

ON STATIC VS DYNAMIC (SWITCHING OF) OPERATIONAL POLICIES IN AIRCRAFT TURNAROUND TEAM ALLOCATION AND MANAGEMENT

Sudipta Saha
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

Silvia Padrón

TBS Business School
20 Boulevard Lascrosses
31068, Toulouse, FRANCE

ABSTRACT

Aircraft turnaround operations represent the fulcrum of airport operations. They include all services to be provided to an aircraft between two consecutive flights. These services are executed by human operators, often organised in teams, who employ some related equipment and vehicles (e.g. conveyor belts, trolleys and tugs for baggage loading/unloading and transportation). In this paper, we focus on the real-time management of turnaround operations, and assess the relative merits and limitations of so-called dispatching rules that originate from the manufacturing literature. More precisely, we focus on the real-time allocation, on the day of operation, of teams of ground handling operators to aircraft turnarounds. This is pursued from the viewpoint of third-party service providers. We employ simulation, in conjunction with deep reinforcement learning, and work on the case of a real airport and the entirety of its turnaround operations involving multiple service providers.

1 MOTIVATION AND RELATED LITERATURE

Aircraft turnaround operations represent services provided to aircraft, while parked in areas of an airport known as aprons, for the aircraft to be ready for the following take-off (Gök et al. 2020). Example services include passenger disembarking/boarding, baggage unloading/loading, refueling, cabin cleaning, catering, toilet and potable water servicing, etc. Many aircraft turnarounds end with the aircraft being pushed back into the nearest taxiway, so it can make its way to the designated runway for take off.

In aircraft turnarounds, many types of equipment and vehicles are employed, such as: conveyor belts for the loading/unloading of baggage, trolleys and tugs for the transfer of bags between the terminal and the aircraft, fuelling trucks, push-back tractors and tow bars, etc. These go by the name of ground support equipment (GSE). GSE are operated by human resources (ground handlers) employed by some service provider organisation (Tomasella et al. 2019). Handlers are often organised in teams of variable size, depending on the specific service, related GSEs, aircraft model/make, airline, etc.

Uncontrolled growth of air traffic in the past decades has put the global air transport network under unprecedented pressure, with airport capacity as well as operational efficiency being two central problems

far from being resolved (Guimarans et al. 2019). Airport apron operations are the points where the flows of passengers and baggage/cargo interchange between the land and the air transport modes, hence they are a particularly critical component in the air transport value chain. Evidence exists (Eurocontrol 2017) of the considerable delays experienced at airports, especially around the peak times that characterise most airports' schedules (Gelhausen et al. 2013). In 2017, 44.4% of European flights were more than 5 minutes late on departure and the average departure delay was estimated at 12.4 minutes per flight (Eurocontrol 2017). This brings with it an enormous impact in terms of knock-on effects across the global air transport network, with the associated costs (EU 2004).

The ground handling market, meanwhile, has evolved into an extremely fragmented one. Both legislators and regulators have played a primary role in this change process. Before the EU Directive 96/97/EC was introduced in 1996, only 7% of ground handling was undertaken by independent ground handling companies. In the years to follow, it has risen to a remarkable 45% (CAPA 2011). This fragmentation is likely to have impacted on the efficiency of apron operations, as more and more players operate in what is a relatively small area (the apron). At most airports, each ground handling service organisation provides certain services to certain airlines, and more often than not a single aircraft turnaround sees the involvement of multiple service providers.

Several macro-factors concur in determining the punctuality and efficiency of apron operations at any airport. First, air traffic control and management impact directly the punctuality of the aircraft itself, both on arrival and departure. Second, the actual duration of turnaround activities is often uncertain, as the demand for some services (e.g. refueling, baggage loading/unloading) may not be known with certainty until the service is executed. Hence, unforeseen delay in a given operation may impact punctuality of completion of other operations, such as the already mentioned aircraft push-back. Third, capacity, availability and operational management of both turnaround teams and GSE drives much of the dynamics in apron operations. Of these three macro-factors, service providers can exert (partial) control only to the latter.

In practice, the management of turnaround teams and that of GSE are most often handled separately. In fact, information on turnaround teams and GSE are traditionally stored in separate information systems, as well as organised and executed through separate processes and by separate functions of the service provider organisations (e.g. Human Resources for the teams and Maintenance for GSE).

The present paper focuses on the operational management of turnaround teams, currently a topic of growing interest to the scientific community— see Padrón et al. 2016, Tomasella et al. 2019, Hur et al. 2019 and Gök et al. 2020— within an even richer landscape of works dedicated to aircraft turnaround operations— see Schultz et al. (2020) and Evler et al. (2021). However, we do not know of any attempt at studying the management of turnaround teams at the operational level, with all major efforts so far having been spent on tactical-level decision making (the operational level is targeted by the GSE literature though, see for instance Mao et al. (2008) or Ansola et al. (2013)).

The present article is the first to study and compare various operational rules for the management of turnaround teams in real-time, on the day of operation. Two principal settings are studied (Section 2): (a) a static one, in which a given operational rule is applied throughout the entire day of operation; and (b) a dynamic setting, where over the day it is possible to switch from one rule to another, driven by the latest knowledge of the airport's operating status, in a way to recoup cumulative delays being experienced as well as, more in general, to recover from operational disruptions. Our approach blends simulation with reinforcement learning (Section 3). We adopt a real airport as a case study (Section 4), in order to quantify the likely relative pros and cons connected with both the static and dynamic settings (Section 5). We then conclude and present future developments of this line of research (Section 6).

2 DISPATCHING RULES

Whilst scholars may have not been looking much into rules for the real-time assignment of turnaround teams to the tasks to be executed on the day of operation, our talks with some of the global providers of airport ground services have confirmed that a good deal of on-the-spot assignments are made daily, as soon

as the pre-developed tactical plans succumb to the unpredictability of the reality that is airport operations. This often happens as early into the day as during the first wave of morning flights, which for many airports represents the absolute peak of the day in terms of air traffic movements.

These real-time assignments are based on various blends of experience and expertise on the side of the decision maker. In none of our conversations with service providers we were able to identify the logics behind these assignments *exactly* with any specific operational rule of the type one encounters in manufacturing settings. Therein, so called *dispatching rules* are applied to prioritise the jobs awaiting to be processed by the same resource (e.g. a machine), with a decision being made every time a resource is set free, and the job having at that time the highest priority — based on some pre-defined logic— being selected for processing.

The literature on manufacturing dispatching rules is rich and well established (Panwalkar and Iskander 1977), so we turned to it to look for rules for use in our aircraft turnaround setting, with two main aims:

- By working on a case study based on a real airport, we aimed to assess the relative benefits of the selected small set of rules, assuming the same rule is applied throughout the whole day of operation.
- We looked at comparing the previous *static* setting vs one where, as the day unfolds and more up-to-date data about the status of the airport and related flights is collected, one is able to dynamically switch among the given rules. This second scenario aimed to estimate, in cases comparable to the airport reality of our case study, the likely pros and cons from a more dynamic management of airport resources, under various conditions.

To the best of our knowledge, this is the first time either the static setting or the dynamic switching of rules have been studied in relation to airport apron operations. In the remainder of this section, we present our own selection of rules, as well as the mechanism through which we implemented the dynamic switching.

2.1 Dispatching Rules for Turnaround Team Allocation

Given the focus of this paper, a complete review of works on dispatching rules in manufacturing would be clearly out of scope. The interested reader may want to consult a few selected such works. In Panwalkar and Iskander 1977, a survey of in excess of one hundred rules is presented in details. Additional rules are presented in, to name but a few: Blackstone, Phillips, and Hogg 1982; Holthaus and Rajendran 1997b; Holthaus and Rajendran 1997a and Kaban, Othman, and Rohmah 2012. From this literature, it is easy to cluster subsets of dispatching rules around some recurring ideas. For instance, many of the rules are based on the known or expected processing time of a job — in our case, a whole aircraft turnaround, or its constituent tasks. Other rules are based on the number of operations of a job — say the number of tasks connected to the turnaround of a given aircraft. Further rules consider the arrival time of a job — in our case, the aircraft time of arrival at the parking stand after landing, while others take into account the due date of a job — say the time of departure from the parking stand. A quick look at the above cited works will reveal even more examples, as well as many cases of composite rules that take ideas from across multiple clusters. In the following, we list the selected few rules we investigated in our study.

Our first rule is called **Earliest Departure Time** (or **EDT**) and lists aircraft turnarounds in ascending order of scheduled departure time. Each time a turnaround team is set free from their previous turnaround, it is allocated to the turnaround with the most immediate scheduled departure time. EDT is equivalent to the *Earliest Due Date* (EDD) rule of manufacturing fame.

The next rule we selected is simply named **First-In First-Out** (or **FIFO**). Here, higher priority is given to aircraft the arrivals of which was notified earlier to the airport. By this we mean that as soon as the forthcoming arrival of an aircraft is confirmed, through the usual air traffic control channels, a notification is shared with ground service providers, so the allocated parking stand and expected time by which the aircraft will be there can be used to arrange ground services as needed.

We also defined a rule called **Estimated Operation Time** (or **EOT**), giving higher priority to aircraft expected to have longer turnaround times. EOT is opposite to the well known *Shortest Processing Time* (SPT) rule from manufacturing. The *exact* times required to perform many operations are hardly known ahead of the aircraft opening its doors after landing, hence some sort of estimation is involved in determining the value of each operation time. Our predictive analytics approach to estimate what services might be required by an aircraft rests on an artificial neural network (ANN) of five input layers, corresponding to: aircraft model/make, airline, scheduled time of arrival, number of inbound passengers and time of day. There are two hidden layers, of 25 neurons each, and three output layers to predict fueling, cleaning and catering services, respectively. Our choices of inputs and architecture for the network follow the principles suggested by Heaton (2008)— e.g. that one should always aim for a ‘minimalist’ model that produces acceptable results *vs* more complex models leading to over-fitting. Our model produced an accuracy of 90% and 78% on validation and test data respectively with a decent AUC of 76%. Data to be employed in such a procedure are widely available (FlightStats 2019). Assuming data are available on the duration of past occurrences of all the services involved, it then becomes possible to fit distributions to the available data. By bringing together predictions on services and best-fit distributions for the duration of each service, it is then possible to attach values to the expected operation times of all involved aircraft turnarounds.

We then devised a rule called **Least Slack Time** (or **LST**). This, as the name suggests, gives higher priority to turnarounds having lower *slack time* (ST), which in turn we define as:

$$ST = STD - EOT - LT$$

where STD is the scheduled time of the subsequent departure of the same aircraft, and LT is either the actual time of arrival — if the aircraft has already landed, or the current time in the opposite case.

We also considered a slight variation on LST, which we named **Least Slack Time with Look-Ahead** (or **LSTLA**). The idea in this case is to avoid assigning a team immediately to an aircraft which has a considerable slack (e.g. it may be a long haul flight), when other aircraft, still to begin their final approach to landing, may in fact have tighter schedules to meet. We then look ahead in time, for the subsequent 30 minutes, looking for such cases among all aircraft that are still en-route and the landing of which is expected to be notified within the same window.

Finally, in order to facilitate comparisons, we considered a ‘base case’ rule which we named **Status Quo** (or **SQ**). This rule assigns ground handling teams to aircraft turnarounds following exactly the same sequence that was predetermined ahead of the day of operation, with no room for flexible re-assignments as the operational conditions on the day unfold. Given that this tactical-level decision is made differently by different service providers, in our modelling and implementation we settled for the logic adopted by one of the global ground service organisations we have been working with. This allocation rule is a slight variation on the basic idea that assigns teams to turnarounds following their expected arrival time to the parking stand. Optimal or sub-optimal plans developed following some of the tactical-level works cited in Section 1 are likely to perform better than this very simple idea. However, in defining our SQ rule we targeted resemblance to practical approaches already adopted in the industry as opposed to more academic work that is still to be widely applied.

2.2 Dynamic Switching of Dispatching Rules

A major learning point from the literature on dispatching rules in manufacturing is that the same rule can behave very differently under different scenarios — see Holthaus and Rajendran (1997a), Rajendran and Holthaus (1999) and Kaban et al. (2012). A second major learning point states that dynamic switching of rules over time tends to perform well, and that its benefits (*vs* static rules) are case specific. We see no reason for these two major points not to hold true even with respect to aircraft turnarounds. In our experiments, we sought evidence in support of/against this thesis.

In this research, we considered dynamic switching among the six dispatching rules introduced above. This was pursued by integrating Deep Q-Learning — see Arulkumaran et al. (2017) and Kaelbling et al.

(1996), with simulation. More specifically, we deployed Deep Q-Learning with Experience Replay — see Lin (1993) and Sutton and Barto (2018). We aim to maximise the reward collected by going from state s to s' by choosing an action a . The α is the learning rate and γ is a discount factor. In symbols, our objective function to be optimised can be written as:

$$Q^*(s, a) = (1 - \alpha)Q(s, a) + \alpha(r + \gamma \max_{a'} (Q(s', a')))$$

Q-learning is an off-policy epsilon-greedy method which learns by evaluating the action that yields the most reward, independently of the policy (Watkins and Dayan 1992). As such, this makes Q-learning very effective for finding an optimal path in deterministic and discrete action spaces (e.g. in choosing 1 rule out of 6). The idea is to let the machine learn the patterns of switching rules from numerous simulation replications. We trained the model with penalties, i.e. any strictly positive waiting time for an aircraft (measured at half an hour intervals, a fairly common choice in airport performance measurement), for each such rule switch. This will inherently reinforce the machine to select the rule with least accumulated waiting time when presented with a given state of the system. We represent the state as a one-dimensional array of length 15. Each element of the array is binary that states whether or not there is an aircraft awaiting the assignment of a team of handlers. It depicts how busy the airport is at any one time. The ϵ dictates the exploration versus exploitation trade-off (Mnih et al. 2013). The model generates random (model independent) moves with a probability of ϵ and consequently the value of ϵ decreases over time to allow the model decisions on its own. The advantage of Experience Replay it intuitively acts as a memory bank. As such, in addition to having new generated samples, it also helps to converge quickly by having training samples that have been already explored.

3 MODELING AND IMPLEMENTATION

3.1 A Glimpse of the Simulation Model

Our study of both the static and dynamic switching of rules was centred around extensive simulation of entire days of operations of an airport, following actual flight schedules. Our simulation model had been originally developed and validated for a real airport of comparable size, as part of a separate sequence of projects. As such, the level of detail in the model includes key elements of the layout of an airport's apron, with actual distances between aircraft parking stands, permitted routes for ground movements of support vehicles and aircraft, etc. The actual number of service providers operating at the airport was considered. The related capacities in terms of both GSE and teams was inferred from a decade of projects with multiple ground handling service providers, which included two of the service organisations from our case study. With these organisations, a more direct validation of the same capacities was possible.

In terms of modelling stance, a hybrid simulation approach (Brailsford et al. 2019) with elements of both agent-based and discrete event simulation was adopted. Anylogic was chosen as the platform for implementation. The core of the model is agent-based, with six main agents interacting with each other to mimic various key aspects of the interactions among the many organisations that coexist within apron operations. These agents are:

- The *Aircraft* agent, to embed the *life cycle* of an aircraft at the given airport, from the beginning of the final approach to landing until its subsequent departure, all the way through its turnaround process.
- The *Stands* agent, which represent the logic based on which parking stands (and boarding gates) are allocated and managed at the given airport. This is normally a separate function in airport operations management, with its own logics and optimization methods.
- The *Vehicles* agent, to represent the various mobile GSEs employed at the given airport.
- The *Stores* agent, to represent the areas of the apron where GSEs have to be parked/stored between successive uses. In some airports many or all of the GSE needed have to be kept in stores physically

located around the parking stand they service. In other cases, GSEs are moved around the apron, which results in them being displaced in many different locations over the same day of operation.

- The *Handler* agent, to embed the characteristics, behaviours, logics and decision making employed, as well as communication mechanisms with which turnaround teams or individuals from the many service organisations involved go about performing their duties.
- The *StandPreparation* agent, which represents a prominent element of apron operations, i.e. all that happens to the teams/individuals and related GSEs that, having to turn around a given aircraft, set everything ready for the related operations to take place at the allocated parking stand, ahead of aircraft arrival.

State Charts model the behaviours of all these agents and the related interactions. Action Charts model other key aspects such as the logics representing all six dispatching rules. Finally, the discrete event components of our model deal with relevant details of all turnaround tasks to be executed, e.g. its key events, as well as the movements related to the more passive entities (e.g. hold bags and cargo) and resources that do not require being modelled as agents (e.g. certain ‘simple’ GSEs such as tow bars).

3.2 Overall Approach

Figure 1 presents our overall approach to modelling and analysis. The two main components are the ‘Simulation Module’ and the ‘Process Manager’. The former coincides with the hybrid ABS/DES model, while the latter consists of Python code, the main element of which is the RL algorithm. The Process Manager guides the overall control flow.




Figure 1: Conceptual representation of the integration between simulation and RL in our approach

The remaining elements of our approach include four additional components (Reinforcement Learning Algorithm, Artificial Neural Network, Set of Static Rules, Trained Model) and two pseudo data base systems (Flight Schedule Database, Experience Replay Memory). The symbols used in Figure 1 to indicate the additional four components have no particular meaning. The notation used to display the dependencies

between the two main modules and the remaining elements does have a specific meaning though. Dashed lines represent simple one-way dependencies, e.g. the Process Manager makes use of the RL algorithm, and the simulation model reads the flight schedule for the particular day of operation under study. Continuous lines represent instead somewhat more developed links, e.g. the Trained Model gets updated at the end of every episode, while the Experience Replay memory gets updated with new memory as things proceed. The total implementation of Reinforcement Learning is facilitated in Anylogic by DL4J. Additionally, as anticipated in Section 3.1, the dispatching rules are also implemented directly in Anylogic using State Activity Charts.

Lastly, we use Double DQN, as opposed to single DQN, to overcome maximisation bias Sutton and Barto (2018). So, here we use two separate models - primary(Q) and target(Q'), where the primary model weights are copied to the target model after a fixed number of episodes (5 in our study), following (Mnih et al. 2013). Therefore, the updated equation with the overall working principle is as follows:

Algorithm 1: Dynamic Switching training Algorithm using DQN and Simulation

```

Initialise replay buffer  $D$ ;
Initialise primary  $Q_P$  and target  $Q_T$  network;
Initialise  $T$  as 0;
for each Episode do
  for each Interval do
    Observe state  $s_t$  and select  $a_t$  with probability  $(1 - \epsilon)$ ;
    Execute  $a_t$  and observe  $s_{t+1}$  and reward  $r_t = R(s_t, a_t)$ ;
    Store  $(s_t, a_t, r_t, s_{t+1})$  in replay buffer  $D$ ;
    Randomly sample batch  $e_t \leftarrow (s_t, a_t, r_t, s_{t+1}) \sim D$ ;
     $Q^*_P(s_t, a_t) \leftarrow \alpha (r + \gamma Q_P(s_{t+1}, \text{argmax}(Q_T(s_{t+1}, a_{t+1}))))$ ;
    Perform Gradient Descent  $\sim (Q^*_P(s_t, a_t) - Q_P(s_t, a_t))^2$ ;
     $T \leftarrow T + 1$ ;
    if  $T$  is equal to 5 then
       $Q_T \leftarrow Q_P$ ;
  end
end

```

The above algorithm is executed by the Process Manager, where the first loop that is iterating over the episodes is analogous to simulation replications. This is followed by a loop that iterates over every half an hour interval. The next step is to generate an action (a_t) by observing state (s_t) with the probability of $(1 - \epsilon)$. This is followed by randomly selecting a batch of moves from the Experience Replay and training the model using traditional gradient descent methodology. The argmax outputs the index of the action that yields the maximum value. Lastly, as we have implemented Double DQN, we update the target model after every fifth iteration.

4 CASE STUDY

4.1 Selected Airport

With a traffic just short of 17 million passengers throughout year 2018, single-runway Luton Airport ranked at the time as the fifth busiest commercial airport in the UK. We chose it for two reasons: (1) because of our exposure to its operations, thanks to previous projects we carried out at this base, involving both airlines and ground handling service providers; and (2) because its overall (pre-Covid-19) traffic figures were representative of a fairly busy international origin and destination airport. Both the flight schedules for the data days included in the study as well as the information on the airport's layout were gathered from various publicly available resources accessible via the internet.

Our data for Luton Airport in 2018 shows that traffic ranges from about 50 flights/day, up to 62 flights/day and about 74 flights/day for low, normal and high traffic respectively.

In 2018, three ground handling service providers operated at Luton Airport, which for simplicity we name ‘A’, ‘B’ and ‘C’. At the time, they served one, five and four airlines respectively. In a typical high traffic day, they would provide services to 34, 11 and 29 flights each, respectively.

We assumed that, for the related flights, each of the three service providers were responsible for the majority of ground handling tasks, such as aircraft marshalling and parking, passenger boarding/disembarking and baggage unloading/loading. Luton Airport hosts, largely, operations from low-cost carriers, hence the aircraft models/makes to be turned around at its facilities tend to be relatively homogeneous (e.g. different versions of B737s, A319/320/321, etc.). We further simplified our case study by considering a common structure (gant chart) for the turnaround of the different aircraft types handled. We assumed that any aircraft turnaround would need a team of four employees from the related service provider. Given the different number of customers per service provider, as well as the different workload in terms of flights/day and related distribution over the daily schedule, we assumed that four, one and three teams were on the daily roster for respectively ‘A’, ‘B’ and ‘C’.

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, 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 were: 11 fueling trucks, 16 catering trucks, 16 cleaning trucks, 5 ambulifts and 12 passenger coaches respectively.

Finally, a note on our modelling of GSEs. Given the focus on team allocation, in our experiments we fixed the factor related to GSE management, in such a way to minimise its impact on our results and the related variability. More precisely, we assumed a scenario where all needed GSE are always available, at each stand, prior to any aircraft turnaround. We also assumed that, after use, the same GSEs are left in the designated storage locations around the same stand. This scenario effectively eliminates the GSE allocation and management problem outright. The scenario itself is quite realistic, and is the case at many of the busiest airports.

4.2 Performance Metrics

We now explore the different metrics that we chose to assess the performance levels of apron operations, by each service provider and overall for the airport, under both the static and dynamic settings involving the six dispatching rules discussed at the outset.

In the given operational setting, two delays are of particular concern, and hence should be minimised, both in the number of times they occur over the day of operation, and in their average duration. They are *turnaround delays* and *team delays*. With regard to the former, we deem a flight to be delayed when the actual completion of its last operation (mostly, its push-back into the nearby taxiway) overshoots the scheduled version of the same time by three minutes or more. With respect to the latter, a team delay is a positive difference between a flight’s notification time and the time it is assigned a team. For both these metrics, we recorded both the average duration and the number of occurrences over the day.

Given the hourly variations in the daily schedules that most airports experience due to commercial reasons, we also chose to collect the same four metrics at half-hourly intervals. This would ultimately enable a more fine-grained analysis of the dynamics through which each day of operation unfolds under different rules and settings.

In addition to the averages of the two delays over the day of operation, we also collected their cumulative versions over the same time span. Assuming the figures used in the industry to quantify the monetary value of each time unit of delay are sensible, with the assessment of cumulative delays it is possible to put a representative £-value against the punctuality in the system, or lack thereof.

5 NUMERICAL RESULTS

When comparing each of our five static rules against the pragmatic form of tactical allocation (Status Quo), a number of findings can be drawn. Our overall results suggest, although space limitations do not allow us to show this extensively, that the relative merits/demerits of each of the rules play out more or less markedly depending on which performance metrics one is looking at. This is only a confirmation that the intrinsic multi-criteria nature of the problem under study indeed requires appropriate approaches, if one is in search of a “true winner” among the rules investigated. As expected, rules also rank differently on the basis of how busy the daily schedule is. Statistical significance in the differences between some of the rules and the Status Quo is clearly detectable (see 99% confidence intervals in Figure 2), more clearly so for EDT (for low and normal traffic) and FIFO (for high traffic). Where statistically significant, the practical significance in the differences is around 10% of what is the typical duration of the turnaround for the flights investigated.




Figure 2: Average turnaround delays over the day of operation (Normal traffic days, in minutes)

Dynamic Switching (DS) does learn to make effective decisions. We found that in low and normal traffic days, it fails to overtake the best rule (EDT), but performs better than FIFO, LST and LSTLA (Figure 2). In high traffic days, the situation repeats, with DS unable to perform any better than the best (FIFO) but faring better than EDT. Either these results partly point to an over-fitted RL model, or in fact the model does take the best possible decisions but, given the limitation of being able to switch only among the given static rules considered, it cannot outperform the best in any given scenario. Again, even the ‘best’ decisions are not the best with respect to all performance metrics anyway.

Where the contribution from DS becomes more visible is, as expected, at busier times of the daily schedule. A more granular analysis, at half-hourly intervals as opposed to over the entire day of operation, confirms that DS does indeed learn, and keeps choosing the rule that best fits the particular peak situation (Figure 3).

In our case study, despite the non negligible overall traffic, there clearly are fairly long periods in the day where the airport is not very busy at all. Around these times, the performance levels of all rules tend to overlap. This also confirms previous results, as it does a plot of the cumulative turnaround delays (Figure 4).




Figure 3: Average turnaround delays —half-hourly analysis (Low traffic days)

The latter also helps to get an idea of the magnitude of the delays involved, with the related monetary value, when things go really bad. It also shows that, by the end of the day of operations, differences from having adopted one rule vs another can be quite important (about 250 minutes, in the case shown). Figure 4 also shows something we found consistently across all scenarios investigated, i.e. that, in our case study, the cumulative delays accrued over the first wave of morning flights alone, the single most problematic part of the day, make up for about 50% of the cumulative delays for the overall day, almost irrespective of the rule being adopted. Similarly, our graphs confirm that choosing the right rule for this daily peak alone yields to sizeable potential savings (of up to 200 minutes in the case shown).

6 CONCLUSIONS

We have shown that, when assigning teams/employees to turnaround tasks in real-time over the day of operation, the choice of the ‘right’ dispatching rule can indeed make a difference. We have also shown that DS clearly has potential. Limitations in our study include the relatively small set of static rules considered, which likely impacted the potential of DS, as the training of the RL algorithm only featured the given small set of static rules. Removing these limitations may help to foster acceptance of the proposed approach by practitioners, in a sector where this sort of decisions is largely made by specific personnel on the basis of expertise and experience.

We see the potential for investigating multi-criteria decision analysis methods to support the choice of the ‘best’ dispatching rule more thoroughly than it was pursued in this paper. Integration of GSE operational management in future investigation is very much advised, given that in many airport operational settings GSE allocation and management by teams/employees of ground service providers is often far from the simplified scenario assumed in this paper. Finally, controlling coordination both across the multiple stakeholders implicated in turning an aircraft around and within a handling organisation providing services to different airlines appears critical to enhancing airport apron management.




Figure 4: Cumulative turnaround delays —half-hourly analysis (High traffic days)

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

- 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.
- Arulkumaran, K., M. P. Deisenroth, M. Brundage, and A. A. Bharath. 2017. "Deep reinforcement learning: A brief survey". *IEEE Signal Processing Magazine* 34(6):26–38.
- Blackstone, J. H., D. T. Phillips, and G. L. Hogg. 1982. "A state-of-the-art survey of dispatching rules for manufacturing job shop operations". *International Journal of Production Research* 20(1):27–45.
- 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.
- CAPA 2011. "Handle with care – Europe's airport ground handling business to be liberalised further". <https://centreforaviation.com/analysis/reports/handle-with-care--europes-airport-ground-handling-business-to-be-liberalised-further-60660>, Last accessed on 2021-05-13.
- EU 2004. "EU Regulation 261". <https://eur-lex.europa.eu/eli/reg/2004/261/oj>, Last accessed on 2021-05-13.
- Eurocontrol 2017. "Cancellations to Air Transport in Europe 2017". <https://www.eurocontrol.int/sites/default/files/publication/files/coda-digest-annual-2017.pdf>, Last accessed on 2021-05-13.
- Evler, J., E. Asadi, H. Preis, and H. Fricke. 2021. "Airline ground operations: Optimal schedule recovery with uncertain arrival times". *Journal of Air Transport Management* 92. 102021.
- FlightStats 2019. "Airport on-time performance reports". <https://centreforaviation.com/analysis/reports/handle-with-care--europes-airport-ground-handling-business-to-be-liberalised-further-60660>, Last accessed on 2021-05-13.
- Gelhausen, M. C., P. Berster, and D. Wilken. 2013. "Do airport capacity constraints have a serious impact on the future development of air traffic?". *Journal of Air Transport Management* 28:3–13.
- 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.
- Guimarans, D., P. Arias, M. Tomasella, and C.-L. Wu. 2019. “A review of sustainability in aviation: A multidimensional perspective”. In *Sustainable Transportation and Smart Logistics*, 91–121. Elsevier.
- Heaton, J. 2008. *Introduction to neural networks with Java*. Heaton Research, Inc.
- Holthaus, O., and C. Rajendran. 1997a. “Efficient dispatching rules for scheduling in a job shop”. *International Journal of Production Economics* 48(1):87–105.
- Holthaus, O., and C. Rajendran. 1997b. “New dispatching rules for scheduling in a job shop—An experimental study”. *The International Journal of Advanced Manufacturing Technology* 13(2):148–153.
- 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.
- Kaban, A., Z. Othman, and D. Rohmah. 2012. “Comparison of dispatching rules in job-shop scheduling problem using simulation: a case study”. *International Journal of Simulation Modelling* 11(3):129–140.
- Kaelbling, L. P., M. L. Littman, and A. W. Moore. 1996. “Reinforcement learning: A survey”. *Journal of artificial intelligence research* 4:237–285.
- Lin, L.-J. 1993. “Reinforcement learning for robots using neural networks”. Technical report, Carnegie-Mellon Univ Pittsburgh PA School of Computer Science.
- 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.
- Mnih, V., K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. 2013. “Playing atari with deep reinforcement learning”. *arXiv preprint arXiv:1312.5602*.
- 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.
- Panwalkar, S. S., and W. Iskander. 1977. “A survey of scheduling rules”. *Operations research* 25(1):45–61.
- Rajendran, C., and O. Holthaus. 1999. “A comparative study of dispatching rules in dynamic flowshops and jobshops”. *European journal of operational research* 116(1):156–170.
- Schultz, M., J. Evler, E. Asadi, H. Preis, H. Fricke, and C.-L. Wu. 2020. “Future aircraft turnaround operations considering post-pandemic requirements”. *Journal of Air Transport Management* 89. 101886.
- Sutton, R. S., and A. G. Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- 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.
- Watkins, C. J., and P. Dayan. 1992. “Q-learning”. *Machine learning* 8(3-4):279–292.

AUTHOR BIOGRAPHIES

SUDIPTA SAHA is an alumnus of The University of Edinburgh Business School. At the time of this research he was a Master’s student in Business Analytics where his research, encompassed the domains of Simulation, Optimisation and Reinforcement Learning. His email address is saha.sudipta.495@gmail.com.

MAURIZIO TOMASELLA is with the Management Science and Business Economics Group at the University of Edinburgh Business School. His research interests include simulation, multi-criteria decision analysis, and their applications to airport operations as well as manufacturing and supply chains. 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.

SILVIA PADRÓN is Assistant Professor at TBS Business School (Toulouse, France). Her research is focused on metaheuristics and the hybridisation of optimisation methods with simulation techniques to solve stochastic and complex combinatorial problems in application areas such as transportation, aviation, and supply chain management. Her e-mail address is s.padron@tbs-education.fr.