids modelling gate-level geometry that the data does not support.

4.3 Aircraft Types and Taxibot Eligibility

A second dataset lists the most commonly used aircraft types at Brussels Airport (42 unique ICAO type codes). In my model this reference file serves one purpose: determining which aircraft are eligible for taxibot service. Narrow-body A320- and B737-family aircraft, which dominate the Brussels departure mix, are eligible, consistent with current TaxiBot certification; wide-body types such as the A330, B777 and B787 are not, and appear in my model only as conventionally taxiing traffic.

Since the Skeyes data contains no fuel consumption figures, the fuel comparison in Section 5 uses per-type values from the literature. The standard method takes the taxi/idle fuel flow of each aircraft’s engine from the ICAO Aircraft Engine Emissions Databank, which lists fuel flow at the 7% (taxi/idle) thrust setting, and scales it by the number of engines and the taxi duration (Khadilkar & Balakrishnan, 2012). Representative values for the Brussels narrow-body fleet are given in Table 3.

Table 3: Taxi/idle fuel flow per engine at the 7% thrust setting (European Union Aviation Safety Agency, 2026).

Aircraft family	Typical engine	Per engine (kg/min)
A319 / A320CEO	CFM56-5B	~6.2
A319 / A320 / A321CEO	IAE V2500	~7.7
A320NEO / A321NEO	LEAP-1A / PW1100G	~5.0–5.5
B737NG	CFM56-7B	~6.8

These are certification idle figures and should be read as upper estimates: measured taxi fuel burn is typically lower, because engines spend part of the taxi below the 7% setting, and estimates assuming constant 7% thrust have been shown to overestimate observed taxi fuel consumption substantially (Nikoleris et al., 2011). For consistency, the working baseline of the fuel comparison is the aggregate narrow-body value of Camilleri and Batra (2021), with the per-type values of Table 3 used for differentiation between types.

4.4 Airport Surface Network and Traversal Times

The taxiway network is constructed from 199 nodes identified manually on the official aerodrome chart of Brussels Airport (AIM Belgium, 2026), with their positions recorded as pixel coordinates. The nodes comprise apron-level nodes, taxiway intersections, runway entry and holding nodes, the runway itself, and intermediate nodes placed along curved taxiway segments purely for accurate distance calculation. Connections between nodes are derived from the recorded taxiway paths: every consecutive pair of nodes appearing in any path defines a connection, bidirectional by default, yielding 740 directed connections after removing a small number of physically unrealistic direct-to-runway links. Nodes that lie physically close together are grouped into 27 mutual-exclusion zones, used by the conflict constraints described in Section 5.

Each connection receives a physical length: the straight-line distance between its end-point coordinates, converted from pixels to metres with a scale factor calibrated by measuring the known length of runway 25R (3,638 m) on the chart (AIM Belgium, 2026). The traversal time of a connection is then its length divided by a speed, computed separately for each connection and each mode of movement: the engine taxi speed for conventionally taxiing aircraft, the towing speed for taxibot-towed aircraft, and the engine taxi speed again for the empty drive-back. The towing speed is derived from the calibrated engine speed through a speed factor, whose value is set in Section 5.4.

The engine taxi speed is not assumed but calibrated from the data. For each flight, the rolling time is computed as $\text{END TAXI} - \text{TAXI CLR}$. TAXI CLR is used rather than BLK TIME because the interval from BLK TIME to TAXI CLR covers the pushback manoeuvre and engine start, during which the aircraft is not yet taxiing under its own power; including it would inflate the duration and underestimate the true rolling speed. Dividing each flight’s geometric path length by its rolling time gives an empirical speed per flight, and the median across flights is used as the engine taxi speed, the median providing robustness against outliers. Calibrated on the 5,735 usable departures of July 2025, this yields an engine taxi speed of approximately 7.3 m/s (14.1 knots), consistent with published taxi speeds for narrow-body aircraft. Because the node coordinates are extracted manually, a small geometric inaccuracy is unavoidable, but the calibration partly absorbs it: any systematic bias in the measured distances is compensated in the empirical speed.

Connections that cross an active runway receive an additional penalty time reflecting the wait for crossing clearance, applied to all three modes.

Formally, this network is a directed graph $G = (N, E)$, where the node set N contains the apron, intersection, holding, runway and intermediate points, and the edge set E contains the taxiway segments connecting them (Atkin et al., 2010). Each edge $e \in E$ carries three traversal times, one per mode of movement: t_e^{engine} for a conventionally taxiing aircraft, t_e^{taxibot} for a taxibot towing an aircraft, and t_e^{back} for an empty taxibot driving back. Each is the edge's physical length divided by the corresponding speed. Edges that cross an active runway carry an additional penalty on all three modes, representing the wait for crossing clearance.

Table 4: Symbols used in the network representation.

Symbol	Definition
$G = (N, E)$	Directed graph representing the airport surface: nodes N , edges E
N	Set of nodes (apron, intersection, holding, runway and intermediate points)
n	A single node, $n \in N$
E	Set of directed edges (taxiway segments)
e	A single edge, $e \in E$
i	Index of an aircraft
t_e^{engine}	Traversal time of edge e for a conventionally taxiing aircraft (s)
t_e^{taxibot}	Traversal time of edge e for a taxibot towing an aircraft (s)
t_e^{back}	Traversal time of edge e for an empty taxibot driving back (s)
P_i	Route of aircraft i : ordered node sequence $(n_1, \dots, n_m)$ from the data
$E(P_i)$	Edges along the route of aircraft i
τ_i^{tow}	Total towing time of aircraft i along its route (s)

A route for aircraft i is the ordered sequence of nodes

$$P_i = (n_1, n_2, \dots, n_m)$$

taken from the recorded taxiway path. Let

$$E(P_i) = \{(n_1, n_2), (n_2, n_3), \dots, (n_{m-1}, n_m)\}$$

be the edges along it. The towing time of aircraft i is the sum of the towing-mode traversal times over these edges:

$$\tau_i^{\text{tow}} = \sum_{e \in E(P_i)} t_e^{\text{taxibot}}. \quad (1)$$

Because the route is taken as given (Section 3.3), this sum is a constant computed once in the data-preparation stage, which is why τ_i^{tow} enters the scheduling model as a parameter rather than a variable. Different routes between the same apron and runway entry can yield different towing times, so the value is genuinely route-specific and cannot be replaced by a single average. Figure 5 illustrates the graph structure on a small schematic section.




Figure 5: Simplified illustration of an airport surface graph, with gate nodes, taxiway nodes and runway entry. Created with Canva.

4.5 Data Limitations

Three limitations of the data should be noted. First, the dataset contains only departures; the influence of arriving traffic is not represented in the records and is approximated in this model through the penalty times described in Section 4.4. Second, no fuel consumption data is available, so the fuel comparison rests on published values, as described in Section 4.3, and is an approximation. Third, a small share of recorded

taxiway paths (approximately 8.4%) reference segments that could not be reconciled with the constructed network, mostly routes crossing the runway system via less common intersections; where a shortest-path repair could not restore them, these flights are excluded from the simulation. This exclusion is unlikely to bias the fleet-size comparison, as the affected routes occur throughout the operational day, but the reported totals refer to the usable subset of departures.

The filtering chain leaves the following instance sizes. Of the 9,589 departure records in July 2025, 6,596 use runway 25R, 6,315 of these have a pushback clearance time, and 6,261 also have a usable taxiway path; this is the clean dataset. Of those, 4,409 are narrow-body movements eligible for taxibot service and 1,852 are wide-body or other types that appear only as conventionally taxiing traffic, while 4,189 of the eligible flights follow a route through one of the two start-up zones and are therefore also eligible for the pool model.

A further group of flights cannot be served under any fleet size, because attachment and towing do not fit within their time window even at the earliest opportunity; these are removed in preprocessing and taxi conventionally throughout. On 9 July, this leaves 211 taxibot-eligible departures, of which 205 are pool-eligible and 66 are individually infeasible, giving an upper bound of 145 aircraft that any fleet could serve.

4.6 Derived Variables

Several model inputs are computed from the raw fields. Table 5 summarises them; the time unit throughout is seconds.

Table 5: Variables derived from the ground movement dataset. All times are in seconds.

Variable	Definition	Purpose
Rolling time	END TAXI – TAXI CLR	Speed calibration (Section 4.4)
Empirical speed	Path length / rolling time	Engine taxi speed (median over flights)
Towing time τ_i^{tow}	Sum of connection traversal times along the path of aircraft i at towing speed	Known before scheduling, since routes are given
Release time r_i	PUSH CLR (fallback: BLK TIME)	Earliest start of service for departure i
Deadline d_i	Derived from T/O TIME by allowing for detachment and engine warm-up	Latest time by which towing must be completed
Service window	Time available between release and the towing deadline after accounting for attachment and towing	Determines whether an aircraft can be served

The deadline of aircraft i is derived from its observed take-off time by reserving sufficient time for detachment and engine warm-up after towing. It therefore represents the latest moment by which towing must be completed.

Together with the aircraft’s release time and towing duration, this deadline determines whether the aircraft can in principle be served by a taxibot. Aircraft for which attachment and towing cannot be completed within the available time are classified as individually infeasible and are excluded from taxibot scheduling. The formal definitions of these time-window conditions are introduced in Section 5.2, after the complete model notation has been defined.

5 Methodology

5.1 Solution Approach Overview

The scheduling problem described in Section 3 is approached with three models, which differ in how they trade solution quality against computation time.

First, a mathematical formulation of the problem is given as an integer linear programme (ILP), which defines precisely what a valid schedule is and what “optimal” means: the assignment of aircraft to taxibots that serves the largest possible number of departures while respecting all time-window and resource constraints. This formulation is solved with Gurobi, a commercial solver for mixed-integer programmes that searches the solution space with branch-and-bound techniques and either proves a solution optimal or, when stopped at a time limit, reports the best solution found together with a bound on how far it can be from the optimum (Gurobi Example Tour, 2026). The ILP solution provides the benchmark for each problem instance.

Second and third, two variants of a greedy construction heuristic are tested, which build a schedule in a single pass: one processes aircraft in order of their release time, the first-come-first-served discipline that air traffic control applies in practice, and one in order of their deadline (earliest deadline first), the classical priority rule for meeting time windows. Both rules are standard in the scheduling literature and are discussed in Section 2.4; comparing them shows how sensitive throughput is to the ordering choice. The heuristics do not guarantee optimality, but they run in milliseconds where the ILP requires minutes to hours, and comparing their schedules against the ILP benchmark quantifies exactly how much is lost by using them, the optimality gap reported in the results.

The reason for this three-model design is computational. The problem is a variant of the vehicle routing problem with time windows, for which even finding a feasible solution with a fixed fleet is NP-complete (Savelsbergh, 1985). Exact solution is therefore only practical within a time limit, and the solver reports for each instance whether the returned solution is proven optimal or only the best found within the limit. A heuristic that closely tracks the ILP at realistic instance sizes is the more useful tool in practice, since real operations are dynamic and require fast re-computation; the ILP’s role is to certify how close the heuristic comes. All times in my model are expressed in seconds, and timestamps are handled as absolute date-times of the operational day.

The implementation is organised in a sequence of notebooks: data preparation and network construction (Section 4), the scheduling models, and the scenario comparison,

supported by a shared module containing the heuristic and the ILP so that both solve identical instances with identical parameters.

5.2 Mathematical Formulation

5.2.1 Sets, parameters and decision variables

This section presents a mathematical formulation of the scheduling problem as an integer linear programme. Table 6 defines all sets, parameters and decision variables used in the objective and constraints; the subsequent text uses these symbols without redefinition. The table gives the symbols only; the values assigned to the parameters are set out in Section 5.4.

Table 6: Sets, parameters and decision variables of the optimisation model.

Symbol	Definition
Sets	
A	Departing aircraft, indexed i, j
B	Taxibots, indexed k
Z	Mutual-exclusion zones, indexed z ; each zone groups physically adjacent nodes (Section 4.4)
$\text{Slots}(z)$	Slots of start-up zone z , one per unit of its capacity
A_z	Aircraft whose route passes through mutual-exclusion zone z
Parameters	
K	Number of taxibots, $K = B $
r_i	Release time of aircraft i : pushback clearance (Section 4.6)
$t_i^{\text{T/O}}$	Observed take-off time of aircraft i
d_i	Deadline of aircraft i : latest time by which towing must be completed,
	$d_i = t_i^{\text{T/O}} - \tau_{\text{det}} - \tau_{\text{warmup}}$
	(Section 4.6)

Table 6 – continued from previous page

Symbol	Definition
τ_i^{tow}	Towing time of aircraft i along its recorded route, including runway-crossing penalties; known before scheduling
τ_{att}	Attachment time
τ_{det}	Detachment time
τ_{back}	Drive-back time
τ_{warmup}	Engine warm-up time after detachment
C_z	Capacity of start-up zone z , in number of aircraft
T_i^{busy}	Time a taxibot is occupied by aircraft i :
	$T_i^{\text{busy}} = \tau_{\text{att}} + \tau_i^{\text{tow}} + \tau_{\text{det}} + \tau_{\text{back}}$
det_off_i	Offset from service start to start of warm-up:
	$\text{det_off}_i = \tau_{\text{att}} + \tau_i^{\text{tow}} + \tau_{\text{det}}$
M	Big- M constant, large enough to deactivate a disjunctive constraint when it is inactive
α_{iz}	Offset from the service start of aircraft i to its entry into zone z
β_{iz}	Offset from the service start of aircraft i to its exit from zone z
z_i	Start-up zone assigned to aircraft i , determined by its route
Decision variables	
y_i	1 if aircraft i is served by a taxibot, 0 otherwise
x_{ik}	1 if aircraft i is assigned to taxibot k , 0 otherwise
u_{ij}	1 if aircraft i is served before aircraft j on the same taxibot, 0 otherwise

Table 6 – continued from previous page

Symbol	Definition
q_{ic}	1 if aircraft i occupies slot c of its start-up zone, 0 otherwise
o_{ijc}	1 if aircraft i completes its warm-up in slot c before aircraft j begins its own, 0 otherwise
s_i	Start of service (start of attachment) for aircraft i , $s_i \geq 0$
w_{ijz}	1 if aircraft i clears zone z before aircraft j enters it, 0 otherwise

The deadline requires one remark, as its definition matters for correctness. The deadline d_i is the latest moment by which towing must be finished so that detachment, engine warm-up and line-up still fit before the observed take-off time.

Because each aircraft's route is given, its towing time τ_i^{tow} is known before scheduling. The requirement that towing finishes by d_i therefore translates into a condition on the start of service:

$$s_i + \tau_{\text{att}} + \tau_i^{\text{tow}} \leq d_i. \quad (2)$$

An aircraft is schedulable at all only if this condition can already be met at the earliest possible opportunity. Therefore,

$$r_i + \tau_{\text{att}} + \tau_i^{\text{tow}} \leq d_i. \quad (3)$$

Aircraft failing the condition in Equation 3 taxi conventionally regardless of fleet size.

5.2.2 Objective function

The objective is to maximise the number of departing aircraft served by the taxibot fleet:

$$\max \sum_{i \in A} y_i. \quad (4)$$

The motivation for this objective over direct fuel minimisation is given in Section 3.1.

5.2.3 Constraints

Constraint 1: Assignment consistency. Each served aircraft is assigned to exactly one taxibot, and the served indicator equals the number of assignments:

$$y_i = \sum_{k \in B} x_{ik} \leq 1, \quad \forall i \in A. \quad (5)$$

The inequality permits leaving an aircraft unserved; there is no obligation to serve any particular flight.

Constraint 2: Time window. For a served aircraft, service cannot start before release, and attachment plus towing must be completed by the deadline:

$$\begin{aligned} s_i &\geq r_i - M(1 - y_i), \\ s_i + \tau_{\text{att}} + \tau_i^{\text{tow}} &\leq d_i + M(1 - y_i), \quad \forall i \in A. \end{aligned} \quad (6)$$

Detachment and warm-up do not appear here because they are already accounted for in the definition of the deadline,

$$d_i = t_i^{\text{T/O}} - \tau_{\text{det}} - \tau_{\text{warmup}}. \quad (7)$$

The time-window constraint therefore requires attachment and towing to finish by the moment that leaves exactly enough room for the remaining phases before take-off. The relaxation term with M deactivates the deadline for unserved aircraft. Aircraft that cannot meet it even at their release,

$$r_i + \tau_{\text{att}} + \tau_i^{\text{tow}} > d_i, \quad (8)$$

are removed in preprocessing, as they cannot be served under any schedule.

Constraint 3: No overlap on a taxibot. Two aircraft assigned to the same taxibot cannot be served at the same time. Whichever aircraft is served first, the other cannot

start before the full service cycle of the first, including the drive-back, is complete. For every pair $i < j$ and every taxibot k :

$$\begin{aligned} s_j &\geq s_i + T_i^{\text{busy}} - M(1 - u_{ij}) - M(2 - x_{ik} - x_{jk}), \\ s_i &\geq s_j + T_j^{\text{busy}} - Mu_{ij} - M(2 - x_{ik} - x_{jk}), \end{aligned} \tag{9}$$

$$\forall i < j, \forall k \in B.$$

The second big- M term deactivates both inequalities unless aircraft i and j are both assigned to taxibot k . The first big- M term selects which of the two orderings applies, so exactly one of the inequalities binds for a pair that shares a taxibot.

The condition is imposed on every such pair rather than only on consecutive services, which sequences a taxibot's entire workload without requiring successor variables. Ordering variables are created only for pairs whose service intervals can interact, that is, when

$$r_j < d_i + T_i^{\text{busy}} \quad \text{and} \quad r_i < d_j + T_j^{\text{busy}}. \tag{10}$$

For all other pairs, no ordering constraint is generated. Here M is the global big- M constant of Table 6. A tight value of M strengthens the linear relaxation and reduces solver time; setting a smaller constant per constraint rather than one global value is a possible refinement noted as future work.

Constraint 4: Start-up zone capacity. Each start-up zone z has a fixed number of slots equal to its capacity C_z (Section 3.5), and every aircraft towed to that zone occupies exactly one slot while it starts and warms up its engines:

$$\sum_{c \in \text{Slots}(z_i)} q_{ic} = y_i, \quad \forall i \in A, \tag{11}$$

where $\text{Slots}(z_i)$ denotes the slots of the start-up zone on the route of aircraft i .

Two aircraft assigned to the same slot may not warm up at the same time. An aircraft occupies its slot from the moment towing and detachment are complete, at offset det_off_i from its service start, until its warm-up ends τ_{warmup} later. For any two aircraft i and j sharing slot c , one warm-up must finish before the other begins:

$$\begin{aligned}
s_j + \text{det_off}_j &\geq s_i + \text{det_off}_i + \tau_{\text{warmup}} - M(1 - o_{ijc}) - M(2 - q_{ic} - q_{jc}), \\
s_i + \text{det_off}_i &\geq s_j + \text{det_off}_j + \tau_{\text{warmup}} - Mo_{ijc} - M(2 - q_{ic} - q_{jc}), \\
&\forall i < j \in A \text{ with } z_i = z_j, \quad \forall c \in \text{Slots}(z_i).
\end{aligned} \tag{12}$$

Modelling each zone as a set of capacity-many slots, rather than as an aggregate count, keeps the constraint a standard pairwise no-overlap constraint of the kind used for the taxibots. The second big- M term deactivates the constraint unless both aircraft actually share slot c . As with the taxibot ordering variables, o_{ijc} is created only for pairs whose warm-up windows can overlap.

Constraint 5: Mutual-exclusion zones. Nodes lying physically close together are grouped into zones $z \in Z$ (Section 4.4) that cannot be occupied by two aircraft simultaneously. A zone is treated as a single shared resource: aircraft i occupies zone z from its entry into the first zone-member node on its route until its exit from the last, at the known offsets α_{iz} and β_{iz} from its service start.

For every zone z and every pair $i < j$ of aircraft routed through it, one aircraft must have cleared the zone before the other enters:

$$\begin{aligned}
s_j + \alpha_{jz} &\geq s_i + \beta_{iz} - M(1 - w_{ijz}) - M(2 - y_i - y_j), \\
s_i + \alpha_{iz} &\geq s_j + \beta_{jz} - Mw_{ijz} - M(2 - y_i - y_j), \\
&\forall z \in Z, \forall i < j, i, j \in A_z.
\end{aligned} \tag{13}$$

The form is identical to Constraint 4, with zone-entry and zone-exit offsets in place of slot offsets, and the second big- M term restricts the condition to pairs that are both towed. As there, the constraints are generated only for pairs whose occupation windows of the zone can intersect.

Constraint 6: Domains. The binary assignment, ordering and occupation variables satisfy

$$y_i, x_{ik}, u_{ij}, q_{ic}, o_{ijc}, w_{ijz} \in \{0, 1\}, \quad (14)$$

for all applicable $i, j \in A$, $k \in B$, $z \in Z$ and $c \in \text{Slots}(z)$, while service start times satisfy

$$s_i \geq 0, \quad \forall i \in A. \quad (15)$$

5.2.4 Implemented model and solution protocol

Two elements of the formulation are not enforced in my implemented models. The route-wide mutual-exclusion zones (Constraint 5, Equation 13) are formulated but left as an extension, as their effect at the start-up locations is already captured by the zone-capacity constraint. Segment-level conflicts between individual aircraft along the taxi routes are omitted by design, for the reasons given in Section 3.7. Everything else is enforced by both the exact model and the heuristic, which is what makes the optimality gap between them meaningful.

Both models are run in two configurations, with and without the start-up zones, so that the contribution of the zones can be isolated in the results. In the start-up zone configuration the same formulation is applied with τ_i^{tow} and d_i redefined to the start-up zone rather than the runway holding point: towing ends at the zone, and the deadline additionally accommodates the aircraft's own taxi from the zone to the runway after warm-up.

5.3 Scheduling Heuristics

Because my exact model is costly to solve, the problem is also addressed with a greedy construction heuristic that builds a schedule in a single pass over the aircraft. Two variants are tested, differing only in the order in which aircraft are considered.

5.3.1 Algorithm

The heuristic processes aircraft one at a time in a fixed order and assigns each to the taxibot that can serve it earliest, if any can. For each taxibot $k \in B$, the heuristic maintains a state variable free_k , defined as the earliest time at which taxibot k is available to start a new service. At the start of the algorithm, all free_k values are initialised to the earliest release time in the instance.

For each aircraft i , taken in the chosen order:

- its earliest possible start on taxibot k is the later of its release time and that taxibot's free time,

$$\max(r_i, \text{free}_k);$$

- taxibot k is a feasible candidate if towing can still finish by the deadline from that start,

$$\max(r_i, \text{free}_k) + \tau_{\text{att}} + \tau_i^{\text{tow}} \leq d_i;$$

- if no taxibot is feasible, aircraft i is left unserved and taxis conventionally;
- otherwise, aircraft i is assigned to the feasible taxibot with the earliest start, its service start s_i is set to that time, and that taxibot's free time is updated to

$$s_i + T_i^{\text{busy}} = s_i + \tau_{\text{att}} + \tau_i^{\text{tow}} + \tau_{\text{det}} + \tau_{\text{back}}.$$

The procedure never revisits an assignment, which is what makes it fast: each aircraft is placed in a single step, giving a total running time that is linear in the number of aircraft for a fixed fleet.

5.3.2 Ordering variants

The two variants differ in the sort order:

- **Release-time order (first-come-first-served).** Aircraft are processed by their release time r_i , matching the chronological flow of operations: decisions are made in the order aircraft become ready, and no aircraft is deferred in the hope of a better later assignment.
- **Earliest-deadline-first (EDF).** Aircraft are processed by their deadline d_i , so those whose windows close soonest are placed first. This tends to protect aircraft

with tight windows, which a purely chronological order might lose to a taxibot occupied by a less urgent aircraft.

Both variants use the identical feasibility and assignment logic of Section 5.3.1; only the processing order changes. Reporting both, alongside the ILP, shows how sensitive throughput is to this ordering choice and which variant to prefer in practice.

5.3.3 Constraints enforced and comparison protocol

In its base form my heuristic enforces the same conditions as my exact model: single assignment, the release-time and deadline window, and no overlap on a taxibot (Constraints 1 to 3), together with a pool variant that additionally respects the start-up zone capacity (Constraint 4).

The route-wide mutual-exclusion zones and segment-level conflicts are not enforced, consistent with my exact model (Section 5.2.4). Because both the heuristic and the ILP enforce the same set of constraints in each configuration, their solutions are directly comparable: for each fleet size and parameter set the two are run on the same instance, and the difference between the ILP’s served count and the heuristic’s is the optimality gap.

A heuristic that stays close to the ILP across fleet sizes, at a fraction of the runtime, is then the more practical tool for the scenario analysis, with the ILP serving to certify how close it comes.

5.4 Parameter Setting

The formulation of Section 5.2 and the heuristics of Section 5.3 define the problem in symbols; this section assigns them values. All times are expressed in seconds, and timestamps are handled as absolute date-times of the operational day. Table 7 lists the base values used throughout the fleet-size sweep, together with their source. Three of them are literature values, three were validated with Skeyes, and one is calibrated from the data itself; the parameters whose values are estimates rather than measurements are varied in the sensitivity analysis of Section 6.5.

Table 7: Base parameter values used in the fleet-size sweep.

Symbol	Value	Source
K	1 to 8	Scenario input (Section 5.5)
τ_{att}	60 s	Salihu et al. (2021)
τ_{det}	60 s	Salihu et al. (2021)
τ_{back}	600 s	Validated with Skeyes (F. Claude, personal communication, 2026)
τ_{warmup}	180 s	Stettler et al. (2018); confirmed by Skeyes
Speed factor	1.5	Salihu et al. (2021)
C_{B1}	2	Validated with Skeyes (F. Claude, personal communication, 2026)
C_{W42}	4	Validated with Skeyes (F. Claude, personal communication, 2026)
M	Day horizon + longest service and warm-up duration	Set to deactivate an inactive disjunctive constraint

The attachment and detachment times are taken from the towing simulation of Salihu et al. (2021) and treated as constants, since their variation between routes is small relative to the towing time itself.

The drive-back time is the parameter where literature and operational practice diverge most. Computing it from the network gives roughly 300 seconds for a return from the runway area to the apron, but Skeyes assessed this as optimistic, since a returning taxibot must be fitted into live ground traffic and handled by the ground controller; 600 seconds is used as the realistic base value and the network-derived 300 seconds is retained as an optimistic case in the sensitivity analysis.

The warm-up time of 180 seconds sits within the two to five minutes that current narrow-body engines require depending on engine type and state (Stettler et al., 2018).

The towing speed is not set directly but derived from the calibrated engine taxi speed through a speed factor, so that both modes share the same distance measurements and any systematic bias in those measurements affects them equally. Published operational towing speeds range from 4 m/s (Salihu et al., 2021) to 11.8 m/s (van Baaren & Roling, 2019; van Oosterom et al., 2023), against a design maximum of about 23 knots for the TaxiBot itself (Lukic et al., 2019), with the operational restriction imposed to limit nose landing gear fatigue (TLD Group, n.d.). A factor of 1.5 corresponds to a towing

speed of approximately 4.8 m/s at the calibrated engine speed of 7.3 m/s, consistent with the lower end of that range, and it is varied between 1.0 and 2.0 in the sensitivity analysis.

The capacities of the two start-up zones follow the physical space available at each location, as described in Section 3.5: the W41/W42 platform accommodates four aircraft and the B1 zone two. Both are varied in the sensitivity analysis, which shows them not to be binding.

The exact model is solved with Gurobi under a fixed time limit of 3600 seconds per instance, and every result is labelled with its solver status: optimal when optimality was proven within the limit, feasible when the reported solution is the best found. The heuristic requires no such limit, since it terminates after a single pass.

The solver is given the greedy EDF schedule as an initial incumbent, which shortens the search substantially at the fleet sizes where the assignment problem is hardest. A consequence is that the exact model can never report fewer served aircraft than the heuristic it starts from, so the gap between the two should be read as how much the solver can improve on the heuristic within the time limit, rather than as an independent comparison of the two methods.

5.5 Scenario Design

This model is run across a range of fleet sizes to trace how the number of aircraft served depends on the size of the taxibot fleet. The number of taxibots is varied over $K = 1, 2, \dots, 8$, chosen to span the full behaviour of interest: from severe scarcity, where a single taxibot serves only a small fraction of departures, through the operationally realistic range, to the point where adding taxibots yields little further gain because the binding constraint has shifted from fleet size to the aircraft time windows (Section 2.4). Both configurations, with and without start-up zones, are run, and for each value of K the greedy heuristic and the ILP are solved on the same set of departures with the same parameters, so that my models are compared like for like.

Each scenario is evaluated against the conventional baseline in which every aircraft taxis on its own engines, using the observed taxi durations and routes from the data. For each fleet size the following are reported: the number of aircraft served by taxibots and the corresponding proportion of eligible departures covered, the number continuing to taxi conventionally, and the resulting fuel and CO₂ savings relative to the baseline (Section 5.6). Each ILP result additionally carries its solver status, optimal or best-

found within the time limit.

Because several of these parameters are estimates rather than measurements, their influence is examined separately in a sensitivity analysis, reported with the results, covering the towing speed factor, the drive-back time, the warm-up time and the start-up zone capacities (Section 6.5).

5.6 Fuel and Emissions Comparison

The fuel and emissions comparison quantifies the environmental effect of the schedules produced by this model. It is not a separate optimisation but a calculation applied on top of each schedule: every aircraft that a taxibot serves has its engines off during towing, so the difference between its conventional taxi fuel and its taxibot-service fuel is a saving, and summing these gives the effect for the fleet. Because more served aircraft means more savings, this quantity follows directly from the throughput this model maximises (Section 3.1).

Since the dataset contains no measured fuel consumption, the comparison uses published values, and its results are therefore approximations of realistic magnitude rather than measured savings. For a conventionally taxiing aircraft, the taxi fuel is estimated from its observed taxi duration and a per-aircraft fuel-flow rate; the aggregate narrow-body figure of Camilleri and Batra (2021) provides the working baseline, with the per-type values of Section 4.3 used where a finer breakdown is needed. This approach is known to overestimate somewhat, as constant idle-thrust assumptions exceed observed taxi burn (Nikoleris et al., 2011); the estimates are therefore treated as an upper reference, which is acceptable because the comparison is relative, the same method being applied to both the baseline and the taxibot scenarios.

For an aircraft served by a taxibot, its engines are off during the towed portion of the taxi and on for the remainder, so its engine fuel is counted only over the taxi time not covered by towing; the taxibot’s own energy use is then added, over both the tow and the drive-back, and depends on its propulsion type. This is where the distinction between a diesel and an electric taxibot enters, not in the scheduling, which is identical for both, but only in this comparison. A diesel taxibot burns fuel during towing and the return drive; an electric taxibot produces no direct emissions, leaving only the upstream emissions of the electricity used to charge it, estimated from the Belgian grid emission factor. The auxiliary power unit, which runs during towing, and the brief additional burn at engine start are treated as part of the aircraft’s engine-on time rather than as separate terms, a simplification that slightly understates the

taxibot-service fuel and therefore keeps the savings estimate conservative.

The total fuel saving of a scenario is the baseline fuel of all departures minus the fuel of the mixed operation, in which served aircraft use taxibot service and the rest taxi conventionally; CO₂ savings follow by applying a standard jet-fuel emission factor. These totals are reported per fleet size alongside the throughput metrics, giving a combined picture of the operational and environmental effect of taxibot deployment at Brussels Airport.

The estimates carry the limitations of the underlying assumptions, aggregate rather than per-phase fuel rates, no differentiation by weather, and a fixed grid factor for the electric case, so they indicate the order of magnitude of the savings rather than precise figures.

6 Results

6.1 Experimental setup

All results in this section are computed with the base parameters of Table 7, in which the drive-back time of 600 seconds and the start-up zone capacities were validated with Skeyes. The headline results use 9 July 2025, one of the busiest days of the month, with its 211 taxibot-eligible and 205 pool-eligible departures (Section 4.5). To verify that this day is representative, the fleet-size sweep is repeated across the first week of July, whose days carry comparable departure volumes. The exact model is solved with a fixed time limit of one hour per instance, and every ILP result is labelled with its solver status: optimal when optimality was proven within the limit, feasible when the reported solution is the best found. Aircraft whose time window cannot accommodate attachment and towing even at the earliest opportunity are removed in preprocessing and taxi conventionally regardless of fleet size; the number of served aircraft is therefore bounded above by the eligible subset, not by the total number of departures.

This section proceeds from the central throughput result (Section 6.2) through the comparison of the scheduling models (Section 6.3) and the start-up pool analysis (Section 6.4) to the sensitivity of the results (Section 6.5) and the fuel and CO₂ effect (Section 6.6).

6.2 Throughput versus fleet size




Figure 6: Aircraft served against fleet size, pool model, 9 July 2025. Open markers indicate ILP solutions that reached the time limit.

Figure 6 shows the central result of the thesis. With the parameter set of Section 5.4, the exact model of the pool configuration serves 43 aircraft with a single taxibot, 82 with two, 110 with three, 127 with four, 136 with five, and 141, 144 and 146 with six, seven and eight, out of 205 pool-eligible departures. The greedy heuristic in its earliest-deadline-first variant, plotted alongside it, tracks these figures closely at every fleet size; the two are compared in Section 6.3, and the figures quoted below are those of the exact model unless stated otherwise. The pattern of diminishing returns is pronounced: each of the first three vehicles adds between 28 and 39 served aircraft, the fifth adds nine, and the eighth adds two. Beyond roughly five vehicles the binding constraint is no longer the fleet but the time windows of the aircraft themselves, and the curve flattens toward its ceiling.

This shape acquires operational meaning through the assessment of Skeyes. Independently of this model, Skeyes indicated that about five taxibots is a strict practical maximum at Brussels Airport, because every additional vehicle adds return traffic on the taxiways and workload for the ground controllers (F. Claude, personal communication, 2026). The practical ceiling and the economic knee of my model’s curve therefore roughly coincide: at $K = 5$ the fleet serves 136 aircraft, 66% of the pool-eligible departures, and nearly everything the remaining three vehicles would add lies in the operationally questionable range. The full curve is computed to $K = 8$ to characterise the saturation behaviour, but the reading of the results above $K = 5$ is theoretical rather than operational.

Table 8: Share of pool-eligible departures served, per day and fleet size, first week of July 2025. N is the number of pool-eligible departures that day, an absolute count; all other columns are percentages of N .

Date	N	$K = 1$	$K = 2$	$K = 3$	$K = 4$	$K = 5$	$K = 6$	$K = 7$	$K = 8$
1 July	205	20.5	36.6	47.3	54.1	57.6	59.5	60.5	61.0
2 July	212	20.3	35.4	46.2	54.7	58.5	60.4	61.8	63.2
3 July	199	20.6	37.7	47.7	54.3	57.8	59.3	60.8	61.3
4 July	218	18.3	33.0	43.1	50.5	54.6	56.4	57.8	58.3
5 July	162	21.6	37.7	48.8	56.8	61.1	63.6	64.8	64.8
6 July	177	20.3	39.0	50.3	58.2	62.1	65.5	66.7	67.2
7 July	183	18.6	34.4	45.9	53.0	58.5	61.2	62.3	62.8
Mean		20.0	36.3	47.0	54.5	58.6	60.8	62.1	62.7
Std. dev.		1.2	2.1	2.3	2.5	2.5	3.0	2.9	2.9

Table 8 confirms that the pattern is not an artefact of the chosen day. Across the first week of July, the served share rises steeply to around $K = 4$ – 5 and saturates on every

day, with day-to-day variation driven by the traffic volume and its peaks rather than by this model.

On average across the week, the fleet serves 47.0% of eligible departures at $K = 3$ and 58.6% at $K = 5$, with a standard deviation of only two to three percentage points at every fleet size, so the fleet-size effect is far larger than the day-to-day variation. The two ordering rules are practically equivalent: on the headline day, in the configuration without start-up zones, they differ by at most one aircraft at any fleet size (107 against 106 at $K = 3$, 134 against 133 at $K = 5$, identical elsewhere), and on other days the difference occasionally falls the other way. The results above therefore use the earliest-deadline-first variant; the per-day figures for both are given in Appendix A.2 and A.3.

6.3 Heuristic versus exact model

Table 9: Greedy heuristic against the exact model, pool configuration, 9 July 2025, ILP time limit 180 s per instance.

K	Greedy EDF served	ILP served	Gap (ILP – greedy)	ILP status
1	43	43	0	Optimal
2	82	82	0	Feasible (time limit)
3	109	110	1	Feasible (time limit)
4	127	127	0	Feasible (time limit)
5	135	136	1	Feasible (time limit)
6	138	141	3	Optimal
7	139	144	5	Optimal
8	139	146	7	Optimal

The comparison between the two scheduling approaches is unambiguous. Table 9 compares the greedy heuristic with the exact model on the headline day in the pool configuration, at a solver time limit of 180 seconds per instance. The greedy heuristic matches the exact model to within a single aircraft at every fleet size up to $K = 5$, serving 109 against the ILP’s 110 at $K = 3$, and across the first week the daily gap never exceeds a handful of aircraft. The comparison is one-sided by construction:

the solver is started from the greedy schedule (Section 5.4), so it can only equal or improve on it, and the gap therefore measures how much branch-and-bound adds to the heuristic within the time limit rather than an independent contest between the two. What it adds is little. Several instances stop at the limit with status feasible, meaning the reported value is the best solution found rather than a proven optimum; where optimality was proven, the greedy matches it exactly or nearly so. The runtimes differ by orders of magnitude, the heuristic completing in milliseconds against minutes for the exact model.

One deliberately constructed exception qualifies this picture and is worth recording: when the B1 start-up zone is restricted to a single slot (Section 6.5), the greedy falls nine aircraft behind the exact model, because under severe slot scarcity the order in which aircraft are committed interacts with slot availability in a way a single-pass procedure cannot anticipate. At all realistic capacities the gap disappears. For operational use, where schedules must be recomputed as conditions change, the heuristic therefore sacrifices essentially nothing while being effectively instantaneous.

To check whether the time limit constrains the comparison, the sweep was repeated on a separate day in the configuration without start-up zones with a limit of one hour per instance; the full results, including solver runtimes, are given in Appendix A.1. The picture is unchanged: at every fleet size the exact model returns the greedy solution or one aircraft more, and at the four mid-range fleet sizes it exhausts the full hour without improving on the schedule it started from or closing the optimality gap, while the easy instances at one, two, seven and eight vehicles solve to proven optimality in under a minute. The limited value of a longer time limit is therefore not an artefact of the three-minute budget.

6.4 The start-up pool




Figure 7: Taxibot service cycles over the operational day, greedy pool model, $K = 3$, 9 July 2025.

Figure 7 shows how the service cycles distribute over the operational day and makes the weight of the return drive visible: with the validated ten-minute drive-back, the time a taxibot is occupied per service is dominated not by the towing but by the overhead around it, and the fleet saturates during the morning departure peak while idling in the quieter early afternoon.

The pool changes the picture asymmetrically. On the headline day at $K = 3$ the pool model serves 109 of 205 eligible aircraft against 106 of 211 for towing to the holding point, a modest throughput gain. Underneath, however, the two zones behave very differently. Detaching at the W42 zone frees the taxibot on average 147 seconds earlier per service, while detaching at the B1 zone saves only 20 seconds, because the B1 detachment point lies close to the holding point it replaces. For the aircraft the asymmetry is sharper still: at W42 the aircraft’s total journey is on average 113 seconds shorter, whereas at B1 the pool adds around 12 seconds, since the self-taxi after warm-up exceeds the small towing reduction. The pool concept, in other words, works where the geometry gives it room, and the W41/W42 platform is exactly the location Skeyes assessed as operationally workable for taxibot returns, whereas the B1 flow is where returns would require infrastructure works (F. Claude, personal communication, 2026). This model and the operational assessment point at the same place.

6.5 Sensitivity analysis

Several model parameters are assumptions or literature-derived values rather than quantities measured at Brussels Airport. As set out in Section 5.4, the towing speed factor comes from published towing studies at other airports (Salihu et al., 2021), the warm-up time from engine-manufacturer requirements, and the attachment, detachment and drive-back durations partly from the taxibot literature and partly from operational judgement validated with Skeyes. A one-at-a-time sensitivity analysis around the base case therefore serves two purposes: it shows how strongly the conclusions depend on each assumption, and it identifies which parameters would repay precise measurement in a real deployment. Four parameters are varied: the towing speed factor, the drive-back time, the warm-up time, and the capacities of the start-up zones.




Figure 8: Sensitivity of aircraft served to the towing speed factor, $K = 3$, 9 July 2025.

The towing speed factor is by far the most influential parameter. At $K = 3$, moving the factor from 1.0 (towing at engine-taxi speed) to 2.0 (towing at half that speed) reduces the served count from 138 to 63, more than a factor of two, with the base value of 1.5 sitting mid-range at 107 (Figure 8). The sensitivity is structural: a slower tow lengthens every service cycle and simultaneously tightens every deadline, so its effect compounds. Since no towing-speed measurements exist for Brussels, this parameter is the largest source of uncertainty in the absolute results, although the shape of the fleet-size curve and the location of its diminishing returns are stable across factors.




Figure 9: Sensitivity of aircraft served to the drive-back time, pool model, greedy EDF, 4 July 2025.

The drive-back time matters most for small fleets. Comparing the network-derived driving time of 300 seconds with the Skeyes-validated 600 seconds, which includes congestion and controller handling, the cost of realism at $K = 3$ is 17 aircraft per day (114 versus 97 served), shrinking to 8 aircraft at $K = 5$ and 2 at $K = 8$ (Figure 9). With few vehicles, every minute of unproductive return time translates directly into unserved aircraft, whereas a larger fleet absorbs it. Precisely the small fleets an airport would start with are the most exposed to return-drive logistics.




Figure 10: Sensitivity of aircraft served to the engine warm-up time, $K = 8$, 9 July 2025.

The warm-up time has a steady, almost linear effect: each additional 60 seconds of required warm-up costs roughly six to seven served aircraft at $K = 3$, from 130 at zero warm-up down to 107 at the 180-second base. This quantifies the potential of an operational refinement discussed in Section 3: starting the engines during the final towing segment rather than after detachment would recover part of this margin.

The zone capacities, finally, are not binding. Doubling the W42 capacity from two to four changes nothing at any fleet size, and the B1 capacity shows a threshold: one slot costs roughly ten aircraft, while two or more give identical results. The bottleneck of the system is the taxibot fleet and the aircraft time windows, not the space at the start-up zones. The corresponding sensitivity tables are reported in Appendix A.4, A.5 and A.6.

6.6 Fuel and CO₂ effect




Figure 11: Net CO₂ saved against fleet size, in kilograms per day (left) and as a percentage of the eligible baseline (right), 9 July 2025.

The environmental effect follows directly from the throughput. On the headline day, conventional taxi-out of the 211 eligible departures emits approximately 71,700 kg of CO₂. With three diesel taxibots the net saving is 23,696 kg per day, a 33.1% reduction of that baseline; with five vehicles it reaches 30,405 kg. In fuel terms the $K = 3$ saving corresponds to roughly 7,500 kg of jet fuel per day, in the order of €1.66 million per year at €0.85/kg over 260 operating days. Across the first week of July the daily $K = 3$ saving ranges between roughly 17,500 and 22,000 kg, confirming the order of magnitude beyond the busiest day. As explained in Section 5.6, these figures are estimates built on published consumption values rather than measured fuel data, so they indicate the scale of the effect rather than exact quantities; the vehicle's own energy use is accounted for over its full operating time, towing and the ten-minute drive-back included.

Because the taxibot's own energy use is small against the avoided jet fuel, the choice of propulsion type barely affects the daily CO₂ balance: at $K = 3$ the electric variant saves 24,083 kg per day against 23,696 kg for diesel, a difference of 387 kg, under two percent of the saving. The environmental case for taxibots is therefore carried overwhelmingly by switching the aircraft's engines off; the propulsion type of the vehicle is a second-order refinement in this day-level CO₂ accounting, although it also determines local air quality and noise, which this comparison does not capture.

The start-up zones add a trade-off. At $K = 3$ the configuration with zones serves three more aircraft than towing to the runway holding point, yet saves slightly less CO₂ (an estimated 22,624 against 23,696 kg for diesel), because detaching at the zone shortens the engine-off towed segment and the aircraft covers the remainder on its own engines. The zones thus buy throughput at a small emissions price, paid almost entirely at

the B1 zone, where the self-taxi exceeds the towing reduction. Which configuration is preferable depends on whether the operator weighs aircraft served or CO₂ avoided more heavily, a choice that echoes the stakeholder discussion of Section 3.

6.7 Summary of findings

Four findings summarise this section. A small fleet already delivers most of the achievable effect: three taxibots serve about half of the eligible departures and avoid roughly a third of their taxi-out CO₂, and diminishing returns set in exactly around the five-vehicle mark that Skeyes independently identifies as the operational maximum. The greedy heuristic is operationally sufficient, matching the exact model to within one aircraft at realistic settings while computing in milliseconds. The dominant uncertainties are the towing speed and the return-drive logistics, both of which would repay measurement in a real trial. Finally, the choice between propulsion types barely affects the CO₂ balance, while the choice between pool and holding-point delivery trades a few served aircraft against a few hundred kilograms of CO₂ per day.

7 Discussion

7.1 Interpretation of the main findings

The central pattern of the results is that a small taxibot fleet already captures most of the achievable effect. Three vehicles serve roughly half of the eligible departures on a busy day and avoid about a third of their taxi-out CO_2 ; five vehicles bring this to two thirds of departures, after which each additional vehicle adds only a few aircraft per day. The explanation lies in the structure of the service cycle. With the validated ten-minute drive-back, a taxibot is occupied far longer per aircraft than the towing itself, so during the morning departure peak the fleet is the binding resource, while outside the peak the aircraft time windows are. Adding vehicles relieves the first constraint but not the second, which is why the curve saturates well below the number of schedulable aircraft.

The comparison between the exact model and the heuristic supports a practical conclusion. Up to five vehicles, exactly the range an airport would realistically operate, the single-pass heuristic matches the exact model to within one aircraft while computing in milliseconds. The gap of up to seven aircraft that opens at seven and eight vehicles occurs where many near-equivalent assignments exist and is of little operational relevance, since that fleet range is unlikely to be deployed. For a setting in which schedules must be recomputed whenever delays occur, the heuristic is the appropriate tool, and the exact model's role is to certify its quality.

The environmental effect follows the throughput almost mechanically, which is precisely what the problem framing of Section 3 intended: emissions savings are an effect of serving aircraft, not a separate objective. Two nuances qualify this. First, the choice of propulsion type changes the day-level CO_2 balance by less than two percent, because the vehicle's own energy use is small against the avoided jet fuel; the case for the electric variant rests on local air quality and noise rather than on the CO_2 accounting. Second, the start-up pool buys a few served aircraft at a small emissions price, since detaching earlier shortens the engine-off segment. Which configuration is preferable is therefore not a purely technical question but depends on how the operator weighs throughput against emissions, the same stakeholder trade-off identified in Section 3.

7.2 Position relative to the literature

The results can be read against the towing-fleet studies reviewed in Section 2, which pose the inverse question: they require all aircraft to be towed and minimise the fleet needed, finding for Amsterdam Schiphol that on the order of 38 vehicles would be required (van Oosterom et al., 2023). The present thesis fixes a small fleet and asks what it delivers, and the answer, half of eligible departures with three vehicles, suggests that the gradual-introduction perspective is informative precisely where the full-coverage perspective is silent: in the steep part of the curve, where each vehicle still pays for itself operationally. The mixed-operation setting that existing models treat as a boundary case is, in this range, the whole story. The finding that eligibility and time windows, rather than fleet size, become binding beyond five vehicles also echoes the eligibility figure of Postorino et al. (2017) at Bologna, where only 61% of movements qualified for towing at all: at Brussels the schedulable subset plays the same ceiling role, and no fleet size can tow past it.

7.3 Operational feasibility at Brussels Airport

My model’s numbers acquire their meaning only against operational reality, and the assessment provided by Skeyes sharpens the reading in three ways. First, Skeyes regards about five taxibots as a strict practical maximum, because every vehicle adds return traffic on the taxiways and workload for the ground controllers (F. Claude, personal communication, 2026). My model’s diminishing returns set in at the same point, so the practical ceiling and the economic one coincide; but the reasons differ, and the operational reason applies even at fleet sizes this model finds worthwhile. My model does not represent the taxibots’ return trips as traffic on the network, so the congestion and coordination cost that concerns Skeyes is precisely the cost this model does not see. The true optimum fleet is therefore at most this model’s, and likely at its lower edge.

Second, the two start-up zones are not operationally equivalent. Skeyes assesses the W41/W42 platform as workable, with space for four aircraft and a return path for the taxibot between the aircraft lines, whereas returning from the B1 taxiway would require using the runway, which is operationally unacceptable, or infrastructure works (F. Claude, personal communication, 2026). My model’s own results point in the same direction: relative to towing all the way to the runway holding point, detaching at W42 frees the taxibot on average 147 seconds earlier per service, whereas at B1 the taxibot gains only about 20 seconds and the aircraft’s own journey is even slightly

longer, because the self-taxi from B1 after warm-up exceeds the small reduction in towing (Section 6.4). Since roughly four out of five departures follow the B1 flow, this convergence of model and operational judgement identifies the main bottleneck for a real deployment at Brussels: the concept works where the geometry gives it room, and most of the traffic is where it does not.

Third, these observations generalise beyond the specific zones. Brussels Airport, like most existing airports, was not designed with separate return paths for towing vehicles, and adapting such infrastructure is not easily done (L. Decock, personal communication, 2026). The saturation levels found here should therefore be read as what the current infrastructure permits, not as a property of the taxibot concept itself; a layout designed for towing operations would relax exactly the constraints that bind at Brussels.

7.4 Limitations

The limitations of this study fall into three groups: simplifications of the operational reality, uncertainty in the parameters, and the scope of the data. For each, it matters whether the limitation threatens the conclusions or merely their precision.

The first group concerns what this model does not represent. Routes are taken as given from the data, and conflicts between individual aircraft on shared taxiway segments are not modelled, on the argument that the recorded schedule was already conflict-free and this model only shifts the towed subset within their windows (Section 3.7). The taxibots' return trips, however, are new traffic that the recorded schedule never contained, and their interaction with other movements is captured only through the fixed drive-back time, not as movements on the network. This is my model's most consequential simplification: it is exactly the effect Skeyes identifies as the practical limit on fleet size, and it biases the results toward optimism at larger fleets. The mutual-exclusion zones are formulated but not enforced in the implementation, and arriving traffic is approximated through penalty times rather than modelled. The departure schedule is treated as deterministic and known, so the results describe planning capacity, not performance under delays and disruptions, which the disruption-management literature shows to be a substantial further challenge (Zoutendijk et al., 2023b).

The second group concerns parameters. The towing speed factor is the single largest source of uncertainty: no measurements exist for Brussels, and varying it over the literature range moves the served count at three vehicles between 63 and 138. The fuel comparison rests on published consumption values rather than measured data,

with idle-thrust figures known to overestimate actual taxi burn, and the taxibot's own consumption rates taken from generic sources; the resulting savings indicate an order of magnitude, not precise figures. The drive-back time, although validated by Skeyes as a realistic average, is a constant, whereas in reality it depends on the return distance and the traffic situation. The attachment, detachment and warm-up durations are literature values, of which the warm-up was confirmed by Skeyes.

The third group concerns scope. The analysis covers one airport, one runway configuration and one month of summer traffic, and about eight percent of recorded taxi paths could not be reconciled with the constructed network and were excluded. The B1 start-up zone is included as a modelling construct although its operational feasibility is doubtful without infrastructure works, so the pool results should be read with the W42 zone as the realistic component. None of these limitations undermines the central qualitative findings, the shape of the fleet-size curve, the sufficiency of the heuristic, and the dominance of the aircraft-side savings in the CO₂ balance; they do mean that the absolute numbers are specific to the assumptions and the summer traffic of Brussels Airport in 2025.

8 Conclusion and future work

8.1 Conclusion

This thesis asked how many departing aircraft a fixed fleet of taxibots can serve in a single operational day at Brussels Airport, and what fuel and CO₂ savings follow, compared with the conventional dual-engine baseline. To answer it, the problem was formulated as a scheduling model that assigns a fixed number of taxibots to departing narrow-body aircraft under time-window and resource constraints, with each aircraft following its recorded taxi route and emissions computed as a consequence of the resulting schedule. This model was built on real ground movement data of July 2025 provided by Skeyes, solved both exactly and with a fast construction heuristic, and validated against operational judgement from Skeyes.

The answer is that a small fleet already delivers most of the achievable effect. On 9 July 2025, one of the busiest days of the month, three taxibots serve about half of the eligible departures and avoid roughly 24 tonnes of CO₂, one third of the taxi-out emissions of the eligible traffic; five taxibots serve two thirds of departures and avoid about 30 tonnes. Beyond five vehicles the returns diminish rapidly, and the fleet stops being the binding constraint: the aircraft time windows and the eligible traffic itself take over. This economic saturation point coincides with the operational maximum of about five vehicles that Skeyes identifies independently, on the grounds of return traffic and controller workload. The pattern is stable across the first week of July, with day-to-day variation of only a few percentage points against fleet-size effects of ten points and more per added vehicle.

Methodologically, this thesis shows that this problem does not require heavy machinery to be solved well. A single-pass greedy heuristic matches the exact optimisation model to within one aircraft across the operationally relevant fleet range while computing in milliseconds, which makes it suitable for a real operation in which schedules must be recomputed continuously. The exact model's role is to certify that quality. The inverse framing of the fleet question, fixing a small fleet and maximising service rather than requiring full coverage and minimising the fleet, proved to be the informative perspective for an airport that would introduce taxibots gradually, and it is the perspective the existing literature does not cover.

For Brussels Airport specifically, the findings are double-edged. The savings are substantial and the fleet needed to realise them is small, but the operational analysis identifies where the concept strains against the existing infrastructure: the start-up

zone concept works cleanly at the W41/W42 platform, while the B1 flow, which carries most of the traffic, offers the taxibot no acceptable return path without infrastructure works. As Skeyes observes, an airport not designed for separate towing-vehicle return paths adapts to this way of working only with difficulty, and older infrastructure such as Brussels is exactly where such adaptation is hardest; newly designed airports could incorporate taxibot operations from the start (L. Decock, personal communication, 2026). The quantified benefits of this thesis are therefore what the current layout permits, and they should be read as a lower bound on what the concept could deliver in an infrastructure designed for it.

8.2 Future work

Several extensions follow directly from the limitations. The most valuable single improvement would be to model the taxibots' return trips as traffic on the network rather than as a fixed duration, since this is both my model's largest simplification and the effect Skeyes identifies as the practical limit on fleet size; a simulation of the shared surface, including arriving traffic, would show how quickly return movements congest the taxiway system as the fleet grows. Second, the deterministic setting should be relaxed: real operations face delays, and rescheduling under disruption, as studied for Schiphol by Zoutendijk, Van Oosterom & Mitici (2023), is the natural next layer on top of the planning model developed here. Third, the towing speed factor, the dominant parameter uncertainty, could be replaced by measurements: the operational TaxiBot trials at Brussels and Schiphol generate exactly the data needed, and a calibrated speed profile per taxiway segment would sharpen every result in this thesis. Fourth, the objective could be inverted into a direct fuel-minimisation formulation, assigning taxibots preferentially to aircraft with long taxi routes; comparing its schedules with the throughput-maximising ones would quantify how much the two objectives actually diverge in practice. Finally, the scope can be widened: other runway configurations, winter traffic, wide-body aircraft as certification expands, engine start during the final towing segment to recover the warm-up margin, and the battery-charging dimension of a fully electric fleet, which the Schiphol literature shows to be a fleet-sizing factor in its own right.

References

- AIM Belgium. (2026, April). Belgium & luxembourg aeronautical information publication.
- Alshaikhli, K., & Alshaikhli, A. (2024). Algorithms for multicriteria scheduling problems to minimize maximum late work, tardy, and early. *Journal of Applied Mathematics and Physics*, 12(02), 661–682. <https://doi.org/10.4236/jamp.2024.122043>
- Atkin, J. A. D., Burke, E. K., & Ravizza, S. (2010). The airport ground movement problem: Past and current research and future directions.
- Bennell, J. A., Mesgarpour, M., & Potts, C. N. (2011). Airport runway scheduling. *JOR*, 9(2), 115–138. <https://doi.org/10.1007/s10288-011-0172-x>
- Boussauw, K., & Vanoutrive, T. (2019). Flying green from a carbon neutral airport: The case of brussels. *Sustainability*, 11(7). <https://doi.org/10.3390/su11072102>
- Bräysy, O., & Gendreau, M. (2005). Vehicle routing problem with time windows, part i: Route construction and local search algorithms. *Transportation Science*, 39(1), 104–118. <https://doi.org/10.1287/trsc.1030.0056>
- Brucker, P., Drexl, A., M., Ohring, R., Neumann, K., & Pesch, E. (1999). Resource-constrained project scheduling: Notation, classification, models, and methods. *European Journal of Operational Research*, 112, 3–41. [https://doi.org/10.1016/S0377-2217\(98\)00204-5](https://doi.org/10.1016/S0377-2217(98)00204-5)
- Brussels airport. (2024). Brussels airport and tui fly are testing the taxibot for sustainable taxiing.
- Brussels Airport Traffic Control. (n.d.). *Preferential runway system (prs)*. Retrieved August 10, 2026, from <https://www.batc.be/en/runways-in-use/how-are-runways-selected/preferential-runway-system-prs>
- Camilleri, R., & Batra, A. (2021). Assessing the environmental impact of aircraft taxiing technologies. *32nd Congress of the International Council of the Aeronautical Sciences*.
- Cordeau, J.-F., Desaulniers, G., Desrosiers, J., Solomon, M. M., & Soumis, F. (2002). The vrp with time windows. In *The vehicle routing problem* (pp. 157–193). <https://doi.org/10.1137/1.9780898718515.ch>
- European Union Aviation Safety Agency. (2026, March). *Icao aircraft engine emissions databank*. Retrieved August 10, 2026, from <https://www.easa.europa.eu/en/domains/environment/icao-aircraft-engine-emissions-databank>
- Farhadi, F., Ghoniem, A., & Al-Salem, M. (2014). Runway capacity management - an empirical study with application to doha international airport. *Transportation Research Part E: Logistics and Transportation Review*, 68, 53–63. <https://doi.org/10.1016/j.tre.2014.05.004>

- Gupta, G., Jung, Y. C., & Malik, W. (2009). A mixed integer linear program for airport departure scheduling. *9th AIAA Aviation Technology, Integration, and Operations Conference (ATIO)*. <https://doi.org/10.2514/6.2009-6933>
- Gurobi Example Tour. (2026).
- Horn, W. A. (1974). Some simple scheduling algorithms. *Naval Research Logistics Quarterly*, 177–185.
- International Council on Clean Transportation. (2022, January). *Performance analysis of evolutionary hydrogen-powered aircraft*. <https://theicct.org/wp-content/uploads/2022/01/LH2-aircraft-white-paper-A4-v4.pdf>
- Khadilkar, H., & Balakrishnan, H. (2012). Estimation of aircraft taxi fuel burn using flight data recorder archives. *Transportation Research Part D: Transport and Environment*, 17(7), 532–537. <https://doi.org/10.1016/j.trd.2012.06.005>
- Liu, C. L., & Layland, J. W. (1973). Scheduling algorithms for multiprogramming in a hard-real-time environment. <https://doi.org/10.1145/321738.321743>
- 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. <https://doi.org/10.1109/TTE.2019.2956862>
- Nikoleris, T., Gupta, G., & Kistler, M. (2011). Detailed estimation of fuel consumption and emissions during aircraft taxi operations at dallas/fort worth international airport. *Transportation Research Part D: Transport and Environment*, 16(4), 302–308. <https://doi.org/10.1016/j.trd.2011.01.007>
- Postorino, M. N., Mantecchini, L., & Gualandi, E. (2017). Integration between aircraft and handling vehicles during taxiing procedures to improve airport sustainability. *International Journal of Transport Development and Integration*, 1(1), 28–42. <https://doi.org/10.2495/TDI-V1-N1-28-42>
- Pyrgiotis, N., Malone, K. M., & Odoni, A. (2013). Modelling delay propagation within an airport network. *Transportation Research Part C: Emerging Technologies*, 27, 60–75. <https://doi.org/10.1016/j.trc.2011.05.017>
- Ravizza, S., Atkin, J. A. D., Maathuis, M. H., & Burke, E. K. (2013). A combined statistical approach and ground movement model for improving taxi time estimations at airports. *Journal of the Operational Research Society*, 64(9), 1347–1360. <https://doi.org/10.1057/jors.2012.123>
- Roling, P. C., & Visser, H. G. (2008). Optimal airport surface traffic planning using mixed-integer linear programming. *International Journal of Aerospace Engineering*, 2008, 1–11. <https://doi.org/10.1155/2008/732828>
- Salihu, A. L., Lloyd, S. M., & Akgunduz, A. (2021). Electrification of airport taxiway operations: A simulation framework for analyzing congestion and cost. *Trans-*

- portation Research Part D: Transport and Environment*, 97. <https://doi.org/10.1016/j.trd.2021.102962>
- Savelsbergh, M. W. P. (1985). Local search in routing problems with time windows. *Annals of Operations Research*, 4(6). <https://doi.org/10.1007/BF02022044>
- Schiphol. (2021, April). *Sustainable taxiing: Taxibot trial*. Retrieved August 10, 2026, from <https://www.schiphol.nl/en/innovation/blog/sustainable-taxiing-taxibot-trial/>
- Sirigu, G., Cassaro, M., Battipede, M., & Gili, P. (2018). Autonomous taxi operations: Algorithms for the solution of the routing problem. *AIAA Information Systems-AIAA Infotech at Aerospace*. <https://doi.org/10.2514/6.2018-2143>
- SKYbrary Aviation Safety Auxiliary Power Unit (APU). (2026).
- Solomon, M. M. (1987). Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research*, 35(2), 254–265. <https://doi.org/10.1287/opre.35.2.254>
- Stettler, M. E. J., Koudis, G. S., Hu, S. J., Majumdar, A., & Ochieng, W. Y. (2018). The impact of single engine taxiing on aircraft fuel consumption and pollutant emissions. *Aeronautical Journal*, 122(1258), 1967–1984. <https://doi.org/10.1017/aer.2018.117>
- TLD Group. (n.d.). Taxibot.
- van Baaren, E., & Roling, P. (2019). Design of a zero emission aircraft towing system. *AIAA Aviation 2019 Forum*. <https://doi.org/10.2514/6.2019-2932>
- van Oosterom, S., & Mitici, M. (2022). Analyzing the impact of battery capacity and charging protocols on the dispatchment of electric towing vehicles at a large airport. *AIAA AVIATION 2022 Forum*. <https://doi.org/10.2514/6.2022-3920>
- 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. <https://doi.org/10.1016/j.trc.2022.103995>
- Zoutendijk, M., Mitici, M., & Hoekstra, J. M. (2023). An investigation of operational management solutions and challenges for electric taxiing of aircraft. *Research in Transportation Business and Management*, 49. <https://doi.org/10.1016/j.rtbm.2023.101019>
- Zoutendijk, M., Van Oosterom, S., & Mitici, M. (2023). Electric taxiing with disruption management: Assignment of electric towing vehicles to aircraft. *AIAA AVIATION 2023 Forum*. <https://doi.org/10.2514/6.2023-4219>

Appendix

Table 10: A.1: Greedy heuristic against the exact model, 4 July 2025, configuration without start-up zones, ILP time limit 3600 s per instance.

K	Greedy EDF served	ILP served	ILP status	ILP runtime (min)
1	40	40	Optimal	0.0
2	72	72	Optimal	0.8
3	94	94	Feasible (time limit)	60.3
4	110	110	Feasible (time limit)	60.1
5	119	120	Feasible (time limit)	60.1
6	123	124	Feasible (time limit)	60.0
7	126	127	Optimal	0.1
8	127	127	Optimal	0.0

Table 11: A.2: Aircraft served by the greedy EDF heuristic per day and fleet size, first week of July 2025, configuration without start-up zones. $K = 0$ is the conventional baseline in which no aircraft is towed.

K	1 July	2 July	3 July	4 July	5 July	6 July	7 July
0	0	0	0	0	0	0	0
1	42	43	41	40	35	36	34
2	75	75	75	72	61	69	63
3	97	98	95	94	79	89	84
4	111	116	108	110	92	103	97
5	118	124	115	119	99	110	107
6	122	128	118	123	103	116	112
7	124	131	121	126	105	118	114
8	125	134	122	127	105	119	115

Table 12: A.3: Aircraft served by the greedy heuristic in release-time order per day and fleet size, first week of July 2025, configuration without start-up zones. N is the number of taxibot-eligible departures that day.

K	1 July	2 July	3 July	4 July	5 July	6 July	7 July
N	205	212	199	218	162	177	183
0	0	0	0	0	0	0	0
1	42	43	42	40	35	36	34
2	75	74	75	72	61	69	62
3	97	97	95	94	80	89	84
4	111	115	109	110	93	103	96
5	118	125	115	119	99	110	107
6	122	128	118	123	103	115	111
7	124	131	121	126	104	117	113
8	125	133	122	127	105	118	115

Table 13: A.4: Sensitivity of aircraft served to the B1 start-up zone capacity, pool model, $K = 3$, W42 capacity 2, 9 July 2025, of 205 pool-eligible departures.

B1 capacity	Greedy	%	ILP	%
1	98	47.8	107	52.2
2	109	53.2	110	53.7
3	109	53.2	110	53.7
4	109	53.2	110	53.7

Table 14: A.5: Sensitivity of aircraft served to the engine warm-up time, $K = 3$, 9 July 2025.

τ_{warmup} (s)	Greedy EDF	%	ILP	%
0	130	61.6	130	61.6
60	124	58.8	124	58.8
120	117	55.5	118	55.9
180	106	50.2	108	51.2

Table 15: A.6: Sensitivity of aircraft served to the drive-back time and to the W42 capacity, pool model, greedy EDF, 4 July 2025, of 211 pool-eligible departures.

K	$\tau_{\text{back}} = 300 \text{ s}$	450 s	600 s	W42 = 2	W42 = 4
1	50	46	42	42	42
2	90	80	76	76	76
3	114	100	97	97	97
4	126	116	114	114	114
5	131	126	123	123	123
6	135	131	129	129	129
7	136	134	133	133	133
8	136	135	134	134	134




Figure 12: A.7: Total departures, taxibot-eligible departures and aircraft served by the greedy EDF heuristic at $K = 3$, per day, July 2025 (top), and the served share of eligible departures (bottom).




Figure 13: A.8: Net CO₂ saved per day at $K = 3$, first week of July 2025, for the pool and no-pool configurations with a diesel and an electric taxibot fleet.

Per-aircraft ILP schedule — 2025-07-04 (K=3)




Figure 14: A.9: Per-aircraft ILP schedule on 04/07/2025 (part 1).




Figure 15: A.9: Per-aircraft ILP schedule on 04/07/2025 (part 2).