

RUNWAY-AWARE CAPACITY ENVELOPES FOR LONG-RANGE ARRIVAL METERING UNDER OPERATING-DIRECTION CHANGES AT SINGAPORE CHANGI

Ehsan Asadi[1], Gabriele Enea[2] & Michael Schultz[1]

[1]Institute of Flight Systems, University of the Bundeswehr Munich, Germany
[2]ATC & Weather Systems Group, MIT Lincoln Laboratory, Lexington, USA

Abstract

This paper proposes a runway-aware capacity envelope for long-range arrival metering at Singapore Changi Airport (WSSS), where weather-driven changes in runway operating direction create short but operationally relevant throughput losses. Routine METAR reports are converted into ATMAP weather-impact indicators, surveillance tracks are used to infer the active operating direction, and inbound demand is represented by rolling-hour entries at a 170-NM reference ring. A direction-stable reference capacity is first obtained from September 2019 operations with predominantly North configuration. Runway-direction transitions are then modeled as local perturbations of this reference, consisting of a pre-flip shoulder, a post-flip trough, and an optional recovery bump. The resulting time-indexed capacity envelope is embedded in a discrete LR-ATFM metering model in which each flight selects one feasible crossing-time option. For a flip-rich day at WSSS, the runway-aware formulation reduces the overload integral from 2085 to 395 flight-minutes, while keeping residual overload localized to the most severe transition cluster. The results show that routinely available weather, track, and traffic data can support lightweight pre-tactical metering that better prepares inbound demand for terminal-area arrival management.

Keywords: Airport capacity, runway configuration change, long-range ATFM, arrival metering, weather impact

1. Introduction

Demand–capacity balancing at major hub airports is strongly affected by short-lived reductions in usable arrival capacity. Even moderate capacity losses can propagate quickly when inbound demand is organized in dense arrival banks and when aircraft are already close to terminal-area sequencing constraints [1, 2]. For this reason, arrival planning requires more than an average airport acceptance rate; it requires a time-indexed estimate of the capacity that the terminal system can realistically sustain during the relevant planning horizon.

This requirement is particularly important for Singapore Changi Airport (Singapore Changi (WSSS)), a major Asia–Pacific hub with a mixture of long-haul, medium-haul, and regional arrival flows. A substantial share of inbound flights remains airborne for several hours before reaching Singapore, which creates an opportunity for Long Range Air Traffic Flow Management (LR-ATFM): moderate speed adjustments during cruise can absorb time before aircraft enter the terminal arrival-management environment [3–6]. However, the benefit of upstream metering depends on the quality of the capacity information used to regulate demand. If capacity is treated as static while the airport is temporarily weakened, metering may shift traffic into periods where the terminal system has the least remaining flexibility.

At WSSS, one important source of time-varying capacity is the runway operating direction. The airport alternates between North and South operating modes in response to surface wind conditions,

convective weather, and local wind shifts. These transitions are operationally necessary, but they disturb the short-term stability of the arrival process. Arrival streams may need to be re-phased, runway-dependent procedures can change, and spacing may temporarily become more conservative [7–10]. As a result, realized throughput can deviate from the level expected under stable operating-direction conditions.

Conventional planning values tend to obscure this behavior. A single capacity value, or even a time-of-day average, may represent the normal operating level but cannot capture the short depressions associated with runway-direction transitions. Conversely, purely tactical detection of overloads occurs too late for long-range metering, where the useful control action is gradual time absorption while aircraft are still in cruise. This paper therefore treats arrival capacity as a runway-aware capacity envelope: a time-indexed trajectory that follows the stable operating baseline and is locally modified around observed or anticipated runway-direction changes.

The proposed approach uses only operationally available data streams. Meteorological conditions are represented through routine Meteorological Aerodrome Report (METAR) reports and converted into weather-impact indicators using the Air Traffic Management Advancement Program (ATMAP) procedure [11]. Runway operating direction is inferred from low-altitude surveillance tracks near the airport, and arrival demand is measured as rolling-hour inbound entries at a 170-NM reference ring. Building on previous work on runway-direction changes and arrival-management implications at WSSS [12], this paper focuses on the construction of a runway-aware capacity envelope and its use in a discrete LR-ATFM metering formulation.

The contributions are threefold. First, a direction-stable reference capacity template is constructed from a flip-free baseline period in September 2019 using high-percentile rolling-hour inbound counts. Second, runway-direction changes are represented through a compact piecewise response consisting of a pre-flip shoulder, a post-flip capacity trough, and a possible recovery bump. Third, the resulting capacity envelope is embedded in a discrete arrival-metering problem in which each flight selects one feasible time-shift option and overload beyond the time-varying capacity is penalized.

The remainder of the paper is organized as follows. Section 2. describes the operational data sources and indicators. Section 3. presents the runway-aware capacity-envelope construction. Section 4. formulates the arrival-metering problem. Section 5. discusses the case-study results, and Section 6. concludes the paper.

2. Operational Data and Indicators

The proposed capacity model is built from routinely available operational data: airport weather reports, surveillance tracks, and time-stamped inbound traffic observations. The aim is not to reproduce tactical controller decisions in detail, but to extract consistent indicators of stable runway operation, runway-direction transitions, and the associated demand response.

All time series are represented on a common 5-minute local-time grid. This resolution captures short runway-transition effects while remaining consistent with rolling-hour demand measures used in flow-management applications. Three indicators are used: a weather-impact signal derived from METAR reports, a track-based runway operating-direction state, and rolling-hour inbound entries at a 170-NM reference ring around WSSS.

2.1 Weather-impact representation

Meteorological conditions are described using routine METAR reports for WSSS. These reports provide surface wind direction and speed, gust information, visibility, present weather, and convective indicators. For the present application, wind and convective activity are particularly relevant because they influence the feasibility of a given runway operating direction [8]. Wind shifts can make the current direction less favorable, while convective cells and gust fronts may introduce abrupt short-term variability.

To convert heterogeneous METAR elements into a uniform operational signal, we use the ATMAP procedure, which maps weather phenomena to normalized impact scores for visibility, wind, precipitation, freezing, and an aggregate danger index [11]. The danger score is not used as a direct capacity value; it serves as a diagnostic indicator for relating adverse weather episodes to runway-direction changes and throughput depressions.

The raw ATMAP series is resampled to the 5-minute grid and smoothed using a short weighted moving average followed by a centered one-hour mean. This filtering suppresses isolated reporting noise while preserving the broader weather episodes associated with convective passages and gust fronts.

2.2 Track-based runway operating direction

Runway operating direction is inferred from surveillance tracks rather than from manually curated runway logs. Only near-airport, low-altitude track segments are used, typically below 4000 ft and within approximately 50 km of WSSS. These restrictions reduce ambiguity from far-field headings and focus the classification on aircraft aligned with the terminal arrival geometry.

For each sliding time window, approach tracks are classified according to their geometry relative to the runway axis. The resulting state is assigned to North, South, or mixed operations. North operations correspond to the use of the 02 direction, with arrivals approaching from the south; South operations correspond to the use of the 20 direction, with arrivals approaching from the north. Mixed-operation windows occur during short transition periods or when the observations do not support a single dominant direction.

A runway-direction change, or flip, is identified when the dominant state changes between North and South operations. Short isolated classifications are filtered to avoid false transitions caused by sparse tracks or temporary ambiguity. The resulting direction timeline provides the temporal reference for aligning weather signals, traffic counts, and capacity changes.

2.3 Rolling-hour inbound entries at the 170-NM ring

Inbound demand is represented by aircraft crossing a circular reference boundary around WSSS. A radius of 170 NM is used because it is consistent with the interface between upstream LR-ATFM and terminal-area arrival management. At this distance, many aircraft are near the transition from cruise toward descent, but still far enough from the airport for moderate in-cruise time absorption to be meaningful [3, 4, 6].

For each inbound flight, the first crossing time of the 170-NM boundary is extracted. The rolling-hour entry count at time t is

$$q(t) = \sum_{i \in \mathscr{I}} \mathbf{1}\{t - 60\ \text{min} < \theta_i \leq t\},$$

where θ_i is the 170-NM crossing time of flight i, and $\mathscr{I}$ is the set of inbound flights. The count is updated every 5 minutes. This representation captures sustained pressure on the arrival system while revealing transient overloads around runway-direction changes.

Figure 1 illustrates the combined indicator set. The runway operating direction provides transition timing, rolling-hour entries describe the demand response, and the ATMAP danger score provides weather-impact context. Their combination reveals repeatable capacity signatures that can be represented as a time-varying envelope.

3. Runway-Aware Capacity Envelope

The operational indicators are used to construct a time-indexed arrival capacity envelope $U(t)$. The envelope combines a direction-stable reference template with local perturbations around runway-direction changes. The reference describes sustained capacity under stable operations, while the perturbation model captures the short transition-related reductions and recovery phases observed around flips.

3.1 Direction-stable reference template

A reference capacity should describe what the arrival system can sustain when runway operations are stable. We therefore select a September 2019 baseline period with minimal operating-direction changes, during which WSSS operated predominantly in North configuration. This block is used to estimate a stable North-operations capacity template.

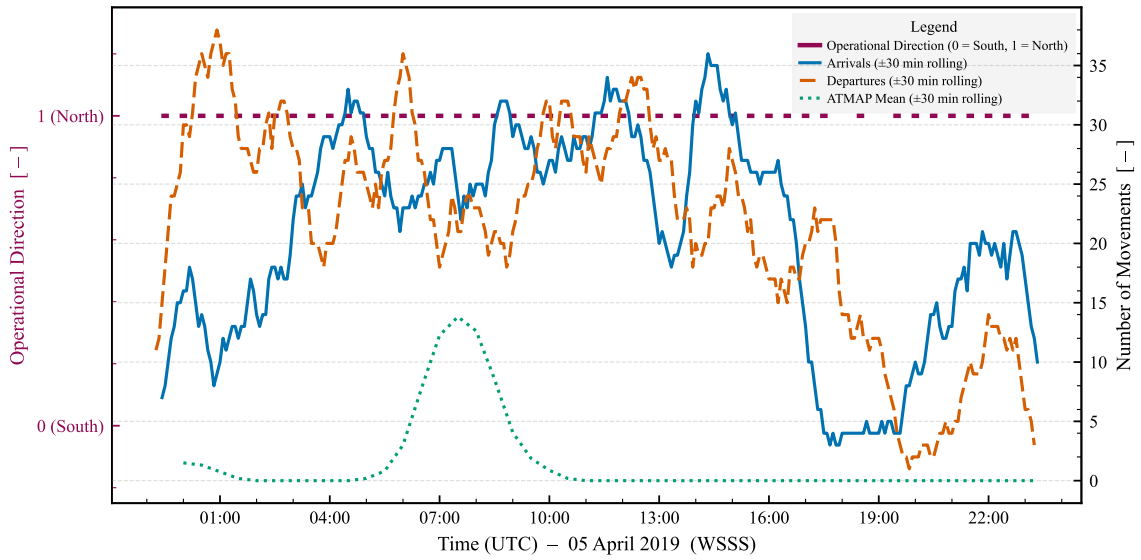


Figure 1 – Coupled time series of runway operating direction (left axis), 1-hour rolling arrival and departure counts, and smoothed ATMAP convective weather index at WSSS on 5 April 2019.

For each reference day d and each 5-minute time stamp t, the rolling-hour inbound entry count $q_d(t)$ is computed at the 170-NM boundary. The reference capacity is obtained from a high percentile of the distribution across reference days,

$$U_{\mathrm{ref}}(t) = S\left(Q_p\left[q_d(t),\, d \in \mathscr{D}_{\mathrm{ref}}\right]\right),$$

where $\mathscr{D}_{\mathrm{ref}}$ is the set of reference days, Q_p is the selected percentile, and $S(\cdot)$ denotes a light smoothing operator. A high percentile is used rather than the mean because the objective is to estimate sustained usable capacity rather than average scheduled demand [13]. The resulting curve preserves the diurnal structure of the arrival system, including high-utilization periods, routine troughs, and the evening reduction.

3.2 Runway-direction transition response

When the runway operating direction changes, realized throughput often deviates from the stable reference. Around many transitions, the response follows a compact structure: a pre-flip shoulder, a post-flip trough, and a recovery phase. The pre-flip shoulder captures the preparation period immediately before the detected change. The post-flip trough represents the dominant capacity effect, when streams are re-phased, runway-dependent procedures are adapted, and spacing may temporarily become conservative [14]. A recovery bump may occur once the new configuration stabilizes and accumulated demand is released.

The transition response for flip r at time τ_r is represented by a local perturbation $\Delta_r(t)$ added to the reference template:

$$\Delta_r(t) = \begin{cases} \Delta_r^{\mathrm{pre}}, & \tau_r - L_r^{\mathrm{pre}} \le t < \tau_r, \\ -\Delta_r^{\mathrm{tr}}, & \tau_r \le t < \tau_r + L_r^{\mathrm{tr}}, \\ \Delta_r^{\mathrm{rec}}, & \tau_r + L_r^{\mathrm{tr}} \le t < \tau_r + L_r^{\mathrm{tr}} + L_r^{\mathrm{rec}}, \\ 0, & \text{otherwise.} \end{cases}$$

Here, L_r^{pre}, L_r^{tr}, and L_r^{rec} denote the phase durations, while Δ_r^{pre}, Δ_r^{tr}, and Δ_r^{rec} describe the corresponding capacity changes relative to the reference curve. The trough parameter Δ_r^{tr} is positive in magnitude but enters the envelope with a negative sign.

The full runway-aware capacity envelope is

$$U(t) = \max\left[0,\, U_{\text{ref}}(t) + \sum_{r \in \mathscr{R}} \Delta_r(t)\right],$$

where $\mathscr{R}$ is the set of detected runway-direction changes. A final light smoothing step avoids artificial discontinuities between adjacent 5-minute bins. When flips occur close together, overlapping perturbations represent the resulting prolonged constrained state.

3.3 Empirical parameterization for the case study

The flip response is calibrated by comparing realized rolling-hour entries on a flip-rich day with the direction-stable reference template. For 25 September 2019, ten runway operating-direction changes are detected. Early-morning transitions generate the strongest depressions: between approximately 04:30 and 06:30, the post-flip deficit reaches about 8–15 movements per hour and typically lasts 45–60 minutes, extending to approximately 90 minutes when flips occur close together. Mid-morning transitions are milder, with deficits of about 3–6 movements per hour over 30–45 minutes. Around midday, some transitions are followed by a short recovery bump of about 2–4 movements per hour for 20–30 minutes.

This behavior motivates the schematic representation in Figure 2. The purpose of the model is not to predict each aircraft movement, but to provide a compact capacity envelope that captures the operationally relevant time scale and magnitude of runway-transition effects.

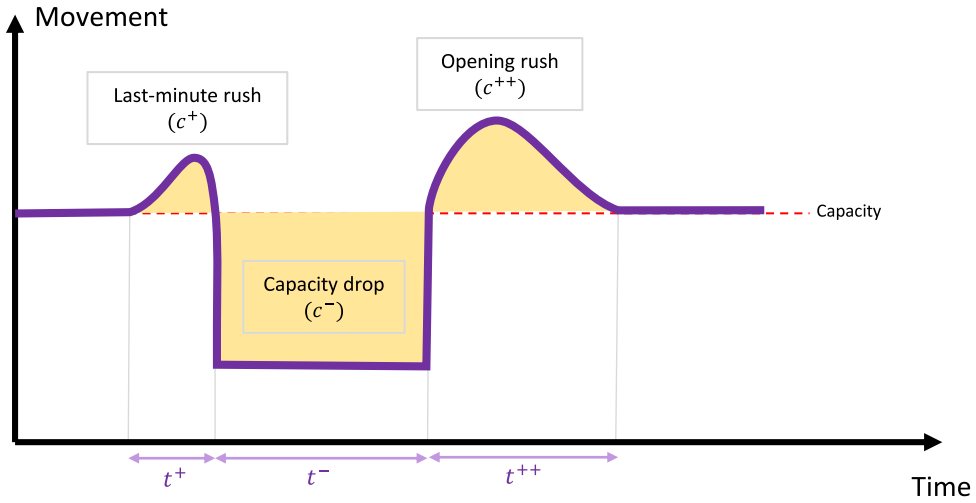


Figure 2 – Schematic runway-flip response relative to the direction-stable reference template.

4. Arrival Metering Formulation

The runway-aware capacity envelope $U(t)$ provides the time-varying capacity input for the arrival metering problem. The objective is to reshape inbound demand at the 170-NM reference ring so that the rolling-hour entry count remains close to the available capacity. The formulation is designed for pre-tactical LR-ATFM: it does not assign runways, build terminal sequences, or model vectoring decisions. Instead, it adjusts planned crossing times within operationally feasible limits before aircraft enter the terminal-area arrival-management environment.

Each inbound flight i has an unregulated reference crossing time θ_i^0. A finite option set $\mathscr{J}_i$ is generated, where each option j corresponds to a feasible time shift δ_{ij} and shifted crossing time

$$\theta_{ij} = \theta_i^0 + \delta_{ij}.$$

The option set represents bounded in-cruise time absorption or recovery through moderate speed adjustment. This keeps the model compact while remaining consistent with LR-ATFM, where small time corrections are accumulated upstream before terminal-area entry [3, 4, 15].

To connect flight options with rolling-hour demand, the binary parameter T_{ijt} is defined as

$$T_{ijt} = \begin{cases} 1, & t - 60\ \text{min} < \theta_{ij} \le t, \\ 0, & \text{otherwise.} \end{cases}$$

Thus, $T_{ijt} = 1$ if option j of flight i contributes to the rolling-hour count at time step t. Let $\mathscr{T}$ denote the planning horizon and $\mathscr{I}$ the set of inbound flights. The binary variable x_{ij} indicates whether option j is selected for flight i, and the nonnegative variable ρ_t measures residual overload above the capacity envelope. The optimization problem is

$$\min \sum_{t \in \mathscr{T}} \Delta t\, \rho_t$$

subject to

$$\sum_{i \in \mathscr{I}} \sum_{j \in \mathscr{J}_i} x_{ij} T_{ijt} \le U_t + \rho_t \qquad \forall t \in \mathscr{T},$$

$$\sum_{j \in \mathscr{J}_i} x_{ij} = 1 \qquad \forall i \in \mathscr{I},$$

$$x_{ij} \in \{0,1\} \qquad \forall i \in \mathscr{I},\ j \in \mathscr{J}_i,$$

$$\rho_t \ge 0 \qquad \forall t \in \mathscr{T}.$$

Here, Δt is the 5-minute time-step length. Since ρ_t represents the number of aircraft above the capacity envelope, the objective measures the overload integral in flight-minutes. This metric captures both magnitude and duration of imbalance and is closely related to the operational risk of downstream holding when demand remains above capacity [16].

The first constraint limits optimized rolling-hour demand to the available capacity except for penalized residual overload. The second ensures that each flight selects exactly one option. Because the option sets are discrete and T_{ijt} is precomputed, the problem is a compact mixed-integer linear program. The resulting metering layer prepares a smoother demand profile at the 170-NM boundary, while detailed sequencing, runway balancing, and tactical spacing remain with the terminal Arrival Manager (AMAN).

5. Case Study Results

The case study evaluates the runway-aware metering approach for WSSS using September 2019 traffic and operating-direction data. The reference capacity template is constructed from the direction-stable period between 5 and 20 September 2019, during which the airport operated predominantly in North configuration. The metering experiment is then performed for 25 September 2019, a flip-rich day with ten detected operating-direction changes.

5.1 Capacity-envelope behavior on the flip-rich day

Figure 3 compares realized rolling-hour inbound entries on 25 September with the direction-stable reference template. The deviations reveal three recurring transition effects. First, throughput deterioration begins approximately 10–15 minutes before several detected flips, indicating that the arrival process is affected already during the preparation phase. Second, the strongest depressions occur during the early-morning cluster between approximately 04:30 and 06:30, where repeated direction changes produce deficits of about 8–15 movements per hour. These depressions typically last 45–60 minutes and can extend to approximately 90 minutes when a second flip occurs before the system has recovered. Third, mid-morning transitions are milder, with reductions of about 3–6 movements per hour over 30–45 minutes, while some midday transitions are followed by short recovery bumps of about 2–4 movements per hour.

These observations are translated into the capacity envelope $U(t)$ used by the metering model. The envelope follows the direction-stable reference during stable periods and introduces local perturbations around detected flips. The deepest troughs are assigned to the early-morning cluster, while

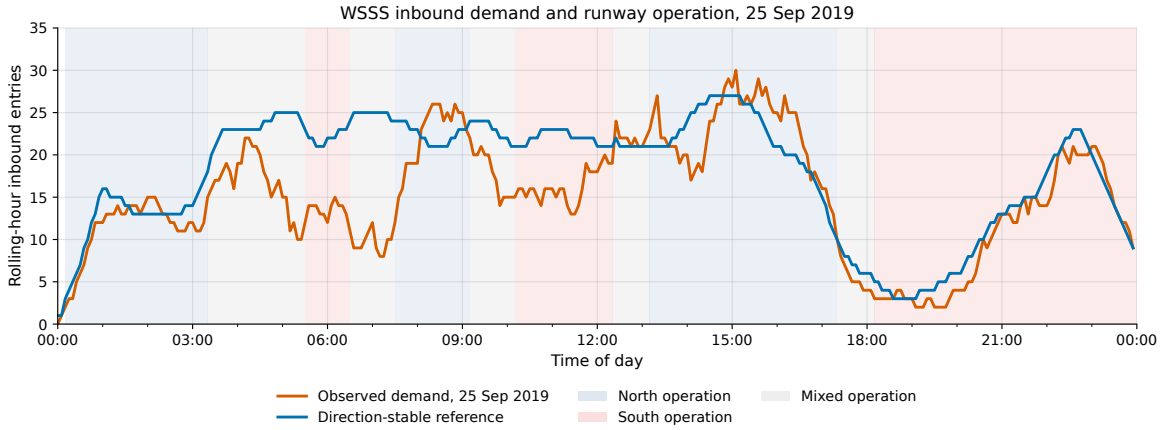
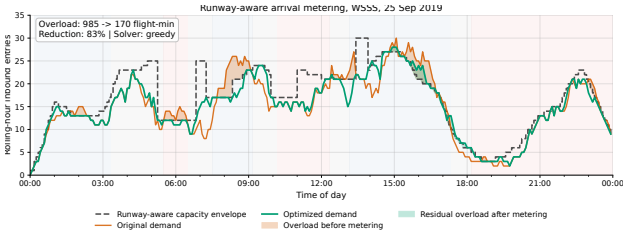


Figure 3 – Realized rolling-hour inbound entries on 25 September 2019 compared with the direction-stable reference template.

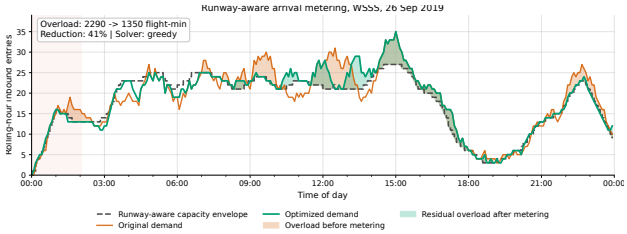
shorter and shallower depressions are used for later transitions. Where supported by the data, a small recovery bump is included after stabilization.

5.2 Metering performance

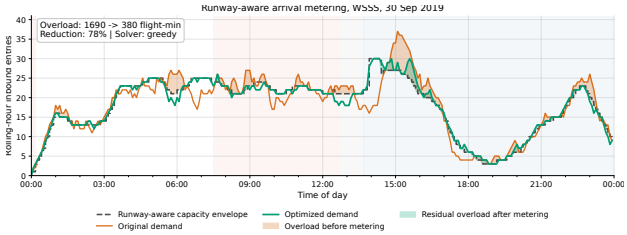
The metering optimization is applied to inbound flights crossing the 170-NM reference ring on 25 September 2019. The original demand trajectory is evaluated against the runway-aware capacity envelope, and the optimized trajectory is obtained by selecting one feasible time-shift option for each flight. Figure 4 shows the original demand, optimized demand, and time-varying capacity envelope.



(a) 25. Sep



(b) 26. Sep



(c) 30. Sep

Figure 4 – Runway-aware metering results for different days in September 2019. The dashed curve denotes the time-varying capacity envelope, the original rolling-hour entries show the unregulated demand, and the optimized entries show the demand after feasible time-shift options are selected.

Without metering, the original demand exceeds the runway-aware capacity envelope in several periods, especially during the early-morning transition cluster. The overload integral is 2085 flight-minutes. After optimization, it is reduced to 395 flight-minutes, corresponding to an 81% reduction. The maximum residual exceedance is limited to six aircraft and occurs during the deepest morning trough.

The optimized trajectory remains close to the available capacity but generally avoids crossing it. Rather than delaying a large number of flights to the end of a constrained period, the model re-phases demand locally, using available shoulders before and after troughs while preserving the broader daily arrival structure. Remaining overload is concentrated in the most difficult part of the day, where high

https://doi.org/10.71945/icas2026_0408

demand, a deep capacity trough, and limited recovery time between consecutive flips coincide. This indicates that runway-aware LR-ATFM can substantially reduce overload with moderate upstream adjustments, although it cannot fully compensate for severe clustered capacity losses without stronger interventions.

6. Conclusions

This paper presented a runway-aware capacity-envelope and arrival-metering approach for Singapore Changi Airport. Instead of treating arrival capacity as a static ceiling, the proposed method represents capacity as a time-varying envelope that follows a stable operating baseline and is locally modified around runway operating-direction changes.

The method combines METAR-derived ATMAP weather-impact scores, track-based runway-direction detection, and rolling-hour inbound entries at a 170-NM reference ring. A direction-stable reference template is first constructed from September 2019 operations using high-percentile rolling-hour entry statistics. Runway-direction changes are then represented through a compact piecewise response consisting of a pre-flip shoulder, a post-flip capacity trough, and, where observed, a short recovery bump after stabilization.

The case study for 25 September 2019 shows that runway-direction changes can produce short but operationally relevant capacity depressions, especially when flips occur in close succession. Embedding the resulting capacity envelope in a discrete LR-ATFM metering formulation reduced the overload integral from 2085 to 395 flight-minutes, corresponding to an 81% reduction. The remaining overload was localized to the most constrained transition cluster.

The approach is intentionally lightweight. It does not require detailed runway assignment, tactical sequencing, or proprietary decision-support data. Instead, it uses routinely available weather, track, and traffic observations to identify periods in which the arrival system is temporarily weakened and to shift demand away from those periods through feasible upstream time adjustments.

Future work should couple the envelope with probabilistic runway-direction forecasts based on predicted wind components, convective-weather indicators, and evolving ATMAP trends. Advanced runway-configuration forecasting could further improve pre-tactical applicability [17]. Additional validation across other months and monsoon regimes, as well as airline-centric cost and equity constraints, should also be included. Nevertheless, the results indicate that runway-aware capacity modeling provides a practical bridge between weather-sensitive airport operations and long-range arrival metering.

7. Contact Author Email Address

If you have any questions, please send an email to: ‘ehsan.asadi@unibw.de’ (corresponding author).

8. Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.

References

[1] Yunxiang Chen, Yifei Zhao, Fan Fei, and Haibo Yang. Optimizing large-scale demand and capacity balancing in air traffic flow management using deep neural networks. *Aerospace*, 11(12):966, 2024.

[2] YY Tee and Zhao Wei Zhong. Modelling and simulation studies of the runway capacity of changi airport. *The Aeronautical Journal*, 122(1253):1022–1037, 2018.

[3] Michael Schultz, Daniel Lubig, Ehsan Asadi, Judith Rosenow, Eri Itoh, Srinivas Athota, and Vu N. Duong. Implementation of a long-range air traffic flow management for the asia-pacific region. *IEEE Access*, 9:124640–124659, 2021.

https://doi.org/10.71945/icas2026_0408

[4] G. Enea, A. Andreeva-Mori, J. Bronsvoort, J. Hochwarth, T. Reynolds, and Onji. Implementation of long-range air traffic flow management at large hub airports: An international perspective. In *Proceedings of the 15th USA/Europe Air Traffic Management (ATM) R&D Seminar*, 2023.

[5] Liping Huang, Sheng Zhang, Yicheng Zhang, Yi Zhang, and Yifang Yin. Aircraft landing time prediction with deep learning on trajectory images. *arXiv preprint arXiv:2401.01083*, 2024.

[6] Eri Itoh, Koji Tominaga, Michael Schultz, and Vu N Duong. Untangling complexity in asean air traffic management through time-varying queuing models. *Aerospace*, 11(1):11, 2023.

[7] Diana C Tascón and Oscar Diaz Olariaga. Air traffic forecast and its impact on runway capacity. a system dynamics approach. *Journal of Air Transport Management*, 90:101946, 2021.

[8] DT Pham, M Ngo, N Tran, S Alam, and V Duong. Changi airport. In *Air Traffic Management and Systems IV: Selected Papers of the 6th ENRI International Workshop on ATM/CNS (EIWAC2019)*, volume 731, page 113. Springer Nature, 2021.

[9] Afaq Khattak, Pak-wai Chan, Feng Chen, and Haorong Peng. Assessing wind field characteristics along the airport runway glide slope: An explainable boosting machine-assisted wind tunnel study. *Scientific Reports*, 13(1):10939, 2023.

[10] Adrian Barea, Raul de Celis, and Luis Cadarso. Integration of airport terminal arrival route selection, runway assignment and aircraft trajectory optimisation. *Transportation Research Procedia*, 47:299–306, 2020.

[11] Eurocontrol. Algorithm to describe weather conditions at European airports, 2011.

[12] Ehsan Asadi, Gabriele Enea, and Michael Schultz. Implications of Runway Direction Changes for improved Arrival Management at Singapore Changi Airport. In *15th SESAR Innovation Days 2025*, page 9, Bled, 2025.

[13] Deepak Kulkarni. Models of maximum flows in airspace sectors in the presence of multiple constraints. In *2017 IEEE/AIAA 36th Digital Avionics Systems Conference (DASC)*, pages 1–8. IEEE, 2017.

[14] Ricardo Rodrigues Pacheco, Elton Fernandes, and Eduardo Marques Domingos. Airport airside safety index. *Journal of Air Transport Management*, 34:86–92, 2014.

[15] Judith Rosenow, Ehsan Asadi, Daniel Lubig, Michael Schultz, and Hartmut Fricke. Long range air traffic flow management with flight-specific flight performance. *Future Transportation*, 2(2):310–327, 2022.

[16] Prasenjit Sengupta, Monish D Tandale, and Padmanabhan K Menon. Risk-hedged traffic flow management under airspace capacity uncertainties. *Journal of Guidance, Control, and Dynamics*, 37(5):1487–1500, 2014.

[17] T. Reynolds, M. Matthews, D. Clark, S. Splawinski, K. Arnold, and M. Williams. Weather considerations for airport capacity decision support development. In *Proceedings of the 16th USA/Europe Air Traffic Management (ATM) R&D Seminar*, 2025.