Proceedings of the 2024 Winter Simulation Conference
H. Lam, E. Azar, D. Batur, S. Gao, W. Xie, S. R. Hunter, and M. D. Rossetti, eds.

SIMHEURISTICS FOR STRATEGIC WORKFORCE PLANNING AT A BUSY AIRPORT

Johanna Wiesflecker[1], Maurizio Tomasella[1], and Thomas W. Archibald[1]

[1]Business School, The University of Edinburgh, 29 Buccleuch Place Edinburgh, EH8 9JS, UK

ABSTRACT

Airport demand frequently exceeds capacity. An airport's capacity bottleneck is often located in its runway system. Other times, it is elsewhere, including the terminal facilities. At one of the busiest UK airports, the subject of this study, the bottleneck is located in the passenger security hall, where staff are employed directly by the airport operator. This airport plans to restructure the current workforce's overall size and types of contracts. Our work proposes a Simheuristic approach that utilizes current workforce data and demand forecasts for the upcoming season to adjust the contractual configurations systematically. We employ simulation to test the expected costs (vs. flexibility needs) of the likely rosters that each contractual configuration will allow. Using the simulation results, the algorithm aims to identify the optimal contractual configuration to minimize costs while ensuring the adaptability required to address unforeseen changes in the flight schedule as each season is underway.

1 MOTIVATION

Capacity planning — including but not limited to workforce planning (Lusby et al. 2012) and slot declaration for future seasons (Zografos et al. 2017) — is constantly at the top of the agenda of most airport operators. Its centrality to airport planning and operation has been extensively reported and studied (Jacquillat and Odoni 2015). In many cases, airport demand exceeds capacity (Jacquillat and Odoni 2018). Oftentimes, the airport's runway system is the bottleneck to achieving higher throughput. Other times, including at one of the Top-10 UK's busiest airports — the subject of our study — it occurs in their security hall (in yet different cases, the bottleneck might be the check-in hall, etc.). This particular airport is extremely keen to improve staffing issues and satisfaction with the used rosters. Furthermore, work demands are changing, with people wanting more flexibility to suit their lives outside of work — the literature has consistently reported on this trend for a while, see Laporte (1999). In addition, the UK's aviation sector is facing extensive security updates with upcoming changes to hand luggage regulations (Department for Transport 2022) — thus requiring new ways to run passenger security at airports. As a result, our partner airport operator is looking to overhaul their current rosters to suit better the needs of its employees (especially in the security hall) whilst accommodating growth and keeping staffing costs at its security hall under control (as this is the single biggest cost center at the airport).

The current method adopted for creating staff rosters and deciding suitable workforce sizes largely depends on the experience and expertise of the chief staff scheduler (CSS). They have been in this position for years and have an intrinsically profound knowledge of all relevant working regulations and the process for producing schedules. However, they are human and lack the computational power of modern modeling software (they use spreadsheets for the task). On a related note, as the CSS moves on to the next stage of their career at the same airport, someone else will have to take over their role, creating a clear case for effective knowledge management.

Currently, the strategic planning decisions for upcoming seasons start with budget decisions. The workforce size is set based on the available budget. Here, the CSS needs to argue whether or not the given budget is sufficient to employ a workforce large enough to cover all upcoming demand, staff holidays, unforeseen absences, unforeseen demand fluctuations, staff training, and much more. All in all, it is an enormous task to perform by hand and spreadsheet!

About 25 years ago, it was observed that algorithmic approaches to the construction of work schedules may be "often too constraining and not sufficiently flexible" (Laporte 1999, p. 1011), especially with regard to accommodating workers' preferences. The CSS at the partner airport still agree with this view today, adding that none of the costly software products the airport has tried over the years has ever provided the requisite flexibility to adapt some of the rules (which is virtually always required) in a way to make the schedule manageable and accepted by the unions who drive the negotiation of each and every roster. Laporte is also completely right that at the basis of staff rotas (*rotating workforce schedules*, such as those employed at the airport), there exists a fairly basic arithmetic. But the complexity of the current working regulations, the frequency with which these are modified, the variety of available contracts, and the size of the workforce (all major issues at the airport) make it impossible for anyone having to carry out workforce planning and scheduling tasks fully by hand to attempt any form of optimization. Additionally, unions require that the finalizations of the roster (consisting of frequent cycles of negotiations between the CSS, airport operator, and unions) for the upcoming season take place as early as possible into the start of the current season. However, the time required to produce a single complete roster manually is so high that it has become virtually impossible to accommodate the pace that the mandated timeline and frequent negotiation cycles require. Therefore, shortcuts are needed almost everywhere along the process. Currently, the CSS relies heavily on past schedules, the existing workforce size (and contractual setup), and their experience to estimate whether or not the budget will suffice for the upcoming season. In summary, the estimates depend on basic full-time equivalent (FTE) calculations and highly educated guesses to fill the gaps.

What is currently not possible is the strategic testing of different workforce sizes and contractual configurations to see which combination of workforce size per available contract will best cope with the uncertainty of the upcoming season. Contracts available to the airport's security staff differ in their work regulations (maximum weekly work hours, allowed shift types, etc.) and team structure (size of teams that follow the exact same roster for the entire season). To facilitate the rostering process (e.g., to devise a lower number of roster lines, thereby ensuring repeatability of the same cycle more times during the same season), the contracts with a larger workforce are split into multiple smaller contracts with the same work regulations and team size. The question we help answer in this paper is: "Given the demand forecast, what is the best contractual configuration with which to head into the next season?"

Based on Simheuristics (Juan et al. 2015), we recently developed a software tool to support the above-described decision-making process at our partner airport. At the time of writing, this is being field-tested directly by the CSS for future adoption. Our Simheuristic takes a starting workforce (e.g., the current contractual configuration) and the demand forecast for the upcoming season. It then adjusts the contracts to find the lowest FTE combination across contracts, one that best passes a series of robustness tests — these being carried out via simulation. Once the best FTE and accompanying contractual setup have been devised, the CSS and HR department will still need to recruit the required staff before the plan can be fully implemented. Between advertising, recruitment, and onboarding/training, a cycle of several months awaits. Our proposed tool is, therefore, a decision support system that helps the CSS test their workforce against the upcoming season, but it does not take over the entire job of strategic workforce planning. The latter will remain with HR and the airport's senior management team.

In the remainder, we will start with an analysis of the related literature on both strategic workforce planning and Simheuristics (Section 2). We will then present our Simheuristic approach (Section 3), delving into our approaches to both heuristic search and the way we help the algorithm move towards better solutions by integrating simulation. We then analyze (Section 4) results from an extensive experimental campaign and discuss the benefits and costs of adopting our approach before we wrap up the major learning points that we believe might apply outwith our partner airport and security hall (Section 5).

2 RELATED LITERATURE

2.1 Strategic Workforce Planning

Strategic workforce planning (SWP) studies workforce dynamics, including attrition, promotion, and retention, paired with recruitment decisions to cover the expected workforce demand for the upcoming seasons. In other words, it "examines the gap between staff availability ... and staff requirements ... over time, and prescribes courses of action to narrow such a gap" (Doumic et al. 2017, p.217). For a general introduction to workforce planning, we refer the reader to De Bruecker et al. (2015).

Oftentimes, studies of the phenomenology of workforce dynamics and decision-making quickly lead to some form of stochastic modeling. Given the above-mentioned aspects (staff availability, to cite but one), this is hardly surprising. The SWP literature broadly differentiates between approaches that integrate simulation (such as the current paper) and those that do not. Examples of the latter subset include, for instance, works grounded in stochastic programming, such as Jaillet et al. (2022).

Still, deterministic approaches remain popular. For example, Zeng et al. (2019) work on a similar problem to the one experienced by our partner airport. The authors' aim is to minimize their workforce mix while providing enough demand coverage vs. a given demand forecast. The formulation of their version of the problem, though, is centered around *hierarchical skills* such that staff with higher level skills are permitted to cover demand for staff with lower levels. However, in airport security halls, staff employed have roughly homogeneous skills.

As another such example, Llort et al. (2019) looks at SWP for management consulting firms. Similar to our partner airport's problem, workforce decisions made now will have long-term consequences, as "consultants are highly qualified workers who need very long learning periods to achieve enough expertise" (Llort et al. 2019, p.497). The paper proposes and validates a decision support tool for strategic staff planning, built upon mixed-inter linear programming, tackling decisions on attrition, recruitment, and promotion.

Turning our focus to approaches that indeed include simulation, the reader should look at Jnitova et al. (2017) for an extensive overview of simulation models for workforce planning in military settings.

Partially aligned with the problem studied in Llort et al. (2019) — but two decades earlier, model the workforce through a personnel flow diagram that encapsulates the modeling of transitions to different positions within the company (pay grade, etc.), attrition rates, and recruitment. Using a simulator, the authors then derive the steady-state flow of personnel, on the basis of which a range of SWP decisions can be supported (minimizing cost, maximizing personnel efficiency, etc.).

The incorporation of recruitment, attrition, promotion, retention, and training — or a selection thereof — in SWP studies is, by our very definition of SWP, a central, recurring aspect of most papers. However, it is generally hard to find works adding workforce scheduling to the mix. In particular, there appears to be a lack of papers looking at ways to ensure that staff rosters will remain feasible throughout the time horizon of interest within a context where the demographics of the workforce evolve as a result of the decisions being made over time. A recent exception is represented by Akl et al. (2022) and their simulation-optimization framework to optimize a maintenance schedule and decide the requisite workforce and skills to accomplish that schedule whilst accounting for backlog costs.

We propose a simple yet effective Monte Carlo simulation paired with a metaheuristic search scheme to test if the proposed workforce structure and contract configuration can handle the expected demand patterns whilst allowing for absences due to sickness or holidays. Decisions about how this structure can be attained (attrition, recruitment, promotion, retention, and training) can then be made jointly by the CSS and the HR team directly on the basis of the information generated as the output of our model.

2.2 Simheuristics

With reference to the taxonomy of Simulation-Optimization problems in Figueira and Almada-Lobo (2014), our problem falls into the category of "Optimization with Simulation-based iterations". Suitable methodologies to address such problems include the joint adoption of metaheuristics with simulation. For an overview of *Simheuristics* and their use, see Juan et al. (2023). We now look at how Simheuristics have been successfully applied to various scheduling problems.

Calvet et al. (2016) use SimILS to solve a Distributed Flow-shop Scheduling Problem. The iterated local search (ILS) strategically searches the solution space, and Monte Carlo simulation is used to evaluate the solution at each iteration. Somewhat similarly, Gök et al. (2020) use Large Neighborhood Search (LNS) paired with a Petri-Net-based discrete-event simulation to schedule airport turnaround staff. Building upon that, Gök et al. (2023) use two nested Simheuristics to produce robust schedules for the same problem. The 'outer loop' targets a resource-constrained project scheduling problem (RCPSP) with an enhanced *SimLoop* algorithm (Guimarans et al. 2015). At each iteration of this algorithm, the 'inner loop' solves a traveling salesman problem with time windows (TSPTW) through an enhanced version of *SimLNS* (Guimarans et al. 2015). In both loops, the metaheuristic is used to produce new solutions to the problem, while simulation is used to evaluate the solution's robustness.

Kızıloğlu and Sakallı (2023) propose two Simheuristics to simultaneously solve Flight scheduling, Fleet assignment, and Aircraft Routing problems. The metaheuristics involved are Simulated Annealing and Cuckoo Search, respectively. Solution evaluation happens through Monte Carlo simulation, which produces an expected objective function that is fed back to the metaheuristic part of the method for further iterations. Ultimately, the final solution is extensively tested through longer Monte Carlo Simulations.

Dehghanimohammadabadi et al. (2023) use a multi-objective metaheuristic, based on particle swarm optimization and paired with a discrete event simulation, to optimize patient appointment scheduling in a healthcare setting. The metaheuristic creates the patient schedule, consisting of patient arrival time and type. The simulation then runs through one full working day to test the schedule and measure patient dynamics in the system. The information is fed back to the metaheuristic so that the schedule may be improved. This follows the same ideas from Gök et al. (2020) and Gök et al. (2023). Building upon this evidence of prior success from this sort of approach, we adapted it to our SWP airport problem, by combining Adaptive Large Neighborhood Search (ALNS) (Ropke and Pisinger 2006) and Monte Carlo Simulation.

While not referred to as 'SimALNS', this combination has been used widely in recent years to improve the effectiveness of ALNS for solving problems under uncertainty. The simulation approach varies in the literature from fuzzy simulation (Zhou et al. 2020), over discrete event simulation (Kisialiou et al. 2018), to Monte Carlo Simulation (Nasri, Metrane, Hafidi, and Jamali 2020). Occasionally (e.g., Nasri et al. (2020)), simulation is used in the scheme in a slightly different way, to create realizations for which ALNS is then used to create solutions.

As the next section will show, in this paper, ALNS is used to alter the contractual setup for the rostering problem (the input values to the rostering algorithm rather than the roster itself), and simulation is used to test whether the resulting rosters are feasible under uncertainty.

3 METHOD

Our method for finding a good contractual combination for the upcoming season builds on existing scheduling technology at the airport. The assumption is that there exists a tool that can produce a feasible roster for a given workforce based on the input demand forecast of one model week that satisfies all working regulations set by unions and contracts (i.e., a deterministic solution to the rostering problem). For the purpose of this paper, the exact details and rules represented in the rostering problem are irrelevant as long as the contractual features that are considered in the ALNS procedure are also used as input parameters in the rostering tool (e.g., weekly hours for each contract). We refer to (Ernst et al. 2004; Van Den Bergh et al. 2013) for detailed reviews on scheduling literature. The approach aims to provide a way to find acceptable solutions for all problem stakeholders (CSS, unions, employees) rather than achieving optimality.

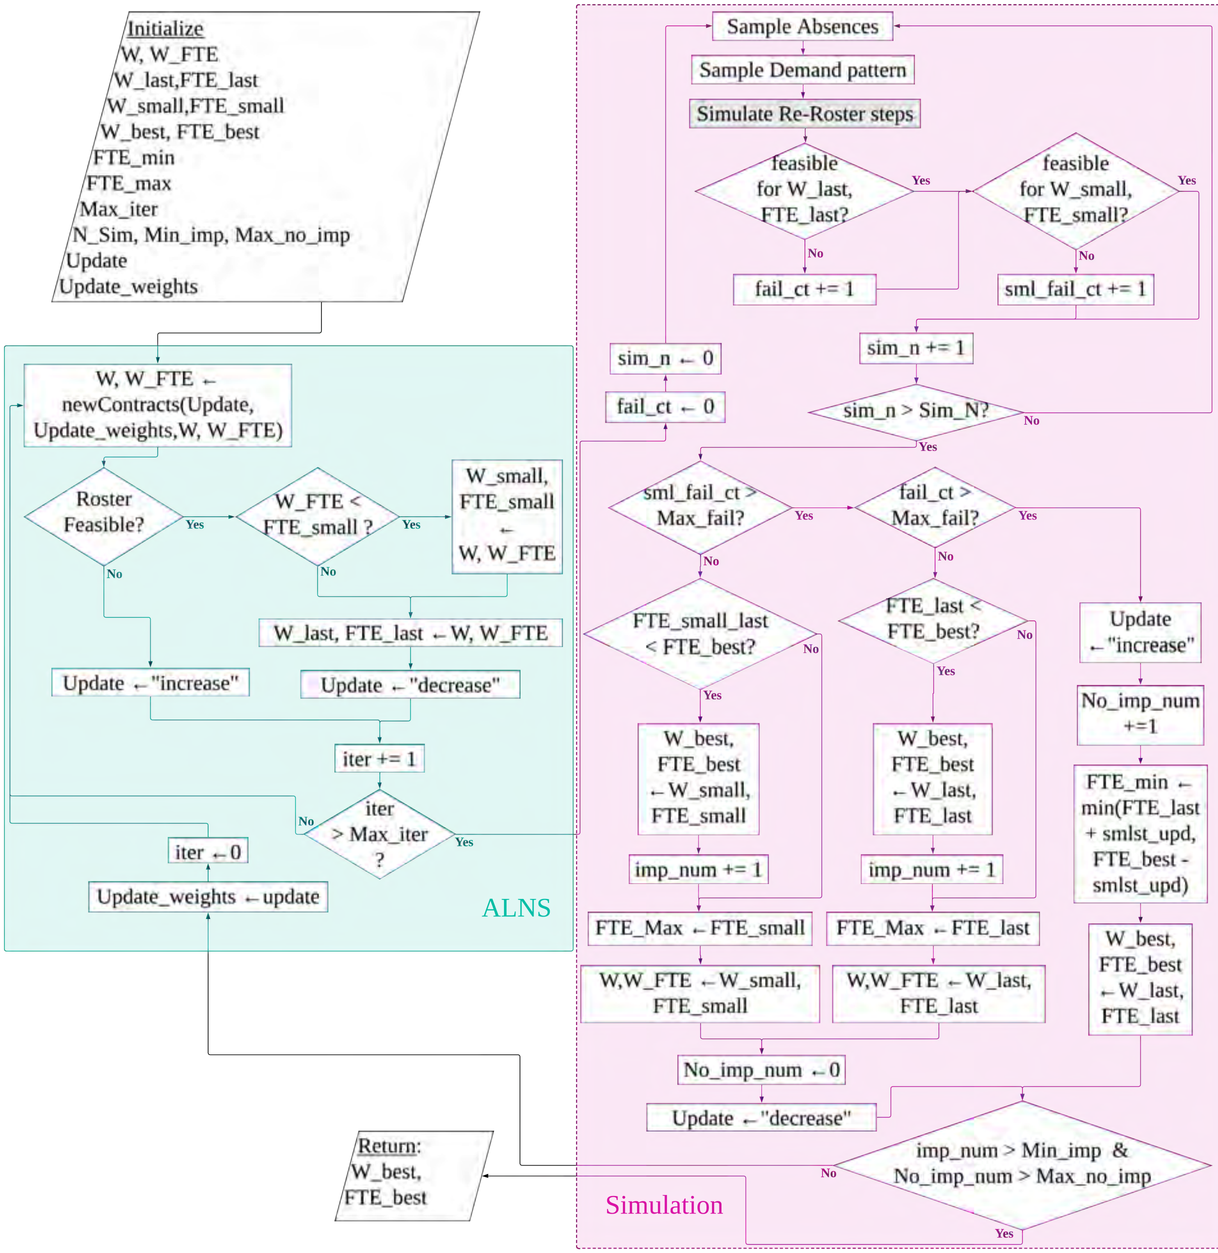

Figure 1: SimALNS for strategic workforce planning.

3.1 Adaptive Large Neighborhood Search (ALNS)

In our approach (Figure 1), ALNS aims to find the workforce with the lowest FTE that enables feasible rosters that also pass all robustness tests performed via simulation.

Starting from an initial feasible workforce (*W*) with associated FTE value (*W_FTE*), each ALNS iteration first updates the current workforce by either increasing or decreasing the associated FTE value through adding or removing (depending on the current setting of *Update_action* — see below) teams, workhours and contracts. The upper and lower limits of *FTE_max* and *FTE_min* are not strict bounds on the solution space but rather guidelines to help lead solutions back to the feasible solutions space. So, once an infeasible solution has been reached, the operators will increase/decrease the contracts until they are back within the limits before exploring further solutions.

Next, we apply the airport's existing roster tool to deterministically test whether a feasible roster (i.e. one that meets the demand forecasts) can be produced for the current workforce. If there is no feasible roster, *Update_action* is set to increase — based on the assumption that higher FTE values more likely lead to feasible rosters than lower FTE values might do. A maximum allowed computational budget (*Max_iter*) ensures the algorithm eventually moves on to the simulation stage.

If the roster is feasible, the associated value could be the lowest in the current ALNS run. In that case, the associated variables (*W_small* and *FTE_small*) are updated before updating the variables associated with the last feasible roster (*W_last* and *FTE_last*). Otherwise, only the latter step is performed. Finally, the *Update_action* is set to 'decrease' to try and improve the current solution further.

Once ALNS has run for *Max_iter* iterations, the lowest (*W_small* and *FTE_small*) and last feasible (*W_last* and *FTE_last*) solutions are sent to the Simulation to test if they are also feasible once the uncertainty inherent with the problem is accounted for.

3.2 Monte Carlo Simulation

At the simulation stage, we aim to replicate the Re-Rostering steps that take place throughout the season to account for and repair infeasibilities in the roster. Regardless of how good demand forecasts are, they can never account for everything that could happen throughout a season. The most common discrepancies happen due to staff absences and unforeseen changes to the demand patterns. To account for this, the simulation runs through multiple rounds of sampling demand patterns and staff absences and then checks if the roster can be repaired without breaking any working regulations. Each time the re-rostering steps fail, the method keeps track of this through a fail counter. This step is repeated for both contract setups that are fed to the Simulation from the ALNS run.

Once the simulations are complete, we first check if the lowest solution (*W_small*) satisfies the pass condition of the simulation. One round of simulation is considered as 'passed' if it is possible to repair the roster via re-rostering — after sampling absences and the new demand pattern — without violating any work regulations. A solution passes the simulation rounds if the re-rostering steps fail no more than the maximum number of fails (*Max_fail*). If it did pass, we check if the solution is better than the previous best solution *W_best* and update *W_best* accordingly before setting the FTE upper limit *FTE_max* to be equal to *FTE_small*. When comparing two feasible contractual combinations that both passed the simulation, we only consider the FTE value where a lower FTE is better (i.e., less expensive). This is due to two main reasons. Firstly, the FTE value is independent of the simulation result — it simply measures the available work hours per week. Secondly, if a contractual combination passes the simulation, we can assume that it will be able to handle the majority of demand patterns and unforeseen absences that might occur during the season. How good the pass was does not matter at this decision stage, as the real demand fluctuations and absences will be different anyway, and a lower FTE value is of higher importance here. The upper limit (*FTE_max*) helps ensure the ALNS operators keep bringing *W* back to feasible solutions throughout the iterations. Finally, the current solution (and starting point for the next round of ALNS) *W* is set to equal *W_small* and the *Update_action* is set to 'decrease' as ALNS will then start from a feasible solution to try and improve. If (*W_small*) does not pass the simulation tests, the same steps are repeated with the last feasible solution *W_last*. Should both *W_small* and *W_last* fail the simulation tests, the current solution *W* is set to equal *W_last*, but the *Update_action* will be 'increase' to try and find a feasible solution. In addition, the lower bound for the operators *FTE_min* will be updated as well. The SimALNS procedure will eventually stop once a minimum number of updates (*Min_imp*) of the best solution have been found and a maximum number of no improvement steps (*Max_no_imp*) have been reached.

4 EXPERIMENTS

Our Simheuristic approach was tested with data provided by the airport encompassing a demand pattern forecast — consisting of staffing requirements for every 15-minute interval of the day — for a 29-week season, paired with a matching contract set-up. The goal of this study was to investigate contractual combinations that can handle the most common disruptions (demand pattern changes throughout the season and staff absences). More detailed stress testing could be performed by adding matching demand patterns to the input forecast, which will then be sampled during the Monte-Carlo simulation or by performing additional tests on the final solution. We consider 15% unforeseen overall absences (based on the CSS's estimate) and 28 days paid leave per FTE employee per year (based on the UK statutory paid annual leave entitlement (UK Government 2024)). With regard to absences, our assumptions are as follows:

- Each individual has a 5% chance of being absent on any particular day
- This is increased to 15% if a team member was absent on the previous day
- Absence duration lies between 1-5 days, with decreasing probability ($\frac{5}{15}, \frac{4}{15}, \frac{3}{15}, \frac{2}{15}, \frac{1}{15}$ respectively)

The modeling of absences is based on rough estimates by the CSS because it was not possible to obtain the exact absence data due to security reasons. The absence patterns were tested to ensure an average absence percentage of 15% over the season. Similarly to the demand patterns, the absence simulation can be adapted to account for a higher or lower absence rate depending on the scenario that is being investigated. The allocation of holidays in our model works as follows:

- Each full-time employee is assigned 16 days off over the 29-week season ($16 \approx \frac{29 \text{ weeks}}{52 \text{ weeks}} \times 28$ days). All part-time employees (i.e. those working 'less than 40-hour' weeks) are assigned a number of days off that is proportional to their work hours (e.g. a 30-hour contract employee would get $16 \times \frac{30}{40} = 12$ days off for the season)
- Holidays are 1-5 days in length with equal probability and based on the availability of days off.
- There is a maximum number of employees that can be off on the same day. We set the maximum to be $2 \times \frac{\text{total number of holidays required for the entire workforce (days)}}{\text{total number of days in the season - assuming a 7-day week}}$.

The SimALNS parameters were set based on trial and error to suit this particular problem setting. However, more detailed methods for metaheuristic parameter tuning have been used successfully in the past Yadav and Tanksale (2023).

- Operator weight updates: +2 for improving the current solution, +4 for improving the best solution
- Decay parameter: 0.95
- Stopping conditions: minimum number of improvements (3) and maximum number of iterations without improvement (100) are both satisfied. These values were based on the outputs obtained after running the SimALNS approach for a time limit of 3 days.
- Number of simulation replications: 10 — This value was chosen to balance the runtime of the re-roster model that needs to be executed during each simulation replication and the need for replications to ensure that significant results can be obtained. Here, 10 replications cover roughly 34% of the demand forecast. In addition, we set three of the 10 replications to cover the three highest demand weeks of the forecast. Paired with a required pass rate of 80%, this ensures that each roster that passes the simulation will be able to handle most weeks during the season.
- The ALNS operators are defined to mimic allowed actions to the contractual set-up of a workforce. The main actions include adding/removing a contract or team and increasing/decreasing weekly hours by 5 hours. The actions are applied to a random contract or to the contract with the highest/lowest number of employees or weekly hours.

All experiments were run on a personal laptop (Apple M1 Pro) running macOS Sonoma 14.4.1. The SimALNS implementation was coded using Python 3.8 and calling Minizinc Version 2.8.3 (Nethercote et al. 2007) with Gurobi 9.5.1 (Gurobi 2024) for the rostering and re-rostering actions.

Figure 2 shows the progression of FTE values for all visited solutions (*W_FTE*), all feasible solutions (*last_FTE*), the lowest feasible solution for each metaheuristic run (*last_smallest_FTE*), and the overall best solution (*best_FTE*) for a 12-hour run of our SimALNS method. For the graph in Figure 2 (a), we randomly created a feasible starting contractual combination with an FTE value double that of the airport's current workforce. The image shows a steep decrease (FTE of 21240 hours down to 10395 hours) of the smallest feasible solution (*last_smallest_FTE*) during the first set of 20 ALNS iterations. This solution is shown to be feasible during the first set of simulations after 20 ALNS iterations (as shown by the start of the *best_FTE* line in orange). The spikes of the blue *W_FTE* and *last_FTE* lines show that as the iterations continue, the approach tests various solutions around the current best solution to avoid getting stuck in a local minimum. The best solution is continuously improved until it reaches an FTE value of 9775 (equivalent to ≈ 245 full-time employees) after 260 iterations. In summary, the metaheuristic first quickly tests a range of solutions and improves the smallest feasible solution with regard to creating an initial roster. Then, the computationally more expensive simulation only tests the best solutions as found

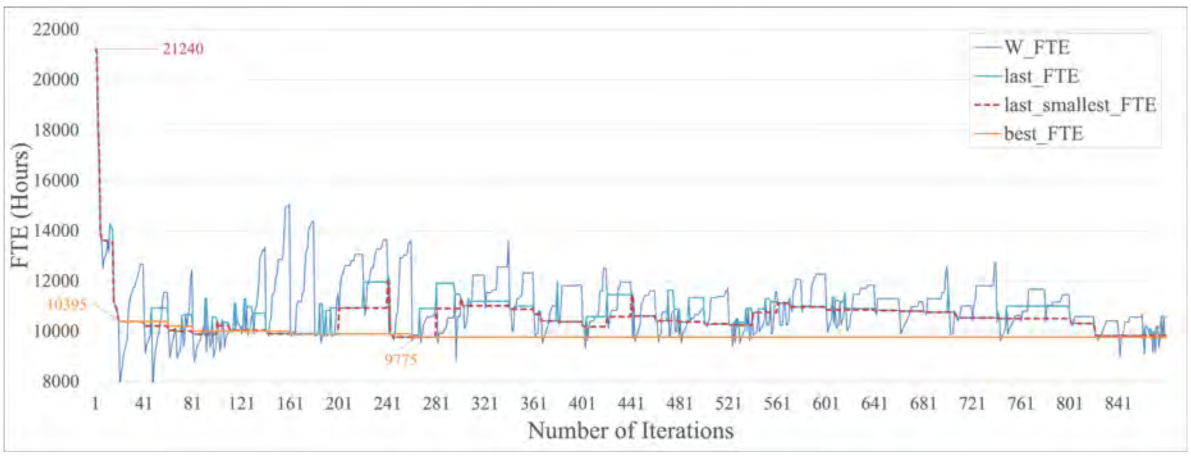

(a) Start with a contractual combination with double the airport's FTE value.

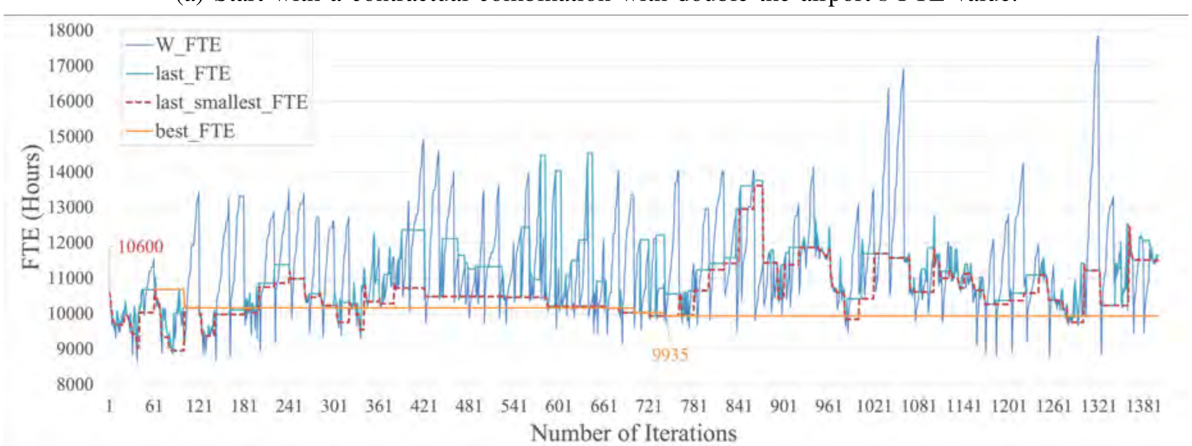

(b) Start with the airport's contractual combination.

Figure 2: FTE change over a 12-hour run of SimALNS.

by the metaheuristic to see if they also stay feasible under uncertain demand patterns, staff absences, and holidays. Figure 2 (b) shows the progression of FTE values while applying our SimALNS approach to the current contractual combination at the airport. Starting from an initial FTE of 10600 hours, we can see that the airport's current set-up is already relatively close to the best possible FTE value. Therefore, we can also not see a drastic decrease in FTE values at the beginning (as seen in Figure 2 (a)). We do, however, see an initial decrease in the best solution between iterations 40 and 120 and a further decrease between interactions 680 and 760. Between these improvements, the graph clearly shows how the method tries to escape the potential local minimum by consistently moving away from the current best FTE and then decreasing again (spikes in blue lines of *W_FTE* and *las_FTE*).

We also tested our Simheuristic by running five random starts for four different scenarios and different random seeds. For the scenarios, we started with a basic run (**basic** demand scenario), where the demand patterns match that of our partner airport, and we model absences and holidays as described at the beginning of the section. We then investigated a scenario with the same demand pattern but without the holiday allocation (**no holiday**), to see how much of the final FTE hours can be associated with holiday allocation. Finally, we tested our approach with two demand forecasts that differ in the number of security lanes needed throughout the season ($\pm$10%, scenarios **10% more** demand and **10% less** demand). For each of the runs, a random combination of contracts was created that satisfied a minimum FTE value (1.6 times the FTE value of the current workforce structure for the 10% more demand case and 1.3 for all other cases) and led to a feasible roster. The results of these runs can be found in Table 1.

Table 1: Ranked best FTE values (hours) for five random starts in various scenarios. Statistics with the basic outlier omitted are given in brackets.

demand pattern	FTE values ranked					Average	% change
	1^{st}	2^{nd}	3^{rd}	4^{th}	5^{th}		
basic	**9355**	10230	10320	10365	10405	10135 (10330)	N/A
no holiday	9835	9870	9945	10185	10235	10014	-1.19% (-3.06%)
10% less	9505	9665	9935	9945	10080	9826	- 3.05% (-4.88%)
10% more	10110	10255	10460	10650	11085	10512	+ 3.72% (+1.76%)

Table 2: Headcounts for different contract features and all runs with the **basic** demand pattern, where the grey column corresponds to the outlier run.

basic demand	1^{st}	2^{nd}	3^{rd}	4^{th}	5^{th}
FTE	9355	10230	10320	10365	10405
headcount - size 3 teams	45	81	45	57	99
headcount - size 5 teams	115	320	195	150	150
headcount - size 8 teams	**184**	0	48	104	88
headcount - 40h	60	0	48	88	0
headcount - 35h	55	10	240	31	158
headcount - 30h	45	81	0	192	80
headcount - 25h	0	250	0	0	99
headcount - 20h	**184**	60	0	0	0
total headcount	344	401	288	311	337

One very noticeable result is that of the best result for the **basic** demand pattern, which produces a significantly lower FTE result than any other run-scenario combination (we use 'scenario' and 'demand pattern' interchangeably). To see what might be causing this drastic difference in FTE value, we take a look at the headcounts for various contract features — Table 2 — of the results associated with the

basic demand pattern. We can see from the table that the majority of the workforce (184 out of 344) employees are assigned to a 20-hour contract with a team size of 8, which does not happen in any other contractual combinations. In addition, we performed additional Simulations (as outlined in Figure 1 to test the performance of this solution further and find a very low rejection rate of $\approx 7\%$. Given that no other results get close to the FTE value of this particular solution (even though the Simheuristic approach consistently trials contractual combinations in this FTE range, see Figure 2), we also present in brackets the results with the data of that run omitted (Table 1).

Looking more generally at the percentage change in FTE for the different demand patterns in Table 1, we find that not modeling holidays leads to an average 3.06% reduction in FTE (hours), suggesting that only just over 316 additional FTE hours ($\approx$ 8 people) are needed per week to cover the holiday requirements of the current workforce. Furthermore, we find that a 10% decrease in demand from the current scenario leads to a significantly higher change in FTE compared to the scenario with 10% additional demand (4.88% decrease vs. 1.76% increase). This can be linked to the set-up of the security hall at our partner airport where all security lanes are paired into so-called 'cells' — two adjacent lanes form a cell — by sharing the costly body scanner between the two. As a result, opening a new cell requires more staff than manning the second lane in an already open cell (or, vice versa, closing a lane in a cell reduces the workforce by less than closing down a cell altogether). Paired with the fact that demand is measured in the number of lanes that are required to cope with the flow of passengers, it is plausible that the discrepancy between demand changes and required FTE hours is mostly due to the cell layout of the security hall.

Table 3: Correlation coefficients for best FTE Hour variation and the number of contracts, teams, and employees in various categories (team size, weekly hours) for the starting and best solution.

categories	contracts		teams		employees	
	Start	**Best**	**Start**	**Best**	**Start**	**Best**
count	/	0.32	0.50	0.43	0.61	0.37
team size 3	0.19	0.42	0.46	0.42	0.46	0.42
team size 5	0.40	0.13	0.00	0.19	0.00	0.19
team size 8	0.17	-0.07	0.19	-0.09	0.19	-0.09
40h	/	0.47	0.50	-0.13	0.61	-0.09
35h	/	-0.03	/	0.08	/	0.05
30h	/	-0.05	/	0.02	/	0.03
25h	/	0.11	/	0.14	/	0.13
20h	/	-0.12	/	0.01	/	-0.13

To better understand the relationship between the solution's contractual features (e.g., the number of employees, teams, and contracts with varying team sizes and weekly hours), we examine the correlation coefficients between these features and the final FTE values. Table 3 displays the correlation coefficients for the various contractual features (number of contracts, teams, and employees for different team sizes and weekly hours) between the starting contractual combination and the final results (i.e. best FTE). A backslash indicates that the contractual features did not appear (20h-35h variations) or had all equal values (number of contracts) in the starting solution, making it impossible to calculate a correlation coefficient in these cases. Overall, the lack of correlation coefficients in Table 3 above magnitude 0.68 indicates that there is no strong correlation between contractual combinations (in the starting or final solutions) and final FTE values. From this, we can draw two main conclusions. Firstly, the starting contractual combination does not affect the final result — a highly important conclusion for our approach as we do not want the final result to depend on the manual input decided by the CSS. Secondly, the final contractual combination also does not appear to affect the final FTE value, meaning that a contractual combination that performs well in one season might not do so in the following season. Connecting this back to the current approach at our partner airport, this result suggests that their approach of rolling old schedules and contracts over to the next season (as a starting combination) may not be the best approach to finding a good contractual combination for the upcoming season. From an HR perspective, the approach makes sense, as it is significantly easier to hire staff for existing contracts that have been approved by union representatives and are tried and tested with the existing workforce. Both are important factors in keeping retention rates high. However, our approach suggests that changes might be beneficial to the performance and cost efficiency of the final workforce to adapt the contracts from season to season — within the possibilities of changing existing contracts, for example, by changing team size rather than weekly hours of contracts.

5 CONCLUSION

Coming back to our partner airport, we can draw conclusions about their work and suggest future changes.

Our results indicate a lack of correlation between contractual combinations and low FTE values. This means that a workforce that works for one season might not be the best choice for the next season. In reality, it will not be possible to revamp the entire contractual set-up of a workforce from season to season, but the CSS should at least try to validate the acceptability of some of the best solutions (perhaps those closer to the current roster) by the workforce and their representative unions.

One interesting result is that of the first-ranked run for the basic demand pattern. It gives the lowest FTE value by a large margin, and the contractual combination is the only one with a high number of 20-hour contracts. From conversations with the airport CSS, we know that they prefer to have lots of employees on these low-hour contracts because they provide more flexibility during the scheduling process, thus making their job much easier. However, high numbers of these contracts are rarely approved by the union representatives — making these contractual combinations infeasible in real life.

REFERENCES

Akl, A. M., S. El Sawah, R. K. Chakrabortty, and H. H. Turan. 2022, 3. “A Joint Optimization of Strategic Workforce Planning and Preventive Maintenance Scheduling: A Simulation–Optimization Approach”. *Reliability Engineering and System Safety* 219:108175.

Calvet, L., A. A. Juan, V. Fernandez-Viagas, and J. M. Framinan. 2016. “Combining simulation with metaheuristics in distributed scheduling problems with stochastic processing times”. In *2016 Winter Simulation Conference (WSC)*, 2347–2357 https://doi.org/10.1109/WSC.2016.7822275.

De Bruecker, P., J. Van Den Bergh, J. Beliën, and E. Demeulemeester. 2015, 5. “Workforce planning incorporating skills: State of the art”. *European Journal of Operational Research* 243(1):1–16.

Dehghanimohammadabadi, M., M. Rezaeiahari, and J. Seif. 2023, 4. “Multi-Objective Patient Appointment Scheduling Framework (MO-PASS): a data-table input simulation–optimization approach”. *Simulation* 99(4):363–383.

Department for Transport 2022. “Passengers to benefit from biggest shake-up of airport security rules in decades”. https://www.gov.uk/government/news/passengers-to-benefit-from-biggest-shake-up-of-airport-security-rules-in-decades, accessed 28th March.

Doumic, M., B. Perthame, E. Ribes, D. Salort and N. Toubiana. 2017, 10. “Toward an integrated workforce planning framework using structured equations”. *European Journal of Operational Research* 262(1):217–230.

Ernst, A. T., H. Jiang, M. Krishnamoorthy, and D. Sier. 2004, 2. “Staff scheduling and rostering: A review of applications, methods and models”. *European Journal of Operational Research* 153:3–27.

Figueira, G. and B. Almada-Lobo. 2014, 8. “Hybrid simulation–optimization methods: A taxonomy and discussion”. *Simulation Modelling Practice and Theory* 46:118–134.

Gök, Y. S., S. Padrón, M. Tomasella, D. Guimarans and C. Ozturk. 2023. “Constraint-based robust planning and scheduling of airport apron operations through simheuristics”. *Annals of Operations Research* 320:795–830.

Gök, Y. S., M. Tomasella, D. Guimarans, and C. Ozturk. 2020. “A Simheuristic Approach for Robust Scheduling of Airport Turnaround Teams”. In *2020 Winter Simulation Conference (WSC)*, 1410–1414 https://doi.org/10.1109/WSC48552.2020.9383947.

Guimarans, D., P. Arias, and M. M. Mota. 2015. “Large Neighbourhood Search and Simulation for Disruption Management in the Airline Industry”. In *Applied Simulation and Optimization*, 169–201. Springer International Publishing.

Gurobi 2024. “Gurobi software”. https://www.gurobi.com, accessed 2nd March.

Jacquillat, A. and A. R. Odoni. 2015. “An integrated scheduling and operations approach to airport congestion mitigation”. *Operations Research* 63(6):1390–1410.

Jacquillat, A. and A. R. Odoni. 2018. "A roadmap toward airport demand and capacity management". *Transportation research. Part A, Policy and practice* 114:168–185.

Jaillet, P., G. G. Loke, and M. Sim. 2022, 3. "Strategic Workforce Planning Under Uncertainty". *Operations Research* 70(2):1042–1065.

Jnitova, V., S. Elsawah, and M. Ryan. 2017, 10. "Review of simulation models in military workforce planning and management context". *Journal of Defense Modeling and Simulation* 14(4):447–463.

Juan, A. A., J. Faulin, S. E. 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.

Juan, A. A., P. Keenan, R. Martí, S. Mcgarraghy, J. Panadero, P. Carroll *et al*. 2023. "A review of the role of heuristics in stochastic optimisation: from metaheuristics to learnheuristics". *Annals of Operations Research* 320:831–861.

Kisialiou, Y., I. Gribkovskaia, and G. Laporte. 2018, 5. "Robust supply vessel routing and scheduling". *Transportation Research Part C: Emerging Technologies* 90:366–378.

Kızıloğlu, K. and Ü. S. Sakallı. 2023, 12. "Integrating Flight Scheduling, Fleet Assignment, and Aircraft Routing Problems with Codesharing Agreements under Stochastic Environment". *Aerospace* 10(12).

Laporte, G. 1999. "The art and science of designing rotating schedules". *Journal of the Operational Research Society* 50(10):1011–1017.

Llort, N., A. Lusa, C. Martínez-Costa, and M. Mateo. 2019, 6. "A decision support system and a mathematical model for strategic workforce planning in consultancies". *Flexible Services and Manufacturing Journal* 31(2):497–523.

Lusby, R., A. Dohn, T. M. Range, and J. Larsen. 2012. "A column generation-based heuristic for rostering with work patterns". *Journal of the Operational Research Society* 63(2):261–277.

Nasri, M., A. Metrane, I. Hafidi, and A. Jamali. 2020, 12. "A robust approach for solving a vehicle routing problem with time windows with uncertain service and travel times". *International Journal of Industrial Engineering Computations* 11(1):1–16.

Nethercote, N., P. J. Stuckey, R. Becket, S. Brand, G. J. Duck and G. Tack. 2007. "MiniZinc: Towards a standard CP modelling language". *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 4741 LNCS:529–543.

Ropke, S. and D. Pisinger. 2006, 11. "An Adaptive Large Neighborhood Search Heuristic for the Pickup and Delivery Problem with Time Windows". *Transportation Science* 40(4):455–472.

UK Government 2024. "Holiday Entitlement". https://www.gov.uk/holiday-entitlement-rights, accessed 2nd April.

Van Den Bergh, J., J. Beliën, P. De Bruecker, E. Demeulemeester and L. De Boeck. 2013, 5. "Personnel scheduling: A literature review". *European Journal of Operational Research* 226(3):367–385.

Yadav, N. and A. Tanksale. 2023, 6. "A multi-objective approach for reducing Patient's inconvenience in a generalized home healthcare delivery setup". *Expert Systems with Applications* 219.

Zeng, L., M. Zhao, and Y. Liu. 2019, 4. "Airport ground workforce planning with hierarchical skills: a new formulation and branch-and-price approach". *Annals of Operations Research* 275(1):245–258.

Zhou, Y., J.-J. Yang, and Z. Huang. 2020. "Automatic design of scheduling policies for dynamic flexible job shop scheduling via surrogate-assisted cooperative co-evolution genetic programming". *International Journal of Production Research* 58(9):2561–2580.

Zografos, K. G., M. A. Madas, and K. N. Androutsopoulos. 2017. "Increasing airport capacity utilisation through optimum slot scheduling: review of current developments and identification of future needs". *Journal of Scheduling* 20:3–24.

AUTHOR BIOGRAPHIES

JOHANNA WIESFLECKER is a final-year PhD student in the Management Science and Business Economics (MSBE) subject group at the Edinburgh University Business School. As part of her research, she applies Simulation-Optimization to workforce planning problems in airport security. Her email address is johanna.wiesflecker@ed.ac.uk.

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 and Politecnico di Milano, Maurizio's work employs Simulation-Optimization and Multi-Criteria Decision Analysis techniques to address applications in airport operations management. His e-mail address is maurizio.tomasella@ed.ac.uk.

THOMAS W. ARCHIBALD is Professor of Business Modelling at the University of Edinburgh Business School and former Head of the MSBE subject group. His research focuses on applications of stochastic modelling to problems in business and management. His email address is: t.archibald@ed.ac.uk.