Minimizing Airplane Boarding Time

Felix J. L. Willamowski[1⋆] and Andreas M. Tillmann[2]

[1] RWTH Aachen University, Lehrstuhl für Operations Research,
Kackertstr. 7, D-52072 Aachen, Germany
[2] Technische Universität Braunschweig,
Cluster of Excellence SE^{2}A – Sustainable and Energy-Efficient Aviation and
Institute for Mathematical Optimization, Universitätsplatz 2, 38106 Braunschweig, Germany

Abstract. The time it takes passengers to board an airplane is known to influence the turn-around time of the aircraft and thus bears a significant cost-saving potential for airlines. Although minimizing boarding time therefore is the most important goal from an economic perspective, previous efforts to design efficient boarding strategies apparently never tackled this task directly. In this paper, we first rigorously define the problem and prove its NP-hardness. While this generally justifies the development of inexact solution methods, we show that all commonly discussed boarding strategies may in fact give solutions that are far from optimal. We complement these theoretical findings by a simple time-aware boarding strategy with guaranteed approximation quality (under reasonable assumptions) as well as a local improvement heuristic and an exact mixed-integer programming (MIP) formulation. Our numerical experiments with simulation data show that for several airplane cabin layouts, provably high-quality or even optimal solutions can be obtained within reasonable time in practice by means of our MIP approach. We also empirically assess the sensitivity of boarding strategies with respect to disruptions of the prescribed boarding sequences and identify robustness against such disruptions as a bottleneck for further improvements.

1 Introduction and Preliminaries

It is folklore knowledge in the aviation industry that a passenger airplane can only generate revenue while in the air, as ground handling operations and the time a plane spends, e.g., at the gate effectively cost airlines money. Thus, airlines wish to minimize the turn-around times of their airplanes, i.e., the times between the last landing and the next takeoff (alternatively, the time an aircraft spends at the gate after arrival and before the subsequent departure). Although there are many steps involved in turning around an airplane (see, e.g., [Horstmeier and de Haan, 2001, Figure 1] or [Ozmec-Ban et al., 2018, Figure 1]), passenger boarding is one of the steps that can most affect the turn-around time [Ferrari and Nagel, 2005]. This is because all other steps either cannot be significantly shortened (e.g., taxiing the aircraft into its parking position), may run parallel to (and finish during) the boarding process (e.g., baggage handling), or necessarily precede it (like refueling, which is typically not allowed while passengers are on board for safety reasons, or cabin maintenance, avoided for passenger comfort reasons). Indeed, Neumann [2019] showed that under normal operational conditions—i.e., in the absence of factors that are uncontrollable by the airline such as bad weather, and provided sufficient aircraft occupancy levels—boarding can be statistically proven to be on the critical path of the turn-around time with high confidence. Moreover, a growing amount of carry-on luggage has been held responsible for an increase of boarding times, further emphasizing the bottleneck-like role of boarding in a time-efficient turn-around process [Marelli et al., 1998].

A boarding time reduction by just one minute per flight can lead to cost savings of tens to hundreds of million dollars per year for a major airline, cf., e.g., [Neumann, 2019, Horstmeier and de Haan, 2001].[3]

⋆ Corresponding author; email: willamowski@or.rwth-aachen.de

[3] We remark that in this paper, we do not take the present challenges due to the global SARS-CoV-2 pandemic into account, and instead refer to pre-COVID air transport operations, as they are widely assumed to reach similar activity levels again after the pandemic will be overcome. Modifications to boarding procedures in response to pandemic-related novel requirements to reduce infection risks—in particular, distancing rules—are discussed in, e.g., [Schultz and Fuchte, 2020, Milne et al., 2021].

Similarly, boarding being on the critical turn-around path also implies that preventing boarding delays avoids extra costs incurred by prolonged turn-around periods. Furthermore, a reduced boarding time is also beneficial for passengers and airport operators. Passengers should generally appreciate earlier boarding completion, and airport operators are possibly able to offer more flights per day per gate. For a more detailed discussion of benefits and challenges, we refer to the fairly recent survey [Jaehn and Neumann, 2015] on boarding methods.

Somewhat surprisingly, it appears that previous works on reducing airplane boarding time have never explicitly formulated the task as an optimization problem with the objective to directly minimize the overall boarding time. Existing strategies are either heuristics or, if based on optimization, target other objectives like passenger interference in aisles or seat rows. The asymptotic analyses in [Bachmat et al., 2009] and related works (cf. discussion in Section 1.1 below) focus on minimum average boarding times of highly randomized boarding schemes, allowing to derive high-level insights into practical boarding behavior but not certifiable instance optimality. Indeed, the need for research into ways of finding boarding sequences that actually minimize boarding time was pointed out in [Jaehn and Neumann, 2015, Sect. 7.1.1].

In this paper, we introduce the (to the best of our knowledge) first rigorous mathematical optimization model, utilizing mixed-integer programming (MIP), for the explicit goal to *minimize the overall boarding time*. We furthermore introduce a novel heuristic—a time-aware variant of the well-known outside-in strategy—and a local-search post-processing routine to further improve a given boarding sequence. Moreover, although the airplane boarding problem has long been recognized as challenging, a formal proof of intractability was lacking. We close this gap by providing several computational complexity and (in-)approximability results, under different assumptions on the input time data (w.r.t. passengers moving through the plane and settling in at their seats) that reflect varying degrees of information resolution from the "uninformed" use of rough estimates identical for each passenger (as may be used in practice) to "fully informed" knowledge of individual-passenger moving and settle-in times (as in simulations). With extensive numerical experiments, we demonstrate, in particular, the competitiveness of our heuristics and that our MIP allows to reliably solve (at least) medium-sized instances to certifiable optimality within reasonable time. Moreover, even if terminating the solution process early, the MIP approach still provides previously unknowable solution quality guarantees and approximation error bounds for heuristics. Indeed, while previous simulation studies to identify "the best" boarding strategies could only deliver *relative* solution quality assessments (i.e., how well a strategy performs relative to others), our approach allows to quantify solution quality in an *absolute* sense (i.e., compared to the actual optimum) for what appears to be the first time.

The paper is organized as follows: In the remainder of this section, we first review the previous literature on airplane boarding strategies (Section 1.1), and then formally define the boarding time minimization problem treated in this paper (Section 1.2) and discuss its underlying assumptions and possible limitations (Section 1.3). In Section 2, we investigate the computational complexity of the problem, deriving strong NP-hardness results as well as approximation guarantees; for readability purposes, all formal proofs have been deferred to an Appendix at the end of the paper. Section 3 introduces a compact mixed-integer programming (MIP) formulation to solve the problem to optimality. The numerical experiments presented in Section 4 demonstrate the practicality of the proposed approaches and the potential gains compared to several previously known boarding strategies that do not explicitly consider the time-minimization objective. In further experiments, we also assess sensitivity/robustness of the boarding sequences w.r.t. time-data perturbations and disruptive passenger behavior. In Section 5, we discuss and evaluate extensions that further increase model realism. Finally, concluding remarks are given in Section 6.

1.1 Related Work

Many different approaches for shortening the boarding time have been presented in the literature; the paper by Jaehn and Neumann [2015] provides a very detailed discussion and survey of many important previous works (see also Coppens et al. [2018]). Although computational intractability had, to the best of our knowledge, not been formally established prior to the present work (cf. Section 2), the majority of methods that have been proposed are simple heuristic schemes that provide easy-to-implement boarding strategies.

Generally, one can distinguish between *by-group* and *by-seat* strategies. By-group boarding rules partition the passengers into several groups and prescribe an ordering for the groups, but passengers within the groups can board the plane in arbitrary order. In contrast, by-seat strategies prescribe a specific position in the boarding sequence for each seat (passenger), usually taking individual instance characteristics such as passenger walking speeds into consideration.

It seems that most if not all airlines currently employ by-group strategies, see, e.g., [Delcea et al., 2018], who nonetheless point out a 2013 pilot study of KLM using a by-seat strategy, the results of which do not appear to have been made public by the airline. The prevalence of by-group strategies can be explained (at least partially) by the fact that such strategies require very little supervision, or control, of the boarding process. While in general, by-group schemes cannot guarantee solution quality for any specific flight, enforcing a by-seat strategy would necessitate much stricter control, which complicates practical realization and is thought to be problematic w.r.t. passenger acceptance—though both control requirements and passengers' willingness to comply might change in the wake of the coronavirus pandemic. On the other hand, clearly, an optimal boarding sequence (for a given specific flight) will always be of a by-seat nature, so there is a general trade-off between economic goals and passenger satisfaction for airlines to consider.

For by-group strategies, the boarding process is of inherently stochastic nature (cf., e.g., [Horstmeier and de Haan, 2001, Bachmat et al., 2009, Schultz, 2018b]). Thus, a large body of works employ computer simulations to compare different boarding strategies over large sets of randomly generated instances to identify those that consistently work better than others, see, for instance, [Marelli et al., 1998, van Landeghem and Beuselinck, 2002, Nyquist and McFadden, 2008, Steffen, 2008, Jaehn and Neumann, 2015, Jafer and Mi, 2017, Kierzkowski and Kisiel, 2017, Ozmec-Ban et al., 2018, Schultz and Schmidt, 2018, Schultz, 2018b]. Extensive simulation allows to draw conclusions on the empirical average performance of different strategies (depending on the distributions of random data) and to relatively easily analyze other aspects such as the impact of cabin configuration changes (e.g., multiple doors) in the same fashion. The most basic by-group strategies that find consideration in many such studies are *random boarding* (all passengers are treated as a single group and can just enter the airplane in an arbitrary order) and the *back-to-front* strategy (filling the airplane from the last row to the first, with passengers seated in one row boarding in arbitrary order among their group), and their variations that further subdivide the passengers into a few larger blocks (e.g., comprising several successive rows) that board together in a certain order (e.g., in back-to-front fashion). Typically[4], notable boarding time reductions compared to these traditional strategies can be obtained by means of more sophisticated strategies such as, in particular, the *outside-in* strategy [Marelli et al., 1998], also known as "WilMA", where passengers with window seats board first, followed by those with middle seats, and aisle seats last, or the *reverse pyramid* scheme [van den Briel et al., 2005, Nyquist and McFadden, 2008], a hybrid of back-to-front and outside-in with "diagonal loading" in several groups so that at the back of the plane, boarding passengers have seats closer to the aisle than those within the same group that are seated in rows toward the front.

Besides the categorization as "by group", the aforementioned boarding schemes may also be called *static* (cf. Zeineddine [2017]), since they prescribe the same respective strategies regardless of individual flight and passenger properties. Most strategies that do take such individual flight/passenger-specific characteristics into consideration indeed fall into the "by seat" category (discussed further below), and vice versa. Notable exceptions are the so-called *Steffen method* [Steffen, 2008] (a static by-seat strategy that boards passengers in back-to-front fashion, starting with passengers who have window seats in odd-number rows, then those with window seats in even-number rows, middle-seats in odd rows, and so on) and the *slow-first* strategy [Erland et al., 2019] (a by-group strategy that utilizes passenger aisle-clearing time estimates, e.g., based on the size of carry-on luggage pieces, to partition the passengers into two groups, and lets the slow group board before the faster one, with arbitrary ordering within the groups). A fairly recent ranking of various boarding strategies including those mentioned above (except slow-first) appears in [Kierzkowski and Kisiel, 2017], see also [Ozmec-Ban et al., 2018]; according to this, the smallest average boarding times are achieved by the Steffen method, followed by the much simpler outside-in strategy and the reverse pyramid scheme.

[4] Simulation results may vary substantially due to the use of different underlying data, though this issue has somewhat dissipated in more recent years as the considered data distributions became more realistic and also more similar.

A modification of the Steffen method that boards passengers in four groups (right side of the aisle, then left side first for odd-numbered rows, then analogously for the even-numbered rows), with arbitrary order within the groups, was meant to be easier to realize in practice, but turned out to be inferior to other strategies. Further simulations have focused on, e.g., the impact of different cabin interior designs [Schultz and Schmidt, 2018] or hand-luggage placement [Coppens et al., 2018].

Complementing these empirical studies based on simulations and occasional field trials (cf., e.g., Nyquist and McFadden [2008], Steiner and Philipp [2009], Schultz [2018a], Hutter et al. [2019]), there is a series of papers that provide asymptotic mathematical analyses of average (or expected) boarding times for different by-group boarding schemes, see, in particular, [Bachmat et al., 2009, Bachmat, 2013, Bachmat et al., 2013, Erland et al., 2019]. This line of work shows that boarding can be viewed as a stochastic process and modeled as a physical system with deep connections to space-time (or Lorentzian) geometry, which yields provable probabilistic bounds on the average boarding completion times in the asymptotic regime, i.e., when the number of passengers tends to infinity. Although asymptotic (and based on some assumptions like infinite walking speeds), these bounds have been found to match empirical findings (with finite numbers of passengers) quite well, and can thus offer a theoretical explanation for simulation results, high-level insights into the relative behavior of different boarding strategies, and estimates for the general order of average boarding completion times. Moreover, the bounds depend on a parameter called *congestion*, defined as the ratio of the total length of a single-file queue of all passengers and the aisle length of the aircraft. Thus, congestion decreases with increased space between rows of seats, or when more passengers are assumed to fit into the space of one row. Analyzing boarding strategies for different congestion factors allows somewhat flexible conclusions, e.g., that a back-to-front policy can notably reduce expected boarding times in comparison to random boarding for congestion parameters at most 1 [Bachmat and Elkin, 2008], that using more than two groups in back-to-front boarding is ineffective [Bachmat et al., 2013], or that the aforementioned two-group slow-first strategy outperforms an analogous fast-first strategy, back-to-front and random boarding [Erland et al., 2019]; in practice, the congestion can often be argued to be about 4, cf. [Bachmat et al., 2009]. It is also worth mentioning that further functional expressions to estimate average boarding times for select strategies have been derived in, e.g., [Frette and Hemmer, 2012, Kaupužs et al., 2019, Hutter et al., 2019].

By-seat strategies have apparently been considered far less in the literature, presumably due to potential difficulties in enforcing them in practice. On the one hand, there are some few variants of more classical (non-deterministic by-group) boarding strategies that prescribe certain (deterministic by-seat) concrete sequences; one example is the aforementioned Steffen's method [Steffen, 2008], another is a hybrid of back-to-front and outside-in boarding that had been considered in, e.g., [van Landeghem and Beuselinck, 2002, Delcea et al., 2018] and which we will refer to as *outside-in back-to-front (outside-in BTF)* throughout this paper. On the other hand, previous attempts to improve boarding have soon begun taking individual flight/passenger data into account to determine a boarding sequence by optimization. However, these works have apparently always circumvented the actual objective of minimizing the boarding completion time. For instance, the paper [Zeineddine, 2017] proposed an adaptive queueing scheme that mainly targets a reduction of the number of on-board (aisle and seat) interferences, i.e., situations where some passenger is kept from proceeding by another who is, e.g., currently loading carry-on luggage into an overhead bin or sitting in the way to the target seat (cf. [Ferrari and Nagel, 2005]). Avoiding such blockages is meant to improve the passengers' perception of the boarding process (as waiting is a "source of annoyance"), and also has an intuitive tangible impact on the overall boarding time. However, it is easily conceived that an interference-minimal boarding sequence is not automatically also a solution that actually minimizes the boarding time. Moreover, in the last few years, mixed-integer programming (MIP) techniques were applied to some related boarding problems. In [Milne and Salari, 2016], the claimed goal is to minimize the boarding time, but differing from the actual boarding time minimization considered here, a seat-assignment MIP problem based on carry-on luggage information is solved so that *the Steffen method* will result in short boarding times (see also [Milne and Kelly, 2014]). A linear MIP to minimize aisle and seat interferences was proposed in [Bazargan, 2007], similar to an earlier nonlinear MIP (MINLP) introduced in [van den Briel et al., 2005]; in the latter work, it is stated that "To make the problem more tractable, we used the minimization of passenger interferences as our objective in lieu of the minimization of boarding time".

Moreover, van den Briel et al. [2005] note their model to be a nonlinear assignment problem, which *in general* is an NP-hard class of problems; however, they do not provide a proof that their concrete problem is still NP-hard itself. Another MIP for interference minimization was put forth in [Soolaki et al., 2012], along with a genetic algorithm that sometimes achieved better results than commercial MIP solvers under solution time limits.

It is also worth mentioning that robustness of boarding strategies has been treated in the literature as well, even in a MIP context (see [Milne et al., 2018]), but, to the best of our knowledge, not in the present setting (actual boarding time minimization) and extent (combining time fluctuations, passengers changing positions in the boarding queue, and/or arriving late at the gate). The most extensive work in this direction we are aware of is [Ferrari and Nagel, 2005], a broad simulation study involving early and late passengers, yet different from how we treat disruptions (cf. Section 4.3). For by-group strategies where passengers within groups board in a random order, random disruptions are naturally less likely to cause significant deterioration, but for by-seat strategies, deviation from the prescribed boarding sequence may cause more severe delays. To the best of our knowledge, no extensive study has yet been conducted to look closely into these effects—especially, due to their unavailability prior to the present work, with respect to *optimal* boarding sequences.

For further, more detailed discussions of the many different (static) boarding strategies and the comparisons and conclusions from the extensive simulation studies, we refer, in particular, to [Jaehn and Neumann, 2015, Nyquist and McFadden, 2008, Bachmat et al., 2009, Coppens et al., 2018] and the references therein. Notably, it seems the (economic) main goal of actually minimizing the boarding time has not been tackled explicitly before, and while it is understandable that, e.g., passenger boarding comfort is taken into consideration—an argument for simple, easy-to-implement by-group strategies—it has already been demonstrated that by-seat strategies can lead to greater time savings. Therefore, in this paper, we focus solely on the economic side, i.e., on determining (by-seat) sequences that minimize the boarding time. In principle, such sequences might be realized by roughly dividing the passengers into groups that are asked to gather in certain areas at the gate, where they are then "sorted" by (or with the help of) gate agents into the intended boarding order, similar to the KLM field test described in [Delcea et al., 2018], or properly construed pre-boarding areas (see also [Steiner and Philipp, 2009]). Nevertheless, we leave practical details of how to enforce by-seat boarding sequences at the gate for future consideration.

Finally, we remark that minimizing airplane boarding time bears some resemblance to certain (permutation) flow shop scheduling problems with makespan (i.e., completion time) minimization objective; we refer to [Brucker, 2007, Pinedo, 2016] for details on machine scheduling, and will occasionally give pointers to related results at the appropriate places in the paper.

1.2 Problem Definition and Notation

We consider the *airplane boarding problem (ABP)* that asks for a sequence in which to board passengers with preassigned seats to an airplane such that the overall boarding time is minimized. We focus on the setting that is most prevalent in the existing literature: there is one entrance at the beginning of the passenger cabin to be used for boarding (through a jet-bridge), and the cabin consists of a single deck with seats to either side of a single aisle. By *(overall) boarding time*, we mean the time between the first passenger entering the airplane cabin and the last passenger sitting down on their given seat. Moreover, we adopt the common assumption of "single-class boarding", i.e., we leave out priority boarding (e.g., first class) and other pre-boarding rounds in our considerations. To formalize things, we need to introduce some notation.

The sets of rational, integer, and natural numbers are denoted by $\mathbb{Q}$, $\mathbb{Z}$, and $\mathbb{N}$, respectively; possible restrictions are indicated by suitable subscripts (e.g., $\mathbb{N} = \mathbb{Z}_{\geq 1}$). For a number $n \in \mathbb{N}$, we write $[n] \coloneqq \{1, 2, \ldots, n\}$. An airplane's *cabin layout* parameters are the (ordered) set of rows $\mathcal{R} \coloneqq [R]$, $R \in \mathbb{N}$ (the row closest to the door being row 1, and all later rows $r \leq R$ being accessible by first passing rows $1, 2, \ldots, r-1$), and the collection $\mathcal{S}$ of seats (given in the form (r, s) for each $r \in \mathcal{R}$ with seat numbers s). We further require a set of *passengers* $\mathcal{P} \coloneqq [P]$, $P \in \mathbb{N}$ with $P \leq |\mathcal{S}|$, unique *passenger-seat assignments* given by an injective $\sigma : \mathcal{P} \to \mathcal{S}$, $p \mapsto (r(p), s(p))$, and for each passenger $p \in \mathcal{P}$, a *settle-in time* $t^s_p \in \mathbb{Q}_{\geq 0}$ (consisting of the time passenger p takes to stow away their carry-on luggage, move within row $r(p)$, and finally sit down at the assigned seat $s(p)$) and *moving times* $t^m_{p,r} \in \mathbb{Q}_{\geq 0}$ (for passing row $r \in \mathcal{R}$) for $r \leq r(p) - 1$.

Now, the task is to find a permutation $\pi \in \Pi_P$ (i.e., a one-to-one mapping describing the passengers' boarding sequence) that minimizes the overall boarding time. Borrowing the common notation for completion times from the scheduling literature (cf., e.g., [Brucker, 2007]), we denote the boarding time induced by a sequence π by $C_{\max}(\pi)$ (or simply $C_{\max}$ if π is clear from the context). Thus, we can abstractly express the ABP of interest as:

Definition 1 (Airplane Boarding Problem (ABP)). *Given an airplane cabin layout $(\mathcal{R}, \mathcal{S})$, a finite set of passengers $\mathcal{P}$ with injective seat assignments $\sigma : \mathcal{P} \to \mathcal{S}$, and moving and settle-in times $t^m_{p,r}, t^s_p \in \mathbb{Q}_{\geq 0}$, find a permutation π of the passengers with minimum $C_{\max}(\pi)$.*

A few further remarks and details on the ABP appear in order. First, consider the cabin layout: For each row $r \in \mathcal{R}$, there are two (possibly empty) ordered sets of seats $\mathcal{S}^1_r := ((r,1), \ldots, (r,k^1_r))$ and $\mathcal{S}^2_r := ((r,k^1_r+1), \ldots, (r,k^1_r+k^2_r))$, with $k^1_r, k^2_r \in \mathbb{Z}_{\geq 0}$; thus, $\mathcal{S} = \dot{\bigcup}_{r \in \mathcal{R}} (\mathcal{S}^1_r \dot{\cup} \mathcal{S}^2_r)$. Each seat $(r,s) \in \mathcal{S}$ is accessible from the aisle by passing seats $(r,k^1_r), \ldots, (r,s+1)$ if $s \leq k^1_r$, or $(r,k^1_r+1), \ldots, (r,s-1)$ otherwise (i.e., for any r, seats (r,k^1_r) and (r,k^1_r+1) are the aisle seats, and $(r,1)$ and $(r,k^1_r+k^2_r)$ are the window seats). Moreover, by $\mathcal{P}(r)$ we will denote the set of passengers with seats in row r.

Furthermore, we presume that each passenger p tries to go directly to their seat without unnecessary "loitering" and that p settles in (in particular, stows away carry-on luggage) at the row $r(p)$, from which the assigned seat $s(p)$ is accessible. Following common modeling practice, we also assume that overtaking in the aisle is not possible, so a passenger can proceed to some row only if it is not presently occupied by another passenger. More precisely, let p be a passenger who has started some action (moving, waiting, or settling in) at row r at time $t \in \mathbb{Q}_{\geq 0}$. If p passes row r or settles in at it (so $r = r(p)$), he/she *occupies*, or *blocks*, the row for the time period $[t, t + t^m_{p,r})$ or $[t, t + t^s_p)$, respectively. Analogously, in case p has to wait a time period $w \in \mathbb{Q}_{\geq 0}$ at r (i.e., p has passed row r but row $r+1$ is blocked for w more time steps), he/she blocks it for the time period $[t, t+w)$.

Finally, while there does not appear to exist a general closed-form expression for $C_{\max}(\pi)$, we can evaluate the boarding completion time induced by any sequence π recursively as follows. The recursion is similar to that for the makespan of a permutation schedule in flow shops with limited intermediate storage, cf. [Pinedo, 2016, p. 162]; a proof by induction is provided in Appendix A.

Lemma 1. *Given an ABP instance $(\mathcal{P}, \mathcal{R}, \mathcal{S}, \sigma, \{t^m_{p,r}\}, \{t^s_p\})$ and a permutation $\pi \in \Pi_P$, it holds that $C_{\max}(\pi) = \max_{r \in \mathcal{R}} C(P,r)$, where $C(P,r)$ is defined by the following recursion over the times $C(i,r)$ at which a row r becomes accessible again after being blocked by a passing or settling-in passenger $p \in \{\pi^{-1}(1), \ldots, \pi^{-1}(i)\}$:*

$$C(i,r) := \begin{cases} \max \left\{ \begin{smallmatrix} C(i-1,r)+t^m_{\pi^{-1}(i),r}, \\ C(i,r-1)+t^m_{\pi^{-1}(i),r}, \\ C(i-1,r+1) \end{smallmatrix} \right\} & \text{if } 1 \leq r \leq r(\pi^{-1}(i)) - 1, \\ \max \left\{ \begin{smallmatrix} C(i-1,r)+t^s_{\pi^{-1}(i)}, \\ C(i,r-1)+t^s_{\pi^{-1}(i)} \end{smallmatrix} \right\} & \text{if } r = r(\pi^{-1}(i)), \\ C(i-1,r) & \text{if } r(\pi^{-1}(1)) + 1 \leq r \leq R \end{cases}$$

for $i \in \{2, 3, \ldots, P\}$, with $C(i,0) := 0$ $(i \geq 2)$ and

$$C(1,r) := \begin{cases} \sum_{r'=1}^{r} t^m_{\pi^{-1}(1),r'} & \text{if } 1 \leq r \leq r(\pi^{-1}(1)) - 1, \\ \sum_{r'=1}^{r-1} t^m_{\pi^{-1}(1),r'} + t^s_{\pi^{-1}(1)} & \text{if } r = r(\pi^{-1}(1)), \\ 0 & \text{if } r(\pi^{-1}(1)) + 1 \leq r \leq R. \end{cases}$$

Consequently, $C_{\max}(\pi)$ can be computed in $\mathcal{O}(PR)$ time.

Remark 1. In the context of simulations to evaluate boarding methods, the time data $(\{t^m_{p,r}\}, \{t^s_p\})$ is typically drawn from a random distribution for each instance of many trials to be averaged. We will distinguish between different levels of "resolution" regarding the time data in our numerical experiments and, in particular, in Section 2, where we will show that this has an impact on the theoretical difficulty of the ABP. Note that in practice, highly detailed time data is likely hard or even impossible to come by, and one has to work with simplifications and estimates that may be available from real-life observations, based on passenger demographics (cf., e.g., [Marelli et al., 1998]), or recent machine learning predictions [Schultz and Reitmann, 2019]. Nevertheless, the distinction between different resolutions of time data is then still crucial to evaluate the empirical (simulation) performance of boarding methods based on simplified data versus the true optimal values achievable with fine-resolution data. Thus, the complexity results for different resolutions (Section 2) indicate, among other aspects, the necessary computational effort to obtain these optimal values, and our MIP formulation (Section 3) will provide a means to do so in practice, even for hard cases corresponding to the most realistic simulation data.

Remark 2. The aforementioned blockages and resulting waiting periods are what has been referred to as *aisle interferences* in the literature, see the earlier discussion of related work. For simplicity, throughout most of the paper, we do not explicitly incorporate further aspects such as *seat interferences* (which are arguably less important than aisle interferences, cf. [van den Briel et al., 2005, Nyquist and McFadden, 2008, Kierzkowski and Kisiel, 2017]) or *groups of passengers wishing to board together* (e.g., couples, or families with small children). Nevertheless, we will later show how to integrate such additional constraints into our exact ABP model and assess the impact on the exact (MIP) and some heuristic boarding schemes in numerical experiments, cf. Section 5.

1.3 Limitations and Scope

Any assessment of boarding strategies comes with typical assumptions that can affect a broad range of aspects, from aircraft cabin layout (such as the single-aisle, single-door setup considered here and in most papers) to simulation/model data (such as infinite passenger walking speeds). The present work is, of course, no exception.

Generally, in practice, the boarding problem suffers from unavailability of true data (moving and settling-in times). Typically, data is randomly generated and average results over many simulations are used to obtain indications for real-world behavior, assuming the data models are sufficiently realistic. Moreover, as mentioned earlier, it may pose new challenges to implement by-seat boarding strategies in practice. The approach in the present paper obviously cannot circumvent these general limitations, since we treat computing optimal boarding sequences—on the one hand, we require instance data to do so, and on the other hand, optimal boarding sequences will be of a by-seat nature. Data aspects are discussed further in the context of boarding problem complexity (see Section 2) and numerical results (see Section 4), respectively; we do not delve into how by-seat strategies might be implemented in practice.

The arguably most restrictive assumption w.r.t. actual boarding practice that is common to very many simulation studies is the disregard of inseparable groups of passengers (like couples or families) that will want to board together. Similarly, our main theoretical results (cf. Section 2) do not take such passenger groups into consideration yet; likewise, seat interferences are presently not explicitly considered there. While our NP-hardness results remain valid for the correspondingly extended ABP variants, generalizing the polynomial-time solvability and approximability results is out of scope for the present work and thus, for now, remains an important open task. Nevertheless, as mentioned earlier, the MIP model we propose (see Section 3) *can* be refined to account for both inseparable passenger groups and seat interferences, as shown and supported by further numerical experiments in Section 5.

Furthermore, the asymptotic analysis in [Bachmat et al., 2009] and follow-up works revealed that the congestion parameter has a tangible influence on the boarding time. In simulation studies, the congestion is typically fixed, e.g., at the value 4 that has been deemed reasonable for practical purposes or the value 6 (both give similar results). Throughout the present paper, we also work with a fixed congestion level: by assuming for convenience that there can only be one passenger per row, we implicitly set the congestion parameter equal to the number of seats per row (for a fully occupied plane). For the medium-sized instances in our experiments, this coincidentally amounts to the common value 4, but it is 2 and 6 for the smaller and larger instances, respectively. In Section 5, we will briefly show how our MIP model could be adapted

Case (for all $p \in \mathcal{P}$, $r \in \mathcal{R}$)	Complexity	Approx./Algo.
$t^m_{p,r} = c^m$, $t^s_p = c^s$.	$\mathcal{O}(P)$ (Thm. 3)	exact scheme outside-in BTF (Thm. 3)
$t^m_{p,r} = c^m$, $t^s_p \in \mathbb{Q}_{\geq 0}$ arbitrary.	strongly NP-hard (Thm. 2)	$\big(1 + \frac{k-1}{k}(\sum_{i=1}^{k} \frac{1}{i})\big)$-approximation max-settle-row (Thm. 6)
$t^m_{p,r} = c^m_p$, $t^s_p = c^s$.	open	k-approximation outside-in BTF (Thm. 4)
$t^m_{p,r} = c^m_p$, $t^s_p \in \mathbb{Q}_{\geq 0}$ arbitrary.	strongly NP-hard (Thm. 2)	MIP
$t^m_{p,r} \in \mathbb{Q}_{\geq 0}$ arbitrary, $t^s_p = c^s$.	strongly NP-hard (Thm. 1)	MIP
$t^m_{p,r} \in \mathbb{Q}_{\geq 0}$ arbitrary, $t^s_p \in \mathbb{Q}_{\geq 0}$ arbitrary.	strongly NP-hard (Thm. 1, Thm. 2)	MIP

Table 1. Overview of main theoretical and algorithmic contributions. Main results of this paper on computational complexity, approximation guarantees and algorithmic approaches to the boarding time minimization problem, under different assumptions on the input time data; c^m, c^s, and $c^m_1, \ldots, c^m_P$ are constants in $\mathbb{Q}_{\geq 0}$. Here, k is the number of seats per row, and MIP refers to our universally applicable exact solution approach. The results for the polynomial schemes in the first three cases are shown to hold for planes fully booked to capacity (i.e., $|\mathcal{P}| = k \cdot |\mathcal{R}|$).

to different congestion parameter values. We will also occasionally comment on the congestion aspect, but a comprehensive theoretical and empirical assessment is well beyond the scope of this work and is thus left for future consideration. Similarly, our approximability results assume fully occupied aircraft.

2 Computational Complexity and Approximation

The availability of detailed passenger-specific moving and settle-in times pertains to an ideal situation that would allow for the most considerate boarding sequence planning. In practice, such "real" data will hardly be available. Reasonable estimates can be obtained from real-life observations (described in several studies, for instance in [van Landeghem and Beuselinck, 2002]; cf. also novel on-plane measurement devices like the sensor network to monitor boarding processes described in [Schultz and Schmidt, 2018] and references therein) and used to derive suitable (probabilistic) model assumptions, see, e.g., the simulation approaches in [Milne and Kelly, 2014, Budesca et al., 2014]. Nonetheless, as noted in [Jaehn and Neumann, 2015], to actually identify the fastest boarding sequence, simplifying assumptions such as, in particular, identical walking speed for all passengers, would have to be done away with. After all, one can only actually judge how well heuristics using simplified time data really work if one can compare to optimal solutions based on detailed data, in simulations.

In the following, we therefore investigate the theoretical complexity of the boarding time minimization problem (i.e., ABP) under a variety of assumptions on the time data. Namely, we discuss all combinations of fully individual or identical-for-all-passengers settle-in times with fully individual (row-dependent), semi-individual (constant for all rows for any given passenger) and identical-for-all-passengers moving times, cf. Table 1. We remind that the technical proofs are all deferred to the Appendix.

2.1 Intractability Results

In general, the ABP turns out to be computationally intractable, even for very simple airplane cabin layouts, as shown by the following two results for the special cases with either constant settle-in times or constant moving times, respectively.

We start off with the case $t^m_{p,r} \in \mathbb{Q}_{\geq 0}$ $\forall p \in \mathcal{P}, r \in \mathcal{R}$ and $t^s_p = c^s$ $\forall p \in \mathcal{P}$ (with $c^s \in \mathbb{Q}_{\geq 0}$ a constant):

Theorem 1. *It is NP-complete in the strong sense to decide whether a given ABP instance admits a boarding sequence π that completes boarding by a given time $T \in \mathbb{Q}_{>0}$, even if restricted to instances with one seat per row, all-zero settle-in times, integer moving times, and integer T.*

In the above setting, the ABP is similar to permutation flow shop scheduling with blockages, which we exploit in the proof (see Appendix B). A different proof (cf. Appendix C) yields the companion result for the case $t^m_{p,r} = c^m \ \forall p \in \mathcal{P}, r \in \mathcal{R}$ (with $c^m \in \mathbb{Q}_{\geq 0}$ a constant) and $t^s_p \in \mathbb{Q}_{\geq 0} \ \forall p \in \mathcal{P}$:

Theorem 2. *It is NP-complete in the strong sense to decide whether a given ABP instance admits a boarding sequence π that completes boarding by a given time $T \in \mathbb{Q}_{>0}$, even if restricted to instances with all-zero moving times, integer settle-in times, and integer T.*

The above two core intractability results for the ABP implicitly (by virtue of the respective reductions) assume that the plane is fully booked, i.e., that $P \coloneqq |\mathcal{P}| = |\mathcal{S}|$. It is, however, easy to show that NP-hardness persists in case some seats are left empty or were filled earlier:

Corollary 1. *The strong NP-hardness results of Theorems 1 and 2 remain valid in their respective settings even if $P < |\mathcal{S}|$, i.e., after pre-boarding and/or if the aircraft is not fully occupied.*

Moreover, clearly, if the congestion parameter (cf. Section 1.1) is considered part of the ABP input, NP-hardness persists, as the instances constructed in the respective proofs show hardness for a specific choice (namely the number of seats per row). However, while it may be possible to modify our proofs accordingly, NP-hardness presently remains open for other *fixed* congestion levels.

Remark 3. Recall that NP-hardness *in the strong sense* implies that, unless P = NP, neither a polynomial-time exact solution algorithm nor a fully polynomial-time approximation scheme (FPTAS) can exists, see, e.g., [Garey and Johnson, 1979] for details. In particular, the "strong sense" assertion may be viewed as evidence that a problem's intractability is not due to possible ill-conditioning of problem data, but reflects its inherent combinatorial nature.

2.2 Approximability Results

In view of the fact that in practice, one will have to work with estimates of the passengers' moving and settle-in times, the following result justifies the use of *pessimistic* such estimates to retain a certain planning robustness:

Proposition 1. *For an optimal ABP solution π with completion time $C_{\max}(\pi)$, should some t^m- or t^s-values actually be smaller than the data used to obtain π, the resulting actual boarding completion time of π still is at most $C_{\max}(\pi)$.*

Moreover, it can make sense to consider certain simplifications of ABP (more precisely, of its data model), in particular in combination with Proposition 1. In the simplest case, one may neglect passenger differences and simply assume identical moving times and identical settle-in times for all passengers (i.e., $t^m_{p,r} = c^m \ \forall p \in \mathcal{P}, r \in \mathcal{R}$ and $t^s_p = c^s \ \forall p \in \mathcal{P}$, with constants $c^m, c^s \in \mathbb{Q}_{\geq 0}$). In that case, and assuming for simplicity that each row has the same number of seats (i.e., $|\mathcal{S}^1_r| + |\mathcal{S}^2_r| = k \in \mathbb{N}$ for all $r \in \mathcal{R}$) and that the aircraft is booked to capacity (every seat is occupied), ABP becomes solvable in polynomial time by the simple boarding strategy that has been called "by seat descending row and then by letter" [van Landeghem and Beuselinck, 2002] or "WilMA-back-to-front by-seat" [Delcea et al., 2018]; we formally define it below under the name *outside-in back-to-front* (*outside-in BTF*, for short). Throughout, we will focus on outside-in BTF boarding and a novel variant of it, since various simulation studies have consistently demonstrated outside-in variants to be among the best boarding heuristics (see also Section 1.1).

Definition 2 (outside-in back-to-front strategy (outside-in BTF)). *Let the passengers board the plane in ordered subsets of size R (one passenger per row per set), ordered first by window-to-aisle (left-right alternating) seat numbers (identifying the sets) and then by decreasing row number (within each set).*

It is important to keep in mind that the outside-in BTF strategy is a by-seat strategy, and hence is not to be confused with the common by-group strategy usually referred to as outside-in (cf. Section 1.1). We will encounter this latter strategy only later, in the computational experiments, where we will call it "randomized outside-in" to distinguish it from the outside-in BTF method. Furthermore, it is interesting to note that the boarding sequence produced by outside-in BTF essentially (up to inconsequential rearrangements) coincides with that prescribed by Steffen's method [Steffen, 2008] if one assumes that each passenger occupies not one but two rows of space; conversely, had Steffen assumed passengers to occupy one row-length of space, one may presume he would have suggested the strategy we now call outside-in BTF.

Assuming w.l.o.g. that $k_r^1 \geq k_r^2 \ \forall r \in \mathcal{R}$, the position in the outside-in BTF boarding sequence of a passenger $p \in \mathcal{P}$ can be directly expressed as

$$\pi(p) = \begin{cases} \big(2s(p)-1\big)R + 1 - r(p), & \text{if } s(p) \leq k_r^2, \\ \big(k_r^2 + s(p)\big)R + 1 - r(p), & \text{if } k_r^2 < s(p) \leq k_r^1, \\ \big(k - s(p) + 1\big)R + 1 - r(p), & \text{if } s(p) \geq k_r^1 + 1. \end{cases} \tag{1}$$

Recall that in our notation, seats are numbered $1, 2, \dots, k_r^1$ from window to aisle on one side, and then continue from aisle to window as $k_r^1+1, k_r^1+2, \dots, k_r^1+k_r^2$ on the other side, in any row r; also, here, $|\mathcal{S}_r^1| + |\mathcal{S}_r^2| = k$ implies $k_r^2 = k - k_r^1$, and the sequence of alternating seat numbers then is $1, k, 2, k-1, 3, k-2, \dots$.

Theorem 3. *A minimum-boarding-time solution to ABP instances with $P = |\mathcal{S}|$, $|\mathcal{S}_r^1| + |\mathcal{S}_r^2| = k \in \mathbb{N}$ for all $r \in \mathcal{R}$, $t_p^s = c^s \in \mathbb{Q}_{\geq 0}$ for all $p \in \mathcal{P}$, and $t_{p,r}^m = c^m \in \mathbb{Q}_{\geq 0}$ for all $(p,r) \in \mathcal{P} \times \mathcal{R}$ can be computed in time $\mathcal{O}(P)$ by the outside-in BTF strategy.*

The proof (see Appendix F) makes use of the following two auxiliary results providing upper and lower bounds on the ABP objective, respectively, that may be of independent interest.

Lemma 2. *For ABP instances with $P = |\mathcal{S}|$ and $|\mathcal{S}_r^1| + |\mathcal{S}_r^2| = k \in \mathbb{N}$ for all $r \in \mathcal{R}$, every boarding strategy producing a sequence π such that*

$$\mathcal{P}(R-i+1) = \big\{\pi^{-1}(i), \pi^{-1}(i+R), \dots, \pi^{-1}(i+(k-1)R)\big\} \qquad \forall i \in \mathcal{R} \tag{2}$$

(i.e., generalized outside-in BTF with arbitrary passengers per row per group, not necessarily according to seating order from window to aisle) yields a boarding completion time $C_{\max}(\pi)$ of at most

$$\sum_{i=0}^{k-1} \max_{j \in \mathcal{R}} \left\{ \sum_{r=1}^{r\left(\pi^{-1}(j+iR)\right)-1} \left(\max_{\ell \in \{0,\dots,R-r-1\}} t^m_{\pi^{-1}(j+iR-\ell), r+\ell} \right) + t^s_{\pi^{-1}(j+iR)} \right\}. \tag{3}$$

Lemma 3. *For arbitrary ABP instances, the minimum boarding time is bounded from below by*

$$\max_{p \in \mathcal{P}} \left\{ \sum_{r=1}^{r(p)-1} t_{p,r}^m + t_p^s \right\}, \tag{4}$$

$$\max_{r \in \mathcal{R}} \left\{ \sum_{\substack{p \in \mathcal{P}: \\ r(p) \geq r+1}} t_{p,r}^m + \sum_{p \in \mathcal{P}(r)} t_p^s + \min_{\substack{p \in \mathcal{P}: \\ r(p) \geq r}} \sum_{r'=1}^{r-1} t_{p,r'}^m \right\}, \tag{5}$$

$$\frac{1}{\min\{R, P\}} \sum_{p \in \mathcal{P}} \left(t_p^s + \sum_{r=1}^{r(p)-1} t_{p,r}^m \right). \tag{6}$$

Note that the previously mentioned similarity of ABP to permutation flow shop scheduling is reflected here, too, as the bounds (4) and (5) are similar to makespan lower bounds derived from maximum job length and machine load, respectively; see, e.g., [Nagarajan and Sviridenko, 2009].

Remark 4. It should be noted that in Theorem 3, we explicitly assume full plane occupancy, i.e., $P = |\mathcal{S}|$. In case some seats are unoccupied (or filled earlier, e.g., during priority boarding), it is currently still unknown whether the quality guarantee for outside-in BTF boarding remains intact. On the other hand, Theorems 1 or 2 do not imply that ABP remains NP-hard in the setting of Theorem 3 with $P \lneq |\mathcal{S}|$ either: While the intractability results can be extended straightforwardly to include seat vacancies, cf. Corollary 1, the data assumptions $t^m_{p,r} = c^m$, $t^s_p = c^s$ of Theorem 3 are more restrictive and might therefore, in principle, yield easier special cases.

It is worth mentioning that Lemmas 2 and 3 provide general instance-dependent approximation ratios for boarding sequences obeying (2), at least for fully occupied airplanes:

Proposition 2. *For ABP instances with $P = |\mathcal{S}|$ and $|\mathcal{S}^1_r| + |\mathcal{S}^2_r| = k \in \mathbb{N}$ for all $r \in \mathcal{R}$, the boarding time of any boarding sequence π with property (2) is at most a factor β worse than the optimal boarding time, where*

$$\beta := \frac{(3)}{\max\{(4), (5), (6)\}}. \tag{7}$$

The main drawback of the general approximation ratio (7) is its dependence on instance-specific time values $\{t^m_{p,r}\}$ and $\{t^s_p\}$. This is unfortunate for two reasons: First, the ratio can conceivably become quite large when an instance contains some settle-in time that is much larger than the moving times—a situation that seems very natural, as walking past some row should take no more than a few seconds but stowing away carry-on luggage and taking a seat might take up to several minutes. Moreover, passenger time data changes for every instance (and is not even known exactly) while aircraft cabin parameters (k and R) are shared by whole fleets of airplanes. This makes it desirable to have approximation bounds that only depend on these instance-size parameters, thus enabling one to judge the quality of a boarding strategy regardless of the actual instance-specific passenger data. In the remainder of this section, we will demonstrate that such approximation results can indeed sometimes be obtained.

For the first such result, we consider generalizing w.r.t. the moving or the settle-in times (i.e., relaxing from identical constant moving times for all rows and passengers to individual moving times, or from constant to arbitrary settle-in times, respectively). Then, the outside-in BTF strategy is no longer necessarily optimal, but nevertheless, it still provides a k-approximation (see Appendix H for the proof):

Theorem 4. *For ABP instances with $P = |\mathcal{S}|$, $|\mathcal{S}^1_r| + |\mathcal{S}^2_r| = k \in \mathbb{N}$ for all $r \in \mathcal{R}$, and either $t^m_{p,r} = c^m_p \in \mathbb{Q}_{\geq 0}$ for all $(p, r) \in \mathcal{P} \times \mathcal{R}$ and $t^s_p = c^s \in \mathbb{Q}_{\geq 0}$ for all $p \in \mathcal{P}$, or $t^m_{p,r} = c^m \in \mathbb{Q}_{\geq 0}$ for all $(p, r) \in \mathcal{P} \times \mathcal{R}$ and $t^s_p \in \mathbb{Q}_{\geq 0}$ (arbitrary) for all $p \in \mathcal{P}$, the outside-in BTF boarding strategy is a k-approximation algorithm.*

Remark 5. For the first case considered in Theorem 4, i.e., for $t^m_{p,r} = c^m_p \in \mathbb{Q}_{\geq 0}$ for all $(p, r) \in \mathcal{P} \times \mathcal{R}$ and $t^s_p = c^s \in \mathbb{Q}_{\geq 0}$ for all $p \in \mathcal{P}$, one can also obtain the approximation ratio $1 + (\max_p c^m_p - \min_p c^m_p)/\min_p c^m_p$ by simplifying (7) (further upper bounding (3) and lower bounding (6)); we omit the proof. While this bound depends on instance time data rather than the size parameters, it may nevertheless be expected to be fairly small, possibly even better than k when k is relatively large. Moreover, if in fact all moving times are identical for all passengers, this provides an alternative proof of optimality (cf. Theorem 3).

As with Theorem 3, it is unclear whether the above approximation guarantees continue to hold if the plane is not fully occupied, cf. Remark 4. Moreover, notwithstanding the general approximation bound (7), the next result (proved in Appendix I) exhibits that outside-in BTF boarding could lead to rather bad results in even more general settings, and regardless of whether all seats are taken.

Proposition 3. *For every $k \in \mathbb{N}_{\geq 2}$ and $\varepsilon \in (0, k(k-1)/(3k-2)]$, there exists an ABP instance with $|\mathcal{S}^1_r| + |\mathcal{S}^2_r| = k$ for all $r \in \mathcal{R}$, and $t^m_{p,r} = c^m_p \in \mathbb{Q}_{\geq 0}$ for all $(p, r) \in \mathcal{P} \times \mathcal{R}$, such that the outside-in BTF boarding time is at least a factor $(2k - \varepsilon)$ worse than the optimal boarding time.*

In fact, Proposition 3 can be generalized to all strategies that do not explicitly take settle-in times into account. In particular, all the classical boarding strategies mentioned in Section 1 (back-to-front, reverse pyramid, etc.) rely solely on the seat assignments and are thus covered by the following result (proved in Appendix J).

Theorem 5. *For every $R, k \geq 2$, there exists an ABP instance for which any deterministic (by-seat) boarding strategy that disregards individual settle-in times leads to a boarding time at least twice as long as the optimal boarding time, and any (by-group) strategy allowing arbitrary passenger orders (within boarding groups) can also yield boarding times twice as long as the optimal one.*

Let us now introduce a modification of the outside-in BTF boarding strategy:

Definition 3 (max-settle-row strategy). *Let the passengers board the plane in subsequences of size R, each with exactly one passenger with seat in row r for all $r \in \mathcal{R}$. The i-th subsequence contains a passenger p of each row with the i-th longest settle-in time of their row, breaking ties in outside-in fashion (seats furthest from the aisle first), and is ordered by decreasing row number.*

Besides generalizing outside-in BTF boarding, the max-settle-row strategy is related to the method proposed in [Milne and Kelly, 2014], which performs seat assignment based on number of carry-on items to achieve fast boarding with the Steffen method [Steffen, 2008] applied afterwards. However, although the number of hand-luggage pieces intuitively correlates with the settle-in times, there is no direct correspondence (and indeed, the recent statistical analysis in Hutter et al. [2019] indicates that the influence of carry-on baggage item numbers on boarding times may have been overestimated in various studies, including Milne and Kelly [2014]). Moreover, here, the seat assignments are fixed a priori and we construct a boarding sequence, not vice versa. The max-settle-row strategy furthermore also bears some resemblance to the "slow first" strategy analyzed in [Erland et al., 2019], which partitions passengers into a slow and a fast group w.r.t. settling-in times and lets the slow group board before the fast one (both in random order). Thus, max-settle-row may be viewed as a by-seat hybrid of slow-first and (one-per-row) back-to-front boarding.

We first note that when assuming identical settle-in times for all passengers, max-settle-row (and, in fact, any by-seat back-to-front strategy that picks one passenger per row) reduces to outside-in BTF and thus immediately inherits the solution quality guarantees from Theorems 3 and 4:

Corollary 2. *If $t_p^s = c^s \in \mathbb{Q}_{\geq 0}$ for all $p \in \mathcal{P}$, the max-settle-row strategy reduces to outside-in BTF boarding. Consequently, in this case, for ABP instances with $P = |\mathcal{S}|$ and $|\mathcal{S}_r^1| + |\mathcal{S}_r^2| = k \in \mathbb{N}$ for all $r \in \mathcal{R}$, max-settle-row gives an optimal solution if $t_{p,r}^m = c^m \in \mathbb{Q}_{\geq 0}$ for all $(p, r) \in \mathcal{P} \times \mathcal{R}$, and is a k-approximation algorithm if $t_{p,r}^m = c_p^m \in \mathbb{Q}_{\geq 0}$ for all $(p, r) \in \mathcal{P} \times \mathcal{R}$.*

As it turns out, the max-settle-row strategy gives an improved approximation result for the case in which we have constant moving times but arbitrary settle-in times (and a full aircraft).

Theorem 6. *For ABP instances with $P = |\mathcal{S}|$, $|\mathcal{S}_r^1| + |\mathcal{S}_r^2| = k \in \mathbb{N}$ for all $r \in \mathcal{R}$, and $t_{p,r}^m = c^m \in \mathbb{Q}_{\geq 0}$ for all $(p, r) \in \mathcal{P} \times \mathcal{R}$, the max-settle-row strategy is a $(1 + H_k(k-1)/k)$-approximation algorithm, where $H_k := \sum_{i=1}^{k}(1/i)$.*

The proof can be found in Appendix K. Note that $1 + H_k(k-1)/k < k$ for all $k \geq 2$, so Theorem 6 (max-settle-row) indeed gives a stronger guarantee than Theorem 4 (outside-in BTF).

Interestingly, in the setting of the above theorems, the asymptotic behavior of both max-settle-row and outside-in BTF boarding reflects the growing influence of moving times on the total boarding time in case of longer airplane cabins, eventually outweighing settle-in time contributions entirely:

Proposition 4. *For ABP instances with $P = |\mathcal{S}|$, $|\mathcal{S}_r^1| + |\mathcal{S}_r^2| = k \in \mathbb{N}$ for all $r \in \mathcal{R}$, $t_{p,r}^m = c^m \in \mathbb{Q}_{\geq 0}$ for all $(p, r) \in \mathcal{P} \times \mathcal{R}$, and (arbitrary) $t_p^s \in (0, T]$, $T < \infty$, both the outside-in BTF strategy and the max-settle-row strategy are asymptotically optimal. More precisely, as $R \to \infty$, the respective boarding times of their solutions tend to the optimal boarding time.*

The discrepancies between the positive and negative results presented in this section exhibit possibly significant potential for improvements to be gained by solving ABP to optimality. Although Theorems 1 and 2 do not directly imply NP-hardness for all the restricted ABP versions (cf. Table 1 and Remark 4), there is also no immediate strategy to exactly solve these problems in polynomial time (nor to obtain results with *guaranteed* approximation bounds in all cases). Therefore, in the following section, we turn to mixed-integer programming to optimally solve ABP in general.

Remark 6. As a final remark on the theoretical groundwork presented in this section, we remind that all of the above theoretical results implicitly assume the congestion level to be fixed at k, by virtue of each passenger occupying aisle space equal to one row. As with changes to input data and plane occupancy (cf. Remark 4), the polynomial-time solvability and approximability results may change under different congestion parameter assumptions. Indeed, it seems conceivable that, for instance, with congestion $2k$, analogous results to Theorem 3 and Proposition 4 should hold for the Steffen method, or that for "cardboard-thin" passengers that yield congestion well below 1, a suitable back-to-front strategy would become optimal (since it then maximizes settling-in parallelism during boarding). A rigorous theoretical analysis of approximability under different congestion assumptions would nicely complement the probabilistic/asymptotic results from [Bachmat et al., 2009] and follow-up papers, but is beyond the scope of this work and thus left for future research.

3 Exact MIP Formulation

In this section, we describe our compact linear MIP model for the ABP. We model the objective by using the index ordering corresponding to the boarding sequence in an *inverse* fashion; therefore, we abbreviate the MIP given below by (8)–(16) as *IPA (inverted position assignment)*. It is worth mentioning that in a preliminary version of this paper, Willamowski and Tillmann [2019], we had also experimented with alternative formulations using direct (non-inverted) assignments, predecessor relations, and a time-expanded multi-commodity flow network, but IPA turned out to be computationally superior to all of those. Thus, we only present IPA here. We note also that IPA is conceptually related to, but different from, MIP models put forth for permutation flow shop scheduling problems, see, e.g., Wilson [1989], Jessin et al. [2020].

The following variables are used to model a boarding sequence π (which corresponds to a permutation, or linear ordering, cf. [Coudert, 2016]) and the associated boarding completion time:

- Let $x_{p,i} \in \{0,1\}$ be the binary *position* variable representing whether a passenger $p \in \mathcal{P}$ is placed at position $i \in [P]$ in the linear order (boarding sequence) π, i.e., $x_{p,i} = 1$ if $\pi(p) = 1$ and $x_{p,i} = 0$ otherwise.
- Let $t^A_{i,r} \in \mathbb{Q}_{\geq 0}$ be the *arrival time* variable representing the time at which passenger $\pi^{-1}(i)$, $i \in [P]$, arrives at row $r \in \mathcal{R}$, and let $t^F_{i,r}$ be the *finishing time* variable representing the time at which passenger $\pi^{-1}(i)$, $i \in [P]$, finishes their action at row $r \in \mathcal{R}$.
- Finally, we introduce the *completion time* variable $C_{\max} \in \mathbb{Q}_{\geq 0}$ representing the boarding completion time of the computed linear order/boarding sequence.

Regarding the input ABP time data, it is helpful to introduce further notation as well: For every passenger $p \in \mathcal{P}$ and each row $r \in \mathcal{R}$, we are given times $\tau_{p,r} \in \mathbb{Q}_{\geq 0}$ that passenger p *must* spend at row r either moving past it ($\tau_{p,r} := t^m_{p,r}$ for $r \leq r(p) - 1$) or settling in in it ($\tau_{p,r} := t^s_p$ for $r = r(p)$). For rows r that a passenger p needs not visit at all, i.e., those behind their assigned seat ($r \geq r(p) + 1$), we may and do assume that $\tau_{p,r} = 0$. Similarly, we neglect what happens before the first passenger has arrived at the first row; by minimizing the final completion time, it is thus automatically ensured that the first passenger will start with the first row action at time zero. Nevertheless, note that times spent walking through jet-bridge could be incorporated easily as lower bounds on $t^A_{i,1}$.

Naturally, every passenger has to be assigned precisely one slot in the sequence, yielding constraints (9) and (10) in IPA. We work directly with indices of the output sequence $\pi := (1, 2, \dots, P)$ and formulate the remaining constraints in terms of the i-th passenger in this sequence, $i \in [P]$. Incorporating the assignment decisions, we can thus model the objective (8) (minimizing total boarding time) via (11), the requirement that rows are traversed consecutively by each passenger via (12) and (13), that passengers cannot be overtaken until they are seated via (14), and the requirements on times spent in rows via (15).

The full IPA model now reads as follows; some further explanations are given immediately afterwards, in particular regarding (13) and the big-M constants $M_{p,r,i}$ occurring therein.

$$\min \quad C_{\max} \tag{8}$$

$$
\begin{aligned}
\text{s.t.}\quad & \sum_{p\in\mathcal{P}} x_{p,i} = 1 && \forall i \in [P] && (9)\\
& \sum_{i\in[P]} x_{p,i} = 1 && \forall p \in \mathcal{P} && (10)\\
& C_{\max} \geq t^F_{i,R} && \forall i \in [P] && (11)\\
& t^A_{i,r+1} - t^F_{i,r} \geq 0 && \forall i \in [P], r \in [R-1] && (12)\\
& t^A_{i,r+1} - t^F_{i,r} \leq \sum_{p\in\mathcal{P}:r(p)\leq r} M_{p,r,i}\cdot x_{p,i} && \forall i \in [P], r \in [R-1] && (13)\\
& t^A_{i+1,r} \geq t^F_{i,r} && \forall i \in [P-1], r \in \mathcal{R} && (14)\\
& t^F_{i,r} - t^A_{i,r} \geq \sum_{p\in\mathcal{P}} \tau_{p,r}\cdot x_{p,i} && \forall i \in [P], r \in \mathcal{R} && (15)\\
& x \in \{0,1\}^{\mathcal{P}\times[P]};\ t^A, t^F \in \mathbb{Q}^{\mathcal{P}\times\mathcal{R}}_{\geq 0};\ C_{\max} \in \mathbb{Q}_{\geq 0} &&&& (16)
\end{aligned}
$$

Note that (12) and (13) also ensure that no times remain uncaptured, by equating the finishing times for any one row up to the one in which the respective passenger's seat is located with the respective arrival times for the subsequent one. Sufficiently large big-M constants $M_{p,r,i}$ in (13) are used to disable equating $t^A_{i,r+1}$ with $t^F_{i,r}$ for rows r beyond the one in which the i-th passenger is eventually seated. This becomes necessary due to working with the *output* index sequence rather than passenger (input) indices: to circumvent that we do not know a priori the data and seat belonging to the i-th passenger *in the boarding sequence*, we let *all* passengers fictitiously pass through to the last row. Here, $\tau_{p,r} = 0$ for $r > r(p)$ ensures no extra time is added anywhere. The big-M term in (13) ensures that passengers who are only "virtually" present in some row after the one in which they are actually settled in cannot block "real" passengers around them. It is not hard to see that the big-M constants can be chosen, for instance, as

$$
M_{p,r,i} = \max_{\substack{\mathcal{P}'\subseteq\{p'\in\mathcal{P}\setminus\{p\}\,:\,r(p')>r\},\\ |\mathcal{P}'|\leq i-1}} \sum_{p'\in\mathcal{P}'} \left(\sum_{r'=r+1}^{r(p')-1} t^m_{p',r}\right) + t^s_{p'}.
$$

This is, for each tuple $(p,r,i) \in \mathcal{P}\times\mathcal{R}\times[P]$, an upper bound on the time passenger p assigned to position i has to (virtually) wait in row r, since only the $i-1$ predecessors of p can induce this (virtual) waiting by blocking some rows $r' \geq r+1$.

Furthermore, note that because of having all passengers fictitiously arrive at the last row eventually, constraint (11) is formulated w.r.t. the last row R, since we cannot directly access the index $r(\pi^{-1}(i))$ of the row in which the i-th passenger in the output sequence is seated.

4 Computational Experiments

We now turn to an experimental evaluation of the MIP formulation IPA and select boarding strategies. We compare the outside-in back-to-front heuristic, Steffen's method [Steffen, 2008], our novel max-settle-row strategy, as well as random boarding and standard outside-in with our exact model.[5] The first two methods (and other outside-in variants) had emerged in previous simulation studies as often being superior to other classical strategies that also do not explicitly take passenger moving and settle-in times into account; we complement these findings by contrasting the behavior of these schemes with the "time-aware" methods proposed in this paper. Random boarding is included as a reference for what would happen if no care at all were to be taken regarding the sequence in which passengers are asked to board. Similarly, standard outside-in boarding can be seen as a randomized variant of outside-in BTF: it mirrors how outside-in boarding is

[5] We also experimented with a by-seat variant of the slow-first strategy, cf. [Erland et al., 2019], but since that turned out to be significantly inferior to the other by-seat approaches and only slightly better than the randomized strategies, we omit reporting the details for by-seat slow-first.

currently implemented in practice by several airlines (e.g., Lufthansa) by replacing the deterministic back-to-front row order by a random one. To distinguish it from outside-in BTF, we will henceforth refer to the standard outside-in scheme as *randomized outside-in*. For any instance, we take the best one of 1000 random realizations as the respective outputs of the random boarding strategies (thus, the stated results for randomized schemes are, in fact, more optimistic than their average behavior, which is reported as well).

Moreover, we will also evaluate a simple time-aware improvement heuristic that can be applied to any boarding method as a post-processing routine:

Definition 4 (2-opt local search). *Given a boarding sequence π of $\mathcal{P}$ with completion time $C_{\max}(\pi)$ for a given ABP instance, repeatedly perform swaps of the positions of a pair of entries in π that lead to a (strict) reduction in the total boarding time, until no further improvements can be achieved by such swaps (one iteration traverses all pairs in lexicographical order of indices, performing all improving swaps along the way).*

We implemented the local search in Cython 0.29.21 and the boarding heuristics in Python 3.7, which was also used as a scripting language and to generate test instances. The MIP model (IPA) was solved with the commercial solver Gurobi 9.1.0 [Gurobi Optimization, LLC, 2019]; for each instance, we provided Gurobi with the best solution obtained by any of the heuristics (plus local search) as a primal incumbent (starting solution). All experiments were carried out on a cluster of 56 Linux machines with Intel Xeon L5630 Quad-Core CPUs (2.13 GHz) and 16 GB memory.

In the following, we first describe our test instances in Section 4.1, before presenting and discussing several computational experiments in Sections 4.2 (main comparison of the considered boarding methods for different levels of "resolution" w.r.t. the passenger time data) and 4.3 (gauging robustness of boarding sequences obtained with inexact data w.r.t. random disruptions).

4.1 Construction of Test Instances

For our test instances, we considered four different cabin layouts: $(R, k) \in \{(10, 2), (20, 2), (20, 4), (30, 6)\}$. We randomly assign each passenger 1, 2, or 3 pieces of luggage, with respective percentages 60%, 30%, and 10%, as proposed by, e.g., van Landeghem and Beuselinck [2002], Schultz et al. [2008], or Audenaert et al. [2009]. The settle-in times were generated as $\ell_p \cdot \lfloor z, 1\rceil + 2.4$, with ℓ_p the number of luggage items of passenger $p \in \mathcal{P}$, $z \sim \mathcal{W}(1.7, 16)$ drawn independently at random from the Weibull distribution with shape parameter 1.7 and scale parameter 16, where $\lfloor \cdot, 1\rceil$ denotes rounding to one significant digit; this scheme was proposed and verified by Schultz [2018b] (with the exception of rounding, which we do for simplicity). The moving times were generated as $\lfloor z, 1\rceil$, with $z \sim \mathcal{T}(1.8, 2.4, 3.0)$ drawn independently at random from the triangular distribution with lower limit 1.8, upper limit 3.0, and mode 2.4, as proposed by, e.g., van Landeghem and Beuselinck [2002], Schultz et al. [2008], Audenaert et al. [2009], or Milne and Kelly [2014].

We created three instance sets *mp-sp*, *m-sp*, and *mp-s*, each with 10 instances for each cabin layout (for a total of 120 instances). These instance sets differ in which times are constant and which are fluctuant: "m" stands for moving and "s" for settle-in times, and a trailing "p" represents that the corresponding times are passenger-dependent (otherwise, the respective times are the same for all passengers). For simplicity, we do not consider passenger- and *row*-dependent moving times $t^m_{p,r} \in \mathbb{Q}_{\geq 0}$ (i.e., only $t^m_{p,r} = c^m_p$ and $t^m_{p,r} = c^m$ are used). Thus, to be precise, when identical constant times are chosen for all passengers, these equal $\lfloor z, 1\rceil$ seconds with a single $z \sim \mathcal{T}(1.8, 2.4, 3.0)$ (per instance) for the moving times and $(1 \cdot 0.6 + 2 \cdot 0.3 + 3 \cdot 0.1) \cdot \lfloor z, 1\rceil + 2.4$ seconds with a single $z \sim \mathcal{W}(1.7, 16)$ (per instance) for the settling-in times.

With these instances, we evaluate the algorithm performances in the different settings, see Section 4.2. Furthermore, we conduct a second set of experiments to assess the robustness of different boarding strategies w.r.t. data perturbations and disruptions of the determined boarding sequences, see Section 4.3. With these simulations, we hope to gauge the influence of using estimated time data as opposed to the "real" time data (that will generally not be available a priori in practice) as well as that of passengers not keeping to the predetermined boarding order or arriving late at the gate, events that are likely to happen in practice. To that end, first, we consider late passengers: for every instance, we independently sample uniformly at random 2, 3, 4, or 5 passengers, respectively, for each cabin layout (the larger the cabin, the more latecomers), and place them at the end of the boarding sequence (in random order). Second, for the purpose of assessing the impact of passengers disregarding their assigned positions, we let a total of 10% of passengers take different places within the boarding sequence. For this, we generate a collection of old/new position tuples (uniformly at random), delete the passengers in descending order of the old positions (w.r.t. the original boarding sequence) and reinsert them in increasing order of their respective new positions. Third, to examine the relevance of precise knowledge of settle-in and moving times, we evaluate the computed boarding sequences also w.r.t. perturbed times, obtained as random numbers drawn from the normal distribution centered at the original respective time values (per luggage item): For each settle-in time t_p^s, a perturbed time is created as $\ell_p \cdot \max\{\lfloor z, 1\rceil, 0\} + 2.4$ with $z \sim \mathcal{N}((t_p^s - 2.4)/\ell_p, 10)$ drawn independently at random, where $\mathcal{N}(\mu, \sigma)$ denotes the normal distribution with mean μ and standard deviation σ. Similarly, for each moving time c_p^m (recall we do not have row-dependency here), we sample $z \sim \mathcal{N}(c_p^m, 0.4)$, clip so that $z \in [1.8, 3.0]$ and round to one significant digit. In the experiments, we evaluate those three disturbances—late passengers, passenger displacements, and perturbed time data—separately as well as all together (passenger displacements performed before repositioning late passengers), the latter representing the most real-world oriented scenario.

Furthermore, in light of Proposition 1, we also study the influence of using overestimated time data to obtain the boarding sequence. To that end, we rounded each moving time c_p^m up to 3, and replace the settle-in times t_p^s by $\ell_p \cdot 30 + 3$, dependent on the respective number ℓ_p of luggage items, which overestimates 94.5% of the time data (per ℓ_p value). The boarding sequences obtained using these "overestimated" time values are then evaluated using the original ("real") values, including the above-described disturbances.

Our complete testset, as well as solution files for all instances and boarding methods, can be downloaded from our project webpage, https://www.or.rwth-aachen.de/airplane-boarding.

4.2 Experiments for Different Data Models

The results for the experiments with different boarding strategies under the three data scenarios are summarized in Table 2. For each scenario, the table provides average values over all instances (10 per cabin layout, i.e., 40 instances per testset), where all time values are stated in seconds rounded to one significant digit (third, fourth, and last column). Averages are calculated as arithmetic means except for the algorithm runtimes, where we used the shifted geometric mean (with a 10 s shift) to reduce the influence of easy instances. The column labeled "objective" gives the average completion times of the boarding sequences obtained with the respective algorithms listed in the second column ("method"); for randomized strategies, the additionally reported objective values in parentheses are the respective average values over the 1000 realizations (the 2-opt routine has always been applied only to the best one of those, per instance). The column "w / 2-opt" shows the average completion times after post-processing the heuristically computed boarding sequences with our local search routine from Definition 4 (we did not apply this post-processing to any MIP solutions). The column labeled "% impr." gives the resulting average improvement percentages (100%·(old value − new value)/(old value), rounded to two significant digits) due to the 2-opt post-processing routine for the heuristic boarding strategies, and the average improvement over the best (post-processed) heuristic solution achieved by Gurobi applied to the IPA model, respectively. The columns "% gap" and "# opt" provide the average MIP optimality gap (100%·(best upper bound − best lower bound)/(best upper bound), rounded to two decimals) and the absolute number (out of 40) of instances that the respective methods (plus post-processing) solved to optimality—as certified by Gurobi/IPA—within a time limit of one hour (3600 s). Finally, the "runtime" column gives the mean runtimes of the algorithms, including local search; the heuristic boarding strategies themselves take less than a second for all instances.

From Table 2, we can draw a variety of conclusions. Starting with the heuristics, we first observe that the max-settle-row strategy yields lower objectives than all other heuristics in all three time-data scenarios, before and after applying the 2-opt local search post-processing routine. The differences between (deterministic) outside-in BTF and max-settle-row boarding are notable across all scenarios and instance sizes. After 2-opt, these differences become much smaller but persist; recall also that for constant settle-in times, max-settle-row and outside-in BTF coincide, so we report their results only once for testset mp-s. The randomized

testset	method	objective	w / 2-opt	% impr.	% gap	# opt	runtime
mp-sp	random	684.7 (837.6)	392.5	60.69	—	0	0.7
	rand. outside-in	687.4 (819.5)	388.9	63.76	—	0	0.7
	Steffen	535.5	386.7	36.16	—	0	0.5
	outside-in BTF	405.7	342.4	16.04	—	1	0.3
	max-settle-row	348.9	335.1	3.05	—	1	0.2
	IPA	332.6	—	0.72	6.55	8	1938.5
m-sp	random	671.4 (820.4)	388.8	60.37	—	1	0.5
	rand. outside-in	669.6 (800.4)	391.4	60.68	—	1	0.6
	Steffen	512.2	372.4	34.99	—	1	0.4
	outside-in BTF	376.9	322.2	15.11	—	8	0.2
	max-settle-row	321.6	317.6	0.85	—	11	0.1
	IPA	315.4	—	0.75	3.08	15	566.4
mp-s	random	648.3 (766.7)	396.4	49.17	—	0	0.5
	rand. outside-in	633.5 (739.1)	399.3	45.38	—	0	0.5
	Steffen	389.2	366.3	6.85	—	0	0.2
	outside-in BTF / max-settle-row	297.1	287.0	2.86	—	9	0.1
	IPA	286.4	—	0.16	4.45	11	1179.3

Table 2. Experimental results on different testsets. Numbers are average values over 40 instances per testset; 10 each for 4 airplane sizes.

outside-in strategy turns out to be much worse than its deterministic (back-to-front) counterpart, with average boarding times roughly twice as large, and still about 13–39% larger after 2-opt post-processing the best of 1000 random realizations. This contrast alone emphasizes the drastic improvements that may be possible by actually enforcing by-seat boarding in practice, even if sticking to simple, structured strategies. Indeed, classical (randomized) outside-in boarding is apparently only slightly better than allowing a completely random boarding sequence (and sometimes even marginally worse after local search refinements). Overall, random boarding performs worst of all strategies, and Steffen's method is also consistently and significantly worse than outside-in BTF and max-settle-row; for both randomized methods and Steffen boarding, large improvements are achievable by 2-opt post-processing, but the end results are still clearly inferior.

Regarding approximation quality, we find that the outside-in BTF and max-settle-row strategies empirically behave more benevolently than guaranteed by our theoretical results from Section 2.2. This conclusion is enabled by our MIP lower bounds on the optimal boarding times, with which we can certify optimality or compute empirical approximation ratios, as the maximum of (heur. $C_{\max}$)/(IPA lower bound) over all instances under consideration; the required values can be obtained from the result files for all boarding methods and all instances that we provide on our project webpage (cf. Section 4.1). For instance, on the m-sp testset (which appears to yield the easiest MIPs, cf. Table 2), max-settle-row is very close to optimal even without local search post-processing, with average deviation from optimum of 5.2% and empirical (worst-case) approximation ratio estimate of about 1.22—much better than the factor $(1+H_k(k-1)/k)$ guaranteed by Theorem 6 already for $k=2$. Note also that for the mp-s instances, where $t^m_{p,r} = c^m_p \in [1.8, 3.0]\ \forall (p,r)$, Remark 5 gives an approximation ratio guarantee of 5/3 for max-settle-row/outside-in BTF, which is better than the ratio k of Theorem 4; on average over all instance sizes, again using the lower bounds derived from the IPA optimality gap, the average deviation from optimality is bounded by 8.1% and the empirical worst-case approximation ratio is about 1.15. Naturally, these ratios can be improved further by applying the 2-opt procedure to the respective solutions, e.g., from about 1.25 to just below 1.21 for max-settle-row on the mp-sp testset (with average bounds on the deviation from optimum decreasing from 11.8% to 8.4%). Note also that for the easier testsets m-sp and mp-s, max-settle-row with 2-opt improvements sometimes already yields an optimal solution, as does outside-in BTF to a slightly lesser extent; for the most general data (testset mp-sp), however, none of the heuristics achieve optimality except for max-settle-row and outside-in BTF on a single instance.

Let us now turn to the MIP formulation, IPA. From Table 2, it becomes clear that IPA/Gurobi can often only marginally improve over the best heuristic (usually max-settle-row + 2-opt) solutions, and apparently (for the larger instance) has trouble increasing the dual bounds in order to certify optimality of the best known solution. While this behavior is not atypical for hard MIP instances, it may nevertheless seem somewhat disappointing. However, the effort is by no means wasted, as the MIP lower bounds provide a hitherto unavailable way to gauge the effectiveness of boarding heuristics. Indeed, the optimality gaps provide a *computational proof* that the solutions produced by max-settle-row (and, to a lesser degree, outside-in BTF) are typically not very far from optimal, especially after 2-opt post-processing, see the previous paragraph. Moreover, Gurobi/IPA can solve a relevant portion of the instances to provable optimality, albeit mostly for the smaller instances. Fixing either moving or settle-in times to the same constant value for all passengers renders all MIPs notably easier: Comparing the results in Table 2 note that for m-sp and mp-s testsets, more instances are solved, and the mean runtimes as well as the optimality gaps are significantly smaller than for testset mp-sp. For completeness, we remark that on small instances, the aforementioned alternative MIP formulation based on predecessor-relations may give slightly better results than IPA on average, but on instances where it fails—in particular, for the (arguably more interesting) large instances—the performance is significantly worse, cf. Willamowski and Tillmann [2019].

Constant moving times (m-sp) appear to generally give the easiest-to-handle instances for the MIP model (or, more precisely, for Gurobi applied to solve it). This observation may have positive practical consequences: Moving times are small compared to settle-in times, unavoidable as they (unlike settling-in operations) cannot be "parallelized", and should not differ drastically for different passengers. Thus, their influence on the overall boarding time is less significant, so if expectedly small fluctuations are neglected (e.g., taking a "pessimistic estimate" standpoint in view of Proposition 1), one may work with constant moving times and more quickly obtain better solutions or similarly accurate empirical approximation bounds for heuristics by solving IPA on m-sp data rather than on the most realistic mp-sp data.

Finally, it is worth mentioning that, on average, the IPA formulation succeeds in producing lower bounds that are better than what we can obtain from Lemma 3: Consistently, the lower bound (5) is significantly better than the other two ((4) and (6)), and yields, for instance, an average value of about 305.4 for the hardest testset (mp-sp); this implies an associated average optimality gap of 7.23%, which is improved upon by IPA (6.55% average gap).

Thus, to summarize, the results discussed in this section demonstrate that our max-settle-row boarding strategy is the best of the considered heuristics across all time-data scenarios, closely followed by its "parent" outside-in BTF. The proposed 2-opt post-processing routine can significantly improve all heuristic solutions, and, combined with max-settle-row, yields boarding sequences that are provably close to optimal, by virtue of our MIP formulation for the ABP and the resulting computational lower bounds on optimal boarding times. IPA generally provides highly useful lower bounds, and the fact that MIP can often only marginally improve the best heuristic solutions gives further (tentative) indication that they often are likely already very close to optimal.

Finally, it is noteworthy that the results in this section can also be seen to confirm the high-level predictions enabled by the analyses provided in [Bachmat and Elkin, 2008, Bachmat et al., 2009, 2013] and related works, or conversely, that the insights from those works can offer an explanation for (parts of) the above observations. As mentioned earlier, we work with a congestion parameter value k here, i.e., the number of seats per row. Since back-to-front boarding strategies can be shown to be (asymptotically) optimal in the sense of the average-case probabilistic analysis of Bachmat and Elkin [2008] (and follow-up works) for congestion levels up to 1, it intuitively makes sense to partition the boarding queue into k groups (so that each group separately boards with congestion 1), and adapt a back-to-front pattern for each such group in order to maximize the number of (settling-in) actions that can take place in parallel. Indeed, this is precisely what the outside-in BTF and max-settle-row strategies do. Moreover, the boarding sequence of Steffen's method essentially consists of $2k$ back-to-front-like groups that board sequentially, and so would not be optimal for the congestion setup considered in this paper. Consequently, Steffen's method can be expected to gain the upper hand over outside-in BTF and max-settle-row if a congestion of $2k$ was enforced in the simulations, i.e., if every passenger blocks two rows instead of one. Similarly, recalling that interpreting Steffen's method with congestion k essentially results in outside-in BTF, modifying the Steffen method by choosing the seats for each row by decreasing settling-in times rather than in outside-in fashion should yield improvements over its standard variant when assuming congestion $2k$. Different congestion parameter values and their effects in simulation experiments with, in particular, by-seat strategies is surely an interesting aspect to consider in the future, to further complement the theoretical analyses instigated by Bachmat and Elkin [2008] and Bachmat et al. [2009].

4.3 Robustness Experiments

For the experiments discussed in the following, we focus on the mp-sp data model (recall that even though such data is generally not available in practice, it allows for the most accurate simulation-based estimation of heuristic approximation ratios, via the MIP lower bounds). Since our 2-opt local search routine improved the output of all heuristics, here, we consider only the post-processed heuristic solutions, and of course the exact method, i.e., (Gurobi applied to) IPA.

To assess robustness aspects of the boarding schemes w.r.t. different possible disturbances as outlined in Section 4.1, the experiments summarized in Table 3 show the impact on boarding times (again in seconds, averaged over 40 instances, i.e., 10 instances each for all four airplane cabin sizes) of placing randomly chosen passengers at the end of the boarding queue (column "late"), of displacing random passengers to random new positions at odds with the predetermined boarding sequence ("displ."), and assuming the actual passenger time data are perturbed w.r.t. those with which boarding sequences were computed ("pert. obj."), respectively, as well as for the combination of all these disruptions (column "comb."). Additionally, the "ref. obj." column lists the average boarding times achieved by letting the respective algorithms run on the perturbed time data, without applying further disruptions afterwards; for each instance, the smallest value here is an upper bound on the respective average ideal possible boarding time (computed with knowledge of the reference perturbed time data and no disturbance of the boarding sequence). Together with the lower bounds from IPA, we can estimate average loss intervals w.r.t. the ideal situation as incurred by the data uncertainty and disruptions, stated in the final column "% comb. loss".

method	obj.	late	displ.	pert. obj.	comb.	ref. obj.	% comb. loss
random + 2-opt	392.5	465.6	496.6	539.7	658.9	403.8	[48.1, 51.4]
rand. outside-in + 2-opt	388.9	461.5	490.6	540.7	655.3	378.7	[47.2, 50.5]
Steffen + 2-opt	386.7	463.1	488.5	533.1	660.5	408.4	[48.5, 51.8]
outside-in BTF + 2-opt	342.4	418.9	444.6	450.1	583.3	368.0	[41.7, 45.4]
max-settle-row + 2-opt	335.1	413.3	439.7	424.2	567.4	355.5	[39.9, 43.7]
IPA	332.6	411.6	435.1	431.3	571.0	352.1	[40.7, 44.4]

Table 3. Results of robustness experiments on testset mp-sp. Considered are disruptions caused by late passengers, passenger displacements, and for perturbed moving and settle-in times as well as all three combined. Numbers are average values over 40 instances; 10 each for 4 airplane sizes.

We begin with a closer look at the computational results of Table 3. It can be observed that, on average across all instance sizes, all three disruption types have a notable negative impact on the boarding completion times. Displacements and time data perturbations seem to lead to stronger deterioration of the solution quality than having a few late passengers (although the numbers are not directly comparable due to the different numbers of disruptions). Combining time-data perturbation and disruptions, the boarding times for all strategies increase drastically—by up to 50% or more—with max-settle-row and IPA yielding the lowest overall times, between about 40% and 44% larger than the ideal baseline/reference solutions on average. For random boarding, randomized outside-in, and Steffen's method, the disturbances lead to final (combined) boarding times are significantly inferior to those of the other three approaches. On the one hand, the overall strong increase in the boarding times when exposing a computed boarding sequence to "realistic" changes

method	pess.	true	ref. true	true, disr.	runtime
random + 2-opt	769.5	651.8	403.8	741.7	0.4
rand. outside-in + 2-opt	770.4	654.8	407.3	740.3	0.4
Steffen + 2-opt	622.7	559.8	408.4	659.5	0.2
outside-in BTF + 2-opt	459.1	437.5	368.0	591.5	0.2
max-settle-row + 2-opt	454.0	422.9	355.5	578.3	0.1
IPA	451.5	423.6	352.1	577.1	225.3

Table 4. Results for robustness experiments on testset mp-sp for pessimistic data assumptions and all disturbances. Numbers are average values over 40 instances; 10 each for 4 airplane sizes.

of environment demonstrates that even then, the newly proposed strategies (max-settle-row, the 2-opt local search routine, and IPA) provide significant improvements over the others, with outside-in BTF being a close follower. On the other hand, the large differences in boarding times suggest that there is still room for improvement by explicitly incorporating robustness aspects into boarding strategy design.

Furthermore, the sensitivity with respect to perturbed time data also raises the question of how much worse things might get if the time data used to compute the boarding sequence is actually quite far off from the true data. In particular, in view of Proposition 1, do we fare much worse if we adopt very rough (over-)estimates of the passenger moving and settle-in times to compute the boarding order? This leads to our next, and final, experiment in this section.

Table 4 presents boarding times that were computed using "pessimistic/overestimated" time data (moving and settle-in times) on the "true" reference time data and after late- and displacement-disturbances; the "true" data here is the perturbed time data from the experiments summarized in Table 3, see Section 4.1 for the details. In Table 4, the columns "pess.", "true", and "true, disr." give the boarding times of the sequences computed using the pessimistic (mostly overestimated) time-data, when evaluated on that same data, and on the true data (perturbed time values from previous mp-sp experiments) without or with late-passenger and position-change disruptions, respectively. The "ref. true" column states the boarding times achieved by the respective algorithms when directly using the true data, which serve as the reference ideal-situation best known solutions. The final column provides the (shifted geometric mean) runtimes of computing the boarding sequences, including local search runtime for the heuristics.

We can observe from Table 4 that the boarding times using pessimistic time estimates are arguably not terribly far off the reference values. Disregarding other disruptions, this indicates that it might not be a huge obstacle that one likely will not have precise knowledge of passenger moving and settle-in times in practice. Indeed, comparing the first two boarding time columns, we see a validation of Proposition 1—replacing the pessimistic data by the smaller true time values but retaining the boarding sequence itself, the boarding times always reduce—though even for the better strategies outside-in BTF, max-settle-row (each post-processed by the 2-opt routine) and IPA, the difference to the reference solution values (column "ref. true") turns out to be around roughly 20% on average. While this may seem a lot, note that, analogously to the experiments from Table 3, once we additionally consider disruptions of the boarding order, the picture turns much worse again, with final estimates of the "real-life" outcome of the present strategies up to over 60% larger than the ideal boarding times without disruptive passengers.

Hence, while a positive take-away message from these last experiments may be that boarding planning (using the "right" strategies) is apparently not strongly sensitive to the accuracy of the utilized time-data, the other large point to be made reiterates the finding of the first set of robustness experiments above: To reduce large fluctuations induced by disruptive passenger behavior, which may be almost unavoidable in real life, one must find a way to *explicitly* robustify by-seat boarding processes (or algorithmic boarding schemes). It will be especially interesting to see if such currently hypothetical robust by-seat strategies can ultimately outperform by-group strategies, which are naturally more robust (on average) to disruptive passenger behavior due to their inherent randomness. Given that congestion has a strong influence on average boarding scheme performance (cf., e.g., [Bachmat et al., 2009, Erland et al., 2019]), it may also be worth looking into devising strategies that are robust w.r.t. uncertainty in the congestion parameter.

5 ABP Extensions

For the sake of simplicity, we have so far not included other possible aspects of the ABP in the theoretical and numerical results in this paper. For the two examples mentioned at the end of Section 1.2—seat interferences and inseparable passenger groups—we will now briefly discuss how they can easily be incorporated into heuristics and the exact IPA model, and evaluate the algorithms' performance with some further computational experiments. In particular, although the theoretical approximation bounds from Section 2.2 do not directly carry over to the modified (outside-in BTF and max-settle-row) heuristics, the lower bounds from the MIP solver applied to the modified IPA model still allow us to obtain empirical estimates for the respective solution qualities. We point out that the extended ABP is also strongly NP-hard in general, since at least the proof of Theorem 1 carries over straightforwardly.

At the end of this part, in Section 5.3, we also show a way to adapt the MIP model to different levels of congestion (recall that we work with a fixed congestion level, k, throughout the paper). Numerical experiments and a further (also, theoretical) analysis of varying congestion parameters are beyond the scope of this paper and are left for future research.

5.1 Incorporating Seat Interferences and Inseparable Passenger Groups

Let us begin by demonstrating how to modify IPA to capture the ABP refinements:

Seat interferences: Suppose we wish to take into account possible delays prompted by a passenger having to get up again to let another one move to a seat in the same row (but further from the aisle). A naïve possibility would be to prohibit such situations altogether by enforcing an outside-in pattern: For any pair of passengers $(p_1, p_2) \in \mathcal{P}^2$, $p_1 \neq p_2$, seated in the same row $r = r(p_1) = r(p_2)$, w.l.o.g. with $s(p_1) < s(p_2) \leq k_r^1$, this can be ensured by adding the constraints

$$1 - x_{p_1,i_1} \geq x_{p_2,i_2} \quad \forall (i_1, i_2) \in [P]^2 : i_1 > i_2. \tag{17}$$

to IPA, i.e., to the MIP model (8)–(16); analogously for passenger pairs on the other side of the aisle. Note that even with such quite restrictive constraints, the MIP still allows for more general boarding sequences than the plain outside-in BTF heuristic.

More generally, we can model additional waiting times for the case that a passenger p_2 boards before a passenger p_1 with $r(p_2) = r(p_1)$ (and again, w.l.o.g. $s(p_1) < s(p_2) \leq k_r^1$). To that end, suppose we are given data $\eta_p \in \mathbb{Q}_{\geq 0}$ ($p \in \mathcal{P}$) for the time passenger p needs to get up (to let someone pass) and settle in again. Let $z_{i,p} \in \mathbb{Q}_{\geq 0}$ for all $i \in [P]$, $r \in \mathcal{R}$, and $p \in \mathcal{P}(r)$ be variables to capture additional waiting times connected to passengers having to get up (and resettle) to let passenger p pass to their seat in row r, conditioned on p being the i-th one to board. The constraints

$$M_r\,(x_{p,i} - 1) + \sum_{\substack{p' \in \mathcal{P}(r):\\ s(p') < s(p)}} \sum_{i' < i} \eta_{p'} \cdot x_{p',i'} \leq z_{i,p} \quad \forall i \in [P], r \in \mathcal{R}, p \in \mathcal{P}(r) : s(p) < k_r^1 \tag{18}$$

ensure (for each (i, r, p)) that $z_{i,p}$ is at least the sum of these seat interference waiting times if $x_{p,i} = 1$, and becomes redundant otherwise (if $x_{p,i} = 0$) for a sufficiently large big-M constant, e.g., $M_r := \sum_{p \in \mathcal{P}(r)} \eta_p$. Additionally, constraints (15) in IPA need to be modified to

$$t_{i,r}^F - t_{i,r}^A \geq \sum_{p \in \mathcal{P}} \left(\tau_{p,r} \cdot x_{p,i} + \mathbb{1}_{\{r=r(p)\}} z_{i,p} \right) \quad \forall i \in [P], r \in \mathcal{R}, \tag{19}$$

where $\mathbb{1}_{\{r=r(p)\}} := 1$ if $r = r(p)$, and 0 otherwise. Then, the overall time minimization brings each $z_{i,p}$ down to the respective bound from (18) (if positive; to zero otherwise). Naturally, analogons of constraints (18) and (19) can straightforwardly be set up for the other side of the aisle as well.

Inseparable passenger groups: Typically, there are various groups of passengers who wish to board together (e.g., couples or families), which can be accounted for by adding constraints that ensure such groups are sequenced consecutively. Suppose first that the order of passengers within a group is fixed a priori (e.g., say, a group of children in some arbitrary but fixed order with chaperones at the front and at the back). Then, such constraints take a quite simple form: For an ordered passenger group $(p_1, \dots, p_\ell) \subseteq \mathcal{P}$, we ensure that its members will be mapped as $p_1 \mapsto i_1, \dots, p_\ell \mapsto i_\ell$ with $i_{j+1} = i_j + 1$ $(j = 1, \dots, \ell - 1)$ by enforcing

$$x_{p_{j+1},i+j+1} = x_{p_j,i+j} \quad \forall\, i \in \{0, \dots, P-k\}, j \in [\ell - 1]. \tag{20}$$

In the more general case that the order within a group $G \subseteq \mathcal{P}$ is itself also to be optimized, let $\mu_i^G \in \{0,1\}$, for $i \in \{0, \dots, P - |G|\}$, be decision variables indicating whether the group starts boarding at position $i+1$ (in the boarding sequence) or not. The following constraints then ensure that the group G boards together:

$$\sum_{i=0}^{P-|G|} \mu_i^G = 1, \tag{21}$$

$$\frac{1}{|G|} \sum_{p \in G} \sum_{j=1}^{|G|} x_{p,i+j} \geq \mu_i^G \quad \forall i \in \{0, \dots, P - |G|\}. \tag{22}$$

Indeed, note that exactly one μ-variable must be 1 (21) and that $\mu_i^G = 1$ forces the $|G|$ consecutive x-variables associated with the group G to 1 (22).

Finally, for both extensions, the big-M in (13) needs to be adapted as well; for simplicity, we may (and later do) use the objective value of the best known (heuristic) solution here.

Regarding the heuristics, note that both deterministic (back-to-front) and randomized (classical) outside-in strategies avoid seat interferences by construction, as does the Steffen method. This is not the case for the max-settle-row strategy, and forcing it would render the boarding scheme equivalent to outside-in BTF. One might consider a modification that weighs estimated seat interference waiting times against settle-in times to decide in each round for each row whether to prioritize avoiding seat interferences or selecting a passenger with maximal settling-in time. However, doing so somewhat destroys the simplicity of the max-settle-row strategy and, assuming seat interferences are less time-intensive than settling-in actions, it seems quite unlikely to actually change the constructed sequence. Finally, note that if seat interferences were to be avoided during random boarding, one would simply end up with randomized outside-in boarding.

Inseparable passenger groups can be incorporated into all boarding schemes in several ways. Given the greedy-like nature of most boarding heuristics, we felt the most natural option was this: While constructing a boarding sequence according to some strategy, as soon as a member of a group is encountered, schedule the whole group (in an order that, if not fixed a priori, best matches the strategy's intent, e.g., outside-in) rather than just the one member. Note that when inseparable groups are considered, seat interferences may occur even for strategies that otherwise avoid them.

Last but not least, the 2-opt improvement heuristic can be adapted to the new setting in the obvious fashion: the local search must not split inseparable passenger groups (instead we can reposition a whole group if one of its members is selected during local search), and possible waiting times due to (new or resolved) seat interferences must be considered when evaluating objective function value changes.

5.2 Computational Experiments

For the sake of brevity, we restrict our evaluation of the modified boarding strategies to random and outside-in BTF boarding as well as the exact Gurobi/IPA method[6]. We conducted experiments with just groups

[6] Randomized outside-in and max-settle-row behave comparably to random and outside-in BTF boarding, respectively, with the Steffen method performing in between the by-group and non-static by-seat strategies.

ext.	method	objective	w / 2-opt	% impr.	% gap	# opt	runtime
(a)	random	729.8 (890.5)	445.4	31.75	—	1	15.6
	outside-in BTF	686.7	415.8	36.32	—	2	0.3
	IPA	403.6	—	2.08	17.58	10	1432.1
(b)	random	752.4 (916.6)	408.1	38.62	—	0	1.7
	outside-in BTF	405.7	349.0	12.57	—	1	1.1
	IPA	341.0	—	0.82	7.86	8	1905.9
(c)	random	783.3 (956.6)	456.5	33.00	—	0	16.0
	outside-in BTF	686.7	425.1	35.44	—	2	0.8
	IPA	412.3	—	2.08	18.49	10	1448.7

Table 5. Experimental results on testset mp-sp with ABP extension to (a) inseparable passenger groups, (b) seat interference consideration, and (c) both. Numbers are average values over 40 instances; 10 each for 4 airplane sizes.

of inseparable passengers (arguably the most important refinement of the ABP in terms of realism), just penalizing seat interferences, or taking both into account, all in both the basic and the disruption setup (cf. Section 4). To that end, we extended the IPA model with passenger group constraints (20) (a priori fixing the order within each group in outside-in BTF fashion, which seemed reasonable) and seat interference constraints (18) and (19). In the robustness experiments here, we do not incorporate time data perturbations and instead focus on assessing the impact of late and displaced passengers (or passenger groups—indeed, besides single passengers, whole groups of passengers were treated as late or displaced if any one of their members was selected as being so affected during instance generation).

The test instances are based on the same mp-sp data as in the previous experiments, with the following additions: Seat interference waiting times (τ_p) were set to 10 seconds for all passengers; more precisely, the time a row is blocked by a passenger p settling in effectively increases by 10 s for each passenger who has to get up to let p pass. This is a slight simplification of the scheme used in [Schultz, 2018a], which essentially adds 7.2–9.6 s for one interference (person needing to get up) and 19.2 s for two. For the groups, one has to consider both the total number and the sizes of groups as well as the associated seat assignments. To that end, we followed [Steiner and Philipp, 2009] and generated instances with 55% of the passengers traveling alone, 38% traveling as couples, and 7% traveling in groups of three persons. The seat assignments for the groups were generated as follows: Start with a free seat drawn uniformly at random, and while the number of chosen seats is smaller than the group size, repeatedly select another seat with minimum total taxicab distance (w.r.t. row and seat index differences) to all seats already chosen for the current group. Distance ties are broken by giving preference to a seat directly beside one of the chosen seats; if those are unavailable, the seat is selected uniformly at random instead.

In Table 5, we summarize the experiments for instances with passenger groups, taking seat interferences into account, and considering both, respectively. Table 6 provides an overview of the results for some robustness experiments with both passenger groups and seat interferences. The columns are labeled analogously to the earlier tables in Section 4.

A first observation that is immediate from comparing Table 5 and (the mp-sp part in) Table 2 is that seat interferences apparently do not have a large influence on boarding times overall. Indeed, the differences are essentially negligible across all cabin layouts, at least for the by-seat strategies, and in the presence of passenger groups also for the randomized schemes, with or without 2-opt post-processing. Thus, these results can be seen to confirm earlier claims from the literature that seat interferences are less important than aisle interferences (see, e.g., van den Briel et al. [2005], Nyquist and McFadden [2008], Kierzkowski and Kisiel [2017]). To be fair, this impression might change if the seat interference waiting times are set to much larger values than the 10 s per passenger having to get up that we assumed here, but presumably not when keeping them at a realistic level (relative to all the other time data in an ABP instance).

Moreover, the overall boarding completion times grow significantly compared to the setting without inseparable passenger groups (regardless of seat interference considerations). Nevertheless, the main earlier

method	obj.	late	displ.	comb.	% comb. loss
random + 2-opt	456.5	514.3	533.9	576.5	[27.7, 40.9]
outside-in BTF + 2-opt	425.1	483.4	522.5	571.5	[26.4, 40.0]
IPA	412.3	471.8	519.0	565.7	[25.5, 39.1]

Table 6. Results of robustness experiments on testset mp-sp with inseparable passenger groups and seat interferences. Considered were disruptions caused by late passengers, passenger displacements, and both combined. Numbers are average values over 40 instances; 10 each for 4 airplane sizes.

conclusions continue to hold: Our 2-opt local search post-processing significantly improves the sequences produced by random and outside-in BTF boarding. Overall, in all cases, outside-in BTF boarding (+ 2-opt) still significantly improves over random boarding. Interestingly, Gurobi/IPA now achieves larger reductions (about 0.8–2.1%) in the boarding completion time compared to the best heuristic solution than we could observe in the experiments that disregard groups and seat interferences (only ca. 0.8%). On the downside, the IPA MIPs do not appear to become easier to solve (for Gurobi) with the additional constraints; indeed, the average gaps now range from about 8% to roughly 18% compared to the 6.55% from Table 2, though the number of instances solved to optimality within the one-hour time limit is comparable. Nevertheless, the lower bounds provided by Gurobi/IPA show that, e.g., (adapted) outside-in BTF with 2-opt refinement produced solutions that never exceeded optimality by more than 58% (31% on average).

Turning to Table 6, we can observe that for the extended ABP, the relative differences in final boarding completion times between heuristics and the exact IPA method (with time limit) are comparable to those in the earlier robustness experiments (cf. Table 3). Due to considering seat interferences and inseparable passenger groups, the boarding times are generally larger overall, but outside-in BTF still outperforms random boarding, and Gurobi/IPA yields some further improvements. The combination of both late-comers and displaced passenger naturally yields the worst overall boarding completion times, which are between 25% and 41% larger than the ideal boarding times (for undisrupted sequences) on average; the sequences computed with Gurobi/IPA are the most stable, incurring the smallest average losses compared to the ideal boarding times. The loss percentage intervals are much wider than in the previous experiments (without passenger groups and seat interference considerations), with notably smaller lower bounds and comparable upper bounds. The spread is likely explained by the MIPs being harder to solve (yielding larger optimality gap estimates when terminating prematurely due to running into the time limit), but the smaller lower ends of the loss intervals are actually good news: If, as may be expected (since it is quite typical MIP solver behavior), the large gaps are mostly linked to the MIP solver having trouble increasing the MIP objective lower bounds but the best known solutions are already (close to) optimal, then the smaller lower loss bounds compared to the basic ABP (see Table 3) indicate that the extended ABP is actually more robust to the disruptions caused by late or displaced passengers. Although larger total boarding times are unavoidable in the more realistic setting with (in particular) inseparable passenger groups, the further loss due to disruptive passenger behavior would be only about 26% here, in contrast to almost 41% for the basic ABP.

5.3 Adapting the MIP to Different Congestion Parameter Values

Throughout this paper, the congestion parameter—i.e., the ratio of the length of the single-file passenger queue over the length of the airplane cabin—is implicitly equal to k, the number of seats per row, for a fully occupied aircraft. However, different cabin layouts and changes in (assumed or practical) personal space requirements of passengers may yield other values. The congestion influences which boarding strategy results in the best expected boarding times on average, cf., e.g., [Bachmat et al., 2009], who also suggest a value of 4 for realistic instances (on larger airplanes). Hence, to enhance its capabilities and comparability in simulation frameworks that allow for varying congestion levels, we now briefly show how our IPA model can be modified to account for quite general congestion values.

Let $c \in \mathbb{Q}_{>0}$ be the congestion parameter. As mentioned, in the current model, we have $c = P/R$, which we now extend to $c = (\ell P)/R$ or $c = P/(\ell R)$ for arbitrary $\ell \in \mathbb{Z}_{\geq 2}$; the two modifications can then be combined to extend the ABP to arbitrary $c \in \mathbb{Q}_{>0}$.

First, consider the case $c = (\ell P)/R$, which amounts to stretching the passenger queue further, e.g., by increasing personal space. Since we scale from "one person per row" by an integer factor, we only need to adjust the constraints (14) to now read:

$$t^A_{i+1,r-\ell+1} \geq t^F_{i,r} \qquad \forall i \in [P-1],\ r \in \mathcal{R}:\ r \geq \ell - 1.$$

Together with (12), this ensures the desired spacing.

For a congestion parameter value $c = P/(\ell R)$, we essentially split every row into ℓ "parts" that can be occupied/blocked separately. To that end, we may introduce dummy passengers ($\ell - 1$ copies for every one) associated with the row parts, i.e., we define $\mathcal{P}' := [\ell P]$ and let each copy of a single passenger i (say, $i, i+P, \ldots, i+(\ell-1)P$) have the same seat number but in a distinct part of row $r(i)$. For convenience, we virtually increase the number of rows by a factor ℓ as well, and so redefine $\mathcal{R} := [\ell R]$. The seat assignments for all (virtual) passengers $i = \ell' P + p \in \mathcal{P}'$ can now be written as $r'(\ell' P + p) := \ell(r(p) - 1) + 1 + \ell'$ for all $(p, \ell') \in \mathcal{P} \times \{0, 1, \ldots, \ell - 1\}$.) Now, to ensure that there is exactly one copy of each original passenger among the P passengers that are boarded, we modify (10) to

$$\sum_{i \in [P]} \sum_{\ell'=0}^{\ell-1} x_{\ell' P + p, i} = 1 \qquad \forall p \in \mathcal{P}.$$

The moving and settling-in time data for all dummy passengers are copied from the respective original passengers as well. Indeed, since we retain that passengers virtually pass every row, we also need to spread the moving-times $t^m_{p,r}$ evenly over the ℓ new rows for each original one, i.e., for each original row $r \in [R]$, we have $t^m_{p',r'} = t^m_{p,r}/\ell$ for each $p \in [P]$ and its copies p', and each new row r' associated with r. Such a split is clearly not necessary for the settling-in times.

By this construction, the extended virtual model now behaves analogously to the original IPA model but indeed only boards one of each passenger copy, ensuring that an optimal boarding sequence is obtained by simply mapping back dummies to the original passengers.

By combining the above model constructions, the case $c = (\ell_1 P)/(\ell_2 R)$ for arbitrary $\ell_1, \ell_2 \in \mathbb{Z}_{\geq 1}$ can also be handled, as would be needed, for example, to move from a congestion level of 6 as in our large instances down to the value 4 (via $\ell_1 = 2$, $\ell_2 = 3$). Note also that the Big-M constants in (13) may need to be adjusted as well; for simplicity, we could again take the boarding completion time of an arbitrary heuristic (cf. Section 5.1).

It goes beyond the scope of the present paper to test the practicality of the congestion-related extensions of our MIP model (note that the second construction increases the number of variables and constraints, possibly significantly). There might also exist other MIP formulations that incorporate other congestion levels more directly than by extending IPA as demonstrated. We leave these aspects for future consideration.

6 Concluding Remarks

The present paper offered several key contributions to the research on time-efficient airplane boarding methods. We provided an (almost) complete characterization of the airplane boarding problem in terms of computational complexity, under mild assumptions: We proved strong NP-hardness for four out of six considered models of passenger moving and settle-in times, and that (assuming a fully booked flight and congestion equal to the number of seats per row) one of the two other cases can be solved to optimality in polynomial time by the back-to-front by-seat variant of the well-known outside-in boarding strategy. In spite of this positive result for outside-in back-to-front boarding, we also proved that all deterministic boarding strategies that neglect passenger time information (encompassing all commonly used ones, including outside-in BTF) can, in general, yield boarding times that are far from optimal. Moreover, we proved theoretical lower bounds for the optimal boarding time that led to the first approximability guarantees for the ABP (again under the assumptions that the plane is fully occupied and that each passenger needs space equal to the aisle-length of one row), which encompass two data models, including the one not covered by our complexity findings. We then developed the first algorithmic approach that directly tackles the problem of minimizing the airplane boarding time, in contrast to previous work focusing on simulations, probabilistic analyses, or related but different goals such as minimizing the number of passenger interferences. In an empirical study, we compared a new MIP model with several heuristics; we were able to solve between a quarter and more than a third of our test instances to optimality within a one-hour time limit, and provide computational certificates of solution quality for the heuristics in any case. Moreover, our computational results showed that the proposed "max-settle-row" boarding strategy outperforms all other heuristics, closely followed by the related outside-in BTF boarding strategy, and that a simple local search can significantly improve heuristic solutions and thus yield better by-seat sequences. We also assessed the robustness of the considered boarding methods with respect to perturbations of the passenger time data and passengers disrupting the planned boarding sequence. Finally, we described how to adapt the exact MIP model and boarding heuristics to the extended airplane boarding problem that also takes inseparable passenger groups and seat interferences into account, along with a computational evaluation; we also demonstrate how the MIP model may be adjusted to different congestion assumptions.

The robustness experiments, in particular, revealed that there is still a large potential for improvements that may be gained by incorporating robustness aspects explicitly. Thus, we believe it to be an interesting topic for future research to tackle the (extended) boarding time minimization problem from the perspective of robust optimization (see, e.g., [Ben-Tal et al., 2009]). Moreover, it may be worth spending some effort to devise a dedicated branch-and-cut solver for the ABP that incorporates bounds like those from Lemma 3 locally, heuristics, and further MIP techniques (branching rules, cutting planes, etc.) still to be investigated for the present problem. Prior research on MIPs and heuristics for related permutation flow shop problems, e.g., [Wilson, 1989, Jessin et al., 2020] and associated references, might be a good starting point for such endeavors. A more efficient MIP algorithm would open the door to more extensive simulation studies, involving further heuristics and larger testsets. Furthermore, it will be interesting to see how further aspects such as allowing passengers to store carry-on luggage at arbitrary locations could be incorporated into the newly proposed MIP framework.

It is also of interest to continue the theoretical investigation of simple (by-seat) boarding strategies. Some possible topics for future research in this direction are resolving the question of possible optimality or approximation guarantees of (say) the outside-in BTF strategy in case the plane is not fully occupied or the congestion parameter value changes, obtaining constant factor approximation algorithms, or answering the open question of the complexity of the ABP in case of passenger-dependent moving times and identical-for-all settle-in times (cf. Table 1).

6.1 Aknowledgement

The authors would like to thank Marco Lübbecke for hosting inspiring discussions in his gORden as well as the two anonymous reviewers who provided valuable and constructive feedback that helped to improve the paper. A. M. Tillmann would like to acknowledge the funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2163/1 – Sustainable and Energy Efficient Aviation – Project-ID 390881007.

7 APPENDIX

A Proof of Lemma 1

To show correctness, it suffices to prove the assertion that $C(i, r)$ is indeed the time row $r \in \mathcal{R}$ becomes accessible again after having been occupied by a passenger $p \in \{\pi^{-1}(1), \dots, \pi^{-1}(i)\}$. To that end, we proceed by induction over $i \in [P]$.

Since the very first passenger $\pi^{-1}(1)$ is never blocked, this clearly holds true for $C(1, 1), \dots, C(1, R)$. For the induction step, consider passenger $p \coloneqq \pi^{-1}(i)$ (and presume the $C(i, r)$-values are computed in increasing row-index order $r = 1, 2, \ldots, R$, i.e., following the progression of p through the plane). For $r \le r(p)$, p can access r only after no predecessor $(\pi^{-1}(i-1), \ldots, \pi^{-1}(1))$ any longer occupies it (i.e., at $C(i-1, r)$) and only if p has already reached row r (i.e., at $C(i, r-1)$). In row r, p then spends $t^m_{p,r}$ time moving on to the next row (if $r < r(p)$) or t^s_p time settling in (if $r = r(p)$). In the former case, if row $r+1$ is still occupied by any predecessor $(\pi^{-1}(i-1), \ldots, \pi^{-1}(1))$ by the time p would be done in row r, p actually has to wait (in row r) until $r+1$ becomes accessible again (i.e., until $C(i-1, r+1)$). Finally, since p does not traverse rows $r > r(p)$, the times $C(i-1, r)$ for those rows may be preserved (note that for passengers boarding after p, it is nevertheless ensured by construction that they cannot arrive at a row $r > r(p)$ earlier than at $C(i, r(p))$, so it does not matter if $C(i-1, r) < C(i, r(p))$). Put together, this concludes the induction, and yields the formula for $C(i, r)$-values as claimed.

Regarding the running time $\mathcal{O}(PR)$ to obtain $C_{\max} = \max_{r \in \mathcal{R}} C(P, r)$, note that by iterating over passengers (in sequence π) and, for each passenger, over the rows (front to back), every required value is readily available (computed earlier or as given data) and computation of any $C(i, r)$ then reduces to a constant number of arithmetic operations. This completes the proof. ◻

B Proof of Theorem 1

For rational data (moving and settle-in times, and T), containment in NP is easy to see: A "yes"-certificate is given by a permutation π of $\mathcal{P}$, for which the boarding completion time $C_{\max}(\pi)$ can be computed in polynomial time by Lemma 1; thus, $C_{\max}(\pi) \le T$ can be verified in polynomial time.

To show hardness, we reduce from $F_2|\, b = 1\, |C_{\max}$, a flow shop scheduling problem that was shown to be strongly NP-hard in [Papadimitriou and Kanellakis, 1980] (for integer processing times). In this problem, we are given a set of n jobs J and two machines m_1 and m_2; each job $j \in J$ consists of tasks j_m to be processed on machine $m \in \{m_1, m_2\}$ with corresponding processing times $t_{j,m}$. Between the machines, there is a temporary job buffer which can hold a single job that finished on the first machine, if the second machine is not yet available; this buffer is of a first-in first-out (FIFO) nature and cannot be bypassed by jobs. The goal is to minimize the makespan, i.e., the completion time of the last task.

Recall that in a flow shop, all jobs share the same technological order, i.e., they must be processed by every machine and always in the same order (w.l.o.g., $m_1 \to m_2$). Consequently, by definition of the FIFO buffer, the processing order of the tasks is necessarily identical on both machines in $F_2|\, b = 1\, |C_{\max}$ (i.e., this is in fact a permutation flow shop problem). Moreover, Hall and Sriskandarajah [1996, Lemma 1] observed that $F_2|\, b = 1\, |C_{\max}$ can be expressed equivalently as an instance of $F_3|\, blocking\, |C_{\max}$ in which the processing times on the second machine are all zero, and "blocking" refers to allowing jobs to remain on a machine after processing if the next machine is still busy, and blocking the machine for other jobs while doing so. We will use a similar idea to transform the FIFO buffer into an extra row in the airplane cabin in our reduction.

Thus, let $C \in \mathbb{N}$, $J = [n]$, $M = \{m_1, m_2\}$ and $\{t_{j,m} \in \mathbb{Z}_{\ge 0} : j \in J,\ m \in M\}$ be an instance of the decision version of $F_2|\, b = 1\, |C_{\max}$ asking for a feasible permutation schedule with $C_{\max} \le C$. We construct an instance of ABP as follows: Set $T := C$, $\mathcal{P} := [n+3]$, $\mathcal{R} := [n+3]$, $\mathcal{S}^1_r := \{(r, 1)\}$ and $\mathcal{S}^2_r := \emptyset$ for each $r \in \mathcal{R}$, and $t^s_p := 0$ for all $p \in \mathcal{P}$. Moreover, let $t^m_{p,1} := t_{p,m_1}$ and $t^m_{p,3} := t_{p,m_2}$ for $p = 1, \ldots, n$, $t^m_{p,r} := 0$ in all other cases, as well as

$$\sigma(p) := \begin{cases} (3+p, 1), & 1 \le p \le n, \\ (p-n, 1), & n+1 \le p \le n+3. \end{cases}$$

We claim that this instance admits a boarding sequence π with boarding completion time at most T if and only if the input instance of $F_2|\, b = 1\, |C_{\max}$ has a feasible solution ρ with makespan at most C. Indeed, π without the three "dummy" passengers $n+1, n+2, n+3$ is in one-to-one correspondence with ρ: Suppose we have a permutation flow shop schedule $\rho = (\rho_1, \ldots, \rho_n)$ with $C_{\max} \le C$. Then, it is easily seen that in the ABP instance constructed above, each task j_m of a job j being processed on machine $m \in M$ with processing time $t_{j,m}$ corresponds exactly to passenger p, $1 \le p \le n$, moving within row r, $r \in \{1, 3\}$, for the same amount of time, where row 1 corresponds to machine m_1, and row 3 to m_2. Row 2 can be seen to take the place of the FIFO buffer: the machine blockages (buffer utilizations) in the flow shop problem are mirrored in row blockages (aisle interferences) by the corresponding "job" passengers passing the "machine" rows on their way to the assigned seats in the dummy rows $4 \leq r \leq n+3$; the dummy passengers seated in the first three rows have no influence on the completion time and merely serve the purpose of assigning one passenger to each seat (which is not necessary and may be omitted; it allows to cover the case of full plane occupancy). Conversely, a boarding sequence π with completion time at most T can be mapped directly to a desired flow shop schedule ρ by simply removing the dummy jobs. □

C Proof of Theorem 2

To show hardness, we reduce from the 3-PARTITION problem, which is well-known to be strongly NP-complete, cf. [Garey and Johnson, 1979]. In this problem, we are given a finite set $A \coloneqq \{1, 2, \ldots, 3m\} = [3m]$ ($m \in \mathbb{N}$), an integer $B \in \mathbb{N}_{\geq 3}$, and numbers $a_i \in \mathbb{N}$ for $i \in A$ such that $B/4 < a_i < B/2$ and $\sum_{i \in A} a_i = mB$, and we wish to decide whether A can be partitioned into m disjoint three-element sets $A_1, A_2, \ldots, A_m$ such that $\sum_{i \in A_j} a_i = B$ for every $j \in [m]$.

Given such an instance of 3-PARTITION, we construct an instance $\mathcal{I}$ of the decision version of ABP with $T \coloneqq mB$, four rows $\mathcal{R} \coloneqq \{1, 2, 3, 4\}$, seats in the first and fourth rows only,

$$k_r^1 \coloneqq \begin{cases} m, & r = 1, \\ 3m, & r = 4, \\ 0, & \text{otherwise}, \end{cases}$$

($k_r^2 = 0$ and $\mathcal{S}_r^2 = \emptyset$ for all $r \in \mathcal{R}$), passengers $\mathcal{P} \coloneqq [4m]$, all-zero moving times $t_{pr}^m = 0$ for all $p \in \mathcal{P}$, $r \in \mathcal{R}$, settle-in times

$$t_p^s \coloneqq \begin{cases} a_p, & 1 \leq p \leq 3m, \\ B, & 3m+1 \leq p \leq 4m, \end{cases}$$

and seat assignments,

$$\sigma(p) \coloneqq (r(p), s(p)) \coloneqq \begin{cases} (4, p), & 1 \leq p \leq 3m, \\ (1, p-3m), & 3m+1 \leq p \leq 4m. \end{cases}$$

Note that we could have included an arbitrary number of seats in the two middle rows along with suitably many dummy passengers with associated times that do not influence the reduction, but for simplicity, we simply defined these rows to be devoid of any seats.

We show that the original 3-PARTITION instance is a "yes"-instance if and only if $\mathcal{I}$ admits a boarding sequence that has completion time *exactly* T $(= mB)$. Note that every boarding sequence has completion time *at least* T, since m passengers with settle-in time B have their seats in row 1. (Thus, any sequence with completion time at most T will in fact have completion time equal to T.)

"⇒": Let $A_1, A_2, \ldots, A_m$ be such that $A_1 \dot\cup A_2 \dot\cup \ldots \dot\cup A_m = A$, and $|A_j| = 3$ as well as $\sum_{i \in A_j} a_i = B$ for all $j \in [m]$. Furthermore, let $A_j = \{i_1^j, i_2^j, i_3^j\}$ (arbitrarily ordered) for every $j \in [m]$. We set

$$\pi(p) \coloneqq \begin{cases} 4(j-1)+k, & p = i_k^j \in A_j \subseteq [3m], \\ 4(p-3m), & 3m+1 \leq p \leq 4m. \end{cases}$$

For any $j \in [m]$, the group of four passengers $\pi^{-1}(4j-3), \ldots, \pi^{-1}(4j)$ would need time exactly B from passenger $\pi^{-1}(4i-3)$ entering the plane until they have all settled in, if they were not blocked by any other passengers. (Note that, by construction, the two middle rows serve as a "buffer" for passengers waiting to get to the last row.) The passengers $\pi^{-1}(1), \ldots, \pi^{-1}(4)$ are indeed not blocked and thus finish settling in at time B. Inductively, since every fourth passenger in the sequence π sits in the first row and thus blocks all subsequent passengers (from getting on board) until having settled in, each passenger group $\pi^{-1}(4j-3), \ldots, \pi^{-1}(4j)$ with $2 \leq j \leq m$ has settled in at time jB, respectively. Consequently, the total boarding completion time is exactly mB.

"$\Leftarrow$": Let π be a boarding sequence for $\mathcal{I}$ with completion time $C_{\max}(\pi) = mB$. We may and do assume w.l.o.g. that $\pi(3m+1) < \pi(3m+2) < \cdots < \pi(4m)$, and for notational convenience, we define $\mathcal{P}' \coloneqq \{p \in \mathcal{P} : \pi(p) < \pi(3m+1)\}$ and $\mathcal{P}'_j \coloneqq \{p \in \mathcal{P} : \pi(3m+j) < \pi(p) < \pi(3m+j+1)\}$ for $j \in [m-1]$. Thus, $\mathcal{P}'$ describes the passengers that, according to π, board before passenger $3m+1$, and $\mathcal{P}'_j$ those that board before passenger $3m+j+1$ but after passenger $3m+j$.

First, we note that the very last passenger to board the plane, $\pi^{-1}(4m)$, must be one of the passengers with seat in row 1—otherwise, the last passenger would be blocked by the last first-row passenger while the latter settles in, and the completion time would trivially be greater than mB (which coincides with the sum of settle-in times of all first-row passengers).

We proceed by contraposition: Assume that either $|\mathcal{P}'| \geq 4$ or $|\mathcal{P}'_j| \geq 4$ for at least one $j \in [m-1]$. Then, there exists an $\ell \in [m]$ such that passenger $3m+\ell$ can *begin* settling in no earlier than at time

$$S \coloneqq \min\{a_p : p \in \mathcal{P}'\}$$

if $\ell = 1$, or

$$S_\ell \coloneqq (\ell-1)B + \min\{a_p : p \in \mathcal{P}'_{\ell-1}\}$$

if $2 \leq \ell \leq m$, respectively. The term $(\ell-1)B$ reflects the sum of settling-in times of all first-row passengers that have boarded before passenger $3m+\ell$. Now, in case $\ell = 1$, passenger $3m+1$ completes settling in no earlier than at $S+B$, and since $S > 0$ and the remaining first-row passengers $3m+2, \ldots, 4m$ each take B time to settle in, the total boarding time exceeds mB. This contradicts the fact that $C_{\max}(\pi) = mB$. Similarly, in case $2 \leq \ell \leq m$, since $S_\ell > (\ell-1)B$, it follows that $S_\ell + (m-\ell+1)B > mB$. Because $S_\ell + (m-\ell+1)B$ is a lower bound on the completion time (analogously to the case $\ell = 1$), this is again a contradiction to $C_{\max}(\pi) = mB$. Thus, it follows immediately that $|\mathcal{P}'| \leq 3$ and $|\mathcal{P}'_j| \leq 3$ for all $j \in [m-1]$. Moreover, if any of these sets had cardinality strictly smaller than 3, a simple counting argument would yield the existence of another set with cardinality exceeding 3, which we have just ruled out. Thus, we have in fact shown that

$$|\mathcal{P}'| = |\mathcal{P}'_1| = \cdots = |\mathcal{P}'_{m-1}| = 3.$$

Next, assume (again by contraposition) that either

$$B < T' \coloneqq \sum_{p \in \mathcal{P}'} a_p$$

or, for at least one $j \in [m-1]$,

$$B < T'' \coloneqq \sum_{p \in \mathcal{P}'_j} a_p.$$

Then, there exists an $\ell \in [m-1]$ such that passenger $3m+\ell+1$ can begin settling in no earlier than at time

$$S \coloneqq \ell B + \begin{cases} T' - B, & \text{if } \ell = 1, \\ T'' - B, & \text{if } 2 \leq \ell \leq m-1, \end{cases}$$

and it holds that $S > \ell B$. By similar arguments as used before, $S + (m-\ell)B$ is an obvious lower bound on the completion time of π, but since $S + (m-\ell)B > mB$, we arrive at a contradiction to the presupposition $C_{\max}(\pi) = mB$. Thus, the settle-in times associated with any three-element set $\mathcal{P}', \mathcal{P}'_1, \ldots, \mathcal{P}'_{m-1}$ sum to at most B; in fact, since $\sum_{p \in [3m]} a_p = mB$, these sums are all exactly equal to B.

This shows that we have identified a "yes"-certificate of the original 3-PARTITION instance by setting

$$A_j \coloneqq \mathcal{P}'_j \text{ for } 1 \leq j \leq m-1, \quad A_m \coloneqq \mathcal{P}',$$

which completes the proof of NP-hardness.

It remains to note that the reduction is clearly polynomial (in particular, retaining boundedness in the problem dimension of all occurring numbers as well as their respective encoding lengths and thus preserving the "in the strong sense" assertion of NP-hardness), and that containment in NP can be obtained analogously to the proof of Theorem 1. □

D Proof of Corollary 1

Since the seat assignments σ are part of an ABP instance, we may simply extend the respective reductions in the NP-hardness proofs by, e.g., appending a row at the back of the aircraft containing at least one seat to which no passenger will be assigned in the constructed instance. (In the instance from the proof of Theorem 2, we could also place unbooked seats in the two middle rows.) Clearly, this has no further implications w.r.t. the rest of those proofs. $\square$

E Proof of Proposition 1

This can easily be seen by considering the boarding sequence π and propagating the changes in t^m- and/or t^s-values: Decreasing any such value either yields no change or leads to earlier termination of some moving and settling-in operations in the given sequence, but never incurs delays that were not present before. Thus, the previously optimal sequence π is a feasible (but not necessarily optimal) solution of the modified ABP instance, which therefore has optimal completion time $C'_{\max} \leq C_{\max}(\pi)$. $\square$

F Proofs of Lemmas 2 and 3 and Theorem 3

Proof. Proof of Lemma 2. The first passenger $p_1 := \pi^{-1}(1)$ in the boarding sequence needs time $\sum_{r=1}^{R-1} t^m_{p_1,r} + t^s_{p_1}$ to finish boarding. The second passenger $p_2 := \pi^{-1}(2)$ can pass a row $r \leq R-2$ either directly in time $t^m_{p_2,r}$, or in time at most $t^m_{p_1,r+1}$ if he/she has to wait (in r) while the next row $r+1$ is still blocked by the preceding passenger p_1 currently passing it. Pursuing this idea, and using (2), we can upper-bound the time that passenger $\pi^{-1}(i)$ in position $i \leq R$ needs to finish boarding as follows:

$$\sum_{r=1}^{r(\pi^{-1}(i))-1} \max\left\{t^m_{\pi^{-1}(i),r}, t^m_{\pi^{-1}(i-1),r+1}, \dots, t^m_{\pi^{-1}(1),R-1}\right\} + t^s_{\pi^{-1}(i)}.$$

With the obvious indexing adjustments, this bound can be adapted straightforwardly to upper bounds for the time that passengers of the other "groups" of R passengers (of which there are k) need for boarding. Clearly, the next group of R passengers can begin boarding *at the latest* when *all* passengers of the previous group have settled in, so the sum of the respective upper bounds over all k groups gives an upper bound on the total boarding time. Combining this summation with maximization of the passenger-specific boarding completion bound within each group yields the claimed upper bound (3) on $C_{\max}(\pi)$. $\square$

Proof. Proof of Lemma 3. Each passenger p has to pass rows $1, \dots, r(p)-1$ before eventually settling in at row $r(p)$. If a passenger is not blocked along the way, he/she needs time exactly $\sum_{r=1}^{r(p)-1} t^m_{p,r} + t^s_p$ to complete boarding. Thus, maximization over all passengers yields lower bound (4).

Passengers p with seat in row $r(p) \geq r+1$ have to pass row r and all passengers with seat in row r have to settle in at r. None of these events can happen at the same time. Thus, row r is blocked for at least $\sum_{p:r(p)\geq r+1} t^m_{p,r} + \sum_{p\in\mathcal{P}(r)} t^s_p$. Additional time is required for the first passenger who will block row r to pass rows $1, \dots, r-1$ (on the way to row r), which can be lower bounded by the minimum of the associated moving time sums. Finally, maximizing over all rows gives lower bound (5).

To see (6), note that in a best case, the sum of all settle-in and moving times is distributed evenly over the maximal number of actions that can take place in parallel. Since this number is obviously bounded from above by $\min\{R, P\}$, (6) indeed provides a lower bound on the minimum achievable boarding time. $\square$

Proof. Proof of Theorem 3. With $t^m_{p,r} = c^m$ and $t^s_p = c^s$ for all p, r, it can easily be seen that the terms in the maximum in the lower bound (5) all reduce to $(P-k)c^m + kc^s$, which hence is the value of this boarding time lower bound itself. Similarly, the upper bound (3) (which is applicable because the outside-in BTF strategy clearly satisfies the requirement (2) from Lemma 2) reduces to

$$\sum_{i=0}^{k-1} \max_{j\in\mathcal{R}} \left\{ c^s + \sum_{r=1}^{r\left(\pi^{-1}(j+iR)\right)-1} c^m \right\} = kc^s + \sum_{i=0}^{k-1}\sum_{r=1}^{R-1} c^m = kc^s + k(R-1)c^m.$$

To see the first equality, recall that by (2), every group of R passengers who board consecutively includes one person with seat in the last row, so for every $i = 0, 1, \ldots, k-1$, the maximum is attained for j such that $r(\pi^{-1}(j+iR)) = R$. Since $P = |\mathcal{P}| = |S| = kR$, this upper bound coincides with the lower bound (5), which immediately shows optimality of the outside-in BTF boarding sequence. For the running time bound $\mathcal{O}(P)$, note that we can compute the position of any passenger $p \in \mathcal{P}$ in the outside-in BTF sequence directly using (1). Therefore, iterating over all passengers once is indeed sufficient to obtain the full boarding sequence. $\square$

G Proof of Proposition 2

Let $C_{\max}(\pi)$ and $C^*_{\max}$ be the boarding completion time of sequence π and the minimum boarding time, respectively. By Lemmas 2 and 3, it holds that

$$C_{\max}(\pi) \leq (3) \quad \text{and} \quad C^*_{\max} \geq \max\{(4), (5), (6)\},$$

whence $C_{\max}(\pi)/C^*_{\max} \leq \beta$, which proves the claim. $\square$

H Proof of Theorem 4

First consider the case $t^m_{p,r} = c^m_p \in \mathbb{Q}_{\geq 0}$ for all $(p,r) \in \mathcal{P} \times \mathcal{R}$, $t^s_p = c^s \in \mathbb{Q}_{\geq 0}$ for all $p \in \mathcal{P}$. In this setting, the upper bound (3) on the completion time of the outside-in BTF boarding strategy collapses to

$$kc^s + \sum_{i=0}^{k-1} \max_{\substack{p \in \mathcal{P}: \\ iR+1 \leq \pi(p) \leq (i+1)R}} (r(p)-1)c^m_p,$$

which can be further estimated from above by $k(c^s + \max_{p\in\mathcal{P}}(r(p)-1)c^m_p)$. The latter corresponds to exactly k times the lower bound (4), which shows (analogous to (7) and its proof in Appendix G) that the outside-in BTF boarding time cannot be worse than k times the minimum boarding time.

Now suppose $t^m_{p,r} = c^m \in \mathbb{Q}_{\geq 0}$ for all $(p,r) \in \mathcal{P} \times \mathcal{R}$, $t^s_p \in \mathbb{Q}_{\geq 0}$ (arbitrary) for all $p \in \mathcal{P}$. In this case, similarly to what we just saw, the upper bound (3) can be further bounded from above by $k \max_{p\in\mathcal{P}}\{(r(p)-1)c^m + t^s_p\}$, which is again equal to k times the lower bound (4). $\square$

I Proof of Proposition 3

First suppose $P = |\mathcal{S}|$. Let $k \in \mathbb{N}_{\geq 2}$ and $0 < \varepsilon \leq k(k-1)/(3k-2)$, let $c^m, c^s \in \mathbb{Q}_{>0}$ be constants to be defined later, and consider the following ABP instance with $R \geq 2k$ (R will also be specified later): For ease of presentation, assume w.l.o.g. that all seats are on the same side of the aisle ($(r(p), s(p)) \in \mathcal{S}^1_{r(p)}$ $\forall p \in \mathcal{P}$). Only the passengers seated in the last row have positive moving times, namely $t^m_{p,r} = c^m$ for all $(p,r) \in \mathcal{P}(R) \times \mathcal{R}$, only k passengers—one per row for the first k rows—have positive settle-in times, namely $t^s_p = c^s$ for $p \in \{\sigma^{-1}(i,i) : i \in [k]\}$ (i.e., the window-seat passenger in the first row, the next-to-the-window passenger in the second row, "diagonally" continued to the aisle-seat passenger in row k), and all other t^m- and t^s-values are zero.

Due to aisle interferences (passengers trailing behind the last-row passengers), a passenger with positive settle-in time in one of the first k rows only ever reaches their target row once the preceding last-row passenger has reached row R. Thus, it is easily verified that for this instance, the boarding time of the outside-in BTF strategy is exactly $k(R-1)c^m + kc^s$. On the other hand, if all passengers from the last row were to board first, followed by the passengers with nonzero settle-in times from the first k rows in decreasing order of their row numbers (back-to-front fashion), and finally all other passengers in arbitrary order, then the boarding time would only be $\max\{(R-1)c^m + (k-1)c^m, (2k-1)c^m + c^s\}$. (The second summand in the first term amounts to the cumulative waiting times of the last-row passengers, and the first summand of the second term to those of the passengers with nonzero settle-in times.)

As the latter gives an upper bound on the actual optimal boarding time, the completion time of the outside-in BTF boarding strategy is at least

$$\frac{k(R-1)c^m + kc^s}{\max\{(R-1)c^m + (k-1)c^m, (2k-1)c^m + c^s\}} \tag{23}$$

times the optimal completion time. Let $R \coloneqq \lceil k(k-1)/\varepsilon \rceil - k + 2$, $c^m \coloneqq 1$, and $c^s \coloneqq (R-k-1)$. Then, $(R-1)c^m + (k-1)c^m = c^s$ and term (23) collapses to

$$k + \frac{k(R-1)}{(R-1)+(k-1)} = 2k - \frac{k(k-1)}{\lceil k(k-1)/\varepsilon \rceil} \geq 2k - \varepsilon.$$

Observing that to extend the result to the case $P < |\mathcal{S}|$, we may simply add, e.g., an additional $(R+1)$-th row with unoccupied seats at the rear-end of the plane (either omitting them from the outside-in BTF sequence, or including dummy passengers with all-zero time data to be "seated" there) completes the proof. $\square$

J Proof of Theorem 5

We start with considering a deterministic strategy.

First, suppose that the given strategy π boards passengers in back-to-front fashion, i.e., for every $(p_1, p_2) \in \mathcal{P}^2$ with $r(p_1) < r(p_2)$, it holds that $\pi(p_1) > \pi(p_2)$. In particular, this implies that all passengers with seats in row R board first. Suppose that $t^s_{\sigma^{-1}(\pi^{-1}(1),R)} = T = t^s_{\sigma^{-1}(\pi^{-1}(P),r(\pi^{-1}(P)))}$ (i.e., the very first and the very last passenger according to the considered boarding strategy have settle-in time T, for an arbitrary constant $T > 0$), and that all other t^m- and t^s-values are zero. Then, the strategy results in boarding time $2T$, since row $R-1$ is blocked by another passenger $p \in \mathcal{P}(R)$ while passenger $\pi^{-1}(1)$ is settling in, so (even if $R = k = 2$) the last passenger necessarily has to wait T time units before possibly starting to settle-in, which then again takes T time steps. Obviously, an optimal boarding strategy would let the two passengers with nonzero settle-in times settle in simultaneously and thus has boarding time T.

Now we consider the case that the given boarding strategy π does *not* follow a back-to-front principle. To that end, assume there are two passengers p_1, p_2 with $r(p_1) < r(p_2)$ but $\pi(p_1) < \pi(p_2)$, i.e., the considered strategy lets passenger p_1 board before passenger p_2, but p_1 has their seat in an earlier row than p_2. Suppose that $t^s_{p_1} = t^s_{p_2} = T$ ($T > 0$ arbitrary), and that all other t^m- and t^s-values are zero. Then, the strategy yields boarding time $2T$, since passenger p_2 is blocked by passenger p_1 (aisle interference) and the two cannot settle in simultaneously. Again, an optimal boarding strategy clearly yields boarding time T.

For (by-group) strategies that allow passengers within boarding groups to be in arbitrary order, it suffices to observe that any sequence that may be produced by such a strategy still adheres to the above case distinction, and so a sequence that yields boarding time $2T$ can very well be the outcome of applying the randomized strategy.

It remains to note that in these constructions, one can obviously always find some passenger p (with $t^m_{p,r} = t^s_p = 0$ for all r) whose removal does not change the described aisle interference behavior, so the arguments go through for both plane occupancy cases $P = |\mathcal{S}|$ and $P < |\mathcal{S}|$. $\square$

K Proof of Theorem 6

Let $t^s_p \in \mathbb{Q}_{\geq 0}$ be arbitrary, and let π be the max-settle-row boarding sequence. To upper-bound the approximation quality, consider the quotient of the upper bound (3) and the lower bound (5), which, in the present setting, reads

$$\frac{\displaystyle\sum_{i=0}^{k-1} \max_{\substack{p \in \mathcal{P}:\\ iR+1 \leq \pi(p) \leq (i+1)R}} \big((r(p)-1)c^m + t^s_p\big)}{\displaystyle\max_{r \in \mathcal{R}} \left(\big(k(R-r) + (r-1)\big)c^m + \sum_{p \in \mathcal{P}(r)} t^s_p \right)} =: \frac{U}{L}. \tag{24}$$

For the remainder of the proof, we may assume that $L \neq 0$, because otherwise, it would necessarily hold that $c^m = 0$ and $t_p^s = 0$ for all $p \in \mathcal{P}$, so any sequence π would in fact be optimal.

Since the passengers from each row are sequenced in non-increasing order of their settle-in times, the settle-in time of each passenger in boarding group i is bounded from above by $(1/i) \max_{r\in\mathcal{R}} \sum_{p\in\mathcal{P}(r)} t_p^s$. Hence, using that $t_{p,r}^m = c^m \in \mathbb{Q}_{\geq 0}$ for all $(p,r) \in \mathcal{P} \times \mathcal{R}$, we obtain

$$U \leq k(R-1)c^m + H_k \max_{r\in\mathcal{R}} \sum_{p\in\mathcal{P}(r)} t_p^s.$$

Furthermore, the denominator of (24) can be bounded from below by

$$L \geq \max\left\{k(R-1)c^m + \min_{r\in\mathcal{R}} \sum_{p\in\mathcal{P}(r)} t_p^s,\ (R-1)c^m + \max_{r\in\mathcal{R}} \sum_{p\in\mathcal{P}(r)} t_p^s\right\}.$$

Consequently, we have

$$\begin{aligned}
\frac{U}{L} &\leq \frac{k(R-1)c^m + H_k \max_{r\in\mathcal{R}} \sum_{p\in\mathcal{P}(r)} t_p^s \Big(+ H_k(R-1)c^m - H_k(R-1)c^m\Big)}{\max\left\{k(R-1)c^m + \min_{r\in\mathcal{R}} \sum_{p\in\mathcal{P}(r)} t_p^s,\ (R-1)c^m + \max_{r\in\mathcal{R}} \sum_{p\in\mathcal{P}(r)} t_p^s\right\}} \\
&\leq \frac{H_k(R-1)c^m + H_k \max_{r\in\mathcal{R}} \sum_{p\in\mathcal{P}(r)} t_p^s}{(R-1)c^m + \max_{r\in\mathcal{R}} \sum_{p\in\mathcal{P}(r)} t_p^s} + \frac{(k-H_k)(R-1)c^m}{k(R-1)c^m + \min_{r\in\mathcal{R}} \sum_{p\in\mathcal{P}(r)} t_p^s} \\
&\leq H_k + \frac{(k-H_k)(R-1)c^m}{k(R-1)c^m} = H_k + 1 - \frac{H_k}{k},
\end{aligned}$$

which establishes the claimed approximation ratio $1 + \frac{k-1}{k} H_k$. $\square$

L Proof of Proposition 4

From the proof of Theorem 4, we know that

$$k \max_{p\in\mathcal{P}} \{(r(p)-1)c^m + t_p^s\} \leq k(R-1)c^m + k \max_{p\in\mathcal{P}} t_p^s \leq k(R-1)c^m + kT \tag{25}$$

is an upper bound on the outside-in BTF boarding time. Since Lemma 2 also applies to max-settle-row, the bound (25) is actually valid for both boarding strategies. The same can be said about the lower bound from the proof of Theorem 6,

$$\max\left\{k(R-1)c^m + \min_{r\in\mathcal{R}} \sum_{p\in\mathcal{P}(r)} t_p^s,\ (R-1)c^m + \max_{r\in\mathcal{R}} \sum_{p\in\mathcal{P}(r)} t_p^s\right\} \geq k(R-1)c^m. \tag{26}$$

Thus, for both outside-in BTF and max-settle-row boarding, the approximation ratio is upper-bounded by

$$\frac{k(R-1)c^m + kT}{k(R-1)c^m} = 1 + \frac{T}{(R-1)c^m} \quad \xrightarrow{R\to\infty} 1,$$

which completes the proof. $\square$

Bibliography

J. Audenaert, K. Verbeeck, and G. Vanden Berghe. Multi-agent based simulation for boarding. In *The 21st Benelux Conference on Artificial Intelligence (BNAIC 2009)*, 2009.

E. Bachmat. Airplane boarding meets express line queues. *European Journal of Operational Research*, 275 (3):1165–1177, 2013.

E. Bachmat and M. Elkin. Bounds on the performance of back-to-front airplane boarding policies. *Operations Research Letters*, 36(5):597–601, 2008.

E. Bachmat, D. Berend, L. Sapir, S. Skiena, and N. Stolyarov. Analysis of airplane boarding times. *Operations Research*, 57(2):499–513, 2009.

E. Bachmat, V. Khachaturov, and R. Kuperman. Optimal back-to-front airplane boarding. *Physics Review E*, 87:Art. No. 062805, 2013.

M. Bazargan. A linear programming approach for aircraft boarding strategy. *European Journal of Operational Research*, 183(1):394–411, 2007.

A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization.* Princeton Series in Applied Mathematics. Princeton University Press, 2009.

P. Brucker. *Scheduling Algorithms.* Springer, Berlin Heidelberg, 5th edition, 2007.

G. C. Budesca, A. A. Juan, and P. Fonseca i Casas. Optimization of aircraft boarding processes considering passengers' grouping characteristics. In A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, editors, *Proc. Winter Simulation Conference*, pages 1977–1988. IEEE, 2014.

J. Coppens, S. Dangal, M. Vendel, S. Anjani, S. Akkerman, S. Hiemstra-van Mastrigt, and P. Vink. Improving airplane boarding time: a review, a field study and an experiment with a new way of hand luggage stowing. *International Journal of Aviation, Aeronautics, and Aerospace*, 5(2):Art. No. 7, 2018.

D. Coudert. A note on Integer Linear Programming formulations for linear ordering problems on graphs. Research report hal-01271838, Inria; I3S; Université Nice Sophia Antipolis; CNRS, 2016.

C. Delcea, L.-A. Cotfas, M. Salari, and R. J. Milne. Investigating the random seat boarding method without seat assignments with common boarding practices using an agent-based modeling. *Sustainability*, 10(12): Art. No. 4623, 2018.

S. Erland, J. Kaupužs, V. Frette, R. Pugatch, and E. Bachmat. Lorentzian-geometry-based analysis of airplane boarding policies highlights "slow passengers first" as better. *Physical Review E*, 100:Art. No. 062313, 2019.

P. Ferrari and K. Nagel. Robustness of Efficient Passenger Boarding Strategies for Airplanes. *Transportation Research Record*, 1915(1):44–54, 2005.

V. Frette and P. C. Hemmer. Time needed to board an airplane: A power law and the structure behind it. *Physical Review E*, 85:Art. No. 011130, 2012.

M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W. H. Freeman & Co., New York, NY, USA, 1979.

Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2019. URL http://www.gurobi.com.

N. G. Hall and C. Sriskandarajah. A Survey of Machine Scheduling Problems with Blocking and No-Wait in Process. *Operations Research*, 44(3):510–525, 1996.

T. Horstmeier and F. de Haan. Influence of ground handling on turn round time of new large aircraft. *Aircraft Engineering and Aerospace Technology*, 73(3):266–271, 2001.

L. Hutter, F. Jaehn, and S. Neumann. Influencing factors on airplane boarding times. *Omega*, 87:177–190, 2019.

F. Jaehn and S. Neumann. Airplane boarding. *European Journal of Operational Research*, 244(2):339–359, 2015.

S. Jafer and W. Mi. Comparative Study of Aircraft Boarding Strategies Using Cellular Discrete Event Simulation. *Aerospace*, 4(4):Art. No. 57, 2017.

T. A. Jessin, S. Madankumar, and C. Rajendran. Permutation flowshop scheduling to obtain the optimal solution/a lower bound with the makespan objective. *Sādhanā*, 45(1):Art. No. 228, 2020.

J. Kaupužs, R. Mahnke, E. Bachmat, and V. Frette. How long does it take to board an airplane? In S. H. Hamdar, editor, *Traffic and Granular Control '17*, pages 415–421. Springer, 2019.
A. Kierzkowski and T. Kisiel. The Human Factor in the Passenger Boarding Process at the Airport. *Procedia Engineering*, 187:348–355, 2017.
S. Marelli, G. Mattocks, and R. Merry. The role of computer simulation in reducing airplane turn time. *Aero Magazine*, 1(1), 1998.
R. J. Milne and A. R. Kelly. A new method for boarding passengers onto an airplane. *Journal of Air Transport Management*, 34:93–100, 2014.
R. J. Milne and M. Salari. Optimization of assigning passengers to seats on airplanes based on their carry-on luggage. *Journal of Air Transport Management*, 54:104–110, 2016.
R. J. Milne, M. Salari, and L. Kattan. Robust Optimization of Airplane Passenger Seating Assignments. *Aerospace*, 5(3):Art. No. 80, 2018.
R. J. Milne, C. Delcea, and L.-A. Cotfas. Airplane boarding methods that reduce risk from COVID-19. *Safety Science*, 134:Art. No. 105061, 2021.
V. Nagarajan and M. Sviridenko. Tight Bounds for Permutation Flow Shop Scheduling. *Mathematics of Operations Research*, 34(2):417–427, 2009.
S. Neumann. Is the boarding process on the critical path of the airplane turn-around? *European Journal of Operational Research*, 277(1):128–137, 2019.
D. C. Nyquist and K. L. McFadden. A study of the airline boarding problem. *Journal of Air Transport Management*, 14(4):197–204, 2008.
M. Ozmec-Ban, R. Š. Babić, and A. Modić. Airplane Boarding Strategies for Reducing Turnaround Time. In *Proc. 18th International Conference on Transport Science (ICTS)*, 2018.
C. H. Papadimitriou and P. C. Kanellakis. Flowshop Scheduling with Limited Temporary Storage. *Journal of the ACM*, 27(3):533–549, 1980.
M. L. Pinedo. *Scheduling: Theory, Algorithms, and Systems*. Springer, 5th edition, 2016.
M. Schultz. Field Trial Measurements to Validate a Stochastic Aircraft Boarding Model. *Aerospace*, 5(1): Art. No. 27, 2018a.
M. Schultz. Implementation and application of a stochastic aircraft boarding model. *Transportation Research Part C: Emerging Technologies*, 90:334–349, 2018b.
M. Schultz and J. Fuchte. Evaluation of aircraft boarding scenarios considering reduced transmissions risks. *Sustainability*, 12(13):Art. No. 5329, 2020.
M. Schultz and S. Reitmann. Machine learning approach to predict aircraft boarding. *Transportation Research Part C: Emerging Technologies*, 98:391–408, 2019.
M. Schultz and M. Schmidt. Advancements in Passenger Processes at Airports from Aircraft Perspective. *Sustainability*, 10(11):Art. No. 3877, 2018.
M. Schultz, C. Schulz, and H. Fricke. Efficiency of aircraft boarding procedures. In *Third International Conference on Research in Air Transportation*, 2008.
M. Soolaki, I. Mahdavi, N. Mahdavi-Amiri, R. Hassanzadeh, and A. Aghajani. A new linear programming approach and genetic algorithm for solving airline boarding problem. *Applied Mathematical Modeling*, 36 (9):4060–4072, 2012.
J. H. Steffen. Optimal boarding method for airline passengers. *Journal of Air Transport Management*, 14 (3):146–150, 2008.
A. Steiner and M. Philipp. Speeding up the airplane boarding process by using pre-boarding areas. In *9th Swiss Transport Research Conference*, 2009.
M. H. L. van den Briel, J. R. Villalobos, G. L. Hogg, T. Lindemann, and A. V. Mulé. America West Airlines Develops Efficient Boarding Strategies. *Interfaces*, 35(3):191–201, 2005.
H. van Landeghem and A. Beuselinck. Reducing passenger boarding time in airplanes: a simulation based approach. *European Journal of Operational Research*, 142(2):294–308, 2002.
F. J. L. Willamowski and A. M. Tillmann. Minimizing Airplane Boarding Time. repORt 2019–56, Chair of Operations Research, RWTH Aachen University, 2019.
J. M. Wilson. Alternative Formulations of a Flow-Shop Scheduling Problem. *The Journal of the Operational Research Society*, 40(4):395–399, 1989.
H. Zeineddine. A dynamically optimized aircraft boarding strategy. *Journal of Air Transport Management*, 58:144–151, 2017.