Optimization of Regulated Airline Arrival Flows via Target Time Management

An Enhanced Slot Swapping Process

Leonardo Caranti, Marie Carré,
Lucas Weber, Dominique Heilmann
Operations Research & Air Traffic Management
Swiss International Airlines Ltd. (SWISS)
Zurich, Switzerland

Richard Stevens
Network Manager, Air Traffic Flow & Capacity
EUROCONTROL
Brussels, Belgium

Abstract*—Growing Air Traffic Flow Management (ATFM) delays emerge within the European Airspace. Solutions to mitigate their impact are researched and implemented by all airspace stakeholders. Target Time Management (TTM) is a proactive form of arrival sequence steering via Calculated Take-Off Times which could improve operational performance. This process is researched at SWISS, with an optimizer which outputs a wished arrival sequence every 60 minutes for Zurich arrival regulated flights. This sequence is sent to EUROCONTROL's slot allocation algorithm (CASA), which takes into account the wished Target Times of Arrival. The optimization's goal for the summer was to improve passenger connections without hampering rotation delays. Over the course of around two months, critical connections were improved by a weighted average of 6.2 minutes, without worsening rotation delays (a negligible improvement of 0.9 minutes was measured). A total of 3,879 flights and 72,730 connecting passengers were affected by the tool during this time. This proves the usefulness of TTM as a tool to improve operations under ATFM regulations. Yet, insights on the behaviour of CASA should be further modelled within the decision-making process. For instance, it was found that requesting a delay typically results in higher acceptance rates than asking for an anticipation (81.6% versus 47.1%). Moreover, it also takes more time to achieve the requested Target Time for anticipated flights (mean of 3.5 minutes versus 14 minutes for delayed flights). This yields interesting points for further work in light of other SESAR 3 supported projects such as HARMONIC.

***Keywords*—Air Traffic Management, Regulations, Air Traffic Flow Management, Demand-Capacity Balancing, Optimization, Mixed-Integer Linear Programming, Airline Operations, Delays, Passenger Connections, Curfew, SWISS International Airlines, EUROCONTROL.**

I. Introduction

In Europe's saturated airspace, a fine balance between demand and capacity rules operational performance. Air Traffic Flow Management (ATFM) delay ensures that, in cases where demand exceeds capacity, safety is still met. The operations of SWISS International Airlines also routinely take a toll due to ATFM delay, often due to the layout of Zurich airport, and otherwise due weather, noise and strikes. In 2019 alone, SWISS was delayed a total of around 4500 hours when only looking at ATFM arrival regulations, as shown in Figure 1. This translates into approximately 12 hours of delay per day. Weather, aerodrome capacity and environmental issues (such as noise constraints), accounted for up to 94% of all ATFM delay for arrival regulations in Zurich airport in the period between 2019 and 2022.

The consequences of ATFM delay, or any kind of delay, are significant. Missed passenger connections, missed departure slots, delayed rotations, missed crew connections, curfew fees and ultimately cancelled flights, taunt all airlines and airspace users. Due to the unpredictable nature of these delays, it is necessary to have solutions for navigating through operationally tough days.

Approaches to mitigate the effects of delay range from reactive ones to proactive ones. The former attempt to decrease the effects of delay once they present themselves - for example, rebooking of passengers, ramp direct services, and increased flight speeds. On the other hand, proactive approaches attempt to act before the delays heavily affect the operations.

Target Time Management is a concept developed by EUROCONTROL which allows for proactive steering of ATFM regulated flights (its applications outlined in [1]). The Target Time Management System is a process which allows airspace users to file ATFM delay change requests to steer their operations in a more desirable way. This process is designed to be fair, transparent, yet still allow for great optimization potential. Its structure is based on the idea that airspace users should be able to request a wished Target Time of Arrival, which EUROCONTROL's slot assignment algorithm can then grant them, if available.

The goal of this paper is to investigate whether TTM is a viable solution to improve operational performance for SWISS's short-haul fleet in case of arrival slot delays. This is done via an operational live trial. In some papers, it was proven that the Target Time Management was a viable solution for long-haul flights [2], while in others, it was proven with shadow mode trials that similar approaches (slot swapping)

SESAR Innovation Days
27-30 November 2023
Seville

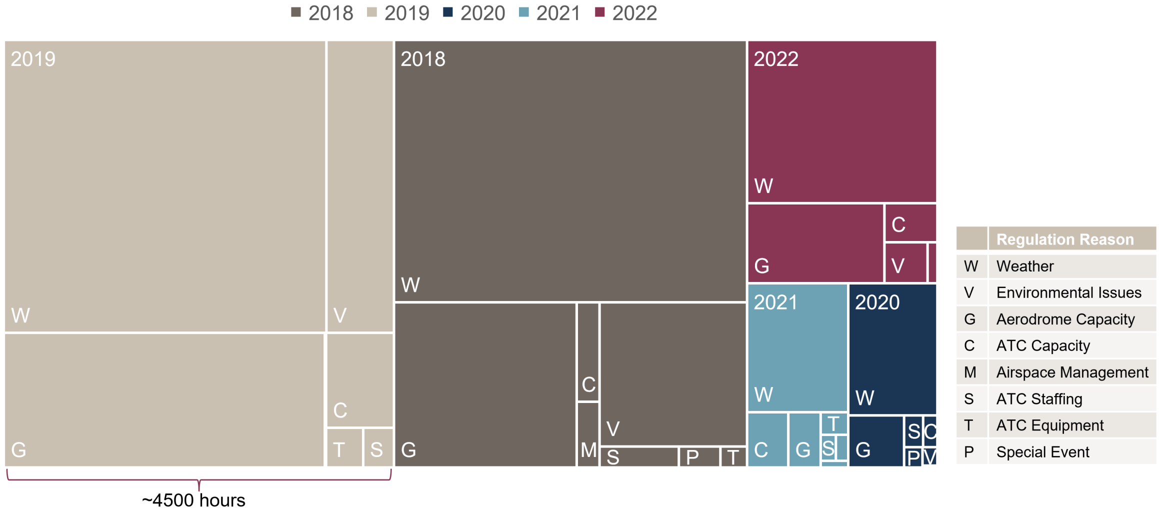

Figure 1. ATFM Delay for arrival regulations in Zurich by year and regulation type.

can improve short-haul performance [3]. This research becomes even more pertinent as SESAR's new project HARMONIC will attempt to investigate this further [4]. As such, a documented operational live trial, even though from the perspective of only one airline, will lay foundations for the HARMONIC project to flourish. HARMONIC is supported by the SESAR 3 Joint Undertaking and its founding members.

It is expected that with this research attempt, SWISS will be able to improve its operational performance by using Target Time Management. It is expected to do so by mainly improving passenger connections, and also rotation delays, if possible. This is due to the business objectives of the summer, which prioritize passenger connections over rotation delays.

The concept of TTM is further described in detail in section II, along with other important attempts to solve similar problems. After this, the optimization process and the operational procedure are better outlined in section III. This is followed by section IV, where the results are outlined, as well as section V, where they are discussed together with recommendations for future work.

II. Literature

TTM, developed by EUROCONTROL, is a concept only narrowly explored within the Air Traffic Management literature branch [1]. It is offered in aid to airlines and Air Navigation Service Providers (ANSPs), which allows for requests to influence the Air Traffic Flow Management delay according to specific needs. Airlines and ANSPs may request a Target Time of Arrival (TTA) for each flight, such that EUROCONTROL can then update the current ATFM delay (by updating the Calculated Take Off Time, CTOT). For airlines, the advantage is to be able to reduce the delay of crucial flights in order to improve operational performance, for instance by improving passenger connections, rotation delay, curfew, crew connections.

In theory, if this method is misused it could lead to exploitation: if an airline would request an improvement on all its flights, the delays of other airlines would drastically increase. This is why EUROCONTROL has only allowed the usage of this tool as a pseudo- slot swapping tool: when an airline requests an improvement, it should match this with a worsened slot on another aircraft, such that the total delay for each airline remains the same.

An example of what a TTA sequence swap could look like is shown in Figure 2. Here, six flights are re-sequenced based on their original Estimated Times of Arrival (ETA). By doing this, it is possible for the airline to maintain the total delay as before the swap. Moreover, this also allows to maintain a minimal amount of other perturbed traffic due to system effects. Both of these elements contribute to maintaining fairness within the system. To also ensure safety in a swap-mechanism, an assessment was made together with SkyGuide and EUROCONTROL, with confirmed system safety.

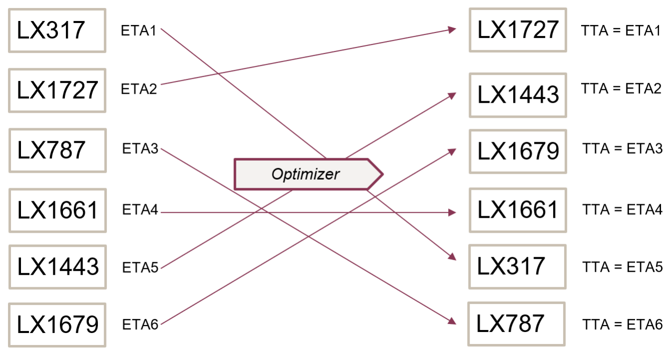

Figure 2. Example of a fictitious TTA swap.

TTM, being a relatively new method, has very limited published research. There are accounts of Heathrow Airport, Charles de Gaulle airport and three airports on the Eastern coast of Spain testing such an approach with local ANSPs within the context of SESAR's Network Collaborative Management Project 24, but no further publications are found [5]. In most cases, only the perspective of an airline sending wished TTAs to a centralized system which then sends them to EUROCONTROL was considered. The authors have not been able to find accounts of TTM being used from purely an airline's point of view on short-haul flights.

In the other hand, there is one example of an operational live trial and system, also developed within SESAR, which uses Target Times of Arrival for long-haul flights. This is the iStream project, which optimized the sequence of early morning long-haul arrivals to avoid unnecessary holdings due to very early arrivals [2]. In this process, airlines submitted estimated arrival times, and received indications to adjust the speed and flight profile accordingly. The results were operationally effective, reducing holdings in the early morning flights, and it was adopted as a standard procedure in many airports, including Zurich. This project, even though very different to the one described in this paper for short-haul flights, is proof of the feasibility of an operational tool based on TTM.

Fortunately from a literary point of view, TTM greatly resembles a more commonly used service: Slot Swapping, which is often regarded as one of the most common methods of reducing reactionary delay [6]. This method swaps the arrival and departure slots in cases where there is a regulation due to demand-capacity imbalance [7]. This is exactly how the approach in the User-Driven Prioritization Process (UDPP) project was taken: here, the airline was able to specify the required swaps to optimize their operations [3]. This project, also part of SESAR's framework, made great steps forward with regard to understanding the limits of an automatic slot-swapping method by an airline. This also was validated with a live, shadow mode trial at SWISS.

Other attempts at Slot Swapping must be mentioned, due to their similarity to TTM. In a research building upon UDPP, ground constraints were added to the slot swapping problem [8]. In a different work, 4D trajectory optimization was attempted to enhance the solution space [9]. Moreover, within the SESAR project SlotMachine [10], researchers were able also to take the perspective of the ANSP, while using airline preferences as an input [11].

No research has been carried out on TTM on short-haul flights, and no operational live trial has been documented regarding automated slot swapping from the perspective of an airline. Although there are similarities between Slot Swapping and TTM, the differences (i.e. a wider window for swapping, increased number of flights to swap around for the latter), warrant research focused on TTM. Papers which investigate Slot Swapping often mention a series of recommendations: (1) to consider swaps beyond pairs of flights, (2) take into account airline-specific business goals. This paper covers precisely this research gap: we investigate the use of a TTM System, from an airline perspective, considering the dynamics between multiple flights.

III. Methodology

The methodology chosen to assign Target Times of Arrival (TTA) to flights was a linear optimization process, to be run automatically every 60 minutes. An optimal result is defined as one which improves passenger connections and rotation delays in the best way possible, for the next wave of flights. Since often there is a tradeoff between these two parameters, a preferred balance can be set for the day by users in the Network Operations Control, if wanted. Once the optimum is found, the result is sent to EUROCONTROL's Computer Assisted Slot Allocation (CASA) algorithm, which then reassigns a new departure slot for the requested Target Time of Arrivals. The methodology has been divided into two perspectives: a procedural one (explained in subsection III-A) and a more technical one referring to the optimization model (subsection III-B).

A. Procedure

From a procedural perspective, it is necessary to include Skyguide (the Swiss ANSP) in the information exchange between SWISS and EUROCONTROL. This is to ensure safety, traceability, and fairness in the changes to the ATFM delay distribution. Moreover, the separation distances in the Terminal Maneuvering Area must be kept.

A visual depiction of the various stages of the process is outlined in Figure 3. Here, one can see how the decision to set a specific TTA is set when the flight is on the way to the outstation, for the flight coming back. The TTA is computed from SWISS side, and is sent to EUROCONTROL via an interface hosted by Skyguide. This procedure can be repeated as many times as required, until the start of the Departure Planning Information (DPI) sequence. By ensuring that after this point, the slot is not affected anymore, it is possible to avoid changing pilot procedures since they would see this as a normal slot change.

B. Model

The decisions taken to affect the process are based on a Mixed-Integer Linear Programming model. This takes into account the business requirements of SWISS airlines, as well as the operational constraints of the Air Traffic Management infrastructure.

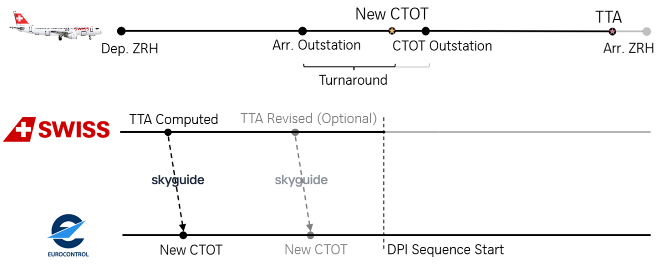

Figure 3. Visualization of the TTA Process. The horizontal lines represent time, with timestamps described on the first from the top.

TABLE I. PARAMETER DEFINITIONS

Parameter	*Definition*
c_{ij}	Criticality of flight i in TTA slot j.
F	Set of flights, indexed with i.
S	Set of TTA slots, indexed with j. Note: $\|F\| = \|S\|$.
M	Chosen large number, $M >> c_{ij}$.
T_i	Scheduled time of arrival of flight i.
t_j	Hours since start of day for slot j.
a	Chosen maximum anticipation limit.
TaT_i	Minimum turnaround time required by flight i.
d	Chosen maximum delay limit.
e_i	Most constraining time (STD and/or ETD) for flight i.
cut	Curfew limit time.

1) Decision Variables: The decision variables of the problem, x_{ij}, represent a binary choice based on whether a flight i is assigned to a specific slot j, as shown in Equation 1. Given that there are only a maximum of 50 or 60 flights for which a decision has to be taken at the same time (given SWISS's fleet size), this formulation still has few enough variables to be solved in very little time ($\leq 1s$) using open source Python solvers [12].

$$x_{ij} = \begin{cases} 1, & \text{if flight } i \text{ is assigned to slot } j \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

2) Objective Function: The objective of the model is to minimize the overall criticality of the flight assignments to slots. This is the responsibility of the first part of Equation 2. Input vector c_{ij} evaluates business parameters to see how critical each flight i is in a slot j. For SWISS, a combination of passenger connecting times, rotation delays, and curfew performance was taken into account. A specific balance between the importance of connecting times and rotation delays can be inputted by the user (with the use of a slider in the front end) to tackle the operational challenges of the day. The second part of Equation 2 ensures that when all flights are equally critical, these are prioritized in terms of minimising the difference to the schedule.

$$\min z = \sum_{i\in F}\sum_{j\in S} \underbrace{x_{ij}\cdot c_{ij}}_{\text{Criticality}} + \underbrace{x_{ij}\cdot \frac{|i-j|}{M}}_{\text{Diff. to Schedule}} \tag{2}$$

The parameters used in Equation 2, as well as in the constraints to follow, are outlined in Table I together with a small explanation.

The criticality parameter, which steers the whole optimization procedure, can be split into four parts. The two most important are the passenger connections and the rotation delay. These evaluate how critical (relative to historical operational performance) a certain flight is. They look several rotations ahead, to see which connections and which rotations will be affected, and output a criticality score. It must also be noted that not all connections are deemed equally important: (1) long haul are given a heavier weight than short-haul ones, and (2) business or first-class are given more weight than economy class. The weights are chosen based on the business benefits that SWISS derives from the various types of passengers. Additionally, the tool also considers two other aspects: the current delay and the number of passengers. These divide flights when all have equal rotation delay and passenger connection criticality. The logic applied is: the higher the delay, and the more passengers a flight has, the higher the criticality score should be.

3) Constraints: The constraints allow for a feasible operational result. Their individual description can be better observed from Table II, yet they can be split up into two groups: Equations 3 and 4 are mathematical constraints, while the rest are more of operational nature. Equations 5 and 8 are written individually to allow direct cause-effect evaluations in case of infeasible results, and the possibility to delete individual constraints.

$$\sum_{j\in S} x_{ij} = 1, \qquad \forall\, i \in F \tag{3}$$

$$\sum_{i\in F} x_{ij} = 1, \qquad \forall\, j \in S \tag{4}$$

$$T_i - \sum_{j\in S} x_{ij} t_j \leq \min\left(a,\ TaT_i\right), \qquad \forall\, i \in F \tag{5}$$

$$-T_i + \sum_{j\in S} x_{ij} t_j \leq d, \qquad \forall\, i \in F \tag{6}$$

$$\sum_{j\in S} x_{ij} t_j \geq e_i, \qquad \forall\, i \in F \tag{7}$$

$$\sum_{j\in S} x_{ij} t_j \leq cut, \qquad \forall\, i \in F \tag{8}$$

IV. RESULTS

The TTM tool was on trial twice: beginning of June (starting on the 6th) and end of June. It then went live officially on July, 1^{st} 2023. The results are based on the trials as well as the operational performance until the 14th of September, when the results were recorded. For all regulated flights for Zurich Arrivals most penalizing regulations, a TTA was sent once per hour. Here, most penalizing regulations entail that the ATFM regulation which allocates the greatest delay to the flight.

TABLE II. CONSTRAINTS EXPLANATION

Constraint	*Explanation*
Equation 3	Each flight must have one TTA slot assigned.
Equation 4	Each TTA slot must be assigned to one flight.
Equation 5	Each flight cannot be anticipated more than the most constraining time between minimum turnaround and a maximum anticipation time.
Equation 6	Each flight cannot be delayed more than the maximum delay time.
Equation 7	Each flight cannot be anticipated more than the most constraining time between STD and ETD if the flight has one.
Equation 8	Each flight cannot be delayed beyond the limit curfew time.

Given that the priority for the summer was passenger performance, the team in charge of operations (i.e. Network Operations Control) advised to shift the balance between delay and passenger connections more towards the latter. This is why, while the tool was not able to significantly improve rotations delay (only by 0.9 minutes on average), it greatly improved passenger connections (up to 6.2 minutes weighted average). As can be seen from Figure 4 and Figure 5, for most days there is a significant improvement for critical passenger connections (only connections with less than 45 minutes are considered in this plot). This comparison was achieved by comparing the final take-off times of the flights (based on the final CTOTs) to the initial ones (initial slots), and looking at how passenger connecting times changed. Weighted connections by business value are considered, because this is also part of the objective of the optimizer, ultimately bringing up business performance. Also unweighted connections are shown for reference in Figure 5, yet the analysis will focus more on weighted connections (Figure 4) due to their similarity to the operational goals at SWISS. The weights were chosen based on the estimated monetary value that each passenger brings to the airline, classified by class and business segment (intercontinental versus continental).

The days with significant improvements in passenger connecting times are mainly due to three different factors. First of all, by manual inspection it was found that with higher total ATFM delay, it was possible to improve the connections more, due to the baseline being worse. Moreover, when there were many flights to swap around with, it was also possible to achieve better results. Finally, the tool performed the best in cases where the arrival sequence was initially inefficiently reordered due to delays, either with respect to the schedule or with respect to planned passenger connections.

It can be noted how there are some dips below zero in Figure 4. This is unfortunately due to the calculation method of improvements of connecting times in a highly dynamic slot allocation environment. The initial status quo is always taken as a reference, and it is assumed that any change to the slots is due to Target Time Management itself. However, this assumption is not valid in two instances: (1)

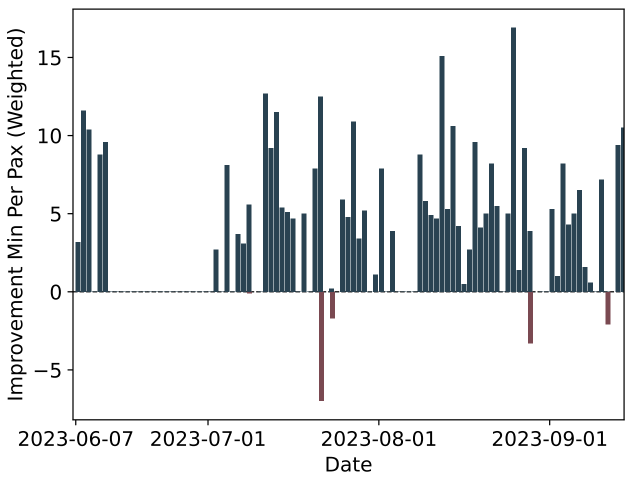

Figure 4. Weighted passenger connection improvement due to the TTM trial. Only passenger connections under 45 minutes have been considered in this plot.

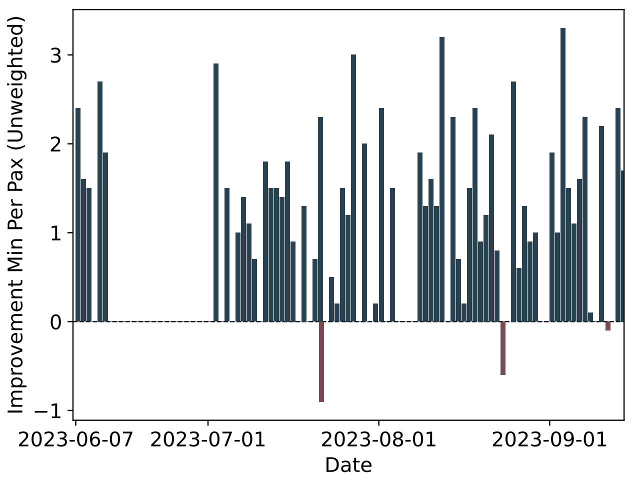

Figure 5. Unweighted passenger connection improvement due to the TTM trial. Only passenger connections under 45 minutes have been considered in this plot.

in the case of slot swapping or slot improvements, (2) when slots delays worsen due to external circumstances, such as worsening weather. The former may be caused by controllers in the Operations Control Center aiming to improve specific regulated flights. However, since this happens for (on average) three flights per regulation, and would only affect around 3.1% of the flights with TTAs, it has no great impact in the analysis. The latter, however, is damaging as it has a direct effect on passenger connection. Slot delays, not Target Time Management, led to the negative dips in Figure 4.

Overall, the average improvement is of 6.2 minutes per weighted connection (weighted by contribution to the business of SWISS), and 1.5 min per unweighted connection. These numbers can also be seen together with other key performance indications in Table III. This was achieved by barely affecting rotation delays: a nearly negligible improvement of 0.9 minutes was measured. It can be seen that in this period, nearly 4,000 flights were affected by a total of nearly 9,000 TTAs, meaning that on average a flight receives 2.3 TTAs throughout its slot developments.

TABLE III. PERFORMANCE KPIS

KPI	*Value*
Flights with TTAs	3,879
Sent TTAs	8,916
Achieved TTAs (± 1 min)	7,026 (78.8%)
Conx. pax. considered	72,730
Conx improvement per pax. (weighted)	6.2 min
Conx improvement per pax. (unweighted)	1.5 min
Rotation delay improvement	0.9 min

It is also interesting to see how the algorithm performs with different regulation types and intensities. From Figure 6 it can be seen how there are two slight trends: when the median delay per flight is higher, the performance of the algorithm tends to increase. Similarly, when there are more regulated flights (aircraft) considered, the performance is also higher. This can be attributed to two factors: first of all, the more flights there are, the more chances there are of finding optimal swaps. Moreover, as the delay worsens so does the baseline that the system is to improve from, hence creating more improvement opportunities for the algorithm.

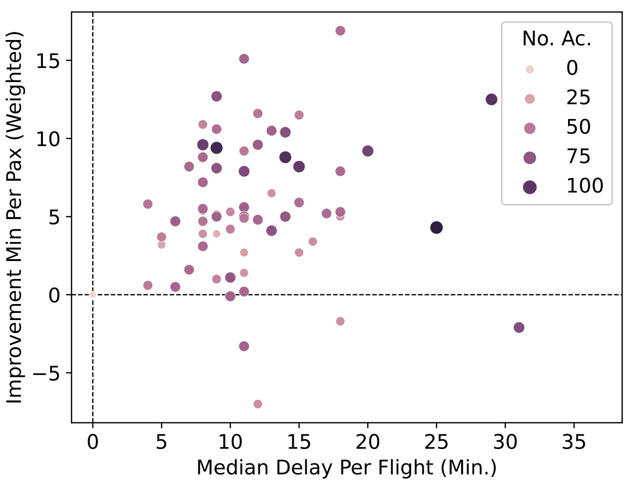

Figure 6. Weighted passenger connection improvements for connections under 45 minutes, as a function of median delay per flight and number of aircraft in the regulation. Here, *No. Ac.* is the number of aircraft considered by the tool in one day.

An interesting metric can be observed in Figure 7: the TTA acceptance rate. This is defined as accepted if and only if sometime in the future, before the next TTA is sent, the slot falls within ± 1 minute of the requested TTA. There is a weak

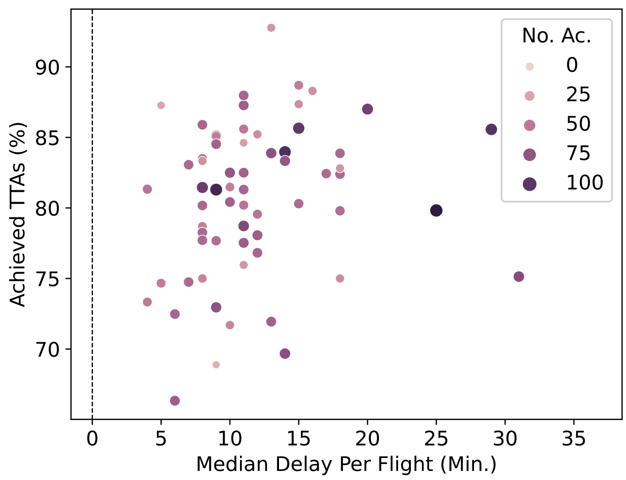

Figure 7. Percentage of achieved Target Times of Arrival within ± 1 min, as a function of median delay per flight and number of aircraft in the regulation.

linear relationship between the regulation severity (median delay per flight) and the percentage of achieved TTAs. And an even weaker relationship (if present at all) between the latter and the number of flights in a regulation.

V. DISCUSSION & FUTURE WORK

The goal of this research was to investigate whether or not Target Time Management could provide a useful operational outcome. The system improved passenger connections without hampering rotation delay performance. The system thus offers operational benefits for an airline: if fewer passengers are delayed for their outbound flight, there will be fewer outbound delayed flights which have had to wait for their connections. Moreover, it also offers advantages from an air traffic control perspective: the optimization objective of the tool can be modified to prioritize rotation delays or curfew performance over connections.

However, there are drawbacks and possible improvements for future research. As shown in Table IV, there is a higher chance that a TTA gets accepted if it is a delay rather than an anticipation request. This may possibly be because there are other flights with higher priority of other airlines which have to be shifted around before the requested flight can achieve its TTA.

TABLE IV. CASA RESPONSE BEHAVIOUR

Parameter	**Delayed Flights**	**Anticipated Flights**
Mean time to reach TTA	3.47 min	14.02 min
Median time to reach TTA	0.55 min	1.54 min
TTA overshoot	14.34%	1.01%

One notable finding was that the slot assignment algorithm by EUROCONTROL (i.e. CASA) would behave differently depending on the TTA requested. More specifically, there are three different types of requests than one can do with TTAs: an anticipation, a delay, and a request to not change the current slot. It was found that for anticipated flights only 47.1% of flights would be accepted, while for delayed flights up to 81.6%. Finally, for flights which asked to remain with the current slot, the acceptance rate was 92.4%. The reasons behind this may be due to the higher priorities of flights which were filed previously to the one with a TTA, as well as the fact that it is significantly faster and easier, in optimization terms, for CASA to accommodate delay requests rather than anticipations.

The fact that delayed flights tend to be more delayed than requested and anticipated flights tend to not achieve their goal has consequences. SWISS tends to have a higher proportion of slot delay than when it is not using Target Time Management. In turn, this means that other airlines are able to take advantage of SWISS's increased delays (assuming that the regulation rate does not change in the process). Figure 8 plots the mean ATFM delay for SWISS and non-SWISS flights under Zurich Arrival regulations. This is computed for both days in which TTAs were sent, as well as past days where the TTA system was not active yet. In general, there is no significant difference between the delays. However, the share of delay for SWISS flights approximately doubled, whilst for the other airlines it roughly shrunk by two times. Based on this plot, SWISS always has more delays than other airlines. Nevertheless, this plot is analysed per unit of time, and SWISS is likely to have more flights than other carriers due to the hub-and-spoke nature of their operations in Zurich.

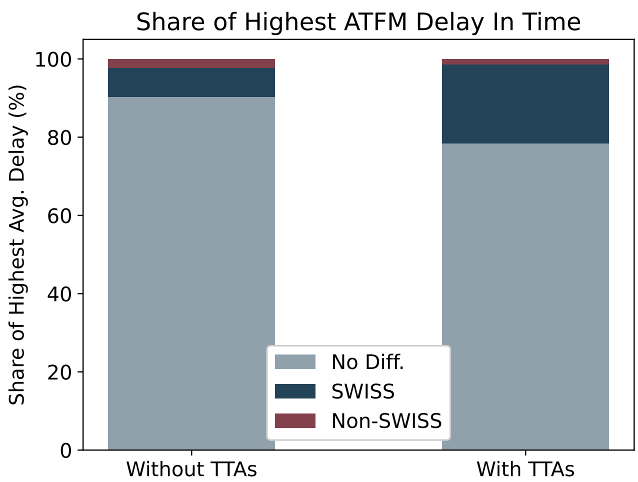

Figure 8. Share of SWISS versus non-SWISS highest average (mean) ATFM delay within Zurich Arrivals regulations over time. Difference found with 2-tailed, 2-σ t-test.

Another CASA behaviour which is reflected in the numbers shown in Table IV derives from the very large difference between the median and the mean times to reach the requested TTAs. This, often referred to as skew, shows that flights can be split in two main groups: one large set of flights which takes very little time to reach the TTA, and one smaller group which takes a very long time. For an optimal use of the system in future research, one would have to recompute TTAs for the latter group after a time threshold, given that the initial TTA is probably either unachievable or unrealistic. For the current usage and dynamics, the authors would recommend a threshold of around 30 minutes. This both ensures enough time for CASA to react while also making sure that the system is not overloaded with information.

Another important number to take note of from Table IV is overshoot. A flight is said to have overshot when its current delay (or anticipation) exceeds the one it requested for. For instance, this could be that a flight asked for 20 minutes of delay but received 30 instead. The reverse is also true: if a flight asks to be anticipated by 20 minutes, but instead is anticipated by 30, it would count within this metric of overshoot. It appears that for delayed flights, in around 14.3% of the cases, overshoot is present, while when asking for anticipations, the likelihood that this happens is as low as 1%. This again creates an imbalance between delayed and anticipated flights due to the combination of CASA behaviour and system effects.

A. *Future Work*

Even though this operational trial yielded a series of positive outcomes, its limitations offer great opportunities for future work. For instance, now that CASA dynamics are better known, it would be interesting to test with asymmetrical TTAs, where the "swap" constraint of the TTA requests is dropped, but the total delay of the airline is still attempted to remain the same. Safety and fairness tests would have to be carried out to ensure that other airlines are not disadvantaged, but the potential of this approach in improving operations is significantly high. Additionally, the algorithm should model CASA dynamics within the request itself: knowing that requesting anticipation has a lower chance of being accepted, it should be reflected in the decisions taken. This, for instance, could be done by modelling CASA as an environment (with techniques such as Reinforcement Learning), instead of assuming that the current slot is static. However, this idea is prone to more constraints on CASA side when dropping swap-style TTAs, hence the achievement rate should be well-monitored and compared.

Finally, other objectives besides passenger connections may be considered. For instance, curfew performance, crew connections, gate allocation, and ground capacity. These inputs should be added as individual objectives, and the system should decide the best trade-off between them, rather than having to specify a priori that passenger connections are the priority. This could be achieved by better financial modelling within the objective function. Moreover, the inclusion of more operational constraints would benefit the system, such as gate allocation and ground constraints. It is expected that by prioritizing the arrival flow into an order which is more operationally valuable, the pressure on the other flight's operations of the airline is reduced. This would also be interesting to test in future research.

VI. Conclusion

The objective of this paper was to investigate Target Time Management as an approach to improve operational performance of an airline's short-haul fleet during ATFM arrival regulations. In the last five years, these have often occurred with a frequency of more than once a week, sometimes up to once a day. Hence, it is of utmost importance for SWISS to mitigate their effects. A case study was carried out, and it was found that it is possible to improve passenger connections and rotation delays, with a trade-off made between the two.

A tool is developed using a Mixed-Integer Linear Programming optimization model, which recomputes solutions automatically every 60 minutes. The testing process lasted from the 6^{th} of June until the 1^{st} of July, when the tool became operational. The data is analysed until the 14^{th} of September. In the trial and during summer operations it was chosen to prioritize passengers: a (weighted) average connecting time improvement of 6.2 minutes is found for all connections below 45 minutes, without increasing rotation delays. This statistic is built on 8916 sent Target Times of Arrival (TTAs), 3,879 flights and 72,730 connecting passengers. The goal is achieved - TTM can significantly improve operational performance for short-haul flights.

Nevertheless, the project also found a number of points of improvement. Firstly, CASA's behaviour on TTA requests has been statistically analysed, and should be modelled in the process of computing the optimal arrival sequence. This is because, due to the hierarchical nature of the slot list, requesting delays results in higher acceptance rates than asking for an anticipations (81.6% versus 47.1%). Moreover, it takes significantly less time to achieve a TTA if the flight was delayed rather than if the flight was anticipated (mean of 3.5 minutes versus 14). Because of this, it would be interesting to explore this topic further with stochastic methods which model the dynamic slot environment. Methods such as Reinforcement Learning would prove to be very interesting for this problem, since they allow for decision-making under highly non-linear, stochastic environments. The authors also recommend to take into account further goals for the optimization such as curfew performance and crew connections.

This paper also lays part of the foundations for HARMONIC, a project which aims to further optimize the regulated streams, from a multi-stakeholder, multi-objective point of view. HARMONIC is supported by the SESAR 3 Joint Undertaking, its founding members and the Swiss State Secretariat for Education, Research and Innovation. This research proves that TTM is a possible solution to optimize arrival streams, and could be expanded to include more stakeholders, objectives and constraints. Examples of these could be gate allocation, ground handling and separation requirements. This will further mitigate the consequences of ATFM regulations, to the benefit of all airspace users. Finally, the effect of multiple TTAs from different airlines should be considered and analysed for future research within the HARMONIC project.

Acknowledgment

The authors would like to thank Stefan Fluck, the SWISS Operations Research & Air Traffic Management team, Anaïs Lacroix, Robert Rudiger, Yannick Carbonneaux, Michael Brügger, Mattes Kettner, and Marta Ribeiro for their innate dedication, hard work and contributions to this project and paper.

References

[1] EUROCONTROL, "ATFCM users manual." [Online]. Available: https://www.eurocontrol.int/publication/atfcm-users-manual

[2] DSNA, SkyGuide, AirFrance, SWISS, Lufthansa, HOP!, ZurichAirport, EUROCONTROL, and AereoportsdeParis, "SESAR joint undertaking | integrated SESAR TRials for enhanced arrival management (iStream)," p. 148. [Online]. Available: https://www.sesarju.eu/node/1551

[3] N. Pilon, L. Guichard, Z. Bazso, G. Murgese, and M. Carré, "User-driven prioritisation process (UDPP) from advanced experimental to pre-operational validation environment," vol. 97, p. 102124. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S096969972100106X

[4] SESAR. SESAR joint undertaking | HARMONIC- harmonised network through smart technology and collaboration. [Online]. Available: https://www.sesarju.eu/projects/HARMONIC

[5] P. Hop, C. Neculae, J. Valle Martinez, M. Garrido, and M. Carré, "Demonstration report PJ24 network collaborative management," p. 639.

[6] B. Campanelli, P. Fleurquin, V. Eguíluz, J. J. Ramasco, A. Arranz, I. Extebarria, and C. Ciruelos, "Modeling reactionary delays in the european air transport network."

[7] J. Schummer and A. Abizada, "Incentives in landing slot problems," vol. 170, pp. 29–55. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S002205311730042X

[8] J. Evler, M. Schultz, and H. Fricke, "Flight prioritization and turnaround recovery - integrating tactical ATFCM slot swapping into resource-constrained ground operations."

[9] L. Delgado, G. Gurtner, P. Mazzarisi, S. Zaoli, D. Valput, A. Cook, and F. Lillo, "Network-wide assessment of ATM mechanisms using an agent-based model," vol. 95, p. 102108. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0969699721000910

[10] SESAR, "SESAR joint undertaking \ SlotMachine - a privacy-preserving marketplace for slot management." [Online]. Available: https://www.sesarju.eu/projects/SlotMachine

[11] C. G. Schuetz, T. Lorünser, S. Jaburek, K. Schuetz, F. Wohner, R. Karl, and E. Gringinger, "A distributed architecture for privacy-preserving optimization using genetic algorithms and multi-party computation," in *Cooperative Information Systems*, ser. Lecture Notes in Computer Science, M. Sellami, P. Ceravolo, H. A. Reijers, W. Gaaloul, and H. Panetto, Eds. Springer International Publishing, pp. 168–185.

[12] L. Perron and V. Furnon, "OR-tools." [Online]. Available: https://developers.google.com/optimization/