

Multi-stakeholder Optimized Arrival Management with the Target Time Management System (TTMS)

Operational Evidence from a Live Trial at Zurich Airport

Laura-Marie Hagemann¹, Marie Carré¹, Michael Brügger², Gregor Sarrazin³, Anaïs Lacroix⁴

¹Operations Research & Air Traffic Management, Swiss International Airlines Ltd. (SWISS), Zurich, Switzerland

²OPF – Performance & Systems Development, Flughafen Zürich AG, Zurich, Switzerland

³Zühlke Engineering AG, Zurich, Switzerland

⁴Capacity Management, Skyguide – Swiss Air Navigation Services Ltd., Geneva, Switzerland

Abstract—We report the first multi-stakeholder, live test of the Target Time Management System (TTMS) for arrival sequencing under Air Traffic Flow Management (ATFM) regulations. TTMS integrates airline, airport, and air navigation service provider (ANSP) priorities by transforming airline cost functions into privacy-preserving weight maps and optimizing a single Target Time of Arrival (TTA) sequence per regulation. In an eight-week trial at Zurich Airport (51 regulations; 3,372 flights, including 2,078 operated by SWISS), TTMS improved SWISS arrival punctuality by 20.2% (+234 on-time flights) and reduced heavy delays (30+ minutes) by 23.5% (−103), while total allocated delay stayed within a band around a Computer-Assisted Slot Allocation (CASA) baseline. Airport-wide, total delay minutes fell by 11% and missed connections by 14%; for SWISS, missed connections decreased by 33.3% and critical connections by 19.8%. Equity analysis found no systematic disadvantage for non-participating airlines (Theil index $T = 0.03072$). We outline operational limitations (low declared rates, short lead times, current Most Penalizing Regulation handling) and the next phases (Q4 2025 extended trial; 2026 year-long evaluation) toward a scalable European rollout.

Keywords—Air Traffic Flow Management (ATFM), Target Time of Arrival (TTA), Target Time Management System (TTMS), arrival sequencing, multi-stakeholder optimization, equity, CASA, SESAR HARMONIC.

I. INTRODUCTION

Air traffic across Europe is again approaching pre-pandemic levels, yet infrastructure, staffing, and environmental constraints limit capacity growth at major hubs such as Zurich. To manage excess inbound demand, Air Traffic Flow Management (ATFM) regulations are frequently applied, resulting in substantial ground delays. In 2019, SWISS alone accumulated about 4,500 hours of ATFM arrival delay at Zurich, with knock-on effects including missed connections, crew disruptions, and curfew violations.

Traditional delay allocation follows the “first planned, first served” (FPFS) principle, which assigns slots according to filed flight plans. While transparent and widely accepted, FPFS does not distinguish between flights with high and low operational impact, often leading to inefficient outcomes and network-wide disruption. Acceptance of FPFS historically reflected the absence of viable alternatives when ATFM regulations were introduced.

The SESAR 3 Joint Undertaking (SESAR 3 JU) HARMONIC project introduces the Target Time Management System (TTMS), a collaborative, optimization-based framework for sequencing arrivals under ATFM regulations. TTMS builds on the ATOM project, where SWISS was the first airline to request Target Times of Arrival (TTAs) for regulated flights. ATOM improved connectivity by reordering slots based on passenger connections and rotation stability but was limited to a single stakeholder.

TTMS extends this approach by transforming flight-level delay cost functions into weight maps that quantify the marginal impact of delay. In this context, a *weight map* is a function that translates each flight’s expected delay cost into a normalized priority weight, thereby linking operational impact to sequencing priority in the optimization. These maps are combined with airport profiles and air navigation service provider (ANSP) priorities in a centralized optimizer, which produces a single TTA sequence maximizing overall benefit. Unlike ATOM, TTMS coordinates multiple stakeholders, enabling system-wide efficiency while preserving fairness.

This paper documents the first operational test of TTMS during a live trial at Zurich Airport (May–June 2025). We describe the generation of flight-specific weight maps, the integration of stakeholder objectives, and the optimizer’s performance under real-world ATFM regulations. Results show measurable improvements in punctuality, passenger protection, and inter-stakeholder coordination, supporting TTMS as a candidate for future European arrival management.

II. LITERATURE

Early European work on collaborative arrival planning demonstrated the feasibility and benefits of Target Time concepts, but without real-time, cost-driven, multi-stakeholder optimization. The iStream project tested locally generated Target Times at Paris–CDG and Zurich, mainly for long-haul arrivals, where higher schedule stability enabled accurate in-flight adjustments and improved pre-departure planning; at Zurich, SWISS reported a 96% reduction in airborne holding [1]. Trials with early-morning short-haul flights were less conclusive, relying on static KPIs rather than dynamic cost functions. In parallel, SESAR’s User-Driven Prioritisation



Process (UDPP) allowed airlines to reprioritize or swap slots within their own fleet during ATFM regulations [2]. While this preserved inter-airline fairness and avoided conflict, it did not support global optimization across all flights in a stream and remained effort-intensive and largely manual.

In the United States, Collaborative Decision Making (CDM) procedures organize arrival metering through Ground Delay Programs (GDPs). Demand is balanced via *ration-by-schedule* (RBS), followed by airline-driven substitution and cancellation actions and centralized compression to recover unused capacity, with controlled cross-carrier trading via Slot Credit Substitution (SCS). These mechanisms and their variants are well documented and have inspired several algorithmic extensions [3]–[6]. However, GDP-based CDM remains largely procedural and market-driven, is primarily departure-side, and typically does not internalize detailed airline or airport cost structures. TTMS is conceptually complementary: it focuses on arrival-side, system-wide optimization based on cost-derived weight maps from airlines, combined with airport and ANSP preferences into a single Target Time of Arrival (TTA) sequence per regulation.

To address transparency and confidentiality in the European context, the SESAR-supported SlotMachine project proposed a privacy-preserving slot management platform using secure multi-party computation and distributed ledgers [7]. Airlines could submit encoded preferences or cost profiles, enabling optimized sequencing and auction- or credit-based slot trading without revealing proprietary data. While conceptually close to TTMS in its treatment of privacy, SlotMachine remained exploratory, validated only at small scale and never deployed operationally.

Recent work has also applied Reinforcement Learning (RL) to ATFM. Chen et al. propose a cooperative multi-agent RL framework in which agents learn delay allocation dynamically through simulation [8]. RL-based approaches outperform heuristic baselines in minimizing total delay under synthetic schedules, but they typically omit airline cost structures, passenger priorities, or airport objectives. They thus remain network-centric rather than genuinely multi-stakeholder, and are best viewed as potential complements to cost- and preference-based optimization frameworks such as TTMS.

A key operational precursor for TTMS is the ATOM project [9], in which SWISS trialed airline-driven optimization at Zurich using a mixed-integer program to generate TTAs every 60 minutes for regulated short-haul flights. The optimizer prioritized passenger connections while maintaining rotation stability. Across 3,879 flights and 72,730 passengers, average connection time improved by 6.2 minutes (weighted by importance and misconnection risk¹), with negligible impact on rotations (+0.9 minutes). However, ATOM was limited to a single stakeholder: SWISS could only optimize allocation within the set of slots already assigned to its flights by Computer-Assisted Slot Allocation (CASA), without considering airport or ANSP inputs or enabling cross-airline optimization. Within

HARMONIC, this work has been extended in [10], which evaluates two TTMS architectures: (i) a trusted, airport-hosted platform using a deterministic assignment algorithm and (ii) a privacy-preserving variant using secure multi-party computation (MPC). For the latter, evolutionary methods such as NSGA-II and SPEA-2 are employed to meet strict runtime limits despite MPC overhead. Synthetic experiments show that both variants can balance airspace-user preferences with airport KPIs, and define realistic input–output structures and equity-weighting mechanisms for benchmarking heuristic optimizers.

Complementary approaches for managing delay and uncertainty during the execution phase exploit in-flight measures such as cruise speed reduction or “linear holding.” Delgado et al. show that modest, controlled cruise speed reductions can absorb delay efficiently while maintaining arrival predictability and reducing fuel burn [11]. Xu et al. extend this idea with a multi-objective optimization framework that balances delay, fuel consumption, and equity considerations through adaptive cruise speed control [12]. Such strategies can complement TTMS by providing additional flexibility once TTAs have been assigned, rather than replacing the need for ground-based arrival sequencing.

Beyond efficiency, several strands of work explicitly address equity in delay assignment and sequencing. Montlaur and Delgado analyze the efficiency–fairness trade-off in ATFM and quantify the “price of fairness” using flight- and passenger-centric metrics [13]. In the U.S. context, Glover and Ball develop stochastic optimization models for GDP planning that balance efficiency and equity objectives [14], while Vossen and Ball formalize a general framework for equity in Traffic Flow Management and propose methods to mitigate exemption bias in GDPs [15]. Churchill and Lovell formulate a coordinated network resource allocation model under uncertainty and demonstrate that system-wide coordination can improve both efficiency and responsiveness compared to independently managed resources [16]. At the network level, Bertsimas and Gupta [17] propose collaborative ATFM models that incorporate fairness among airlines (e.g., controlling overtaking and schedule reversals). More recent European work explores optimization that reduces both primary and reactionary delay while reconciling fairness-related constraints relative to CASA baselines [18], as well as priority-based approaches that manage delay efficiency while limiting displacement and regulating priority levels [19]. These contributions motivate the use of explicit equity diagnostics, including groupwise distributional comparisons and Theil indices, as first applied in ATM by [20].

In summary, prior projects have demonstrated feasibility (iStream), airline-internal reprioritization (UDPP), privacy-preserving coordination (SlotMachine), fairness-aware optimization, and promising AI methods (RL), while U.S. CDM/GDP research offers a mature departure-side collaboration paradigm (RBS, compression, SCS, and automated, privacy-aware slot swapping). TTMS closes a remaining gap by providing a real-time, operational framework that integrates

¹See [9] for details



airline, airport, and ANSP inputs into a single Target Time Management System and evaluates it with live data at Zurich.

III. METHODOLOGY

A. System Overview

The TTMS implemented in the HARMONIC trial is structured as a three-stage coordination process between the airspace users (with SWISS actively participating through the provision of flight-specific weight maps and other airlines passively participating via default profiles), the airport operator (Flughafen Zürich AG), and the air navigation service provider (Skyguide). The final arrival sequence is implemented by EUROCONTROL's Network Manager via the CASA system.




Figure 1. System architecture of the TTMS (adapted from [21])

Airlines, acting as airspace users, provide delay-sensitivity information in the form of weight maps, which capture the marginal cost of delay over time. SWISS computes detailed flight-specific maps, while Zurich Airport supplies standardized default maps for other carriers and its own airport-level preferences. All inputs are collected by the *orchestrator*, which harmonizes airline and airport preferences, triggers the optimizer, and formats the resulting Target Times of Arrival (TTAs) for transmission to EUROCONTROL's Network Manager in the format required by CASA. The centralized optimizer then determines the arrival sequence that maximizes total system benefit while respecting the constraints defined by the local ANSP (Skyguide).

The following subsections detail (i) how default airspace user and airport weight maps are constructed (Section III-B) and (ii) how SWISS generates flight-specific weight maps (Section III-C).

B. Airport and Default Airspace User Weight Maps

For airlines not submitting individual weight maps, Zurich Airport provides standardized default maps to ensure seamless integration of all flights into the optimization process. These default profiles are currently identical across all non-participating airlines and are not derived from historical behaviour or airline-specific preferences, but serve as neutral representations of delay sensitivity. Their construction is based on two components:

- **Scale factor:** For SWISS weight maps, the maximum weight observed within a given regulation is rescaled to 100,000, and all other SWISS weights are adjusted proportionally. Flights without airline-provided weights (i.e., other carriers or airport defaults) are not rescaled and use a scale factor of 1.
- **Priority factor:** The relative importance of a flight in the default map is expressed by a priority factor $p \in [0, 1]$, derived from operational key performance indicators (KPIs).

The maximum weight for a flight without airline-provided data is

$$W_{\max} = 80,000 + 20,000 \cdot p.$$

Thus, a flight with $p = 0$ receives $W_{\max} = 80,000$, while a highly critical flight with $p = 1$ reaches 100,000. Intermediate values provide finer differentiation between operational importance.

Default maps follow a quadratic decay function, initialized at $W_{\max}$ and decreasing smoothly from the estimated time of arrival (ETA) onward. If the ETA is earlier than the scheduled time of arrival (STA) minus taxi-in time, the decay is instead applied from the STA. This ensures consistent delay sensitivity while maintaining comparability with airline-provided maps.

Additionally, Zurich Airport contributes its own weight component reflecting operational constraints. At present, these are primarily based on the scheduled time of arrival (STA), available turnaround time, and available seats derived from the ICAO aircraft code, reflecting the limited data availability. In upcoming trial phases, the airport aims to expand these profiles by integrating additional operational data, such as booked passenger numbers, critical passenger connections, and parking stand constraints, to better capture the true cost impact of delays. These airport-defined weights are linearly combined with airline-provided (or default) maps using a fixed 50:50 ratio to guarantee that both perspectives are equally represented in the optimization.

Figure 2 illustrates how TTMS interprets different weight-map types across three consecutive arrival slots (A = earliest, C = latest). Flight X represents an airspace user submitting its own weight map; the steep decline across slots reflects higher delay sensitivity. Flight Y receives the airport's standardized default map, resulting in a smoother profile with lower overall weights. Flight Z illustrates a flight with operationally constrained or already-fixed arrival options: its feasible slot is given a very large weight (indicating a locked arrival time), while infeasible slots are marked with -1. The optimizer evaluates all flights across all slots and assigns each flight the slot with the highest feasible weight; in the example, these assigned TTAs are highlighted with a red box. Overall, the illustration shows how airline preferences, default behaviour, and hard operational constraints are all expressed in a unified weight-map format used by the optimizer.

Flight	Slot A	Slot B	Slot C
X	91'000	88'500	85'500
Y	88'500	86'800	85'000
Z	1'000'000	-1	-1

Figure 2. Illustrative example of airspace-user, default, and blocked-slot weight maps.

C. SWISS Flight-Specific Weight Maps

While standard weight maps are generated by the airport for all passively participating airspace users, SWISS implements a more granular approach by computing flight-specific cost functions to reflect the real-time operational and economic

impact of delay. These cost functions are constructed individually for each flight and evaluated over a defined time horizon surrounding its scheduled operation.

Specifically, cost evaluations are performed for all SWISS flights with a scheduled time of departure (STD) within the next 15 hours, ensuring sufficient lookahead to capture operational relevance. For each of these flights, potential arrival time slots are assessed from the estimated time of arrival (ETA) up to four hours beyond. This upper limit reflects the operational reality that flights delayed by more than four hours are, in most cases, no longer operationally viable. Cost maps are recomputed every 10 minutes, ensuring that the optimizer works with up-to-date priorities while still respecting these time-based boundaries.

This time-based logic allows the TTMS optimizer to operate with a forward-looking and dynamically updated understanding of flight priorities, while preserving operational stability and decision lead time for dispatch and network control.

The following describes the cost components developed by SWISS to model the impact of arrival delays on passengers, operations, and airport constraints, forming the basis for flight-specific weight map generation in TTMS.

1) *Passenger- and revenue-related costs*: Passenger impacts include (i) *churn risk*, representing long-term revenue losses from delay-induced customer dissatisfaction based on segment yields and SWISS internal churn probabilities; (ii) *connection and baggage costs*, capturing missed connections, reaccommodation, overnight stays, and mishandled baggage; (iii) *EC261 compensation* liabilities as a function of delay thresholds, sector length, and passenger count; and (iv) *F&B voucher costs*, triggered by waiting-time thresholds and distance-based service standards. These elements jointly describe how delay affects both immediate and longer-term passenger-related costs.

2) *Operational propagation costs*: Operational impacts are captured through (i) *speed-up costs*, derived from aircraft-specific performance curves and fuel burn for feasible acceleration margins; (ii) *crew duty limit violations*, which impose cancellation-level costs if projected arrival breaches legal Flight Duty Period limits including buffers; (iii) *rotation disruption*, where inbound delays propagate to subsequent aircraft turnarounds or cockpit crew transfers, allocating downstream costs recursively; and (iv) *curfew penalties*, which impose high costs once an arrival risks infringing Zurich's night-curfew limits, including knock-on effects on later rotations.

3) *Airport- and charge-related costs*: Airport-driven impacts include changes in time-dependent landing and noise charges at Zurich. Delay-dependent marginal costs are computed as the difference between recalculated and originally scheduled charges, internalizing the economic effect of shifting arrivals into less favourable time bands.

These cost components are aggregated into a total delay cost profile for each flight, evaluated across all feasible arrival time slots within the regulation window. The subsequent transformation from cost to weight, detailed in the following section, ensures that time slots associated with higher operational im-

pact receive proportionally lower optimization weights, while lower-cost slots are favored through higher weights. Figure 3 illustrates representative cumulative delay cost profiles for a set of SWISS flights, highlighting the heterogeneity across routes and operational contexts. Steeper curves correspond to flights with higher sensitivity to delay, such as those with critical connections or curfew exposure.




Figure 3. Cumulative cost profiles for a representative set of SWISS flights

D. Cost-to-Weight Transformation Logic

Due to the optimizer's objective of maximizing total assigned weight, the computed cost functions per flight must be transformed into corresponding weight profiles. This transformation not only aligns the input with the optimizer's mathematical objective but also protects sensitive operational and commercial data, allowing airlines to submit preference information in a normalized form without disclosing exact cost values to external parties. At the same time, it ensures that TTMS receives comparable input functions from different airlines, avoiding biases that would arise if raw cost values were used directly—an important consideration for maintaining fairness across diverse business models, including low-cost carriers. These weights serve as input for the Target Time of Arrival (TTA) optimization, which selects delay scenarios that maximize total weight while satisfying operational constraints and the parameters of the applicable ATFM regulation. The transformation ensures that time slots with high delay-related costs receive low weights, discouraging their selection, whereas lower-cost slots are favored through higher weights. The transformation process comprises several stages:

1) *Quadratic Cost Inversion*: For each flight and arrival time slot, costs are transformed into preliminary weights using a quadratic decay function:

$$w_{i,t} = \left(1 - \frac{c_{i,t}}{\max(c_{i,\cdot})}\right)^2 \cdot \left(\frac{\max(c_{i,\cdot})}{\max(c_{\cdot,\cdot})}\right) \cdot (w_{\max} - w_{\min}) \quad (1)$$

where $c_{i,t}$ is the cost of flight i at time t , and $w_{\max}, w_{\min}$ are configurable bounds for the final weight values. This ensures a smooth and monotonic inversion of the cost function, with the steepest penalties applied to high-cost time slots and proportional scaling across flights with different cost magnitudes.

2) *Priority Flight Adjustment*: Flights marked as operationally critical are treated with elevated preference. For these flights, a modified weight profile is applied using an adjusted curve method. This method enhances the weights of the earliest feasible time slots, often just after the scheduled arrival, while preserving the general shape of the original cost-informed curve. As a result, the optimizer is nudged to assign earlier arrival times for priority flights, but at the cost of shifting additional delay to lower priority flights from the same airline, as reflected in their respective weight maps, ensuring that the trade-off remains internally balanced.

3) *Average Weight Correction*: To ensure fairness and influence in the optimization process, each flight's weight profile must meet a minimum average weight threshold. If a flight's average weight falls below this threshold due to a flat or steep cost profile, an additional correction is applied using a quadratic distribution. This adjustment increases the weights of earlier time slots in a structured manner, ensuring the flight remains competitive in the allocation process.

4) *Final Smoothing and Pre-Touchdown Flattening*: As a last step, the weight profiles are smoothed to prevent erratic or non-monotonic patterns that could result from abrupt cost discontinuities. A hybrid smoothing logic combines gradients from the original and idealized curves. Additionally, based on airport-specific requirements, weights before the scheduled touchdown time can be flattened to a constant value, ensuring a stable and predictable optimization landscape. For priority flights, these slots may be assigned the global maximum weight.




Figure 4. Weight profiles for a representative set of SWISS flights

This multi-stage transformation encodes complex operational constraints and business priorities into an interpretable weight structure, a process that can only be carried out by the airlines themselves, as it requires access to sensitive, airline-specific operational data and priorities. This enables the optimizer to perform robust, delay-aware arrival sequencing. Figure 4 shows the resulting weight profiles for the same set of flights as in Figure 3, demonstrating how high delay costs translate into rapidly declining weights, whereas flights with flatter cost functions retain higher weights over a broader range of time slots.

The final weight map is structured as a two-dimensional matrix, with each row representing a flight and each column

a discrete arrival time slot (typically 1-minute intervals). Each cell contains an integer weight that indicates the relative preference for assigning the flight to that arrival time. This map forms the core input to the TTA optimizer and directly drives the outcome of the slot allocation process.

E. Fairness Metric and Visualization

We assess whether TTMS systematically advantages or disadvantages airline groups by comparing assigned TTAs with an operational baseline. For each flight we compute

$$\Delta_{TTA} = \text{assignedTime} - \text{ETA}, \quad (2)$$

which is typically non-negative under regulated operations.

Airline-level equity is evaluated by grouping carriers (hub carriers explicitly, others anonymized) and calculating mean assignment delays μ_g for each group. Let n_g be group size and N the total number of flights. Equity is summarized using the Theil T index:

$$T = \sum_g \frac{n_g}{N} \left(\frac{\mu_g}{\bar{\mu}} \right) \ln \left(\frac{\mu_g}{\bar{\mu}} \right), \quad (3)$$

where $\bar{\mu}$ is the overall mean. The metric equals zero under perfect parity and increases as group means diverge. It is scale-invariant and well suited to comparing groups of different sizes.

To complement this scalar indicator, we visualize group distributions using simple boxplots of Δ_{TTA} . Together, these tools provide a consistent view of between-group fairness without imposing additional assumptions about underlying cost structures.

F. Considerations on Strategic Behavior and Multi-Airline Participation

Normalization of airline inputs into weight maps protects sensitive cost information but raises potential incentives for strategic behavior if multiple carriers actively participate. An airline could, in principle, exaggerate delay sensitivity by assigning high weights across all feasible slots. Several mechanisms mitigate this risk.

First, only *relative* differences within a flight's weight profile matter, as all weights are normalized within each regulation. Second, the airport contributes its own independent weight component, combined with airline inputs at a fixed 50:50 ratio, preventing any single stakeholder from dominating the objective function. Third, fairness monitoring (Section III-E) highlights implausible prioritization patterns, enabling verification and corrective action.

Moreover, inflating weights is inherently self-limiting: improving one flight's slot worsens another flight of the same airline, reducing overall efficiency. Excessive prioritization thus harms the airline's own network performance rather than yielding net advantage.

In future multi-airline deployments, governance and shared validation rules coordinated by the Network Manager will ensure consistent, incentive-compatible use of weight maps. Upcoming trial phases will further refine plausibility checks and stakeholder safeguards.

IV. LIVE TRIAL AT ZURICH AIRPORT

A. Trial Setup and Scope

A live operational trial of the Target Time Management System (TTMS) was conducted at Zurich Airport between 5 May and 30 June 2025. The trial was embedded within the SESAR 3 JU-funded HARMONIC project and involved the collaboration of multiple stakeholders, including Swiss International Air Lines (SWISS), Flughafen Zürich AG (FZAG), Skyguide (ANSP), EUROCONTROL, Frequentis, and academic partners from Johannes Kepler University Linz (JKU) and the Austrian Institute of Technology (AIT).

TTMS was activated selectively during inbound Air Traffic Flow Management (ATFM) regulations declared by Skyguide's Flow Management Position (FMP), specifically those categorized as "Network Cherry Picking" (NCP) regulations and featuring declared arrival rates exceeding 22 aircraft per hour. Activation occurred with a lead time of approximately 120 minutes, during which SWISS submitted their real-time, flight-specific weight maps representing the delay sensitivity of each regulated arrival based on dynamic cost factors.

The optimizer generated an optimized Target Time of Arrival (TTA) sequence for all regulated flights, maximizing the system-wide utility as defined by the aggregated weight maps. The output sequence was required to satisfy all predefined constraints and was validated against Skyguide's operational parameters before implementation. Final slot assignments were transmitted to EUROCONTROL's Network Manager via the Network Manager (NM) TTA API and automatically integrated into the CASA system for execution.

During the trial, SWISS was the only airline providing individual flight-specific weight maps, while flights from non-participating airlines were handled using the standardized weight profiles generated by the airport.

B. Operational Conditions and Constraint Handling

Each TTMS sequence was subject to a stringent set of constraints to ensure compliance with existing safety, regulatory, and procedural requirements. The total delay assigned to all regulated flights was constrained to remain within 10% of a baseline reference. Since EUROCONTROL's ATFM delay metric only accounts for flights where the local regulation is the most penalizing regulation (MPR), a direct comparison with TTMS results was not possible. Instead, we calculated the total delay for all flights in the TTMS solution and compared it to the total delay of all flights in the baseline solution, ensuring a consistent reference. Skyguide's operational rules, in particular the regulation duration and the set throughput rate, were integrated directly into the optimizer's logic, ensuring that the output sequence was conformant by design. As a safeguard, any participating stakeholder retained the authority to request deactivation of TTMS for specific regulations in cases of elevated complexity, such as severe weather disruptions, critical recovery scenarios, or heavy, uncontrollable cascading delays across the network. This included SWISS Network

Operations Control (NOC), Zurich Airport operations, and Skyguide, ensuring that the system could be paused whenever operational circumstances required.

C. Evaluation and Results

During the eight-week trial period, TTMS was applied to a total of 51 distinct ATFM regulations. In total, 3,372 flights were affected by the system, including 2,078 operated by SWISS (60%). For evaluation purposes, each TTMS generated sequence was compared against a counterfactual baseline derived from the CASA system (via the NM B2B Simulation services). While this approach provides a consistent and transparent reference, it does not capture the full dynamic interactions and interventions from real world operations such as actions by airline operations control, airport coordination, Flow Management Positions (FMPs), or the Network Manager, which inevitably influence actual delay distributions. The comparison therefore represents an idealized "all else equal" scenario. Work is already underway to refine the baseline methodology, aiming to capture more of these network dynamics in future evaluations.

The performance results presented below refer primarily to SWISS flights, as the analysis is based on flight- and passenger-level data available for SWISS-operated services during the trial period. Complementary results from Zurich Airport provide an additional system-level perspective.

The introduction of TTMS led to measurable improvements across multiple key performance indicators. For SWISS flights, arrival punctuality (defined as arrival within 15 minutes of the scheduled arrival time published in the flight plan) improved by 20.2%, with 234 more flights arriving punctually. The number of flights with heavy delays (30+ minutes) was reduced by 23.5%, corresponding to 103 fewer severely delayed flights.

From the airport operator's perspective, arrival punctuality improved significantly, with 18% fewer flights arriving more than 15 minutes behind schedule. The number of knock-on delays, i.e. departures affected by late arrivals, decreased by 9%. From the passenger perspective, total delay minutes were reduced by 11%, while missed connections and resulting follow-on delays fell by 14%. For SWISS passengers specifically, the number facing missed connections decreased by 33.3%, while those with critical connections (25–40 minutes) dropped by 19.8%.

The realized gains are very positive but the role of TTMS in the observed increase of total ATFM delay still be quantified. Consequently, the ANSPs' capacity performance targets need to evolve in order to take into considerations the other stakeholders' performance aspects. Beyond the measurable performance improvements, the system also brought about broader qualitative benefits. By reducing the number of missed and critical passenger connections, TTMS contributed to a smoother travel experience and helped strengthen Zurich Airport's position as an attractive and reliable hub for connections. The improved arrival punctuality and reduced variability in arrival times helped stabilize downstream schedules, mitigating delay propagation and fostering greater predictability across






Figure 5. Distribution of baseline-relative assignment $\Delta_{TTA} = \text{assignedTime} - \text{ETA}$ across airline groups (HC = Hub Carrier, A/L = anonymized Airline).

the network. This, in turn, enabled more reliable resource planning for all involved partners, improving operational resilience and coordination. The closer alignment between stakeholder priorities and slot allocation not only optimized outcomes for all airlines but also reinforced trust and collaboration with Zurich Airport and other involved parties. Together, these effects illustrate the broader system-level potential of multi-stakeholder optimization of arrival sequencing, extending the benefits well beyond individual performance gains.

Using the methodology in Section III-E, SWISS conducted an internal equity analysis comparing assigned Target Times of Arrival (TTAs) with ETA across individual airline groups (with hub carriers labeled and non-hubs anonymized). The analysis found no evidence of systematic penalization of any carrier group. Group distributions of Δ_{TTA} were broadly comparable, and the Theil index over group means was low ($T=0.03072$), indicating only modest divergence across groups. Since a value below 0.4 was set as the desired threshold for acceptable equity, the observed result confirms that assignment differences between groups remained well within the target range.

The boxplot in Fig. 5 shows overlapping medians and interquartile ranges, consistent with the fairness evaluation and suggesting that equity principles were preserved despite heterogeneous inputs (e.g., airline-specific weight maps).

At the same time, lower-priority SWISS flights with limited operational impact exhibited higher delays than the average for non-SWISS flights. This pattern indicates that the optimizer intentionally concentrated the delay on flights where it causes the least harm, differentiation within the SWISS network according to priority, thus aligning utility-based optimization with overall network efficiency.

Although no formal feedback was collected from non-participating airlines during the trial period, Zurich Airport reported that no complaints or operational concerns were raised by other carriers regarding the TTMS operation. In addition, the equity monitoring described above provides a quantitative assurance that no airline group experienced disproportionate penalization. Zurich Airport has established internal procedures to analyse potential equity concerns in depth, including flight-by-flight assessments of assigned TTAs and operational impacts. These capabilities ensure that, should any airline raise a concern in future trial phases, the airport and the participating stakeholders will be able to conduct a

transparent and data-driven root-cause analysis. While direct feedback was not available for this initial trial, the absence of complaints, combined with demonstrably fair outcomes, suggests that the system was perceived as operationally neutral by other carriers.

D. Limitations and Observations

The live trial demonstrated the feasibility and operational value of TTMS but also highlighted several limitations and open issues requiring refinement.

First, the safeguard logic applied to ensure compliance with operational timelines required revision. During the trial, the safeguard initially followed an Estimated Off-Block Time (EOBT) – TRS – 10 minutes rule, meaning that a flight would no longer be eligible for further optimization once it approached the point at which its Calculated Take-Off Time (CTOT) could no longer be modified. This logic worked operationally and did not lead to major issues.

However, the safeguard was later changed to a simpler EOBT – 40 minutes cutoff. This change introduced problems, as it imposed the optimization freeze too early, especially for flights with a short flight time—and thus a late EOBT. As a result, these flights were locked out of improvements prematurely and ended up with undesirably poor TTAs, despite still having operational room for optimization. In the final week of the trial, this resulted in noticeable deterioration in slot assignments without a clear operational justification.

For the next live trial, the safeguard will be reformulated so that flights remain eligible for improvement until EOBT minus ten minutes, plus a small operational buffer. This ensures that no deterioration occurs after this point, except in cases of flight plan (FPL) updates, while preventing excessively early freezing of short-haul flights.

Second, the optimizer's performance at very low declared arrival rates was identified as a limitation. In such scenarios, the weighting logic described in Section III-D loses discriminatory power, and optimization should rely primarily on the priority factor from SWISS to ensure that critical flights are less affected. While this logic has been implemented, additional testing under controlled regulation scenarios is required to verify robustness.

Third, issues related to regulation lead-times were noted as operational challenges. On the one hand, FMP actors reported that the time between the publication of a Network Cherry Picking (NCP) regulation and the subsequent slot adaptation by CASA, based on the TTAs provided by TTMS, could feel long. This was partly due to the time required for TTMS computation and subsequent processing by CASA before results became visible in the systems used by controllers. On the other hand, when a regulation was published less than 120 minutes in advance, some flights were already taxiing or airborne and could therefore not be included in TTMS optimization. These flights subsequently appeared in the arrival sequence without TTAs, which was perceived as a more significant limitation of the short lead-time.

Fourth, Zurich Airport’s airport factor, currently based primarily on booking data, was sufficient for the purposes of this trial but requires refinement for broader deployment (see Section III-B). Planned improvements focus on enhancing the integration of passenger-related data, such as more granular booking figures and detailed connection information, while turnaround time is already well captured in the current implementation.

Finally, the integration of Most Penalizing Regulation (MPR) effects could not be fully evaluated. In CASA, when flights are subject to multiple concurrent ATFM regulations, only the one causing the largest delay (the MPR) determines the CTOT, while others are ignored. This meant TTMS optimization at Zurich was effective only when the local arrival regulation (XSZHA) was the MPR; otherwise, potential improvements could not be realized, and in some cases, large TTMS delays risked exposing flights to additional regulations.

V. FUTURE WORK AND OUTLOOK

The next major milestone is a three-month live trial scheduled from October to December 2025. This extended trial will test the system under a wider range of traffic conditions and regulatory scenarios, providing more comprehensive data for evaluation. In particular, it will address several of the limitations identified above, including the adaptation of the ETA-40 minute safeguard rule, procedures for handling short-notice regulations with less than 120 minutes lead time, and lowering the declared rate threshold so that TTMS can be applied to any arrival regulation with rates of 10 aircraft per hour or more. Following this phase, a year-long operational trial is planned to begin in early 2026, marking the transition from experimental testing to sustained operational deployment.

From the airline side, several enhancements are under development to strengthen the fidelity of cost modeling. These include a new cost function for missed operational reserves (with ongoing work to define appropriate cost assumptions), the integration of crew data from wetlease operations, and the inclusion of passenger data from non-ALTEA outstations such as Tirana. Parallel efforts focus on refining the cost-to-weight transformation logic (see Section III-D) to ensure comparability across business models and improve interpretability.

In addition, the current inversion from cost to weight follows a quadratic function, which provides strong delay sensitivity near the scheduled time while maintaining smooth decay for longer delays. However, alternative functional shapes—such as exponential or piecewise-linear transformations—may better reflect the marginal cost dynamics of certain flight types or business models. Evaluating and comparing these alternative formulations is part of ongoing research and will be a central focus of future work.

Further, the inclusion of Edelweiss (EDW) as another Zurich-based hub carrier providing flight-specific weight maps, rather than relying on airport defaults, is being explored. As a hub operator, EDW faces connectivity-driven cost structures similar to those of SWISS, making the investment in

detailed weight maps more beneficial than for predominantly point-to-point carriers.

At the system level, planned improvements include refining the airport factor (Section III-B), optimizing constraint handling for short-lead-time regulations (Section IV-D), and validating operational stability under regulation rectifications. Most critically, the upcoming yearly trial must properly address the current limitation of incomplete Most Penalizing Regulation (MPR) integration. Without a robust solution, TTMS cannot be reliably applied during peak European traffic periods such as Easter and summer, when multiple concurrent regulations are routinely published. Ensuring compatibility with enhanced MPR handling in CASA is therefore essential for large-scale operational use. Together, these developments will provide the robustness and scalability needed for a full European rollout.

The Zurich trial demonstrated that TTMS can be safely integrated into the ATFM framework while delivering measurable performance benefits. The upcoming trials will expand the scope of validation and determine the system’s readiness for large-scale deployment across Europe’s congested airspace.

VI. CONCLUSION

This paper reported the first multi-stakeholder, live deployment of the Target Time Management System (TTMS) during 51 ATFM regulations at Zurich Airport (May–June 2025). TTMS integrates airline, airport, and ANSP priorities through a privacy-preserving cost-to-weight transformation and a unified optimization. Across 3,372 regulated flights (2,078 SWISS), TTMS improved SWISS arrival punctuality by 20.2%, reduced heavy delays (30+ minutes) by 23.5%, and cut missed connections by 33.3% (critical by 19.8%). Airport-wide, total delay minutes fell by 11% and missed connections by 14%, while total simulated delay remained within a predefined band around a CASA-based counterfactual.

Equity analyses found no systematic disadvantage for non-participating airlines, and within the SWISS portfolio delays were intentionally shifted to low-priority flights, reflecting explicit utility trade-offs. Operationally, TTMS proved stable, compatible with existing procedures, and sufficiently transparent for stakeholder oversight.

Limitations remain: reduced discriminatory power at very low declared rates, short lead times for regulation publication, limited granularity of passenger data, and CASA’s current handling of only the Most Penalizing Regulation. Addressing these issues, together with the observed increase in total ATFM delay, requires further refinement and consideration at the European level.

Overall, the results indicate that cost- and preference-aware arrival sequencing can be operated safely at scale and can improve punctuality, reduce severe delays for critical flights, and meaningfully protect passenger connectivity. The incumbent additional ATFM delay triggered by optimization on arrival needs to be addressed at the European level. The European Performance Scheme (EU 2024/3128) needs to be reviewed to allow for innovative solutions that bring direct benefits to





Project funded by



Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Economic Affairs,
Education and Research EAER
State Secretariat for Education,
Research and Innovation SERI

aircraft operators and airports as described above, while not deteriorating the measured ANSP performance. The extended trial in Q4 2025 and year-long evaluation in 2026 will test robustness across traffic regimes, refine airport and default profiles, expand airline participation, and enable more precise assessment and fair apportionment of generated ATFM delay—critical steps toward a European rollout of collaborative arrival management.

ACKNOWLEDGMENT

The authors would like to thank Marie-Sophie Hainzinger, Lucas Weber, and the entire SWISS Operations Research & Air Traffic Management team, as well as Richard Stevens, Mattes Kettner, Christoph Rihacek, Christoph Fabianek, Christoph Schütz, Sebastian Gruber, Tobias Harzfeld and Thomas Lorünser for their innate dedication, hard work and contributions to this project and paper.

This work was conducted as part of the HARMONIC project. This project has received funding from the SESAR Joint Undertaking under grant agreement No 101114675 under the European Union's Horizon Europe research and innovation program. UK participant NATS in Horizon Europe Project HARMONIC receives funding from UK Research and Innovation (UKRI) under the UK government's Horizon Europe funding guarantee [grant number 10091990]. This work has received funding from the Swiss State Secretariat for Education, Research and Innovation (SERI).

REFERENCES

- [1] SESAR Joint Undertaking, "istream: optimising the arrival management at congested airports," SESAR news, Jun. 2015, <https://sesar.eu/news/istream-optimising-arrival-management-congested-airports>.
- [2] J. Evler, M. Schultz, and H. Fricke, "Flight prioritization and turnaround recovery - integrating tactical atfm slot swapping into resource-constrained ground operations," 09 2021.
- [3] T. Kistan, A. Gardi, R. Sabatini, S. Ramasamy, and E. Batuwangala, "An evolutionary outlook of air traffic flow management techniques," vol. 88, pp. 15–42. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0376042116300458>
- [4] Y. Sheng, H. Zhu, and G. Wang, "Slot reallocation for ground delay programs," in *2018 9th IEEE Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON)*. IEEE, pp. 322–327. [Online]. Available: <https://ieeexplore.ieee.org/document/8796573/>
- [5] J. Tian, X. Hao, Y. Yan, J. Huang, and M. Gen, "Multi-objective slot allocation in ground delay program based on estimation of distribution algorithm," in *2019 IEEE International Conference on Smart Manufacturing, Industrial & Logistics Engineering (SMILE)*. IEEE, pp. 178–182. [Online]. Available: <https://ieeexplore.ieee.org/document/8965305/>
- [6] H.-M. Le, D.-T. Pham, and V. N. Duong, "A multi-agent decision-making approach for automated inter-airline slot swapping in ground delay program with airline privacy preservation," in *2025 Integrated Communications, Navigation and Surveillance Conference (ICNS)*. IEEE, pp. 1–7. [Online]. Available: <https://ieeexplore.ieee.org/document/10976865/>
- [7] C. Schütz, S. Ruiz, E. Gringinger, C. Fabianek, and T. Lorünser, "An auction-based mechanism for a privacy-preserving marketplace for atfm slots," 09 2022.
- [8] Y. Chen, Y. Xu, and M. Hu, "General multi-agent reinforcement learning integrating heuristic-based delay priority strategy for demand and capacity balancing," *Transportation Research Part C: Emerging Technologies*, vol. 153, p. 104218, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0968090X23002073>
- [9] L. Caranti, M. Carré, L. Weber, D. Heilmann, and R. Stevens, "Optimization of regulated airline arrival flows via target time management an enhanced slot swapping process." [Online]. Available: <https://api.semanticscholar.org/CorpusID:267388913>
- [10] S. Gruber, P. Feichtenschlager, C. Fabianek, E. Gringinger, and C. G. Schuetz, "Towards a heuristic optimizer for a target time management system in air traffic flow management," in *2024 AIAA DATC/IEEE 43rd Digital Avionics Systems Conference (DASC)*, 2024, pp. 1–10.
- [11] L. Delgado, G. Gurtner, T. Bolić, and L. Castelli, "Estimating economic severity of air traffic flow management regulations," *Transportation Research Part C: Emerging Technologies*, vol. 125, p. 103054, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0968090X21000838>
- [12] Y. Xu, R. Dalmiau, M. Melgosa, A. Montlaur, and X. Prats, "A framework for collaborative air traffic flow management minimizing costs for airspace users: Enabling trajectory options and flexible pre-tactical delay management," *Transportation Research Part B: Methodological*, vol. 134, pp. 229–255, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0191261518310464>
- [13] A. Montlaur and L. Delgado, "Flight and passenger efficiency-fairness trade-off for atfm delay assignment," *Journal of Air Transport Management*, vol. 83, p. 101758, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0969699719303011>
- [14] C. N. Glover and M. O. Ball, "Stochastic optimization models for ground delay program planning with equity–efficiency tradeoffs," *Transportation Research Part C: Emerging Technologies*, vol. 33, pp. 196–202, 2013. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0968090X11001653>
- [15] T. Vossen, M. Ball, R. Hoffman, and M. Wambsganss, "A general approach to equity in traffic flow management and its application to mitigating exemption bias in ground delay programs," *Air Traffic Contr Q*, vol. 11, 10 2003.
- [16] A. M. Churchill and D. J. Lovell, "Coordinated aviation network resource allocation under uncertainty," *Transportation Research Part E: Logistics and Transportation Review*, vol. 48, no. 1, pp. 19–33, 2012, select Papers from the 19th International Symposium on Transportation and Traffic Theory. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1366554511000652>
- [17] D. Bertsimas and S. Gupta, "Fairness and collaboration in network air traffic flow management: An optimization approach," *Transportation Science*, vol. 50, no. 1, pp. 57–76, 2016. [Online]. Available: <http://www.jstor.org/stable/43903249>
- [18] J. Berling, A. Lau, and V. Gollnick, "Constraint reconciliation optimization to minimize primary and reactionary delay in a research network strategic tool," *CEAS Aeronautical Journal*, vol. 15, 04 2024.
- [19] Z. Zhao, Y. Li, X. Liu, J. Zhu, and S. Zhao, "Research on atfm delay optimization method based on dynamic priority ranking," *Aerospace*, vol. 12, no. 9, 2025. [Online]. Available: <https://www.mdpi.com/2226-4310/12/9/793>
- [20] M. Carré, "Multi-Airline Operations Optimization under major Disruptions," Theses, Université Clermont Auvergne, Jan. 2024. [Online]. Available: <https://theses.hal.science/tel-04547756>
- [21] T. Harzfeld, S. Gruber, C. G. Schuetz, M. Carré, M. Brügger, A. Lacroix, C. Fabianek, C. Rihacek, and E. Gringinger, "Towards conformance criteria for ensuring fairness among airspace users in collaborative optimization of flight lists in air traffic flow management," in *2025 Integrated Communications, Navigation and Surveillance Conference (ICNS)*. IEEE, pp. 1–14. [Online]. Available: <https://ieeexplore.ieee.org/document/10976866/>

