Proceedings of the 2021 Winter Simulation Conference
S. Kim, B. Feng, K. Smith, S. Masoud, Z. Zheng, C. Szabo, and M. Loper, eds.

ENHANCED OPERATIONAL MANAGEMENT OF AIRPORT GROUND SUPPORT EQUIPMENT FOR BETTER AIRCRAFT TURNAROUND PERFORMANCE

Siddhanta Saggar
Maurizio Tomasella

Business School
University of Edinburgh
29, Buccleuch Place
Edinburgh, EH8 9JS, UNITED KINGDOM

Giovanni Cattaneo
Andrea Matta

Department of Mechanical Engineering
Politecnico di Milano
Via G. la Masa, 1
Milano, 20156 , ITALY

ABSTRACT

Within the context of airport operations, this work focuses on enhancing the planning and real-time allocation of certain resources that are used to turn around an aircraft between two consecutive flights. This sort of operations takes place in the area of an airport that is called its apron. At peak times in particular, and when resource capacity is really tight, apron operations tend to be affected by either unavailability or late arrival of certain assets at the stand. The key element of this paper is the proposal of a new resource booking system, which operates in real-time, and deals with the related uncertainties. The booking mechanism aims to allow the airlines to book in advance certain resources, in particular ground support equipment. Our choice of a real case study will help us to assess the likely benefits, pros and cons of this system.

1 MOTIVATION AND BACKGROUND LITERATURE

The aircraft turnaround process consists of a collection of so called 'ground handling operations' that are required to set an aircraft ready for its subsequent departure (Ashford et al. 2013). Employed in such a process are a number of resources, which include both human operators and physical/engineering assets. Among the latter, which altogether go by the name of Ground Support Equipment (GSE), relevant examples are shown in Figure 1. For instance, aircraft Ground Power Units (GPUs) allow operators to provide DC/AC power to the aircraft while on the ground. Ambulifts are vehicles needed for the transport, boarding and disembarking of passengers with reduced mobility. Finally, push-back tractors are vehicles equipped with a heavy-load bar, named tow bar, that latches to the aircraft landing gear and pushes it back into the taxiway, on its way to the runway. To avoid delays, each of the turnaround activities requires some specific GSEs as well as the appropriate operators, all at the right time and the right place, to avoid delays.

The two flows of GSEs and human operators are closely intertwined. In practice, the operational management of turnaround teams and GSEs is pursued separately by the ground handling service organisations, going through separate processes and relying on separate information systems. A decade of conversations between us and many of the global players from the ground handling market has suggested that an integrated approach to managing these two types of resources is still far from realistic. The present paper focuses on the operational management of GSEs.

For the past two to three decades, apron operations and the related resource planning and management problems have often been studied by the Operational Research/Management Science community. Most works do target tactical level planning and scheduling problems, especially looking at turnaround teams/employees. Some of the most recent contributions in this direction include Tomasella et al. (2019), Gök et al. (2020) and Hur et al. (2019). Other approaches do exist though.

978-1-6654-3311-2/21/$31.00 ©2021 IEEE

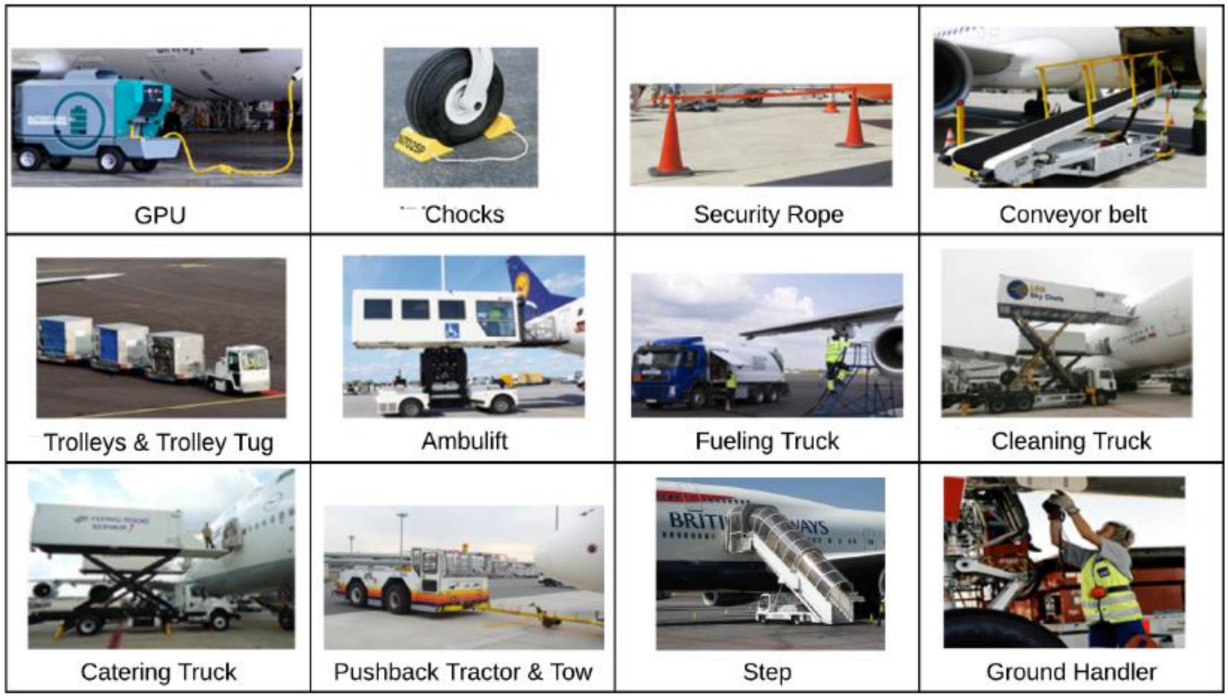

Figure 1: Example GSEs involved in aircraft turnaround operations.

Ter Mors et al. (2007) studied the scheduling of scarce de-icing resources by means of a decentralized approach. The problem was solved through a simple First Come First Served (FCFS) heuristic, with the objective to minimize flight delays. The heuristic algorithm is then improved in a robust way to enable re-planning after incidents/disruptions.

Similarly, Mao et al. (2008) adopted a multi agent based approach, again from a decentralized viewpoint. The authors focused on the lack of a common shared objective among the different stakeholders, and assumed two scenarios. They investigated two types of stakeholder behaviour: cooperative and non cooperative, considering airlines and ground handlers as agents. Results prove that when agents are cooperating, better plans and overall operational performance are achieved. An adaptive version of the proposed heuristic was proposed also including stochastic aircraft arrivals and variable operational times needed by ground handlers to complete turnaround operations. In this way, the scheduling is re-generated after some predefined events.

Andreatta et al. (2014) formulate their own GSE and turnaround team scheduling problem through integer programming. Their solution method was then based on a sequential heuristic, with the objective to improve robustness and keep high operational efficiency by minimizing a cost flow network. Similarly, Andreatta et al. (2014) proposed an equivalent heuristic algorithm for the GSE Allocation Problem (GAP), solved in a robust way also taking into account the statistical distribution of delays, aiming to improve the trade-off efficiency/robustness. However, the main assumption of this method is based on the traceability of ground handling equipment across the airport apron in real-time, which might not be always true in real-cases. Works such as Ansola et al. (2013) study the benefits from tracking and tracing of airport assets through a number of technologies.

The rest of this paper is organised as follows. In Section 2, we discuss in details two very different scenarios of GSE management, for subsequent analysis. One represents the current status quo or most common situation, whereby GSE handling within the context of aircraft turnaround operations tend to happen mostly reactively and in real-time, with ground handling service providers substantially working independently from one another, with little to no coordination at all taking place. The second scenario aims at introducing some form of regularity in the way GSEs are handled, by adding a planning that translates

into the possibility for airlines to book in advance certain GSEs they retain to be key to more punctual turnarounds. For this booking system to take place, a certain degree of information sharing needs to be in place among the airport operator, the airlines and the service providers. This whole concept and the related analysis of possible benefits and risks constitutes a novelty in the literature on apron operations and the operational management of GSEs. Within the extant literature, Padrón et al. (2016) suggested a few years ago that collaboration among the many stakeholders involved in multiple turnarounds at the same airport should be enhanced, for the effectiveness of apron operations to benefit positively from it, although this was not elaborated upon any further. In Section 3, we provide selected details on our simulation modelling activity. In Section 4, we give an overview of the main features of the real airport that we chose as a case study. Section 5 will then analyse our results in details, while Section 6 will draw the main conclusions on the idea for a possible GSE booking system, and will set out possible next steps for this investigation.

2 SCENARIOS AND RELATED RESEARCH QUESTIONS

2.1 Baseline Case

In this section, the two main scenarios we investigated in this work will be presented and discussed in details. The first scenario is quite common in many airport settings, and simple to explain. Service providers either fully control the operations of a number of assets or share assets with one or more other service providers. Service providers either own GSEs outright, or they lease them from a third party. In the latter case, a third party owns a fleet of GSEs, which are made available for several service providers to use. Whenever a turnaround team/employee has to carry out a specific turnaround task for a given aircraft, it searches for a suitable GSE, in a reactive way. This may require quite some time in moving around the apron area to secure the needed asset. After use, the GSE is parked in a dedicated asset store area on the airside of the airport, until the next time someone else needs the same asset. The usage cycle 'fetch— use— store' for each piece of equipment/GSE is very much of an emerging nature, driven by the events of the specific day of operation. Occasionally, the time to fetch many GSEs may be so long that they reach the parking stand later than ideal, thus delaying directly the operation in which the GSE is employed, and possibly impacting other operations downstream. This reactive system is as simple to explain as it can be chaotic as events unfold on the day of operations.

2.2 GSE Booking

The main idea behind this scenario is to add a level of planning to the otherwise fully reactive real-time allocation of GSE to turnaround tasks just described in the baseline scenario. This level consists in letting the airlines choose, if they wish and in exchange of a service fee, whether they want to ensure that the GSEs needed to turn around one or more of their flights (either a subset of assets, or all of them) be present at the parking stand once the aircraft shows up sometime after landing. This is imagined to be happening via a GSE booking system.

Once a particular asset is booked by an airline through this system for a particular flight, it cannot be used, around that time, by any other service provider who may be able to access it. This may of course in turn generate unwanted delays. For instance, when the flight for which the asset was booked is delayed on arrival, and other aircraft, already on the ground, are in need of the same type of asset whilst GSE capacity is constrained and the apron busy, delays may easily propagate downstream. In the latter portion of our experiments (Section 5), we will investigate a slight variation of this basic booking system, in order to counter the possibility of such a delay propagation.

To begin with, our core research objective was to investigate whether the idea itself of adding the pre-booking facility results in improved turnaround effectiveness, i.e. lower turnaround delays in a way that may foster more punctual departures and, as a side effect, help to minimise the penalty costs incurred by the airlines and service providers as a result of undesired delays. Side questions for this research include whether any positive effects from the booking facility require the system to be applied at a certain scale,

i.e. to a large number of the flights on the schedule and/or a large portion of the GSEs operating at the given airport.

Any benefits resulting from turnaround asset booking are likely to depend on many factors. Mostly, different operational settings will be more or less prone to experience such benefits. The extent of the benefits in each setting are likely to depend on the type, configuration, size and layout of the airport and its apron area, as well as the specific daily schedule, the number of airlines and service providers operating concurrently at the airport, and the arrangements in place around GSE operations and management (e.g. whether or not GSE are shared among multiple service providers). To investigate the above questions and compare the two main scenarios we then selected a real airport, representative of relatively good passenger traffic levels (pre-Covid-19 pandemic) for an international origin and destination airport facility. More details on the chosen airport setting will be provided in Section 4.

We now anticipate a few other elements of how we see the basic version of the booking system working. Each ground handling service provider can open-up the resource-booking system for the airlines they service to possibly join, as the schedule is largely known in advance of the day of operation. The service providers may share a chart-sheet of all the resource sets in their control with all the airlines they have a contract with, along with the capacity and possibly the availability of each resource type for each day and each hour slot, theoretically across the entire season to which the overall schedule applies. The associated fees attached to the booking service will clearly have to be shared as well, the extent of which may depend on many factors, including the type of GSE itself. Those airlines joining the scheme will book assets ahead of the day of operation. The closer this is done to the actual day of operation, the more reliable the information around asset availability, e.g. as some of the assets may undergo planned maintenance from time to time and the service provider may have to replace a booked asset with a somewhat equivalent one, in case the asset originally booked became unavailable for the specific day attached to the booking. With each confirmed booking, the chart-sheet gets updated with the new availability status of the booked resource type for the slot it was booked and finally gets re-released to the airlines thereby adhering to the "real-time" aspect of the system. With each asset type having a related unique identification number/code assigned to it, the airlines may get to choose and reserve a specific set of resources for the whole of a turnaround process, or a subset thereof.

The latest time by which a booking can be confirmed for a particular flight by an airline is the *notification time* of the flight, that is the time at which the flight confirms its arrival to the airport— say about 30 minutes before the actual time of arrival, depending on arrival routes and related airborne operations on approach to landing. While reading the information of an incoming flight during the notification process, the associated ground handling service provider confirms the availability of a stand to have the incoming aircraft parked, followed by checking if the concerned aircraft has a booking made for the resources or not. If resources are indeed booked, the concerned service organisation begins searching for the booked resource-ID, while preparing the stand for execution of the turnaround process. Where no booking is found for one or more GSEs, then the service provider proceeds with the default approach described in the baseline scenario, that is a more extensive asset search process takes place around the apron area.

3 MODELING

From a modelling perspective, we adopted a hybrid simulation approach (Brailsford et al. 2019), including elements of both agent-based and discrete event simulation. We implemented the resulting simulation model in Anylogic. The core of the model is agent-based, with agents representing aircraft, parking stands, GSEs, their storage locations, and ground handlers, to name but the most prominent in determining the related emerging dynamics. State Charts, Action Charts and Process Flowcharts all concur in defining agent behaviour and interactions, as well as relevant details of all turnaround tasks to be executed, the corresponding defining events, and the flow of more passive entities such as baggage and cargo.

For instance, when information regarding an incoming aircraft comes in, the assignment of a suitable parking stand is the first step to be taken (Figure 2). The actual stand to be allocated may be the 'preferred'

stand, possibly the one most often allocated, or one which is conveniently located within the terminal, if available. Otherwise, then, another stand that is available gets assigned to this flight, following the logics and procedures for stand and gate assignment that are adopted locally at the given airport.

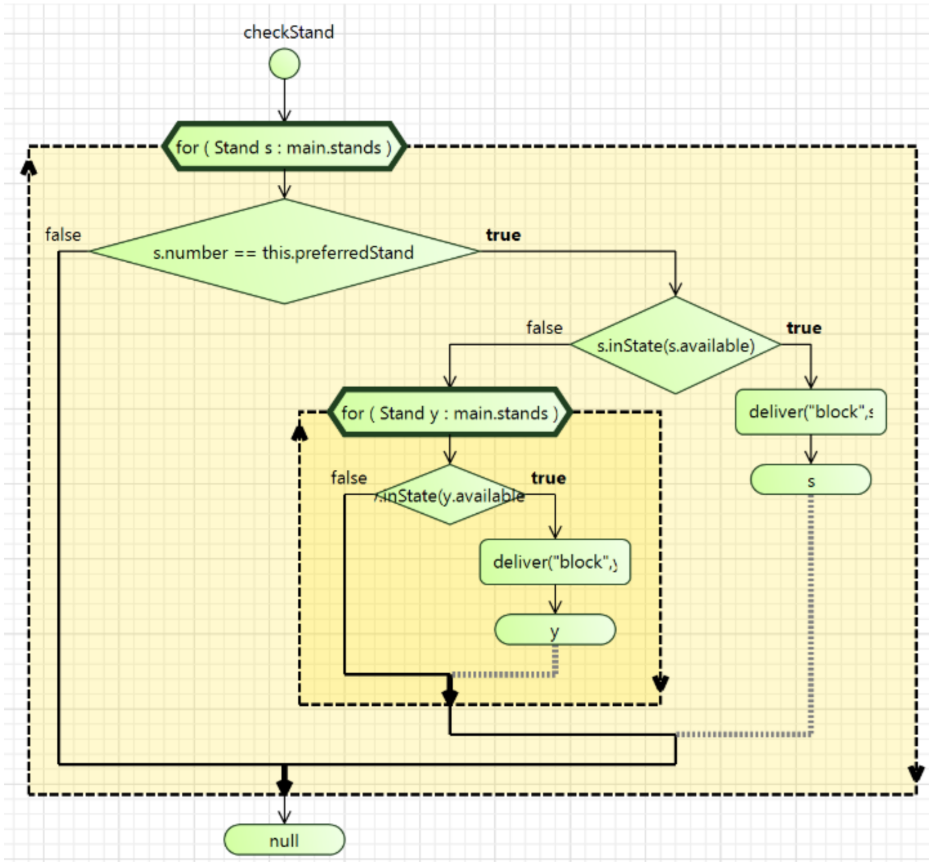

Figure 2: Action Chart: Check Stand.

Figure 3 presents a simple example of a process flowchart showing how a subset of turnaround equipment is setup. This includes the chocks to keep the aircraft from moving while parked, the ground power unit and the safety ropes.

4 CASE STUDY

4.1 Selected Airport

As a case study for our experiments we chose Luton Airport, fifth busiest commercial airport in the UK in 2018, based on passengers figures. In terms of these figures, Luton Airport is fairly representative (pre-Covid-19) of many international origin and destination airports. Figure 4 presents a condensed overview of the main data describing the configuration of the case study airport.

Three international players from the ground handling arena operated at Luton Airport in 2018, which we name ‘A’, ‘B’ and ‘C’. Figure 4 shows, for each of these service providers, how many airlines and flights they service, how many ground handling teams they have on their daily roster, how many individuals make up each team, and how many assets of each type they control directly, for the core turnaround tasks for which they are subcontracted by the 14 airlines operating at the airport.

Additional specialised service organisations provide the remaining services, to all airlines. These services included, at the time: fueling, catering, cleaning, mobility to people who required assistance,

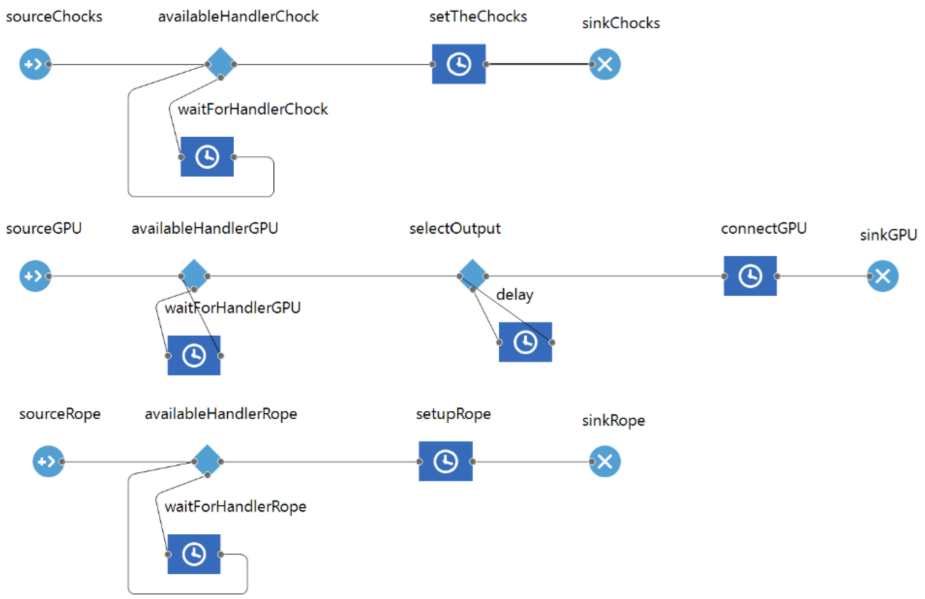

Figure 3: Discrete event process model: setup of chocks, GPUs and safety ropes.

and transfer of passengers between remote stands and the terminal and vice versa. Team allocation and management for these service providers was left out of scope from our study. The capacities in terms of GSEs employed to deliver each such service are also shown in Figure 4.

We collected both the flight schedules for the data days of our experiments and the information on the airport's layout from a host of publicly available resources accessible via the internet. A 'normal traffic' day would involve about 62 flights, as shown in Figure 4. 'Low traffic' and 'High traffic' days would see operations of about 50 and 74 flights per day respectively.

Short-haul, low-cost flights dominate these schedules, with a relatively homogeneous aircraft mix (B737s, A319/320/321, and similar).

4.2 Performance Metrics

In order to compare the relative effectiveness of the two scenarios described above, the following four performance metrics were considered: A) Turnaround Delays; B) Number of Turnaround Delays; C) Asset Fetching Time; and D) Asset Storing Time (all times in minutes).

Turnaround delays are positive differences between the actual and the scheduled departure time of an aircraft from its stand. The fetching time for a particular asset is defined as the time taken by a handler (an operator) to fetch the GSE and bring it to the stand. Asset fetching repeats for all assets, until the stand is fully set for execution of the aircraft turnaround. Finally, the storing time for an asset is the time taken for a handler to start from the assigned stand with the GSE to be stored, find the store which can accommodate it, and store the GSE there. The same happens for all assets.

5 NUMERICAL RESULTS

5.1 Overall results

To start with, Table 1 provides an overview of results, in the form of point estimates for the means of four different performance metrics, from running 650 independent replications of each of our two main scenarios, namely the 'Status Quo' (SQ) and the 'Asset Booking' (AB) settings. In the case of the latter, about 50% was the proportion of flights, randomly selected from the daily schedule, allowed to book their

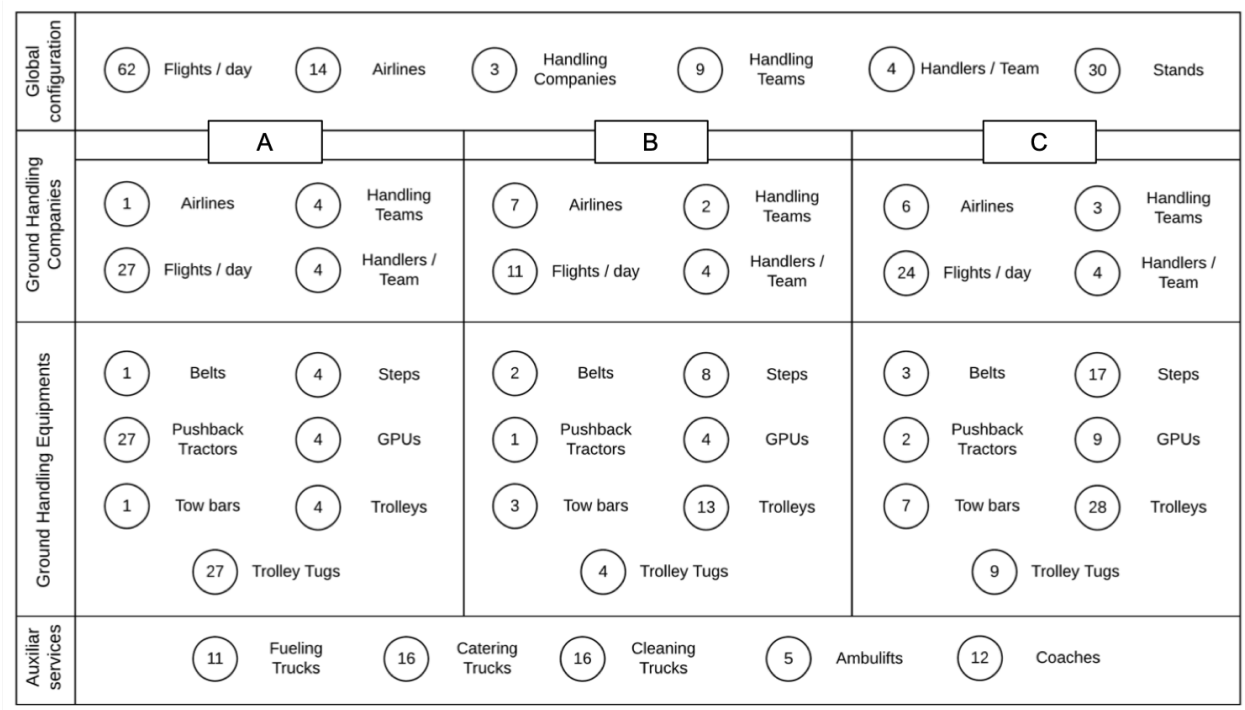

Figure 4: Case study overview.

assets. Scenario AB appears to enhance operations, with respect to the SQ, in terms of three out of the four metrics, with the exception of the number of turnaround delays. The directions of these changes in performance remain the same for each of the three prototype data days investigated.

Table 1: Results overview.

Metrics	SQ (Status Quo)			AB (Asset Booking)		
	Low Day	Normal Day	Busy Day	Low Day	Normal Day	Busy Day
Turnaround Delay Time (mins)	24.73	29.43	32.23	19.7	23.8	25.4
No. of Turnaround Delays	34	51	55	39	53	61
Asset Storing Time (mins)	2.45	2.82	2.89	2.13	2.37	2.45
Asset Fetching Time (mins)	1.7	2.1	3.1	1.02	1.1	1.5

Confidence intervals (95%) on the differences between the means from the same two scenarios confirm the statistical significance of these differences in the cases of both turnaround delays (Figure 5) and asset fetching times (not shown, for reasons of space). We also found the differences in terms of asset storing times to be not very significant, which seems reasonable, given that the dynamics of asset storing are virtually the same in the two scenarios.

Similar confidence intervals built for the differences in the number of turnaround delays (Figure 6) show that these differences are also, although present, not statistically significant.

To understand better the reasons behind these preliminary results, we now reorganise our output data for a more fine-grained discussion, namely two-hourly, flight by flight and asset by asset.

5.2 Time-based Analysis

To investigate any patterns in the differences in mean turnaround delays incurred throughout the working day (7am-11pm), a time-based analysis has been performed over slots of 2 hours each. The overall pattern is similar across the three data days investigated (Low, Normal, Busy), though the magnitude of the differences

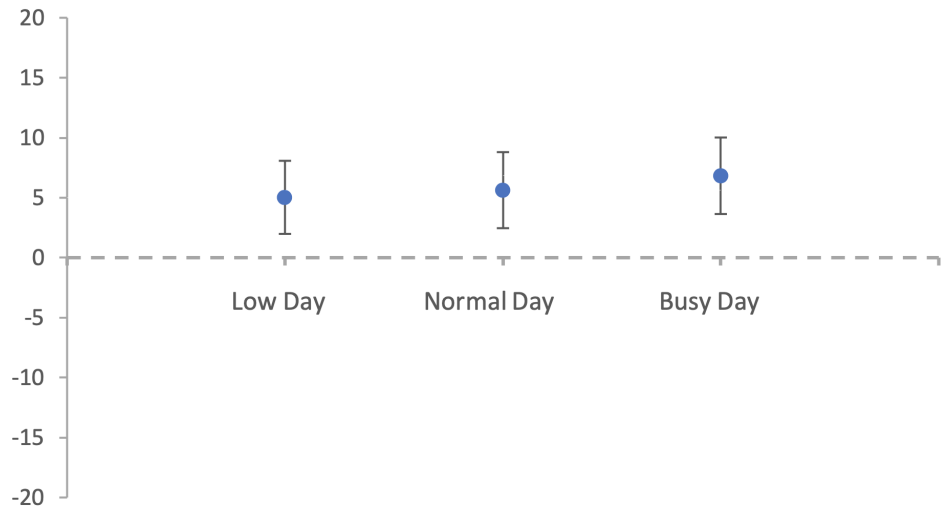

Figure 5: Differences in Mean Turnaround Delays (SQ-AB, in minutes).

change. The resulting 95% confidence intervals (Figure 7) show significant differences in mean turnaround delays for the majority of the slots, with more prominent differences around the daily peaks, particularly the first wave of flights which happens in the morning.

5.3 Impact on Flights with Booking Arrangements

Figure 8 presents, for each of the flights with asset bookings over the Normal Day, the difference between its mean turnaround delay when not booking assets and that when the flight does book its GSE. For the majority of flights, more precisely 27 out of 32, the booking facility indeed helps to keep delays under control. Similar situations happen for both the Low Day (with 20 flights out of 26 with booking improving their delays), and the Busy Day (with 29 flights out of 37). The magnitude of many of the differences displayed in the figure is quite important, including around peak times.

5.4 Asset-by-Asset Analysis

All operational delays and related differences between the two scenarios as discussed so far have been computed either on a flight by flight basis, or over a number of flights covering part or all of the day of operation. When thinking about booking assets for a particular flight, an airline may consider focusing the booking on specific assets, on the basis of some priority ranking. Figure 9 shows the differences in mean operational delays between the SQ and the AB scenario, with regards to GPUs. Clearly, booking GPUs (especially for particular flights, as anticipated above) is likely to considerably improve operational delays for this asset. Differences tend to be both statistically and practically significant in this case. Other GSEs/assets fare differently; for instance, in the case of conveyor belt devices for baggage loading/unloading the differences are not significant at all, in any of the three data days. Factors likely to impact this sort of results for different GSEs/assets include asset capacity, how busy the daily schedule is, the place of the operation in which the asset is used within the overall gantt chart that is the flight's turnaround, and more.

5.5 Enhanced Booking System

Finally, we wish to explore the likely effects from slight variations on two of the factors of interest that have so far been kept constant. The first of these is the proportion of flights for which booking of GSEs is made possible over the day of operation. Instead of the 50% level that has been considered so far, we

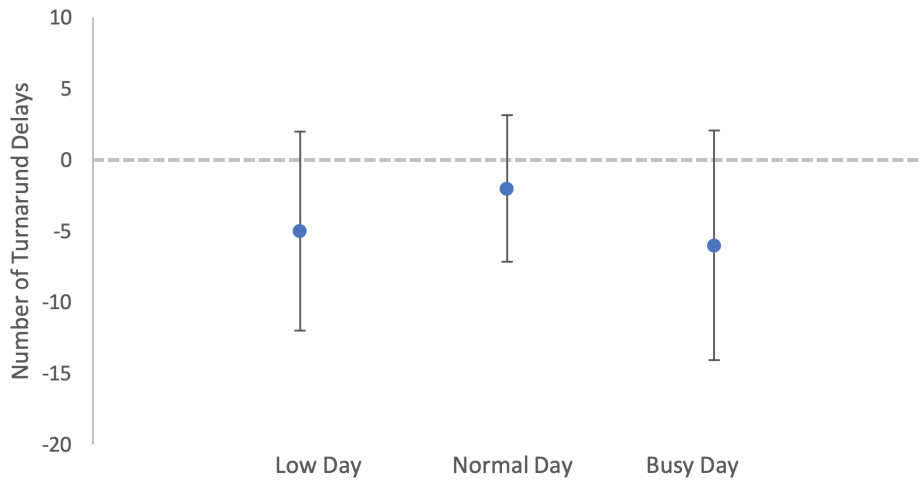

Figure 6: Differences in Number of Turnaround Delays (SQ-AB).

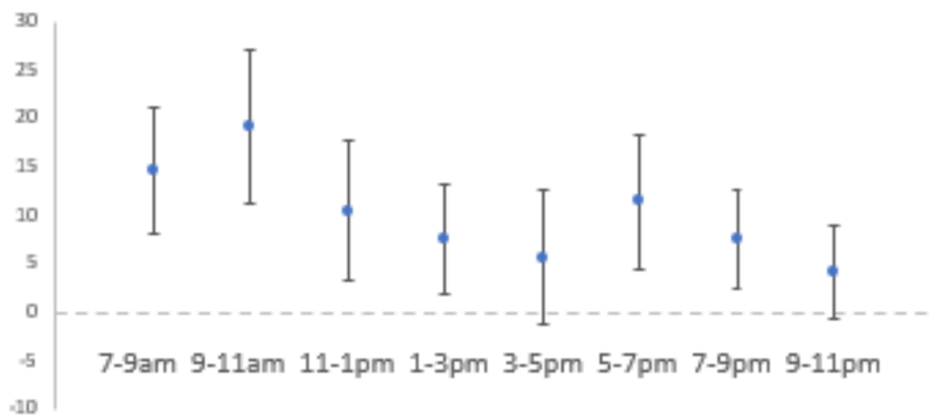

Figure 7: Differences in Mean Turnaround Delays, by Two-Hour Interval (SQ-AB, in minutes, Busy Day).

have investigated what happens if 33% or 66% of flights can be targeted for bookings. We called this first factor 'Book_Prop'. The second factor is the type of booking per se ('Book_Type'), as the one considered so far is very basic indeed. We have named the enhanced version 'Collab', to underline the additional form of collaboration that takes place between flights/airlines. Within the 'Collab' scheme, a partial level of information sharing between the airport stakeholders comes into play, which in scenarios having a flight with booked assets running delayed makes the assets available to other needful flights at the airport, while also ensuring that the assets are made available back to whom had booked them, when the originally booked delayed flight arrives. The baseline version that has been considered so far was renamed *Srch*, to underline that with no redistribution of assets possible, the asset originally booked remains unavailable and other flights which need it at the same time will have to keep searching for alternative assets, as they become available around the apron. We ended up running a replicated 2^2 full factorial experiment.

Figure 10 shows that working on enhancing the type of booking mechanism is likely to pay off way better than extending the possibility of booking to a higher proportion of flights on the day of operation. In fact, the latter may even be slightly counterproductive. No significant interaction seems to be there, at least in the investigated case, between the two factors. Similar results hold for the mean turnaround delay metric, where the booking proportion factor appears even more insignificant, and the type of booking even

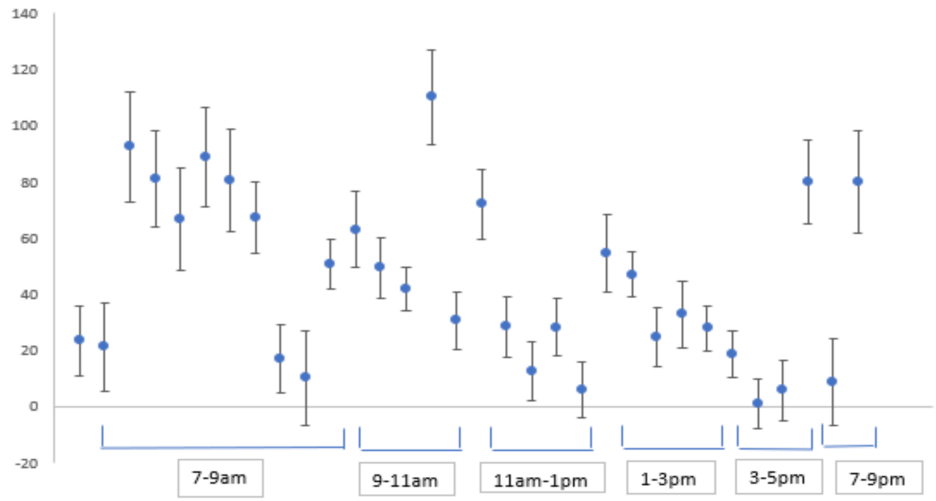

Figure 8: Differences in Mean Turnaround Delay–flights with bookings (SQ-AB, in minutes, Normal Day).

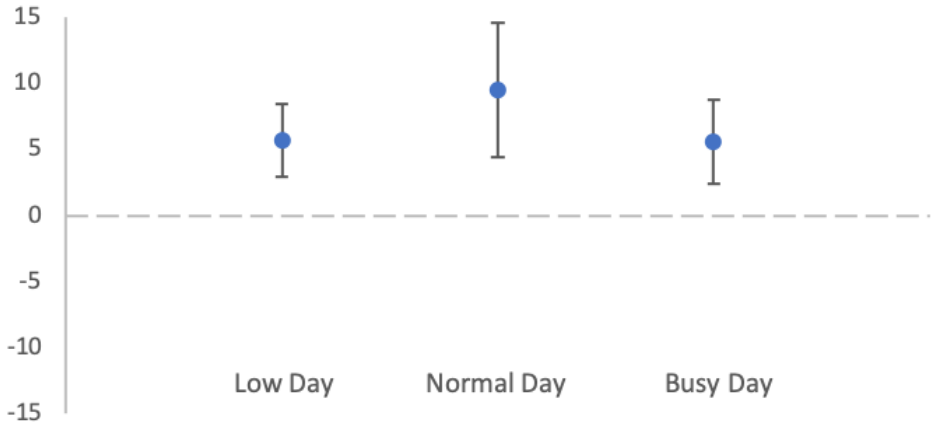

Figure 9: Differences in Mean Operational Delays–GPUs (SQ-AB, in minutes).

more significant. In terms of asset fetching time, both factors appear significant (again with no detectable interaction between the two). The results for the asset storing time confirm the initial findings, with no statistically significant main or interaction effect.

6 CONCLUSIONS

In these pages, we have investigated a novel idea for managing the allocation of GSEs to aircraft turnarounds in apron operations. The idea relates to allowing the airlines operating at the given airport to book GSEs in advance, in a way to secure at least some key assets that, if otherwise delayed on the way to the parking stand, are likely to contribute to generating costly delays to the turnaround as a whole.

In our experiments, we have contrasted such an asset booking scenario with the status quo, i.e. a situation in which GSEs are handled on a first-come-first-served basis, as they are needed over the day of operations. Our results show that, with respect to our case study, the asset booking concept has potential to enhance apron operations, by improving on some key metrics such as the average turnaround delay, as well as by ensuring a more timely delivery of the right asset to the parking stand. The extent to which the latter is true depends on many factors, primarily the type of GSE/asset. We also found that further

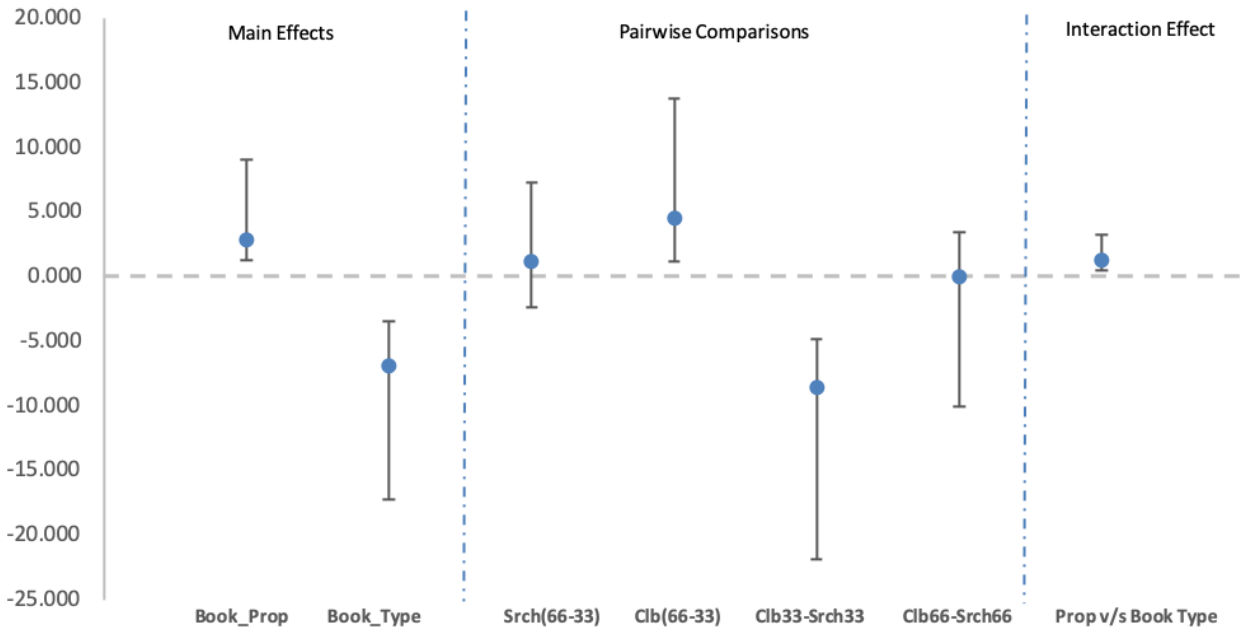

Figure 10: Main and Interaction Effects on Mean Number of Turnaround Delays (Number of Flights).

refinements of the asset booking scenario are likely to be more decisive, in fostering more punctual apron operations, than having a higher proportion of assets booked.

One of the more obvious possible next steps for this research stems from investigating the pros and cons from asset booking in different operational settings, primarily airports of different size and layout, with different mixes of airlines and aircraft model/makes turned around. A second extension regards the development of alternative forms of asset booking mechanisms, especially extending the exploitation, only attempted in this paper at a basic level, of data sharing mechanisms among the interested parties, that is the airport operator, the airlines and the ground handling and other service providers. As more evidence may be collected in favour of the asset booking concept, the investigation of feasible forms (contractual, and else) of implementation of the concept should also be pursued, including an assessment of the likely costs of this service to the airlines, penalties for the service providers that fail to deliver the right asset on time, and more.

ACKNOWLEDGMENTS

The authors wish to acknowledge the many aviation service providers, airport operators and airlines who in the last decade have worked with us and, as such, have directly informed the research investigated in this paper. We also wish to acknowledge the contribution from former University of Edinburgh Business School Master's students Alix Clare, Anton Ryjov and Alqamah Bhutta for their collaboration in earlier versions of our simulation model.

REFERENCES

Andreatta, G., L. Capanna, L. D. Giovanni, M. Monaci, and L. Righi. 2014. "Efficiency and Robustness in a Support Platform for Intelligent Airport Ground Handling". *Journal of Intelligent Transportation Systems* 18(1):121–130.

Andreatta, G., L. De Giovanni, and M. Monaci. 2014. "A fast heuristic for airport ground-service equipment–and–staff allocation". *Procedia-Social and Behavioral Sciences* 108:26–36.

Ansola, P. G., A. G. Higuera, F. J. Otamendi, and J. de las Morenas. 2013. "Agent-based distributed control for improving complex resource scheduling: Application to airport ground handling operations". *IEEE Systems Journal* 8(4):1145–1157.

Ashford, N., P. Coutu, and J. Beasley. 2013. *Airport Operations*. McGraw-Hill Education.

Brailsford, S. C., T. Eldabi, M. Kunc, N. Mustafee, and A. F. Osorio. 2019. "Hybrid simulation modelling in operational research: A state-of-the-art review". *European Journal of Operational Research* 278(3):721–737.

Gök, Y. S., M. Tomasella, D. Guimarans, and C. Ozturk. 2020. "A Simheuristic Approach for Robust Scheduling of Airport Turnaround Teams". In *Proceedings of the 2020 Winter Simulation Conference*, edited by K.-H. G. Bae, B. Feng, S. Kim, S. Lazarova-Molnar, Z. Zheng, t. Roeder, and R. Thiesing, 1336–1347. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Hur, Y., J. F. Bard, M. Frey, and F. Kiermaier. 2019. "A stochastic optimization approach to shift scheduling with breaks adjustments". *Computers & Operations Research* 107:127–139.

Mao, X., N. Roos, A. Salden, and B. Almende. 2008. "Distribute the selfish ambitions". In *Proceedings of BNAIC 2008 Belgian-Dutch Conference on Artificial Intelligence*, 137–144.

Padrón, S., D. Guimarans, J. J. Ramos, and S. Fitouri-Trabelsi. 2016. "A bi-objective approach for scheduling ground-handling vehicles in airports". *Computers & Operations Research* 71:34–53.

Ter Mors, A., X. Mao, N. Roos, C. Witteveen, A. Salden, B. van de Walle, P. Burghardt, and K. Nieuwenhuis. 2007. "Multi-agent system support for scheduling aircraft de-icing". In *Proceedings of ISCRAM 2007-Intelligent Human Computer Systems for Crisis Response and Management*, 467–478.

Tomasella, M., A. Clare, Y. S. Gök, D. Guimarans, and C. Ozturk. 2019. "STTAR: a simheuristics-enabled scheme for multi-stakeholder coordination of aircraft turnaround operations". In *Proceedings of the 2019 Winter Simulation Conference*, edited by N. Mustafee, K.-H. G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, 488–499. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHIES

SIDDHANTA SAGGAR is an MSc Business Analytics graduate from the University of Edinburgh and is currently part of the graduate program in the International Bureau of Aviation (IBA Group). Therein, he is part of the InsightsIQ intelligence team dealing in fleet analysis, carbon emission calculations, valuations and market trend information. His interests lie in analysing data associated with the aviation industry and drawing relevant insights from them. His email address is sidh47497@gmail.com.

MAURIZIO TOMASELLA is with the Management Science and Business Economics Group at the University of Edinburgh Business School. Previously at the University of Cambridge Engineering Department, his research interests include simulation, multi-criteria decision making, and their applications to airport operations as well as manufacturing. His e-mail address is maurizio.tomasella@ed.ac.uk.

GIOVANNI CATTANEO is a Graduate Trainee (Supply Chain) with Tenaris. At the time of this research he was a Master's student at Politecnico di Milano, and worked on his research dissertation while visiting the University of Edinburgh Business School. His email address is cattaneogiovannipegla@gmail.com.

ANDREA MATTA is Full Professor at Politecnico di Milano, where he currently teaches integrated manufacturing systems and manufacturing. His research area includes analysis, design and management of manufacturing and health care systems. He is Editor in Chief of the Flexible Services and Manufacturing Journal. His email address is andrea.matta@polimi.it.