STTAR: A SIMHEURISTICS-ENABLED SCHEME FOR MULTI-STAKEHOLDER COORDINATION OF AIRCRAFT TURNAROUND OPERATIONS

Maurizio Tomasella
Alexandra Clare
Yagmur S. Gök

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

Daniel Guimarans

Faculty of IT
Monash University
Building H, Level 6, 900 Dandenong Road
Caulfield East, VIC 3145, AUSTRALIA

Cemalettin Oztürk

United Technologies Research Centre Ireland Ltd.
4th Floor, Penrose Business Center, Penrose Wharf
Cork City, Co. Cork, REPUBLIC OF IRELAND

ABSTRACT

Aircraft turnaround operations involve all services to an aircraft (e.g. passenger boarding/disembarking, re-fuelling, deicing) between its arrival and immediately following departure. The aircraft, parked at its stand, witnesses a number of service providers move around it to perform their duties. These companies run substantially independent operations, working for different airlines/flights within a confined area where many resources, including physical space itself, have inescapably to be shared. Inter-dependencies among service providers abound, and knock-on effects at disrupted times are rife. Coordination from the side of the airport operator is difficult. We envisage a tactical robust scheme whereby ground handlers and the airport operator cooperate, albeit indirectly, in the development of plans for the next day that are less likely to be impacted by at least the more frequent operational disruptions. The scheme is based on a simheuristic approach which integrates ad-hoc heuristics with a hybrid simulation model (agent-based/discrete-event).

1 MOTIVATION

Demand for commercial air transport continues to grow. Passenger volume grew globally 7.1% in 2016 (IATA 2017) and is expected to more than double by 2035 (IATA 2017), which will require a strong infrastructure and operational efficiency. Additionally, there is a push to move to aircraft 4D trajectory in the US NextGen and Single European Sky initiatives, in which aircraft movement would no longer be restricted to the current node to node concept, but could exist in any *xyz* space coordinate at any time. This would increase the capacity of aircraft in the airspace (Schultz 2018) tremendously, adding to the need to optimize ground operational efficiency for the turnaround of aircraft in between ever more frequent flights.

Ground handling consists of all of the processes that occur between the time that the aircraft reaches the allocated 'block' at the airport (i.e. its parking stand/position) and the time at which the aircraft departs for its next flight. This has been deemed the second most common reason for delays to flights (Van Leeuwen and Witteveen 2009), which means that ground handling service providers (SPs) are in need of support to

optimally plan their operations. The focus of this research is directed at the daily assignment of teams to handle these operations in a tactical and robust nature. It is considered 'tactical' because the assignment is completed ahead of the day of operation. The plan is expected to be 'robust' as the uncertain elements of this process are taken into consideration in the assignment procedure, in a way that will be described later.

2 BACKGROUND

Every day, a multitude of flights arrive and depart from commercial airports. Each flight requires a 'turnaround process' where several operations must be completed, such as baggage loading/unloading, fuelling, and pushing back into the taxiway. Additional activities, such as cleaning, catering, and the transport of passengers with reduced mobility, are also required for many flights. All these operations are completed/supervised by teams of ground handlers, working for different SPs. Recent research on the turnaround process can be found in: (Dohn et al. 2008); (Padròn et al. 2016); (Fitouri-Trabelsi et al. 2013); (Andreatta et al. 2014); (Mota et al. 2017); (Norin et al. 2009) and (Wu 2008).

Airport operations get disrupted all the time, due to e.g. weather, technical faults such as passenger terminal black-outs, terrorism (McDonnell 2014), or simply inefficiencies, including in ground handling. Moreover, there is increasing competition in ground handling due to the now established deregulation, thus lowering the market entry costs(Padròn et al. 2016). To keep up with the dynamic schedule of flights each day, there are two decision points that SPs must consider. First, there is the turnaround process itself, as there are dependencies in the order in which processes can happen. The other decision is the assignment of turnaround teams to flights, together with the needed physical assets, called Ground Service Equipment (GSE), such as passenger stairs, baggage loading/unloading conveyor belts, etc. This paper focuses on team assignments; some related literature also consider the allocation of GSE (Andreatta et al. 2014).

Tactically, SPs will plan their operations up to a month in advance, based on the expected flight schedules; however, daily disruptions require a process to re-plan the allocation of ground handling teams (Padròn et al. 2016) just before the day of operation. Literature on the topic supports single SPs to develop tactical team assignments. What it does not consider in doing so, to our knowledge, are key elements that are out of the control for the single SP, in particular: (1) the tactical plans developed by the other SPs operating at the same airport, and their impact on the dynamics of ground handling on the day of operation; and (2) uncertain aspects of aircraft turnaround operations, such as duration of specific turnaround tasks, uncertain times of aircraft arrival to stand, availability of GSE when needed, etc. This paper presents an idea for how to include these two aspects in tactical planning of turnaround teams by the single SP.

3 PROBLEM STATEMENT AND FORMULATION

In this section we formulate the Turnaround-Team Assignment Problem (T-TAP) - see (1) to (8), inspired by (Andreatta et al. 2014), which however focuses on allocation of GSEs to aircraft turnarounds (GSE Assignment Problem, or GAP). We are at a given airport and look at the tactical planning problem whereby all teams of handlers, simply 'teams', of a given SP, are assigned to aircraft turnaround tasks, 'tasks' for short, for the entire duration of the day of operation. The allocation is devised ahead of the day of operation, say in the night-time gap between two consecutive days.

The model seeks a plan that minimizes (1) the total delays associated with team-task allocations that, at the planning stage, appear unlikely to be able to start on time. A delay $c_{ij}^h > 0$ will be incurred if team h is allocated to task j after completion of task i, that is if $x_{ij}^h = 1$, despite the guesstimate that the aircraft associated with task j is likely to reach its destination stand prior to the moment when team h is available at the same stand, ready to cater for it. This 'moment' is calculated as the sum of the expected time of team h departing from the stand where they completed task i, plus the time for the same team to move to the stand of task j, plus the time to set up the destination stand for the next turnaround to start (11).

The sets of all teams and all tasks are denoted as H and K respectively, while expected start and end times for task $i \in K$ are denoted as s_i and e_i. We adopt symbols A_h and Ω_h to denote 'dummy' tasks, one

each for the beginning and the end of the day/shift respectively, to represent that each team start and end their job from the operational base/office of the SP at the given airport. Our decision variable x_{ij}^h is binary.

The T-TAP can be then formulated as:

$$\min \sum_{h \in H} \sum_{i \in K} \sum_{j \in K, j \neq i} c_{ij}^h x_{ij}^h \tag{1}$$

subject to:

$$\sum_{h \in H} \sum_{i \in K} x_{ij}^h = 1 \ \forall j \in K \setminus \{A_l : l \in H\} \tag{2}$$

$$\sum_{j \in K} x_{A_h j}^h = 1 \ \forall h \in H \tag{3}$$

$$\sum_{v \in K} x_{vi}^h - \sum_{w \in K} x_{iw}^h = 0 \ \forall i \in K \setminus \{A_h : \Omega_h\}, h \in h\} \tag{4}$$

$$\sum_{i \in K} x_{i\Omega_h}^h = 1 \ \forall h \in H \tag{5}$$

$$x_{ij}^h \leq \phi_{ij}^h \ \forall i, j \in K, h \in H \tag{6}$$

$$x_{ij}^h \leq \lambda_{ij}^h \ \forall i, j \in K, h \in H \tag{7}$$

$$x_{ij}^h \in \{0, 1\} \ \forall i, j \in K, h \in H \tag{8}$$

Constraint (2) ensures that all tasks are completed exactly once. Constraints (3)-(5) ensure that each team starts from the related initial dummy task and completes the working shift at the related final dummy task, whilst passing through a feasible sequence of actual tasks in between. Constraints (6) and (7) include binary parameters to restrict the compatibility of team h with specific sub-sequences of tasks $i \to j$. As (9) shows, a sub-sequence $i \to j$ is inhibited from being assigned to team h if either of the following is true: (a) if $i = j$; or (b) if team h cannot feasibly travel from the stand of task i to the stand of task j within a set buffer time ε from the planned start of task j.

$$\phi_{ij}^h = \begin{cases} 0 & \text{if } i = j \ \forall i, j \in K, \\ 0 & \text{if } e_i + t_{ij}^h + \varepsilon < s_j, \\ 1 & \text{otherwise.} \end{cases} \tag{9}$$

All other reasons that inhibit sub-sequence $i \to j$ from being assigned to team h are captured by (10), e.g. all flights of the same airline at a given airport are generally serviced by teams of the same SP.

$$\lambda_{ij}^h = \begin{cases} 0 & \text{if } h \text{ is incompatible with } j, \\ 1 & \text{otherwise.} \end{cases} \tag{10}$$

Finally, (11) defines the cost parameter of the objective function, that is delay c_{ij}^h, where: β_h labels the physical location of 'the office' of the given SP, at the given airport; $t_{\beta_h i}^h$ denotes the travel time for team h to move from its own office to the stand where the same team will execute task i; $t_{i\beta_h}^h$ denotes the travel time for team h to move from the stand where they just completed task i, back to its own office; and δ^h is a parameter representing the average time for team h to be ready to start a task, from the moment when it arrives at the designated stand.

$$c_{ij}^h = \max\left(0, e_i + t_{i\beta_h}^h + t_{\beta_h j}^h + \delta^h - s_j\right) \tag{11}$$

Solving the T-TAP efficiently/quickly is quite hard. This is because the T-TAP is modelled after the GAP, the computational intractability of which is discussed in (Andreatta et al. 2014). Our observations from working with many airports over the past decade, in the UK and the EU, confirm that practitioners do solve the T-TAP by following various forms of ad-hoc heuristic reasoning. A simple rule frequently adopted by ground handling service providers is the FCFS (First-Come-First-Served), whereby teams are allocated to tasks one after another, following the order of arrival of the related flights at the planned stand.

The reality of air side operations at most commercial airports of interest is that the complexity and uncertainty characterising their operational environments are such that, quite quickly into the day of operations, most of the tactical plans developed for the day become unfeasible to follow. This is true for planning turnaround teams as it is for planning GSEs and airport resources in general. Many SPs respond to this by developing a tactical plan only for the start of the day, often covering only until the end of the 'first wave' of morning flights. This makes sense as in many airports the first wave is one of the peak daily periods in terms of traffic intensity. Then, throughout the rest of the day, operations are typically just 'run', with teams assigned as the necessity arises, following some given operational rules. A more structured, approach is to look from the outset at the development of tactical plans that are somewhat 'robust'. A robust tactical plan would be one that performs better at standing the test of time on the day of operation.

The very nature of the T-TAP itself does not, however, allow to develop the intended robust plans, simply because all its main parameters are deterministic, especially those mentioned in (11). In this paper, we devise a scheme that enables the development of robust plans by each SP, where we integrate the solution of the deterministic T-TAP optimization problem with the assessment, through simulation, of the robustness of the developed solution.

4   THE SYSTEMATIC TURNAROUND-TEAM ALLOCATION ROUTINE

In this section, a process is proposed that involves all the SPs working at the given airport and the airport operator (AO) running the same facility. The purpose of this process is to support each SP separately to develop tactical plans for team allocation and assess their robustness. The latter is achieved by quantifying the likely delays to be incurred by each SP on the next day of operation, if their plan was enacted as envisaged. This would in turn lead to an estimation of the likely impact of any such delays on the Service Level Agreements (SLAs) in place between SPs and the airlines/the AO. We call this process STTAR, acronym standing for Systematic Turnaround-Team Allocation Routine - see Figure 1.

Let us assume that, as is often the case, SPs plan and run their operations for the next day independently from one another, and that they employ totally disjointed fleets of GSE. Let us also assume that there exists a 'central computing facility' at the given airport, on which a sufficiently detailed and validated simulation model of ground handling operations at the same airport is available. We assume that this model is perceived as credible by both SPs and the AO. Let us finally assume that, between two consecutive days of operations, there exists a night-time window during which the STTAR process can take place. During this period, the process itself evolves in two subsequent phases.

In 'Phase 1', each SP can submit to the central computer its intended tactical plan for the next (whole) day, in which all of its teams on shift are allocated to some tasks. After the submission by the single SP, which we call SP X, the simulation model is run for a known number of independent replications. The model mimics the allocation plan submitted by the same SP. It also reproduces tactical plans for the remaining SPs, say SPs Y, Z, etc. A range of such plans are simulated, so that interaction between the plan submitted by X and meaningful combinations of the plans to be possibly enacted by Y, Z, etc. can be estimated. The simulation model also uses distributions to model other relevant stochastic aspects such as turnaround task duration, aircraft arrivals/departures, availability of GSEs and teams, etc. Assumptions around the possible tactical plans of SPs Y, Z, etc. can be easily developed based on previous submissions by these SPs to the central system. Assumptions around the statistical distributions that model all relevant stochastic aspects of the given airport can be derived, as we did, from historical data.
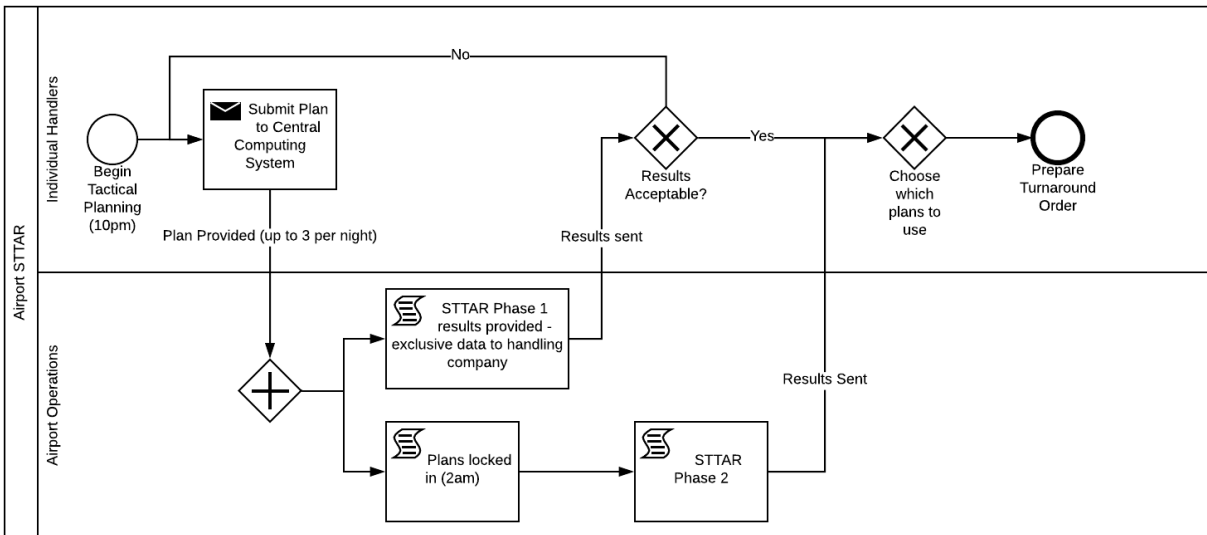
Figure 1: A simplified process map for STTAR, with tentative example timings.

After running Phase 1 simulations, simplified descriptive statistics of the likely performance of each single SP, e.g. ranges for the expected number and duration of the delays occurred in their own turnarounds, are computed and sent back to the SP. The SP can then reflect upon the obtained feedback, and possibly test other alternative plans for the next day. A cap to the total number of submissions by the same SP may apply, e.g. three in the example case of Figure 1. We assume that the SPs cannot access the simulation model directly, but can only upload their own tactical plans, one at a time, independently from one another. We also assume that no visibility exists of other SPs' submitted plans and, quite naturally so, of the results obtained by other SPs' at each submission. Exclusivity of data is vital to this process so that ground handling companies do not share data that is paramount to their competitive advantage.

Each SP benefits from participation to Phase 1 in that they obtain context information, albeit simplified, concerning the likely 'robustness' of the intended tactical plan for their own teams, over the next day of operation. Phase 2 of STTAR will offer benefits that are similar in nature, but will add optimization steps that are intertwined with the simulation runs, in the hope that even more robust tactical plans can be developed and suggested to the SPs for possible consideration. The robustness in Phase 2 will however be computed for the airport as a whole, as opposed to being computed separately for each single SP. As discussed in details in the next section, Phase 2 is meant to be even more highly automated. For this purpose, two alternative simulation-optimization procedures, both of which can be classified as simheuristics (Juan et al. 2015), were developed and will be discussed in the next section.

5   SIMHEURISTICS-ENABLED TACTICAL ROBUST PLANNING

Phase 1 ends at a set time in the considered night-time window - say 2:00am in the example of Figure 1. Then, one tactical plan per SP will be locked-in to be taken into Phase 2 as part of the 'initial solution' - see Figure 2. Locked-in plans can be, according to choices made by each single SP, either their 'preferred plan' among those tested in Phase 1, or a different plan altogether, if they wish so. Because it is SPs, ultimately, that decide what plan of action to follow, when the day starts. The locked-in plans as confirmed by the SPs, either explicitly or implicitly, and are then collated by the central computer, to form the 'initial solution' for the problem of Phase 2. This initial solution includes feasible allocations to all teams of all SPs to all turnaround tasks, throughout the day of operation that is about to start.
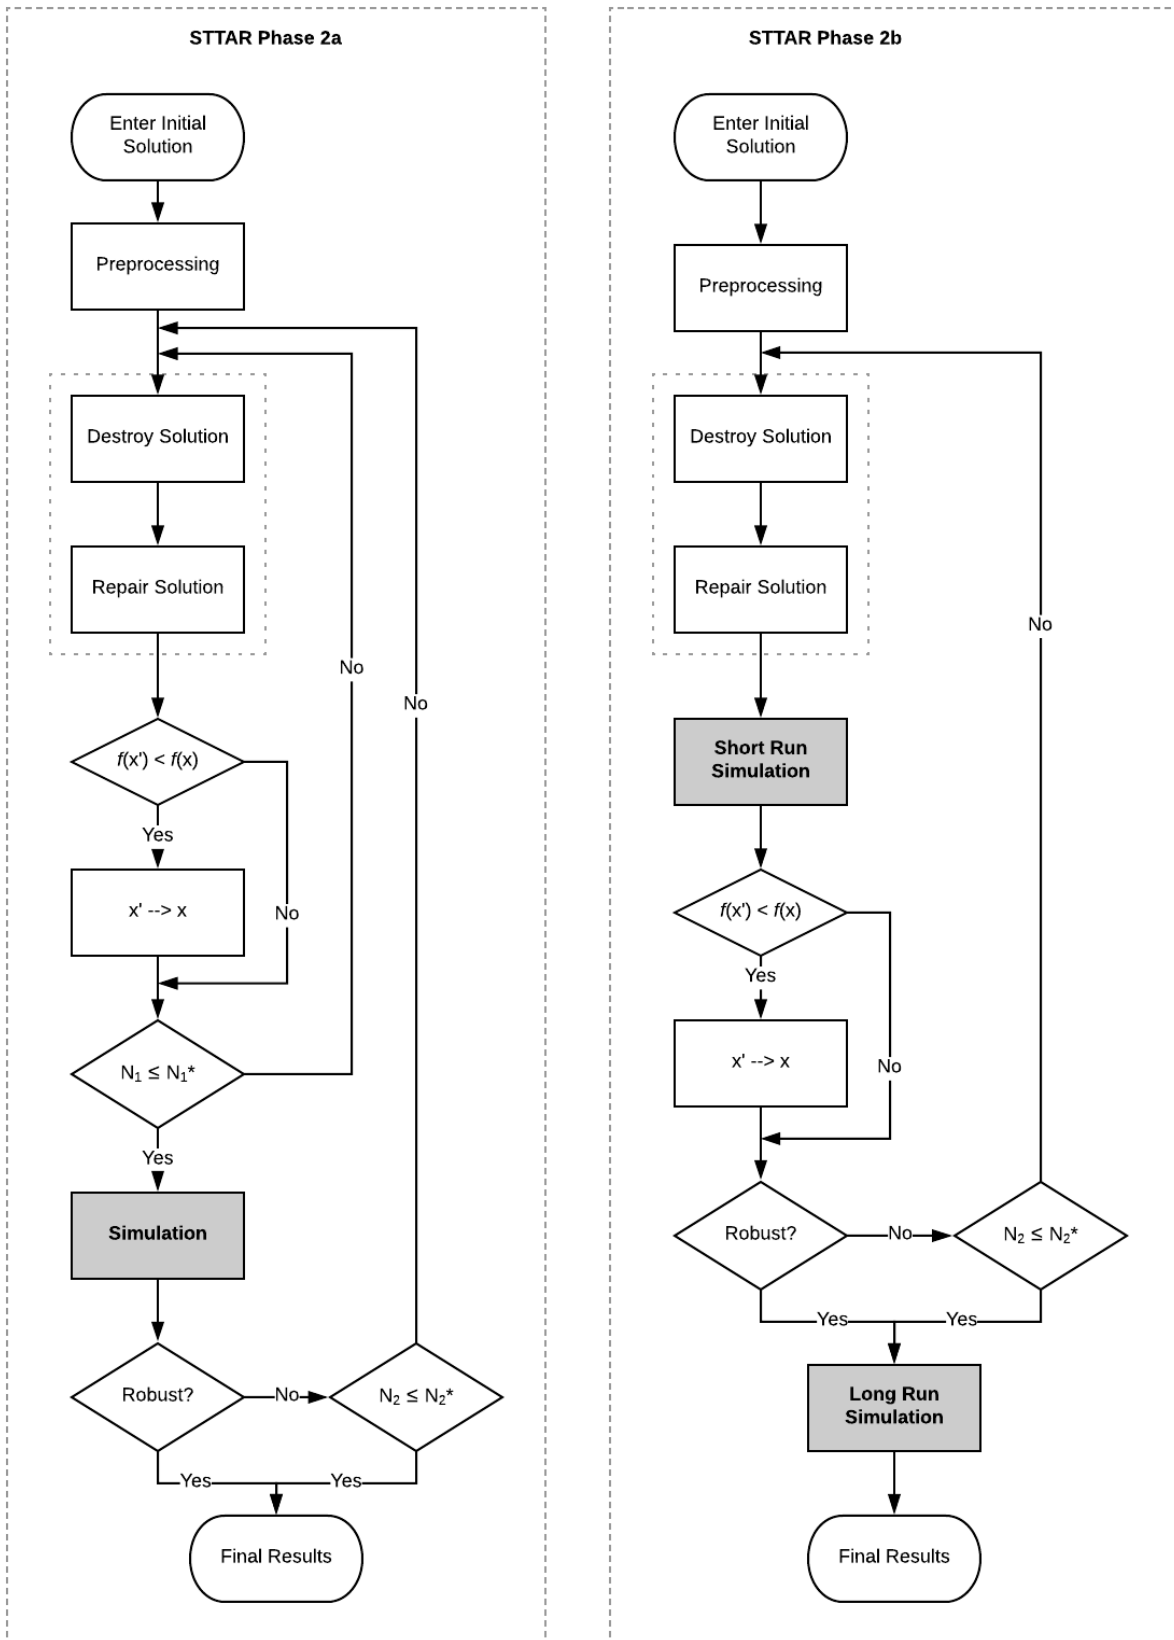
Figure 2: Simheuristic approaches proposed for Phase 2 of STTAR.

During Phase 2, the overall expected delays, from the airport's perspective, will be minimized, whilst ensuring minimum levels of robustness. This is achieved first by solving a (deterministic) T-TAP for the whole airport, through a combination of heuristics developed ad-hoc for the problem. After that, simulation is employed to test the robustness of the overall set of team allocations across all SPs. We propose and test two alternative ways for intertwining the heuristic reasoning component with the simulation component of our approach, shown in Figure 2 as Phase 2a and 2b respectively.

In option 2a, we allow for up to N_1^* iterations of the heuristic search, before we take the best solution found up to that point to the subsequent stochastic simulation stage, where estimates for a set of robustness-related, relevant performance indicators are computed. Should the robustness of the overall plan be considered unsatisfactory, e.g. because a set of SLAs are not met, we allow for up to N_2^* iterations of the whole procedure, thus restarting from the heuristic search once again.

The question of what notion of/metric for robustness should be adopted can be developed in many different ways (Bertsimas and Sim 2004). Here, our robustness metric of choice is the expected total number of turnaround delays over the day of operation, which has to fall below a percentage threshold of the number of flights in a day. The threshold itself is then a problem parameter; different thresholds could be set for different SPs, as well as for the airport overall. Alternatively, one could focus on the magnitude of turnaround delays themselves as another measure of robustness in this context, but other choices can be made as well.

In option 2b, at each iteration of the heuristic search procedure, we run shorter stochastic simulations, the results of which are used to decide on whether the current best solution needs updating. After that, the best solution is tested for robustness, based on the exact same performance indicators employed in option 2a. In the case when the robustness of the plan is considered unsatisfactory, we allow for up to N_2^* iterations of the whole procedure, at the end of which we run a final stochastic simulation to compute a wider range of performance indicators, before these are fed back to the SPs.

At the end of either option, the SPs will then have information at their disposal, to inform their final choice of team allocation plan for the next day of operation.

5.1 Stochastic elements of the problem and the role of simulation

There are four main elements of the problem that, at the end of the day of operation, will have impacted most significantly the operational performance of the system. They are: (1) actual flight arrival times on stand; (2) duration of turnaround operations; (3) team-task allocations; and (4) GSE-task allocations.

While (1) is totally out of the control of SPs, the variability inherent to (2) can only be partially controlled by them. First, because not all the activities in the turnaround task will be executed by the same SP. For instance, while operations such as baggage loading/unloading and passenger boarding/disembarking are executed by the SP, other operations, such as re-fuelling or toilet cleaning, are generally provided by other SPs that specialize in such services. Coordination among SPs involved in the turnaround of the same aircraft, and beyond, is then paramount. Other aspects of the problem that impact the variability in the duration of turnaround activities and that is out of the control of the main SP are e.g. the actual number of outbound/inbound passengers, of outbound/inbound pieces of baggage, etc., which vary even for the same flight on different dates. Finally, additional elements of variability in processing times are directly connected to the way operations are actually executed by the different teams of the many SPs.

Element (3) above is more controllable by SPs and is in fact the focus of the T-TAP as defined above. Finally, element (4) is very often the subject of its own planning problem, similar to the T-TAP, as in (Andreatta et al. 2014). In Phase 2 of STTAR, elements (1), (2) and (4) are modelled as uncertain factors and considered during the simulation runs in which a number of alternative team-task assignment plans are tested, as explained above.

The main novelty in our approach is neither in the use of heuristic search to solve the T-TAP (Andreatta et al. 2014), nor in the way simulation and heuristic search are intertwined in 2a or 2b (Juan et al. 2015) in a scheme going by the name of 'Simheuristics'. On the other hand, and to the best of our knowledge, this

paper is the first applying Simheuristic schemes to the tactical planning of turnaround teams, and one of a very small number of works looking at turnaround team planning more in general (Ho and Leung 2010).

5.2 On developing an initial feasible solution

Both algorithms 2a and 2b make sure they can start from a feasible initial solution to the T-TAP, or in general, with any feasible tactical plan for turnaround team allocations across all SPs. To ensure feasibility of the initial solution considered, we developed a construction heuristic (a simplified version of the fast construction heuristic in (Andreatta et al. 2014)), which is based on assigning tasks to teams as long as they meet basic compatibility requirements, particularly in terms of maximum tolerable delays. Our ad-hoc heuristic also provides a fast solution but, similarly to (Andreatta et al. 2014), not always a feasible one. This means that our heuristic does not always provide a team assignment for every aircraft turnaround/task. This happens especially when the time compatibility criterion is set to be too strict. In such a case, if there are, for instance, five flights that are within a similar time window, with only four teams to allocate among these flights, there would be one that would be left unassigned, which is not acceptable in practice. For this reason, we developed an ad-hoc insertion heuristic. Our 'plan repair' tactic in this algorithm seeks to ensure feasibility whilst targeting a lower overall delay incurred over the day. For space reasons, we do not provide the pseudo-code of either the construction or the insertion heuristic in this paper, but will make them available on request.

5.3 Destroying and Repairing

As shown in Figure 2, and similarly to many existing approaches in the vast field of heuristics-/metaheuristics-based search (Pisinger and Ropke 2010), a sequence of two steps is helpful to first 'destroy' and then 'repair' a solution already examined and transform it into a new solution to be taken to numerical examination. As for the generation of the initial solution, our approach in both destroy and repair steps is very pragmatic. The 'destroy' portion of our heuristic is simply based on choosing a turnaround task out of a limited number of options and the removing the corresponding assignment to the team. The result of this process would leave an infeasible solution, which can then undergo the same insertion process already mentioned above. Our 'destroy' tactic identifies the task for removal by randomizing out of three options: (a) the task with the largest delay; (b) the task directly preceding the one with the largest delay; and (c) a task at random. Option (b) is similar to (a), but follows the idea that if the 'preceding task' can be assigned to another team, the total delay could be lowered overall, by helping to avoid at least a portion of the knock-on effects of earlier delays.

6 NUMERICAL INVESTIGATION

6.1 Case study

For our numerical experiments, we studied the case of London Luton airport. From publicly available data, we took Luton's flight schedules of 2018, from which we 'extracted' three daily schedules, that we believed were representative of three very different workload levels for the airport (especially in terms of total number of flights per day), and we labelled them 'Low', 'Normal' and 'Busy' Day respectively. Also from publicly available data, we took the aerial map of the airfield and worked out the positions of all aircraft parking stands, as well as ground vehicle routes, etc. Again from public data sources, we got to know that three SPs were operating in Luton at the time of the study, and the airlines whose flights each of them serviced. Then, based on our previous simulation and optimization projects involving ground handlers, we estimated the likely number of teams and GSEs for each SP at a general airport of the size and operational characteristics of Luton. Also based on experience from previous projects at real airports and working with ground handling SPs was the definition of the typical operational policies run by SPs for an airport comparable to Luton. With this, we wish to acknowledge that ours is a realistic case study,

whilst at the same time that none of the numbers we used in our simulation-optimization investigation were provided directly by either the actual AO or SPs operating at Luton. As a result, the numbers discussed below may not be representative of the actual performance of either the airport overall or of any of the SPs operating at Luton in 2018.

6.2 Heuristics

In our Phase 1 experiments, in absence of SP-specific team allocation policies from the SPs operating at the case study airport, we implemented the following three options for heuristic search:

- Option *H*: This is the simple FCFS policy that many SPs implement in their operational planning;
- Option *AA*: This simplified form of heuristic search runs (once) our construction heuristic, followed by a single run of the insertion heuristic, to ensure feasibility of the plan;
- Option *AA+*: This algorithm looks very much like algorithm 2a, without either the simulation runs or the robustness check - it basically represents a number of iterations of option *AA*.

In our Phase 1, we first ran each of the three options for each SP and the related teams/tasks, to determine the related tactical plan. All $3^3 = 27$ possible combinations of plans were then simulated to determine the operational performance. We also assumed that each SP would lock-in one of the three options for use at the start of Phase 2.

6.3 Technical details

The construction, insertion and destroy and repair algorithms were implemented in the Eclipse IDE and run on a laptop with 4 cores at 2.7GHz with 8GB of RAM, with the Apache POI library used for I/O operations involving spreadsheets. The simulation model was developed as an agent-based/discrete-event model in the Anylogic 8 University Edition software, and run on a workstation equipped with an Intel Xeon processor at 3.4 GHz with 32GB of RAM. Algorithms 2a and 2b were implemented within Anylogic itself, exploiting the Custom Experiment functionality of the software.

6.4 Parameter Settings

Through standard techniques of output data analysis (Law 2015), we determined that we needed 95 replications to build 90% confidence intervals for the performance indicators (KPIs) to be estimated during the study, with a considered absolute error of $\beta = 1$ delay per day. The KPIs of choice were: number of turnaround delays, duration of turnaround delays, number of team delays (i.e. number of times in a day teams of the SP get late to the designated stand) and their duration. Of these KPIs, we computed averages from the simulation replications and determined confidence intervals on the mean. We set the number of replications equal to 95 in the experiments of Phase 1 as well as in the stages of Phase 2 labelled 'Simulation' (2a) and 'Long Run Simulation' (2b), while we set an arbitrary number of replications equal to 15 for the 'Short Run Simulation' of Phase 2 - see Figure 2. (Juan et al. 2015) makes a similar choice in relation to the number of replications for the 'short run'.

We set $N_1^* = 100$, $N_2^* = 10$ in 2a and $N_2^* = 50$ in 2b. We did not attempt to fine tune these parameters, for the purpose of this paper. There can be a number of improvements added to future implementations. For instance, in 2a one may add a stopping criterion that forces the exit from the inner loop if a better solution is not found in a given amount of iterations, thus avoiding having to run the heuristic for N_1^* times. Other improvements concerning the outer loop - which might not be obvious from Figure 2 - include e.g. considering that if the next solution generated through destroy/repair has already been evaluated in the past, its simulation should be skipped and a new destroy operation should be run.

Finally, the robustness threshold we set for our experiments was of 50% of the daily tasks being delayed. This means that if a given solution taken to the robustness check leads to turnaround delays on

Table 1: Sample Output from STTAR Phase 1

Options	Num T/Round Delay (Min)	Num T/Round Delay (Max)	T/Round Delay (Min)	T/Round Delay (Max)
SP1 - Low Day				
Option AA	13.7	15.84	25.51	28.06
Option AA+	13.85	15.7	25.65	28.13
Option H	14.41	16.43	23.11	26.29
SP1 - Normal Day				
Option AA	16.99	19.36	30.81	34.17
Option AA+	16.53	19.18	31.63	35.57
Option H	16.88	18.97	28.84	33.59
SP1 - Busy Day				
Option AA	24.04	26.51	47.51	51.2
Option AA+	23.51	26.1	46.88	51.55
Option H	23.75	26.45	45.11	48.63

50% or less of the daily tasks at the given airport, then 2a will stop, while 2b will enter the 'Long Run Simulation' stage. Normally, in real airports, this threshold would be set to a considerably lower value, say 20% or lower. We chose 50% because of the characteristics of our case study, where we intentionally chose to 'stress-test' the system by setting a number of teams per handler that is much lower than what one would normally find in an airport with the considered daily schedules.

6.5 Results

Table 1 shows a simplified mock-up of the results we envisage that could be presented to service provider SP1 (fantasy name, for confidentiality reasons) at the end of each run of Phase 1 computations, focusing on turnaround (T/Round) delays ('Num' for 'number of'). All statistics shown in the table were computed by the airport's central computer by combining the single option shown on the left, supposed to have been chosen by SP1, in combination with all other $3^2 = 9$ options possibly chosen by the two competing SPs. While the table could be augmented with additional KPIs (e.g. team delays and number thereof) and, for each KPI, with additional descriptive statistics, we believe this output should be kept as simple as possible, to aid the SP in its selection and understanding of the likely operational (and, ultimately, commercial) consequences from choosing an option for the next day.

Table 2 shows the running time and number of iterations in the outer loop for both algorithms 2a and 2b, under different operational workloads. Algorithm 2a seems slightly quicker than 2b in 'Low Day' conditions; the difference in running times is not statistically significant otherwise. Looking at 2a data, we notice that, sometimes, robustness is reached before the N_2^* threshold. This is not the case however for 2b, where, ultimately, the condition leading to exiting the outer loop seems to be always the reaching of N_2^*.

Finally, our Phase 2a/2b results also gave us confidence intervals for the difference between the means of KPIs, estimated as $\hat{\mu}_{2a} - \hat{\mu}_{2b}$, for all possible combinations, something that for reasons of space is not shown here. In terms of number of both types of delay, differences were not statistically significant for most combinations and, when they were, they were of the order of ± 1 either way. Differences in the magnitude of both types of delay were more marked instead, and appeared to be always to the advantage of 2b. These results are preliminary and difficult to generalise, at the present time, without a more extensive study where parameters are better fine-tuned and where more varied operational conditions for experimentation are investigated. However, even within the above limitations, especially on problem parameters, we believe the same results are credible, with reference to the specific case. Overall, the airport case being investigated

Table 2: Runtime and Performance Comparison of STTAR Phase 2

	Low Day	**Normal Day**	**Busy Day**
Num Flights	50	62	74
Num Flights (Peak Hour)	9	11	14
Avg Runtime (s) (2a)	538.629	2064.438	2362.902
Avg Runtime (s) (2b)	643.958	1983.009	2125.742
Conf Interval	[-118.904, -91.756]	[-113.138, 275.998]	[-210.514, 684.834]
N_2 Min (2a)	7	7	7
N_2 Max (2a)	10	10	10
N_2 Min (2b)	50	50	50
N_2 Max (2b)	50	50	50

is characterised by challenging operational conditions, with GSE and team capacities are extremely tight, especially at daily peak times.

7   CONCLUSION

In this paper, we have shown the concept and application of a not so futuristic approach to enhance the overall ability of SPs operating at the same airport to plan next day turnaround team operations. A first element of novelty is that this is the first paper to employ a Simheuristic approach (Phase 2a/b of the STTAR scheme) to problems of tactical resource allocation in aircraft turnaround, especially for the allocation of turnaround teams of ground handlers. The detailed discussion of this element of novelty has been developed throughout this paper. A related level of novelty regards the fact that ours is the first Simheuristic approach to employ a hybrid agent-based/discrete event simulation model. For reasons of space, the discussion of this aspect was not developed in the previous pages, and will become the main topic of a separate paper.

The implementations of the algorithms of both Phase 1 and 2a/2b leaves plenty of room for improvement. In Phase 1, work should be undertaken on all the heuristics involved, while in Phase 2 more work need to be done on understanding how the parameters would be set, given a specific airport where the potential use of STTAR is being investigated. More interesting developments we are currently working on include aspects such as making sure the simulation components in Phases 2a/b provide more constructive feedback to the heuristic reasoning component, every time the outer loops need to be run again. For instance, results of the simulations may be exploited to refine the very definition of the T-TAP and/or to dynamically generate cuts of various nature that would ultimately improve the methodology itself and the results reached.

Finally, we believe the main advantage for SPs in adopting the STTAR scheme, with respect to the current situation, resides in the additional information that, between Phase 1 and Phase 2a/b, SPs will have at their disposal, to support them in the development of more robust tactical plans. Phase 2a/b also adds advantages in the sense that an overall 'global' optimum for the airport as such is devised, without being binding to SPs. SPs will ultimately enact the tactical plans they feel are best for themselves, hence this non binding nature of the STTAR scheme should help to overcome potential barriers to adoption. Other barriers may still exist, such as the need to ensure continuing credibility of the simulation models from the perspectives of all SPs, and this is something which requires further study.

REFERENCES

Andreatta, G., L. Capanna, L. De Giovanni, M. Monaci, and L. Righi. 2014. "Efficiency and robustness in a support platform for intelligent airport ground handling". *Journal of Intelligent Transportation Systems: Technology, Planning, and Operations* 18(1):121–130. cited By 3.

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.

Bertsimas, D., and M. Sim. 2004. "The price of robustness". *Operations research* 52(1):35–53.

Dohn, A., E. Kolind, and J. Haywood. 2008. "Optimizing manpower allocation for ground handling tasks in airports using column generation". *Technical University of Denmark, Department of Management Engineering*.

Fitouri-Trabelsi, S., C. Nunes Cosenza, and F. Mora-Camino. 2013. "Ground handling management at airports with fuzzy information". In *IFAC Proceedings Volumes (IFAC-PapersOnline)*, Volume 6, 373–378. cited By 0.

Ho, S. C., and J. M. Leung. 2010. "Solving a manpower scheduling problem for airline catering using metaheuristics". *European Journal of Operational Research* 202(3):903–921.

IATA 2017, 12. "Annual Review 2017". Web Page. An optional note.

Juan, A., J. Faulin, S. Grasman, M. Rabe, and G. Figueira. 2015. "A review of simheuristics: Extending metaheuristics to deal with stochastic combinatorial optimization problems". *Operations Research Perspectives* 2:62–72. cited By 65.

Law, A. M. 2015. *Simulation modeling and analysis*. New York: McGraw-Hill.

McDonnell, D. 2014, 11. *A-CDM Concept of Operations*. 15 ed. Brisbane Airport, Sydney Airport, Melbourne Airport, Perth Airport, Airservices Australia.

Mota, M., G. Boosten, N. De Bock, E. Jimenez, and J. de Sousa. 2017. "Simulation-based turnaround evaluation for Lelystad Airport". *Journal of Air Transport Management* 64:21–32. cited By 0.

Norin, A., T. A. Granberg, P. Värbrand, and D. Yuan. 2009. "Integrating optimization and simulation to gain more efficient airport logistics". In *Eighth USA/Europe Air Traffic Management Research and Development Seminar*.

Padròn, S., D. Guimarans, J. Ramos, and S. Fitouri-Trabelsi. 2016. "A bi-objective approach for scheduling ground-handling vehicles in airports". *Computers and Operations Research* 71:34–53. cited By 2.

Pisinger, D., and S. Ropke. 2010. "Large neighborhood search". In *Handbook of metaheuristics*, 399–419. Springer.

Schultz, M. 2018. "Fast aircraft turnaround enabled by reliable passenger boarding". *Aerospace* 5(1). cited By 1.

Van Leeuwen, P., and C. Witteveen. 2009. "Temporal decoupling and determining resource needs of autonomous agents in the airport turnaround process". Volume 2, 185–192. cited By 1.

Wu, C.-L. 2008. "Monitoring aircraft turnaround operations - Framework development, application and implications for airline operations". *Transportation Planning and Technology* 31(2):215–228. cited By 8.

AUTHOR BIOGRAPHIES

MAURIZIO TOMASELLA is a Chancellor's Fellow and Lecturer in Management Science at the UEBS. Prior to that, he was Research Associate at the Institute for Manufacturing of 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.

ALEXANDRA CLARE is currently a Manager, Risk Analytics/Modeling at Charles Schwab. During the time of this study, she was an MSc student in Business Analytics at the UEBS. The work presented in these pages was partially developed as her Masters dissertation. Her email address is alix.clare2@gmail.com.

YAGMUR SIMGE GÖK is a Doctoral Student at the UEBS. Yagmur's work focuses on the developement of hybrid simulation-optimization techniques applied to airport ground operations. At the time of writing, Yagmur is working on a hybrid Simheuristic/Constraint Programming technique for GSE planning and operations. Her email address is Yagmur.Gok@ed.ac.uk.

DANIEL GUIMARANS is Research Fellow in the Data Science and Artificial Intelligence Group of the Faculty of IT at Monash University. Prior to that, he was Senior Researcher - Lecturer at the Aviation Academy of the Amsterdam University of Applied Sciences. Daniel works on the hybridisation of optimisation methods, and its combination with simulation techniques, looking at problems in various areas, including logistics, manufacturing and aviation. His email address is Daniel.Guimarans@monash.edu.

CEMALETTIN OZTÜRK is Senior Research Scientist at United Technologies Research Center, and is based in Cork, Ireland. Prior to that he was a Post-Doctoral Researcher at the Insight Centre for Data Analytics, also in Cork, and Lecturer at Izmir University of Economics. Cemalettin is enthusiastic about industrial applications of combinatorial optimisation and simulation in various domains, including manufacturing and aviation. His email address is OzturkC@utrc.utc.com.