Analysing Delay Propagation and Stability in Aviation Networks Using Max-Plus Linear Systems

MSc Thesis Report

Jelmer de Haan

Delft University of Technology

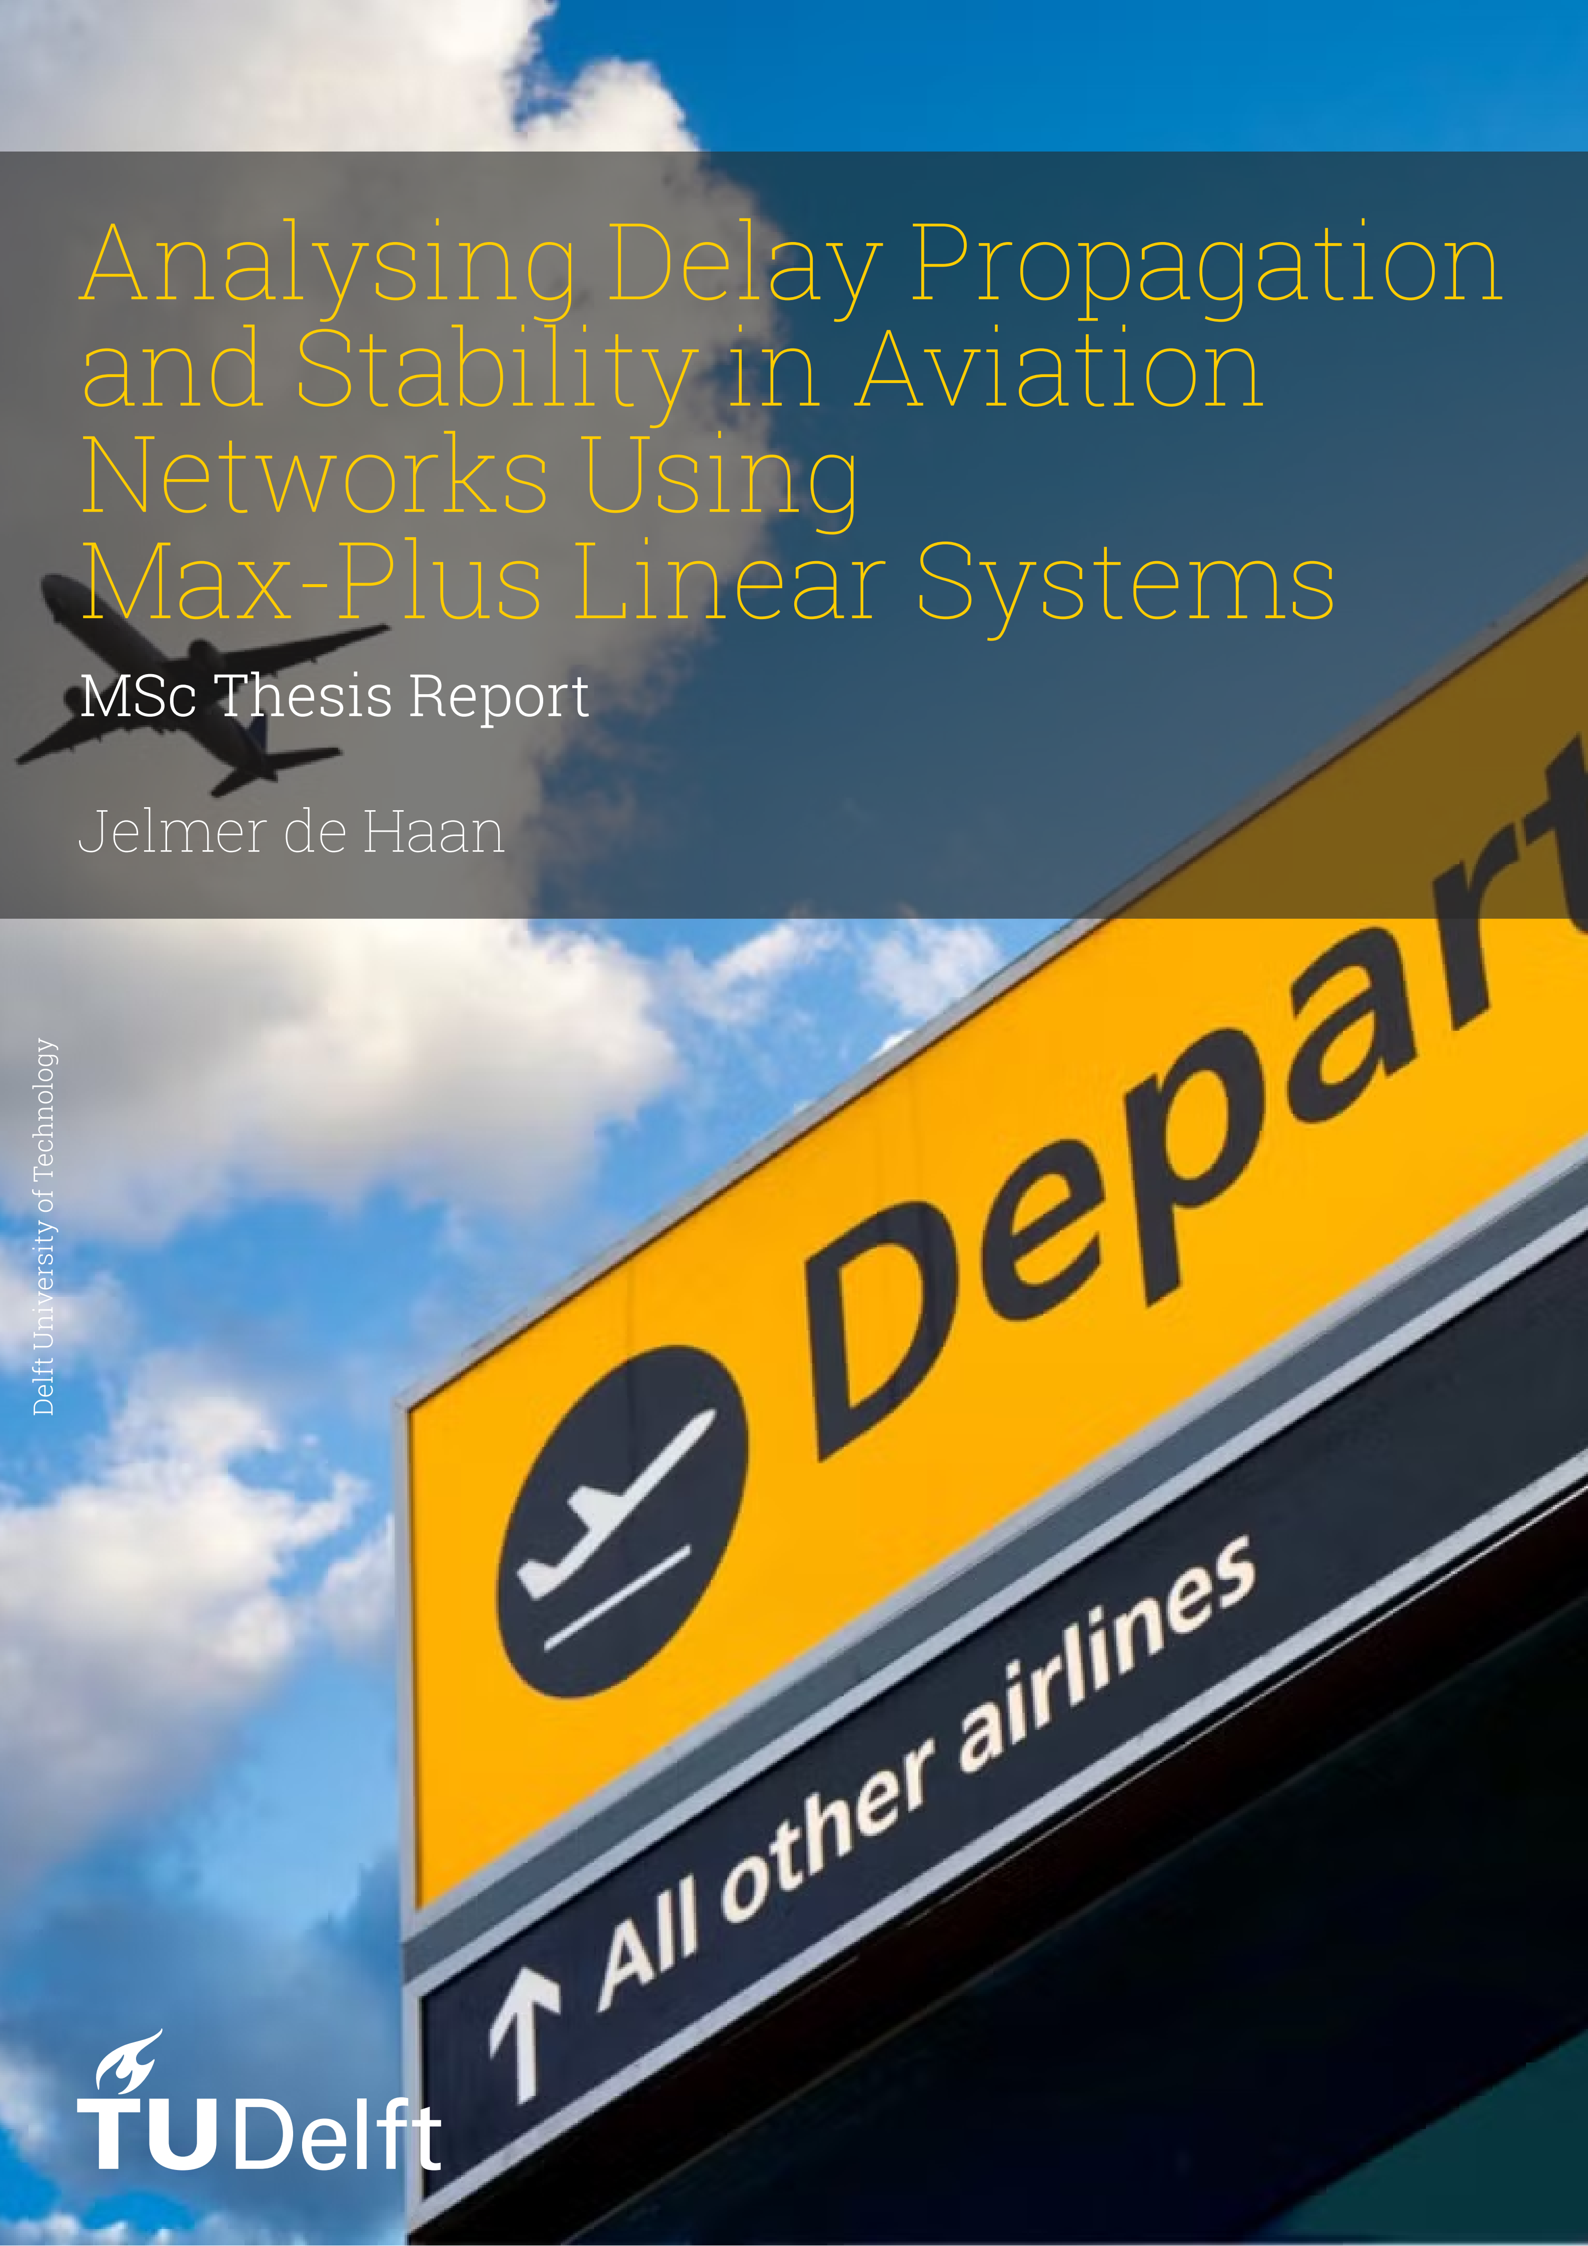

Analysing Delay Propagation and Stability in Aviation Networks Using Max-Plus Linear Systems

MSc Thesis Report

by

Jelmer de Haan

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Thursday August 28, 2025 at 14:00.

Student number:	4864131	
Project duration:	October 23, 2024 – August 28, 2025	
Thesis committee:	Dr. P.C. Roling,	TU Delft, chair
	Dr. A. Bombelli,	TU Delft, supervisor
	Prof. Dr. R. M. P. Goverde,	TU Delft
	Dr. S. Theodoulis,	TU Delft

Cover: ”Airport Departure and Arrival sign at Heathrow, London” © alice-photo, Shutterstock
Style: TU Delft Report Style, with modifications by Daan Zwaneveld

TU Delft

Preface

Here we are, at the end of the road. The end of the road of 10 months of thesis work, the end of the road of 7 years of studying Aerospace Engineering. Lots has happened in those seven years. Who would have thought, back in 2018, that we would join an association, face a pandemic, stay with the in-laws, take a gap year, meet new friends, sail the Mediterranean, work at Heineken, hike across Nepal, move to Rotterdam, and now hand in this thesis? ”Let’s do it all again” is a bit of a stretch, but what a ride it has been.

The idea for this thesis started over two years ago, and I am happy to finally see it complete. Diving into max-plus algebra hasn’t always been easy, especially when the aviation networks you’re interested in aren’t as periodic as you would have hoped or your dataset isn’t as neatly formatted as you would have liked. Should anyone ever wish to save themselves a few weeks of work converting BTS data to UTC time, I do consider myself somewhat of an expert at this point and am available at a competitive rate.

A great number of thanks are in order, for I could not have reached this point alone. First, I would like to thank my supervisors: Alessandro, for your enthusiasm to support a far-fetched idea from the very early start nearly 20 months ago. Rob, for your unwavering attention to detail and reliable guidance throughout this process. Algebra is an exact art.

To my thesis partners in crime, Sebas, Rijk, Mats: Thanks for helping me drag myself to the EWI high-rise, thanks for all the cups of tea and thanks for being (somewhat) excited about my visualisations. To Nynke, Xando, Lil, Sophie for kicking me out of bed, to Ilse, Isa, Jona for keeping me sane, to Zoë, Sanne, Giulia for easing me into adulthood and to Puk, Tim, and Jona for keeping me young.

Finally, thanks to André, Anne, Eline: I wager those past 24 years haven’t been the easiest of your lives, what with all my late nights and odd plans. Still, let’s hope we can cross one worry of your list now.

To new adventures.

Jelmer de Haan
Rotterdam, August 2025

Contents

Nomenclature

Table 1: List of symbols and abbreviations used in this thesis.

Symbol	Description
	Max-plus algebra definitions
$\oplus$	Max-plus addition: $a \oplus b = \max(a, b)$.
$\otimes$	Max-plus multiplication: $a \otimes b = a + b$.
$x^{\otimes r}$	Max-plus power: $x \otimes x^{\otimes(r-1)}$.
ε	Zero element $-\infty$ in max-plus algebra.
e	Unit element 0 in max-plus algebra.
$\mathcal{E}_{m\times n}$	Zero matrix in max-plus algebra.
E_n	Identity matrix in max-plus algebra.
$\mathbb{R}_\varepsilon$	Set of extended real numbers $\mathbb{R} \cup \{-\infty\}$.
$\mathbb{N}_0$	Set of extended natural numbers $\mathbb{N} \cup \{0\}$.
	Precedence graph
$\mathcal{G}(A)$	Precedence graph corresponding to matrix A.
n	Number of events/nodes in the precedence graph.
i, j	Event (node) indices in precedence graph.
a_{ij}	Element of max-plus matrix A, weight from event j to i.
μ_{ij}	Period shift between events i and j.
	Scheduled max-plus linear system
$x(k)$	Event time vector in period k.
$d(k)$	Timetable vector containing scheduled event times.
$z(k)$	Delay vector in period k.
k	Discrete time period index in max-plus system.
p	Order of the system.
T	Scheduled cycle time.
γ	Backward shift operator.
	Stability analysis
A^+	Critical path matrix: $A \oplus A^{\otimes 2} \oplus \ldots$.
A_0^*	Transitive closure matrix: $\bigoplus_{l=0}^{n-1} A_0^l$.
t_0	Time instant defining known events at simulation start.
λ, η	Eigenvalue (maximum cycle mean) of max-plus system.
v	Eigenvector of max-plus system.
ρ	Network throughput.
Δ_1	Average slack on the critical circuits.
Δ_2	Stability margin.
R, r_{ij}	Recovery matrix R with values r_{ij}.
	Abbreviations
DAG	Directed Acyclic Graph
DM	Delay Multiplier.
MAE	Mean Absolute Error between predicted and observed delays.
OTP	On-Time Performance (flights within 15 minutes of schedule).
SCC	Strongly Connected Components
TAT	Turnaround Time.

Introduction

Flight delays remain a significant operational and economic challenge for the aviation industry worldwide. In 2023, only 71% of flights in European Civil Aviation Conference (ECAC) states and 78% of United States domestic flights arrived on time (Bureau of Transportation Statistics, 2024; Walker, 2024). This results in substantial operational disruptions and high associated costs, with annual global estimates ranging into the tens of billions of dollars (AirHelp, 2023; Ball et al., 2010; Scheelhaase et al., 2023; Value Group, 2023).

While many factors contribute to flight delays, delay propagation, where delays spread through interconnected flights, crews, or passengers, is recognised as a key driver. Unlike primary delays caused by factors such as weather or technical issues, propagated delays emerge from the interdependencies within airline operations. Despite its importance, delay propagation remains less thoroughly understood than primary delays, particularly at the network level.

Existing studies often focus on short-term forecasting or the effects of individual delay events, using approaches such as statistical models, machine learning, simulation, or queueing theory. While useful, these methods often struggle to model delay propagation at a high level of operational detail across entire networks, or they require high computational resources (C. Li et al., 2024).

This thesis explores the use of *max-plus linear systems* to model delay propagation and analyse network stability in aviation. Scheduled max-plus linear systems provide an efficient means of representing cyclic, event-based systems with precedence constraints, making them suitable for analysing periodic transport networks. While max-plus methods have been widely applied in railway contexts for timetable stability and delay propagation analysis, their application in aviation has thus far been limited.

The research aims to answer the following research question and subquestions:
To what extent can scheduled max-plus linear systems be leveraged to model, analyse, and predict delay propagation in aviation networks?

1. What data is available and how must it be processed to prepare it to model delay propagation in aviation networks?
2. To what extent can scheduled max-plus linear systems be adapted and applied to model delay propagation in aviation networks?
3. To what extent can and should the interdependencies between aircraft and passengers be integrated into a scheduled max-plus linear system, and what different roles do they play in the propagation of delays according to this system?
4. To what extent can stability indicators for scheduled max-plus linear systems (e.g., cycle time, stability margin, and recovery matrix) be effectively calculated for aviation networks, and to what extent do they provide insights about network robustness?
5. To what extent can an efficient delay propagation algorithm based on scheduled max-plus plus linear systems be developed for aviation networks?
6. How scalable is a max-plus algebra model for large-scale aviation networks, and how well does it capture the real-world complexities of delay propagation compared to smaller test networks?

The core of this thesis is a scientific article, in which max-plus linear systems are applied to the Hawaiian Airlines network to evaluate their suitability for modelling delay propagation in aviation networks. The article develops a modelling framework that incorporates key operational dependencies such as aircraft rotations and passenger transfer connections. It applies this framework to compute structural stability indicators, simulate delay propagation under various scenarios, and assess the trade-offs between model complexity, computational efficiency, and predictive accuracy.

The main findings demonstrate that max-plus linear systems can be effectively applied to aviation networks to compute the recovery matrix and to identify structural weak points in the network. They can also be used to evaluate the impact of including passenger transfer precedence relations. However, the results also show that the predictive accuracy of the model is limited and that the stability characteristics which give useful insights in railway networks do not necessarily do so in aviation networks.

In addition to the scientific article, this thesis includes a literature review and research proposal. These components provide further background on the state of the art in delay propagation modelling, including methods from aviation, railway, and systems theory domains. They also outline the research questions and methodology that guided the development of the scientific article. Note that as the literature review was completed a length of time before the scientific article, the contents of the article can deviate from those of the literature review. The article should be considered the most recent and accurate version.

The remainder of this thesis is structured as follows. Part I presents the scientific article, containing the main research contributions of this thesis. Part II contains the literature review and research proposal that provided the foundation for this research.

Part I

Scientific Article

Analysing Delay Propagation and Stability in Aviation Networks Using Max-Plus Linear Systems

J.S. de Haan
Under the supervision of Dr. A. Bombelli & Prof. Dr. R. M. P. Goverde
Operations and Environment, Faculty of Aerospace Engineering
Delft University of Technology, The Netherlands

Delay propagation is a significant driver of flight delay in aviation networks, yet modelling it at a network-wide scale remains challenging. This study investigates to what extent scheduled max-plus linear systems, as used in railway delay modelling, can be applied to aviation networks. Using the Hawaiian Airlines network as a case study, a methodology is developed to model aircraft rotation and passenger transfer precedence relations within a max-plus linear system. The approach enables the calculation of stability indicators such as maximum cycle mean, recovery times, and network slack, as well as the simulation of delay propagation under various initial delay scenarios.

Results show that the recovery matrix is a valuable tool for identifying structurally vulnerable parts of the network and for assessing the impact of holding aircraft for transferring passengers. However, predictive accuracy of delay propagation for individual flights is limited, primarily due to uncertainties in process time estimation and incomplete knowledge of precedence relations. The 24-hour periodicity of aviation timetables, combined with large overnight buffers, further limits multi-day delay propagation modelling. These limitations are partly specific to the case under study and partly inherent to the deterministic, periodic structure of scheduled max-plus systems.

The study concludes that max-plus linear systems can provide meaningful insights into structural robustness and the systemic impact of schedule design choices, but their use for precise short-term delay prediction in aviation is constrained without high-quality operational data. Future work should explore integration of stochastic max-plus models, application to networks with shorter periodicity, and validation using airline-provided operational datasets.

I. Introduction

Flight delays continue to pose a major challenge for the aviation industry. In 2023, 71% of flights in European Civil Aviation Conference (ECAC) states and 78% of United States domestic flights arrived on time [1, 2]. Per industry convention, a flight is considered on time if it arrives within 15 minutes of its scheduled time [1–4]. Evidently, more than 20% of flights were delayed, with an estimated annual cost of tens of billions of dollars worldwide [5–8]. While many factors contribute to flight delays, delay propagation is a key driver. Propagated delay refers to secondary delays caused by late arrivals of aircraft, crew, passengers, or cargo from previous delayed flights, in contrast to primary delays such as weather, technical issues, or airspace congestion.

Although the financial and operational impact of delay propagation is significant, research on this topic has historically lagged behind that on primary delays [9–11]. Existing studies often focus on short-term forecasting or the effects of individual delay events, using approaches such as statistical models, machine learning, simulation, or queueing theory. While useful, these methods often struggle to model delay propagation at a high level of operational detail across entire networks, or they require high computational resources [12].

This paper explores the use of *max-plus linear systems* as a novel approach to model delay propagation in aviation networks. Max-plus algebra provides a computationally efficient framework for modelling cyclic, discrete event-based systems with precedence constraints. Max-plus linear systems have been applied in railway networks to analyse timetable stability, predict delay propagation, and compute system stability characteristics such as the maximum cycle mean and stability margins [13–16]. However, their application to aviation networks remains limited, with only a conceptual study conducted on an artificial network [17]. No prior study has applied scheduled max-plus linear systems to real-world aviation networks for stability analysis or network-wide delay propagation modelling.

The main contribution of this paper is to explore to what extent max-plus linear systems can be used to model, analyse, and simulate delay propagation in aviation networks. Using the Hawaiian Airlines network as a case study, this work:

- Develops a detailed methodology to model aviation networks as max-plus linear systems, including aircraft rotation and passenger transfer precedence relations.
- Explores the value of calculating stability indicators such as the maximum cycle mean, recovery times, and stability margins to assess the structural robustness of an aviation network.
- Simulates delay propagation under various initial delay scenarios and evaluates the model's predictive accuracy and sensitivity.
- Evaluates the impact of including passenger transfer precedence relations between flights in the system.

This paper is structured as follows. Section II provides an overview of relevant literature on delay propagation modelling and max-plus algebra. Section III presents the modelling methodology, including the formulation of the max-plus linear system and the calculation of stability indicators. Section IV introduces the Hawaiian Airlines case study and describes the dataset preparation. Section V presents the main results, including stability analyses and delay propagation simulations. Section VI discusses the broader implications of the findings. Finally, Section VII presents the main conclusions of the paper.

II. Related Work

Delay propagation in aviation networks has been addressed through a variety of modelling approaches, each with its own strengths and limitations. This section first provides an overview of these approaches and their key characteristics in Section II.A. It then identifies the remaining gaps in the literature and introduces max-plus linear systems as a novel framework with the potential to address them in Section II.B.

A. Overview of Delay Propagation Modelling Approaches

A recent review [12] classified modelling approaches for delay propagation into seven categories: economic models, statistical models, epidemic-spreading models, machine learning models, simulation-based models, queueing models, and network representation models. Each category has distinct strengths and limitations.

Economic, statistical, and epidemic-spreading approaches are computationally efficient but often oversimplify the complex, non-linear dependencies between flights. Economic and statistical models frequently focus on total impact estimation [18–20] or on regression-based relationships between delay causes and effects [21–23]. A relevant indicator from economic approaches is the *delay multiplier* (DM) [18], which shows the impact of an initial delay on propagated delay and is defined as $\text{DM} := \frac{\text{total delay minutes}}{\text{primary delay minutes}}$. Epidemic-spreading approaches [24, 25] model either airports or flights as nodes across which a delay 'infection' can spread.

Simulation-based and machine learning models are more flexible in representing complex relationships but tend to be computationally intensive. Machine learning techniques such as random forests, support vector machines and neural networks have been used to predict delay propagation for single flights or entire networks [26, 27]. However, their 'black-box' nature limits interpretability, and retraining is often required when schedules change. Simulation-based models have been used to analyse complete networks [28–30] or to investigate specific delay propagation mechanisms [31, 32]. Some works incorporate crew and passenger connectivity [28, 29], but without using real connectivity data.

Queueing models represent airports [33, 34], runways [35], or specific airport and airspace regions [36] as servers for which aircraft must queue. These models are well-suited to analysing delay propagation at bottlenecks but can become inefficient for network-wide studies.

Most closely related to the technique presented in this work are the network representation models. Within this category, several subtypes can be distinguished. Some studies use complex network theory to compute network characteristics such as degree and betweenness, or to construct a delay causality network [37–39]. Bayesian networks [20, 40–42] model delay propagation as a set of probabilistic cause-effect relationships between various factors, captured in a directed acyclic graph.

A notable study models the network itself as a directed acyclic graph (DAG) of event precedence [43]. In this study, slack times between events are computed using a topological sort combined with a shortest-path algorithm. Building on this concept, delay propagation trees [20, 42, 44, 45] extend the DAG event precedence approach by tracing the full set of downstream flights affected by a primary delay and quantifying how that delay propagates through successive connections in the network. This representation allows the identification of critical flights or connections that contribute disproportionately to overall delay spread, and has been extended to include crew and passenger dependencies [46]. Closely related are Petri net models [47, 48], which represent the system as a network of places and transitions connected by arcs. The state of the system is indicated by a distribution of tokens across the places. Each transition is associated with a fixed or bounded firing time, allowing explicit representation of event sequencing, concurrency, and timing constraints. The granularity, scope, and computational demands of network representation models vary widely. However, there is always a trade-off between these three elements.

B. Research Gap and the Role of Max-Plus Linear Systems

As evident from the previous section, a wide variety of delay propagation modelling approaches exist. However, no model currently provides network-wide coverage at high granularity while maintaining low computational cost. A potential solution is to represent an aviation network as a max-plus linear system.

Max-plus linear systems offer a computationally efficient framework for modelling discrete event systems with cyclic behaviours and precedence relations. Their application has been well established in railway networks, where they are used to model periodic timetables, compute system characteristics such as cycle time and stability margins, and simulate delay propagation [13–16]. The key advantage of modelling scheduled systems in max-plus algebra lies in the linear formulation of otherwise complex scheduling constraints. Its close link to graph theory enables the use of efficient graph algorithms to compute system properties, making the calculation of delay propagation effects and stability characteristics highly efficient.

Applications of max-plus linear systems in aviation networks remain limited. One example computed the maximum cycle mean and established a realisable timetable for a small artificial hub-and-spoke network [17], but without stability analysis or a delay propagation algorithm. Other studies have recognised the usefulness of the max() operation to model delay propagation [21, 23, 46, 49, 50], but did not use max-plus algebra. Max-plus algebra has also been applied to model aircraft queuing [51], but only for a single airport.

Thus, there remains a gap in the literature for scalable, network-wide models of delay propagation in aviation that can simultaneously capture detailed operational precedence relations and provide system-wide stability insights. Max-plus linear systems, as applied in railway networks, offer an underexplored approach to address this gap.

III. Methodology

This research models a real-world aviation network as a scheduled max-plus linear system. In a scheduled max-plus linear system, a network is modelled as a set of arrival and departure events and precedence relations between them. The precedence relations show that an event cannot take place before a preceding event has occurred and a given *process time* has passed. This section aims to outline this methodology.

Throughout this section, the example network shown in Fig. 1 will be used for illustration. The network consists of a hub airport B connected to two spoke airports A and C. Each day, one flight arrives at B from A and one flight arrives from C. Passengers can transfer between these inbound flights at B. Later the same day, one flight departs from B to A and one flight departs to C.

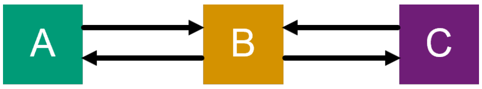

Figure 1. Example network consisting of hub airport B connected to spoke airports A and C.

This section is structured as follows. First, Section III.A will introduce max-plus algebra and Section III.B will show its relation to precedence graphs. Afterwards, Section III.C presents scheduled max-plus linear systems. A delay propagation model based on these systems is presented in Section III.D. Section III.E then outlines how to calculate the maximum cycle mean. Section III.F highlights how stability characteristics can be calculated using max-plus linear systems. Finally, Section III.G outlines how the precedence constraints in a scheduled max-plus linear system can be set up for an aviation network.

A. Introduction to Max-Plus Algebra

Max-plus algebra is somewhat of a misnomer, as it is not formally an algebra in the strictly mathematical sense. Rather, max-plus algebra refers to the idempotent semiring of extended real numbers $\{\mathbb{R}_\varepsilon, \oplus, \otimes\}$ [52–55]. This section aims to outline the basic operations and matrix operations of max-plus algebra.

1. Basic Operations in Max-Plus Algebra

In max-plus algebra, the standard algebraic addition and multiplication operations are replaced by maximum and addition operations. In mathematical notation, the following operations are defined [52, 53]:

$$a \oplus b := \max(a, b),$$

$$a \otimes b := a + b.$$

In conventional addition, the zero element 0 is absorbing: Adding 0 to a value returns the initial value. In max-plus addition, $\varepsilon := -\infty$ is defined as the zero element and is similarly absorbing: $a \oplus \varepsilon = a = \varepsilon \oplus a$ and $a \otimes \varepsilon = \varepsilon = \varepsilon \otimes a$ for all $a \in \mathbb{R}$ [52, 53]. ε can be used to define the numbers used in max-plus algebra as $\mathbb{R}_\varepsilon := \mathbb{R} \cup \{-\infty\}$ [54]. Note that as ε is the zero element, a non-zero element automatically refers to a finite element [13].

Max-plus addition is commutative, associative and idempotent ($a \oplus a = a$). Due to the idempotency, max-plus addition has no reverse operation. Max-plus multiplication is associative and has a unit element $e := 0$. Multiplication is distributive over addition, meaning that $\otimes$ has priority over $\oplus$. $a \otimes b$ can therefore also be written as ab, as is common in conventional algebra [13, 52, 53].

The max-plus-algebraic power operation is defined as $x^{\otimes r} = x \otimes x^{\otimes r-1}$ for an $x \in \mathbb{R}_\varepsilon$ and $r \in \mathbb{N}_0$, where $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$ [52, 53]. Note that $x^{\otimes 0} = 0$ and $x^{\otimes -1} = -x$. For the special case $\varepsilon^{\otimes r}$, $\varepsilon^{\otimes r} = \varepsilon$ if $r > 0$, $\varepsilon^{\otimes r}$ is undefined if $r < 0$, and $\varepsilon^{\otimes 0} = 0$ by definition [54]. If it is clear max-plus algebra is used in an equation, the denotion $\otimes$ is often not included in the power and only x^k is shown [13, 52].

2. Matrix Operations in Max-Plus Algebra

Max-plus algebra is particularly useful to translate a system of non-linear evolution equations in conventional algebra to linear equations in max-plus algebra. This allows for the use of linear algebra techniques and for representing the system in a state-space notation [55]. As such, it makes sense to also define matrix operations for max-plus algebra. If $A, B \in \mathbb{R}_\varepsilon^{m\times n}$ and $C \in \mathbb{R}_\varepsilon^{n\times p}$ then for all i, j [52, 53]:

$$(A \oplus B)_{ij} := a_{ij} \oplus b_{ij} = \max(a_{ij}, b_{ij}),$$

$$(A \otimes C)_{ij} := \bigoplus_{k=1}^{n} a_{ik} \otimes c_{kj} = \max_k(a_{ik} + c_{kj}).$$

The max-plus-algebraic zero and identity matrices, $\mathcal{E}_{m\times n}$ and E_n, are defined as follows. $(\mathcal{E}_{m\times n})_{ij} := \varepsilon$ for all i, j; $(E_n)_{ii} := 0$ for all i and $(E_n)_{ij} := \varepsilon$ for all i, j with $i \neq j$ [53]. Note that $\mathcal{E}_{m\times n}$ has dimensions $m\times n$ and E_n has dimensions $n\times n$. Max-plus algebraic matrix power is defined as $A^{\otimes k} := A \otimes A^{\otimes k-1}$ for $k = 1, 2, \ldots$ and $A \in \mathbb{R}_\varepsilon^{n\times n}$. Additionally, $A^{\otimes 0} := E_n$ [53]. Note that max-plus algebraic power is undefined for $k < 0$ [53].

B. Max-Plus Algebra in Precedence Graphs

There exists a strong connection between max-plus algebra and graphs. A max-plus matrix $A \in \mathbb{R}_{\varepsilon}^{n \times n}$ can be represented as a precedence graph $\mathcal{G}(A)$. In this graph, nodes represent events, and directed arcs represent precedence relations between events. An arc (j, i) exists with weight a_{ij} if and only if $a_{ij} \neq \varepsilon$, indicating that event i cannot occur until at least a_{ij} time units after event j [13, 53].

This correspondence allows for mutual benefits between max-plus algebraic and graph-theoretical methods. Properties defined algebraically can often be interpreted structurally in the associated graph, and vice versa. An example of this is the calculation of whether a max-plus algebraic matrix A is *irreducible*, which plays an important role in determining the connectivity and eigenvalues of a system.

Algebraically, irreducibility ensures that all variables in the system are coupled through the max-plus dynamics. In graph theory, this indicates that the precedence graph $\mathcal{G}(A)$ is *strongly connected*, meaning that a path exists between any two nodes. This condition is often assumed in the analysis of max-plus systems, where it guarantees the existence of a unique eigenvalue. A matrix $A \in \mathbb{R}_{\varepsilon}^{n \times n}$ is said to be irreducible if:

$$A^{+}_{ij} \neq \varepsilon \quad \text{for all } i, j \in \{1, \ldots, n\}, \tag{1}$$

where A^{+} is the *critical path matrix*. The critical path matrix collects, for each node pair, the maximum-weight path.It can be calculated algebraically: For each $k \in \mathbb{N}_0$, the entry $(A^{\otimes k})_{ij}$ gives the weight of the longest path of length k from node j to node i, or ε if no such path exists. By summing these entries over all relevant path lengths, one obtains the critical path matrix [13, 52]:

$$A^{+} := A \oplus A^{\otimes 2} \oplus \ldots . \tag{2}$$

However, the calculation in Eq. (2) is computationally intensive. In graph theory, many more efficient critical path algorithms exist. In this work Johnson's algorithm [56, 57] is used. Like many other path algorithms in this work, Johnson's algorithm has initially been defined as a shortest path algorithm but can be adapted to find the longest path. The algorithm first applies the single-source Bellman-Ford algorithm from a dummy node to reweight the graph to only contain non-positive values. It then repeatedly applies Dijkstra's label-setting algorithm from each node to calculate the critical path and afterwards reverts the graph to its original weights. Johnson's algorithm has been shown to run in $O(n^2 \log(n) + nm)$, with n the number of nodes and m the number of arcs [57]. For large sparse graphs, it will often converge even faster [16].

C. Scheduled Max-Plus Linear Systems

Max-plus algebra and precedence graphs can be used to model transport networks. In this work, an aviation network will be modelled as a *scheduled max-plus linear system*. This modelling approach assumes networks are cyclical: The same set of events and precedence relations repeat at some cycle time [13]. This is the case for transport networks which operate under a periodic timetable.

In the scheduled max-plus linear system, arrivals and departures are modelled as discrete periodic events and represented as nodes in a precedence graph. An arrival is the moment an aircraft arrives at its block As each event occurs once every kth period, the actual event times can be collected in a vector $x(k)$. Similarly, the scheduled event times can be collected in a vector $d(k)$. The occurrence of an event $x_i(k)$ is then dependent on the scheduled event time $d_i(k)$ and the occurrence of events which precede it $x_i(k - \mu_{ij})$. Here, the *period shift* μ_{ij} represents how many periods ago the preceding event occurred.

$a_{ij\mu}$ is the process time of the precedence relation from event j to event i across period shift μ. If there is no precedence relation, $a_{ij\mu} = \varepsilon$ per definition. $a_{ij\mu}$ can show a precedence relation with an event within this cycle ($\mu = 0$) or from preceding cycles ($\mu = 1, 2, \ldots$). For a given event pair (i, j), a different $a_{ij\mu}$ can exist for each μ. However, this does not occur in this research and $a_{ij\mu}$ will thus be simply denoted as a_{ij}.

Combining all terms and defining n as the number of events, the following max-plus algebra recursive equations can be defined [13]:

$$x_i(k) = \bigoplus_{j=1}^{n} (a_{ij} \otimes x_j(k - \mu_{ij})) \oplus d_i(k) \qquad \text{for all } i = 1, \ldots, n. \tag{3}$$

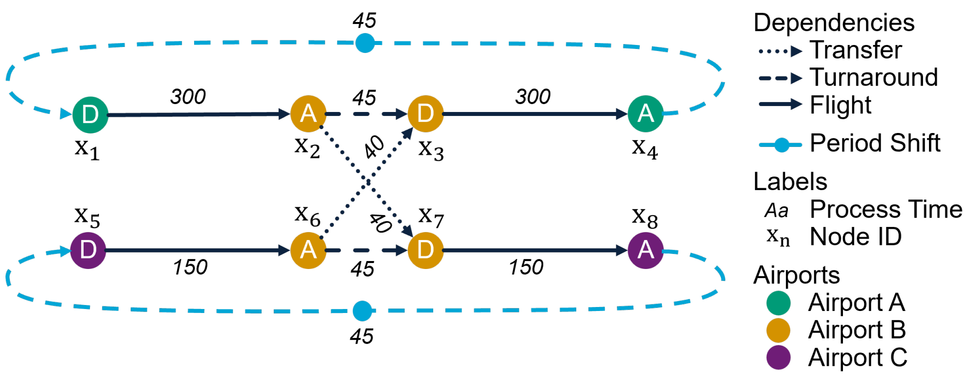

Figure 2. Precedence graph of the example network of Fig. 1. The letters in the event nodes denote the distinction between arrival (A) and departure (D) events. Process times are given in minutes.

For $l \in \mathbb{N}_0$, all process times a_{ij} where $\mu_{ij} = l$ can be collected into matrices A_l, with $[A_l]_{ij} = \varepsilon$ if no such term exists. Equation (3) can then be written in vector notation as follows [13]:

$$x(k) = A_0 x(k) \oplus \cdots \oplus A_p x(k-p) \oplus d(k) = \bigoplus_{l=0}^{p} A_l x(k-l) \oplus d(k), \quad (4)$$

for some $p \in \mathbb{N}$, which is the *order* of the system. The scheduled event time vector $d(k)$ is formally defined as $d(k) = d_0 \otimes T^k$, where $d_0 = d(0)$ and T is the scheduled cycle time.

Finally, Eq. (4) can also be written in a polynomial form. This is done by using the backward-shift operator γ to combine the matrices A_l into a single polynomial matrix $A(\gamma)$. The presence of γ^l is best understood to represent a precedence relation with an event from period $k-l$, or mathematically $\gamma^l x(k) = x(k-l)$. This results in the following system [13]:

$$x(k) = \bigoplus_{l=0}^{p} A_l \gamma^l x(k) \oplus d(k) = A(\gamma)x(k) \oplus d(k). \quad (5)$$

As an example, consider the network of Fig. 1. Adapting this into a precedence graph results in Fig. 2, assuming the network is operated by two aircraft and arbitrarily choosing process times for transfers, turnarounds and flights. The max-plus linear system is then given by:

$$x(k) = A_0 x(k) \oplus A_1 x(k-1) \oplus d(k),$$

with corresponding process time matrices:

$$A_0 = \begin{bmatrix} \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon \\ 300 & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon \\ \varepsilon & 45 & \varepsilon & \varepsilon & \varepsilon & 40 & \varepsilon & \varepsilon \\ \varepsilon & \varepsilon & 300 & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon \\ \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon \\ \varepsilon & \varepsilon & \varepsilon & \varepsilon & 150 & \varepsilon & \varepsilon & \varepsilon \\ \varepsilon & 40 & \varepsilon & \varepsilon & \varepsilon & 45 & \varepsilon & \varepsilon \\ \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & 150 & \varepsilon \end{bmatrix},$$

$$A_1 = \begin{bmatrix} \varepsilon & \varepsilon & \varepsilon & 45 & \varepsilon & \varepsilon & \varepsilon & \varepsilon \\ \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon \\ \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon \\ \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon \\ \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & 45 \\ \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon \\ \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon \\ \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon & \varepsilon \end{bmatrix}.$$

D. Delay Propagation Model

The scheduled max-plus linear system of Eq. (4) can be adapted to form a complete delay propagation model. Assume all event times up to a time $t_0 \in [T, 2T)$ are known. Then, the initial conditions to the system are all $x(k)$ of the last p periods $x_{1-p}, \ldots, x_0 \in \mathbb{R}^n_\varepsilon$ such that $x(l) = x_l$ for $l = 1-p, \ldots, 0$. The event times for events with scheduled event time $d_i(1) \in [T, t_0]$ are assumed to be known. The still unknown event times for events with scheduled event time $d_i(1) \in (t_0, 2T)$ can be initially estimated as $x_i(1) = d_i(1)$. Combined, these event times form the partially unknown initial state vector of period 1, which is denoted as $x(1^-) = x_1 \in \mathbb{R}^n_\varepsilon$. Given a timetable vector $d_0 > \varepsilon$ with entries in a timetable period of length T, and defining an output vector $z(k)$ denoting the delays in period k, the scheduled max-plus linear system becomes [13, 14]:

$$\begin{cases} x(k) = \bigoplus_{l=0}^{p} A_l x(k-l) \oplus d(k), \\ d(k) = T \otimes d(k-1), \\ z(k) = x(k) - d(k), \\ x(0) = x_0, \ldots, x(1-p) = x_{1-p}, \\ x(1^-) = x_1,\ d(0) = d_0. \end{cases} \tag{6}$$

The presence of zero-order terms $A_0 x(k)$ in Eq. (6) introduces a complication: the value of $x(k)$ is required to compute itself. This implicit dependency makes the direct computation of $x(k)$ non-trivial, as $x(k)$ is not fully defined until the equation is resolved. In other words, due to intra-period precedence relations, the scheduled max-plus update equation cannot be evaluated in a purely recursive manner.

One solution is to eliminate these zero-order terms and reformulate the system into an explicit (purely recursive) form. However, doing so would prevent modelling intra-period delay propagation and exclude initial delays occurring within the current period. To preserve these effects and maintain a complete representation of propagation dynamics, the implicit equation involving $x(k)$ is solved using the *transitive closure matrix* A_0^*, under the assumption that the precedence graph $G(A_0)$ is acyclic [13].

The transitive closure matrix is defined as $A_0^* = \bigoplus_{l=0}^{n-1} A_0^l$. Equivalently, $A_0^* = E \oplus A_0^+$ and A_0^* can thus be calculated using efficient algorithms to calculate the critical path matrix A_0^+ [52]. As $G(A_0)$ is acyclic, the critical path matrix A_0^+ can be computed in linear time $O(n+m)$ using a topological sorting algorithm and a dynamic programming algorithm [57, 58].

This yields a well-defined and computable expression for $x(k)$ despite the zero-order precedence relations. The first equation in the system of Eq. (6) then becomes [13]:

$$x(k) = \bigoplus_{l=1}^{p} A_0^* A_l x(k-l) \oplus d(k) \text{ for all } k \geq 2, \tag{7}$$

and $x(1)$ is given by:

$$x(1) = A_0^* \otimes \left(x_1 \oplus \bigoplus_{l=1}^{p} A_l x_{1-l} \right) \oplus d(1). \tag{8}$$

E. Calculating the Maximum Cycle Mean

The polynomial form of the scheduled max-plus linear system given in Eq. (5) is valuable for calculating the *maximum cycle mean* of the associated precedence graph. The maximum cycle mean is defined as the largest average weight per arc over all circuits in the graph, where the weight of a circuit corresponds to its total process time and the mean is obtained by dividing by the number of arcs in the circuit. In the context of periodic event scheduling, it represents the system's cycle time, i.e., the minimal achievable period for repeating the schedule without conflict.

There is a conceptually straightforward way to calculate the maximum cycle mean using the precedence graph $\mathcal{G}(A(\gamma))$, which will be discussed first. Afterwards, a more computationally efficient method using the policy iteration algorithm will be presented.

1. Defining the Maximum Cycle Mean

Max-plus algebra can be used to calculate the maximum cycle mean of a system. This calculation is in fact a max-plus algebraic eigenvalue problem. Generally, for an $A \in \mathbb{R}^{n\times n}_\varepsilon$, if there exists $\lambda \in \mathbb{R}_\varepsilon$ and $v \in \mathbb{R}^n_\varepsilon$ with $v \neq \mathcal{E}_{n\times 1}$ such that $A \otimes v = \lambda \otimes v$, then λ is a max-plus-algebraic eigenvalue of A and v is the corresponding eigenvector. It has been shown that every square matrix with entries in $\mathbb{R}_\varepsilon$ has at least one eigenvalue [54].

For a square polynomial matrix $A(\gamma)$, the generalised eigenvalue problem is to find a nonzero scalar $\lambda \in \mathbb{R}$ and a nonzero vector $v \in \mathbb{R}^n$, such that $A(\lambda^{-1}) \otimes v = v$ [13].

Finding the solution to this eigenvalue problem algebraically can be complex. However, the correspondence of the max-plus linear system to a precedence graph allows for a graph-theoretical approach to calculating the maximum cycle mean. If $A(\gamma)$ is a polynomial matrix with acyclic subgraph $G(A_0)$, then the eigenvalue λ is equal to the the maximum cycle mean η and can be calculated as follows [13]:

$$\eta = \max_{\xi \in C} \frac{w(\xi)}{\mu(\xi)}, \tag{9}$$

where C is the set of all elementary circuits in $\mathcal{G}(A(\gamma))$, $w(\xi)$ is the weight of circuit ξ and $\mu(\xi)$ is the number of period shifts in circuit ξ.

For sake of simplicity, a shorthand notation of the polynomial matrix evaluated at a scalar is defined as $A(\nu^{-1}) = A_\nu$ for a scalar $\nu \in \mathbb{R}_\varepsilon$. If λ is a generalised eigenvalue of $A(\gamma)$, A_λ has eigenvalue e, as $A_\lambda \otimes v = v = e \otimes v$. In this case, the critical circuits in $G(A_\lambda)$ will have mean circuit weight e. A *critical circuit* is a circuit with maximum cycle mean. For each node i on a critical circuit, the diagonal values of the critical path matrix $[A^+_\lambda]_{ii} = e$. The associated columns of A^+_λ are in fact the generalised eigenvectors of $A(\gamma)$ associated with λ. This is useful, as eigenvectors can be interpreted as initial timetable vectors $d_0 = v$ such that the scheduled max-plus linear system exhibits periodic behaviour with minimal cycle time $T = \lambda$ [13].

If the polynomial matrix $A(\gamma)$ is irreducible, there is a unique eigenvalue. If this is not the case, there is not necessarily a unique eigenvalue λ. Rather, if there are multiple strongly connected components in the network, each component can have its own eigenvalue. Note that in this case, Eq. (9) can be applied to each component. The eigenvalue which corresponds to the most critical strongly connected component of the system is denoted as λ_0.

In the example of Fig. 2, the upper circuit of nodes 1-2-3-4-1 is the critical circuit, with a corresponding maximum cycle mean $\lambda_0 = 690$ minutes.

2. *Calculating the Maximum Cycle Mean Using the Policy Iteration Algorithm*

While the calculation of the maximum cycle mean using Eq. (9) is conceptually straightforward, enumerating all cycles in a network is computationally intensive, especially for large or dense graphs. An efficient alternative is the *policy iteration algorithm*, which solves the generalised max-plus eigenvalue problem by iteratively refining a selection of arcs until convergence [16, 59].

In this algorithm, a *policy* is a selection of one incoming arc per node (i.e. one predecessor per event), forming a subgraph of the original precedence graph. The initial policy is chosen based on a heuristic. In this work, the arc with the highest weight adjusted for period shift is initially selected, i.e. $\frac{a_{ij}}{\mu_{ij}}$ if $\mu_{ij} > 0$ and a_{ij} otherwise. The algorithm then alternates between two steps: policy evaluation and policy improvement.

The *policy evaluation* step begins by identifying all strongly connected components in the subgraph defined by the current policy. For each component, a circuit is traced, and its cycle mean is calculated as the total process time divided by the total number of period shifts. This value is then propagated along the rest of the component using a breadth-first search. For each event i, the potential v_i is updated as:

$$v_i = w_{ij} - \mu_{ij}\chi_j + v_j,$$

where (j, i) is the selected arc, w_{ij} is the process time, μ_{ij} is the period shift, and χ_j is the cycle mean of the component.

In the *policy improvement* step, the algorithm evaluates whether replacing a selected arc with an alternative arc from the original graph would improve either the cycle mean or the node potentials. Two forms of improvement are considered:

- *First-order improvement:* If a different predecessor provides a higher cycle mean, the policy is updated accordingly.
- *Second-order improvement:* If the cycle mean remains unchanged, but a different predecessor yields a higher potential v_i, and both nodes are in the same component, the policy is still updated.

This process of evaluation and improvement is repeated until no further updates occur. The final policy defines the set of critical circuits, and the corresponding cycle means provide the eigenvalues λ_c of the system. The node potentials v_i form the associated eigenvectors. In the case of a reducible matrix A, multiple eigenvalues can exist. The maximum among these corresponds to the maximum cycle mean λ_0 [16, 59].

F. Calculating Stability Characteristics of Max-Plus Linear Systems

Various stability characteristics of the max-plus linear system can be efficiently calculated. This section will first discuss those general indicators which can be calculated using the maximum cycle mean. Afterwards, it will introduce the recovery matrix.

1. Stability Analysis Using the Maximum Cycle Mean
The maximum cycle mean λ_0 is significant as it represents the minimal attainable cycle time. If there are multiple eigenvalues, λ_i represents the minimal attainable cycle time of each strongly connected component i of the system. The maximum cycle mean can be used to calculate various other useful stability indicators [13].

Obviously, if the scheduled cycle time T is less than the maximum cycle mean, the timetable will always generate and propagate delay. In other words, a timetable is unstable if $\lambda_0 > T$ and, conversely, a timetable is stable if $\lambda_0 < T$. If $\lambda_0 = T$, the system is critical.

A simple indicator of the network robustness is the *network throughput* $\rho = \lambda_0/T$. If $\rho = 1$, the network is saturated and critical. Naturally, it is infeasible to operate a network at $\rho = 1$: Delays will always occur, and some buffer is necessary to absorb these delays. However, very low values of ρ can indicate that the network is underutilised. A similar indicator is the difference $\Delta_1 = T - \lambda_0$, which gives the average amount of slack on the critical circuits [13].

While the indicators above give insights based only on the critical circuit, sometimes more should be taken into account. The stability margin Δ_2 is the maximum simultaneous increase of all process times to still operate the network within cycle time T. It is the solution to the eigenvalue problem $(\Delta_2 \otimes A(T^{-1})) \otimes v = v$. Similarly to Eq. (9), Δ_2 can then be calculated in a conceptually simple way based on the elementary circuits C in $\mathcal{G}(A(\gamma))$ [13]:

$$\Delta_2 = -\max_{\xi \in C} \frac{w(\xi) - \mu(\xi) \cdot T}{l(\xi)}, \qquad (10)$$

where $l(\xi)$ is the length of the elementary circuit ξ. However, it is more efficient to apply the policy iteration algorithm to the reweighted graph $A(T^{-1})$.

2. Stability Analysis Using the Recovery Matrix
While the previous stability indicators describe stability margins in the complete system, sometimes more granular information is desired. The *recovery time* r_{ij} between events j and i indicates how how much slack exists in the sequence between these events. It thus characterises the maximum delay at an event that will not propagate to a downstream event. All recovery times can be collected in the *recovery matrix* R.

R is defined such that each element r_{ij} denotes the minimum total slack time across all paths from event j to event i. Formally, given a stable scheduled max-plus linear system, R is defined as follows [13]:

$$r_{ij} = d_i^0 - d_j^0 - \left[A_T^+\right]_{ij}. \qquad (11)$$

By convention, if no path exists from j to i, $r_{ij} = -\varepsilon = \infty$.

G. Setting Up Precedence Constraints for Aviation Networks

Precedence constraints indicate that an event i cannot occur before event j has occurred and a process time a_{ij} has passed.

Generally, the process time should be the minimum time to complete the process that is achievable in most circumstances. This excludes the actual minimum process times which were achieved under unusually favourable circumstances, such as strong tailwinds. It is often desirable to take a percentile of all process times to exclude these very low values.

In aviation networks, the same timetable often repeats every 24 hours. Period shifts were thus calculated based on a fixed cycle time of 24 hours. For each 24-hour threshold crossed by a precedence relation between two events according to the schedule, the period shift increases by 1. Note that although scheduled events typically repeat daily, their precedence relations often will not. In this work, these relations are assumed to be identical each day.

The three types of considered precedence constraints are precedence constraints per flight, turnaround precedence constraints and passenger transfer precedence constraints. This section will discuss each type of precedence relation and the calculation of their process time.

1. Precedence Constraints per Flight
The most straightforward precedence constraints are those for flights. These show that the arrival of a flight cannot occur before its departure. In the example of Fig. 2, the arc from node 1 to node 2 represents such a precedence constraint, expressed mathematically as $x_2(k) \geq x_1(k) \otimes a_{21}$.

For the process time, airline block times are used. The block time is the time of a flight measured from the start of movement out of the parking position to the end of movement. Airlines typically have limited influence on block times. The only factor which airlines are in direct control of is airspeed. However, flying at a higher airspeed typically only reduces flight times by 3 to 5 minutes per hour. Its cost-effectiveness is limited and it is not always technically feasible [60]. Airlines can also request shorter flight routes or more efficient gate assignment, but this is very dependent on air traffic control operations. Besides, these are desirable requests to improve operational efficiency even when not delayed.

As such, the minimum generally achievable block time will lie only slightly below the minimum block time. For this reason, the 40^{th} percentile block times of all flights in the dataset between an origin and destination were used for flight precedence constraints.

2. Turnaround Precedence Constraints

Turnaround precedence constraints are based on the aircraft rotations and indicate that the departure of an aircraft at an airport cannot occur before the arrival of that aircraft at that airport. In the example of Fig. 2, the arc from node 2 to node 3 represents a turnaround precedence constraint, expressed mathematically as $x_3(k) \geq x_2(k) \otimes a_{32}$.

Aircraft rotations are reconstructed based on flight schedules and aircraft assignments. If all flights for all tail numbers are available, these can be used directly to build real flight sequences. However, if certain flights or tail numbers are missing, it is not possible to set up full flight sequences. For these flights, certain precedence relations were assumed based on a greedy aircraft assignment algorithm.

The objective of the greedy algorithm was to assign each outbound flight a plausible predecessor flight based on temporal feasibility and the inbound aircraft at the airport. This was done per airport by matching inbound and outbound flights within the same period. Inbound flights were first sorted by their scheduled arrival time and assigned a *ready time*, calculated as the arrival time plus the minimum turnaround time (TAT). Outbound flights were sorted by scheduled departure time. As the algorithm iterated over outbound flights, it assigned the earliest inbound flight whose ready time was no later than the outbound departure. If no suitable inbound flight was available, flights from the previous cycle were assigned. This created a cyclic system.

The process time of the turnaround precedence constraint is the minimum achievable turnaround time between arrival and departure at the stand, which is highly dependent on the airport, the type of aircraft and the airline.

3. Passenger Transfer Precedence Constraints

Passenger transfer precedence constraint indicate that the departure of an aircraft at an airport cannot occur before the arrival of an aircraft containing transferring passengers at that airport. In the example of Fig. 2, the arc from node 2 to node 7 represents a turnaround precedence constraint, expressed mathematically as $x_7(k) \geq x_2(k) \otimes a_{72}$.

To ensure only meaningful constraints were included in the model, only transfers used on average by more than a certain number of passengers were included.

The process time of transfer constraints is the *minimum connection time* (MCT). The MCT between flights can vary significantly between airports, airlines and flights. Generally, connections between domestic flights are often possible within 30 minutes, particularly at smaller airports [3, 61].

IV. Case Study and Data Preparation

To properly study the effectiveness of scheduled max-plus linear systems, this work uses the Hawaiian Airlines network as a case study. This section outlines the selection of the case study and the preparation of the datasets used to construct the scheduled max-plus linear system.

First, Section IV.A outlines the selection of a suitable case study, which is described in Section IV.B. Then, the main datasets are introduced in Section IV.C. Section IV.D describes which timetable was selected to set up the events. The setting up of the precedence structure is described in Section IV.E. Finally, Section IV.F outlines the 7 initial delay scenarios to which the case study will be exposed.

A. Case Study Selection

Much research into delay estimation is not very reproducible [9]. Accordingly, this study prioritises the use of publicly available databases to improve reproducibility. Two databases including actual arrival and departure times were publicly available: EUROCONTROL [62] and the US Bureau of Transportation Statistics (BTS) [63]. The BTS database was selected as it included scheduled arrival and departure times as well as delay causes.

The BTS database is restricted to domestic flights of US carriers with at least 0.5% of total US domestic scheduled-service passenger revenues. Consequently, many networks are either very large or incomplete due to missing international legs. Hawaiian Airlines (HA) was selected as a case study for its small and largely domestic network. Prior to its 2024 acquisition by Alaska Airlines, HA operated as an independent carrier in the US [64].

Data from 2017 was selected as a passenger transfer dataset was only available for 2017. In this year, February was selected as it is a month where daylight saving time does not apply and there are no seasonal timetable adjustments due to Christmas holidays. The selected data included all US domestic Hawaiian Airlines flights in February, 2017, totalling 5789 flights serving 17 airports.

B. Description of Case Study

HA operates flights between Hawaiian islands, to the mainland US and to international destinations in the Asia Pacific region. Its inter-island network is operated by a narrow-body fleet of Boeing 717 aircraft. Other flights are operated by wide-body A330 or B767 aircraft. HA operates the Honolulu (HNL) and Kahului (OGG) airports as hubs from which passengers can fly to the mainland. Note that the BTS only provides data on US domestic flights, and these are thus the focus of this research. The network is shown in Fig. 3.

C. Description of Datasets

The BTS database consists of multiple datasets. The two main datasets are the flights on-time performance dataset (ONTIME) and the passenger connection dataset (PAX). Note that the latter is not publicly available. These are supported by auxiliary datasets containing airport, airline and aircraft specifications such as coordinates, carrier codes and aircraft types respectively.

The ONTIME dataset contains all scheduled flights within the selected period. It includes the origin, destination, carrier, flight number and assigned tail number, as well as the scheduled departure and arrival times as indicated in the central reservation system (CRS). The main columns of interest are those indicating the actual departure, take-off, landing and arrival times. The dataset also includes how many minutes of delay can be attributed to each delay cause the BTS recognises. The BTS causes of delay are: carrier, weather, national aviation system (NAS), security, and aircraft arriving late.

The PAX dataset contains all possible passenger connections between two flights to which more than one passenger is assigned within the selected period. It includes the operating carrier, CRS departure and arrival times for each flight and how many passengers are assigned to a certain connection.

To prepare the datasets for use in this work, the ONTIME and PAX datasets were augmented with the auxiliary datasets. Additionally, times were transformed to Coordinated Universal Time (UTC) including dates. Finally, the first and last day of the month were excluded, as inspection of the dataset highlighted some unexpected cropping at the start and end of the month, skewing data analysis.

D. Setting Up the Timetable and Events

Scheduled max-plus linear systems are based on periodic timetables. Inspection of the data showed no clear periodicity within a day. However, a clear daily periodicity could be recognised. As such, 24 hours was set as the cycle time. The start of each period was set to 00:00 Hawaiian local time (UTC-10), with all flights scheduled to depart between 00:00 and 23:59 UTC-10 included in that period. Note that some flights may arrive in the following period.

Even though many flights are operated daily, not all are. Additionally, the exact arrival and departure times of each flight are not the same on each date. For this research, the timetable of February 21, 2017 was assumed to occur daily. This date was selected as it includes the frequent but non-daily flights to SJC and JFK airports. The day also contains numerous delay events, which is useful for validation purposes. As flights HA466 and HA551 lacked return flights on this date, they were removed to balance the number

Figure 3. Hawaiian Airlines network under study. Note that flights to Pago Pago are not shown.

of flights in and out of each airport.

All arrival and departure events could then be created based on this assumed timetable. The scheduled event time vector $d_0(k)$ was set up using the CRS scheduled arrival and departure times.

E. Setting Up the Precedence Constraints

The precedence constraints for the HA network were set up using the methodology described in Section III.G. This was straightforward for the precedence constraints per flight. However, setting up the turnaround precedence constraints and passenger transfer precedence constraints, and particularly the calculation of their process times, was not as simple. Both types of precedence constraints will thus be given further attention in this section.

1. Turnaround precedence constraints

The turnaround precedence constraints were set up differently for the inter-island flights and the mainland flights. For the inter-island flights, the ONTIME dataset includes all flights for all aircraft. As such, the real aircraft sequences could be used to set up the turnaround precedence constraints. For the first flight of each aircraft, the precedence relation with the flight of the previous day was assigned using the greedy algorithm.

For mainland flights, the data was incomplete as there are no international HA flights in the dataset. To set up these turnaround precedence constraint, the greedy algorithm was applied to all flights. A representative subset of the thus constructed aircraft sequences can be found in Fig. 4.

For turnaround precedence constraints' process times, HA indicates it can achieve TATs of under 30 minutes for its inter-island flights and TATs of 90 minutes for its mainland flights [65]. The actual turnaround times can be calculated by taking the difference between departure and arrival times for flights where the ONTIME dataset indicates there is delay due to an aircraft arriving late. For these flights, there is no turnaround buffer time remaining.

The minimum turnaround time in the dataset can then be calculated to equal 8 minutes, which is not generally achievable. The 25$^{\text{th}}$ percentile of actual turnaround times more correctly approximates the minimum achievable times [23]. These are shown per airport in Table 1 and corroborate the values given by HA. A distinction is made between the turnaround times for inter-island and mainland flights, which are operated by narrow-body and wide-body aircraft respectively. The process times for turnaround times are set to the values in Table 1 if applicable or set to 90 minutes otherwise.

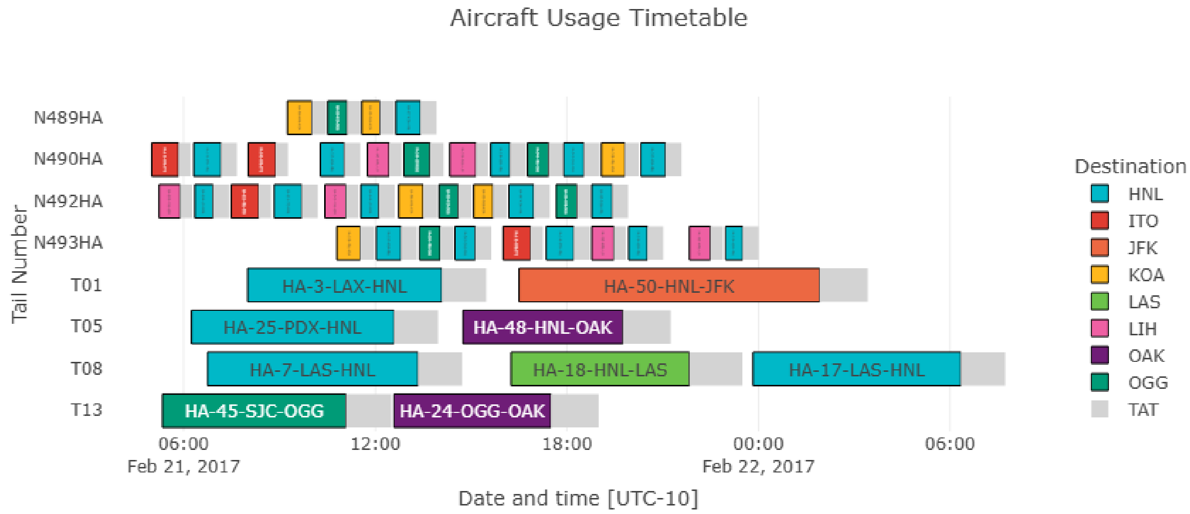

Figure 4. Representative subset of the timetable under study. Note the assumed tail numbers for the mainland flights (T*xx*). Turnaround times are indicated in gray.

Table 1. 25th percentile TAT per airport in minutes.

Airport	Inter-island	Mainland
HNL	30	82
ITO	23	–
KOA	26	–
LIH	23	–
OGG	25	85

2. Passenger Transfer Precedence Constraints

Passenger transfer precedence constraints were added from the PAX dataset. To ensure only meaningful constraints were included in the model, the data was first aggregated by flight pair using the mean number of passengers across the month. Note that certain flights operate at different times across the month. For these flights, only the transfers which matched in time with the selected timetable were used. Transfers with fewer passengers than a selected threshold could be excluded.

For the process times, HA recommends its passengers an MCT of 25 minutes for its inter-island flights and an MCT of 65 minutes (HNL) or 55 minutes (OGG) for transfers between its inter-island and mainland flights [66]. However, it is likely lower MCTs are feasible if the need arises. As such, MCTs were set to 80% of the values specified above.

F. Initial Delay Scenarios

To evaluate the model's behaviour under different operational conditions, a series of delay scenarios were defined. These scenarios serve two primary purposes: (1) to verify and validate that the model behaves in accordance with expectations based on its structural definitions, and (2) to analyse how delay propagates through the network in response to more complex perturbations. The scenarios are divided into two groups: verification and validation scenarios (V1-V4) and analytical scenarios (S1-S3).

V1) *Single flight arrival delay:* A 45-minute delay is applied to the early morning arrival of mainland flight HA17 from LAS to HNL.

V2) *Cross-period flight departure delay:* A 45-minute delay is applied to the departure of the same flight as in V1 (HA17). The precedence constraint between the arrival and the departure events of this flight includes a period shift.

V3) *Inter-island flight delay:* A 45-minute delay is applied to the departure of flight HA392 from HNL to ITO. This flight and its successors are very tightly scheduled, and allow clear inspection of any structural delay propagation effects.

V4) *Real-world delay replication:* Actual initial delays are taken from the BTS ONTIME dataset and applied to all arrival events. Propagated delays are removed from the initial delay as much as possible by subtracting the late aircraft arrival delays. Note that as the BTS groups crew and passenger delay propagation effects under carrier delay, it could be that delay caused by crew and passenger delay propagation is included in the initial delay.

S1) *Delay shock in most vulnerable flights:* The ten flights with the highest historical 75$^{\text{th}}$ percentile initial arrival delays are each delayed by that 75$^{\text{th}}$ percentile value. The propagated delay is also removed here by subtracting the late aircraft arrival delays. This scenario represents a targeted shock to flights statistically prone to delay.

S2) *Morning mainland arrival disruption:* All mainland flights arriving into HNL or OGG between 10:30 and 15:00 are delayed by 20 minutes. This represents the effect of a system-wide disruption (e.g., weather or congestion) affecting the morning inbound wave.

S3) *Temporary closure of KOA airport:* The Kona (KOA) airport is assumed to be closed between 11:00 and 13:00. All arrival and departure events scheduled during this window are rescheduled to occur sequentially from 13:00 onward with a minimum spacing of 4 minutes.

V. Results

This section aims to outline the results of this work. Section V.A will first present an overview of the resulting scheduled max-plus linear system. Section V.B then presents the system stability characteristics. Afterwards, Section V.C discusses the efficacy of the model's delay prediction. Section V.D will discuss the model's sensitivity to process time assumptions. Finally, Section V.E shows the impact of adding passenger transfer precedence constraints to the network.

A. Overview of the Resulting Model

The resulting model represents the Hawaiian Airlines flight network as a scheduled max-plus linear system. The final system consists of 404 events (202 flights) and up to 1122 precedence constraints. These constraints include 404 flight and turnaround constraints and up to 804 passenger transfer constraints, depending on the selected passenger threshold as discussed in Section IV.E. Note that 86 passenger transfer constraints overlap with the aircraft turnaround constraints. The maximum period shift between preceding events was 1, indicating a first-order system.

The model captures approximately 98.8% of the daily average of 206.75 flights operated by HA. It excludes those flights specified in Section IV.D and those not operated on February 21, 2017. Of the 202 flights, 164 were inter-island flights and 38 were flights to the mainland US. In the model, the flights were operated by 16 narrow-body aircraft and 19 wide-body aircraft. When including only flight and turnaround precedence constraints, the network contained four strongly connected components, each a circuit. This indicates that certain events will never influence each others' delays. This system forms the foundation for the delay propagation simulations and stability analyses discussed in the subsequent sections.

B. System Stability Characteristics

As described in Section III.F, scheduled max-plus linear systems allow for the computation of various stability characteristics that quantify the robustness of a given timetable structure. These indicators are presented in this section.

Table 2 summarises the system-wide metrics. Both the maximum cycle mean λ_0 and the derived indicators Δ_1 (the slack on the critical circuit) and ρ (the network throughput ratio) primarily show the network has a lot of nighttime buffer. As such, they offer limited insight into delay resilience in the operational core of the schedule.

The stability margin Δ_2 provides a more meaningful indicator in this context. It reflects the maximum uniform increase in all process times that the system can tolerate before becoming unstable. With a value of 35 minutes, Δ_2 indicates that the system is sensitive to widespread increases of turnaround and block times. However, it is still largely influenced by the major

Table 2. Key stability characteristics of the scheduled max-plus system.

Indicator	Value
Maximum Cycle Mean λ_0	17:50
2nd Cycle Mean λ_2	14:35
3rd Cycle Mean λ_3	11:54
Critical Slack $\Delta_1 = T - \lambda_0$	06:10
Stability Margin Δ_2	00:35
Network Throughput $\rho = \lambda_0 / T$	74.3%

night buffer and 35 minutes is a high value for all process times to increase by in an aviation network.

In addition to global indicators, more granular insights can be drawn from the recovery matrix R. Many entries in R span multiple days, due to the many (indirect) inter-day precedence relations which accumulate the night-time buffers. Practical stability concerns are then better reflected in the share of values below operationally relevant thresholds.

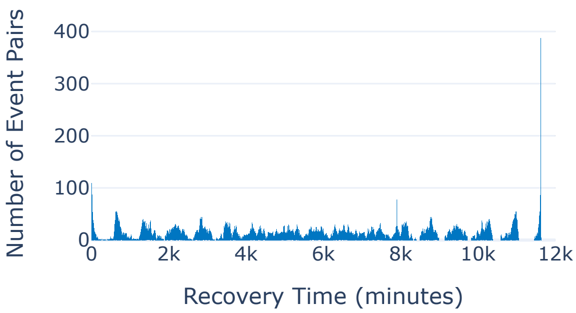

Figure 5. Histogram of recovery times matrix excluding passenger transfer precedence constraints.

Figure 5 presents a histogram of recovery times. The distribution shows a periodic pattern, with a repeated peak at the slack of the third critical circuit ($T - \lambda_3 = 24{:}00 - 11{:}54 = 12{:}06$). This arises because λ_3 corresponds to the cycle time of the critical circuit comprising inter-island flights, which are far more numerous than the mainland flights in critical circuits 1 and 2. Only a small fraction of event pairs are connected by short recovery times: fewer than 1% (1593 entries) fall below 20 minutes, and just 1.8% (2974 entries) fall below 60 minutes. This indicates that short-term structural delay propagation is possible only between a very limited subset of event pairs. Overall, the system is highly buffered at a global level but exhibits localised vulnerability to delay propagation in specific structural bottlenecks.

C. Efficacy of Delay Prediction

Using the delay propagation model described in Section III.D, the system's response to a given initial delay can be evaluated. This section focuses on validating the model's use for delay prediction using only flight and turnaround precedence constraints. As it is unknown whether Hawaiian Airlines delays aircraft to wait for connecting passengers, passenger transfer precedence constraints are excluded for now.

1. Verifying the Model's Basic Functioning

The simplest test of the model's functioning is provided by scenario V1, which applies a 45-minute arrival delay to flight HA17 from Las Vegas (LAS) to Honolulu (HNL). Figure 6 shows the resulting delay propagation, which is restricted to a limited subset of downstream flights. Visual inspection confirms that the model behaves as expected. Similar inspection of the verification scenarios V2 and V3 yields consistent results.

2. Adjusting for Structural Delay

Even in the absence of any initial delay, the model predicts 3 hours and 56 minutes of delay. This *structural delay* arises when the scheduled time between two subsequent events is less than the required process time. Figure 7 shows the distribution of these delays.

A total of 100 arrival or departure events are structurally delayed. One particular structural delay case is that of aircraft N485HA operating flight HA213 and its return flight HA214 only 14 minutes after completing HA190. The delays related to these flights reflect unrealistic assumptions in the schedule and are also observed in the actual delays. Other cases, particularly for inter-island flights, arise because the scheduled block time is shorter than the 40th percentile block time. These small discrepancies are within five minutes in all but one case.

The presence of structural delay can pollute the results. It makes sense to calculate the non-structural delay to accurately study the impact of an initial

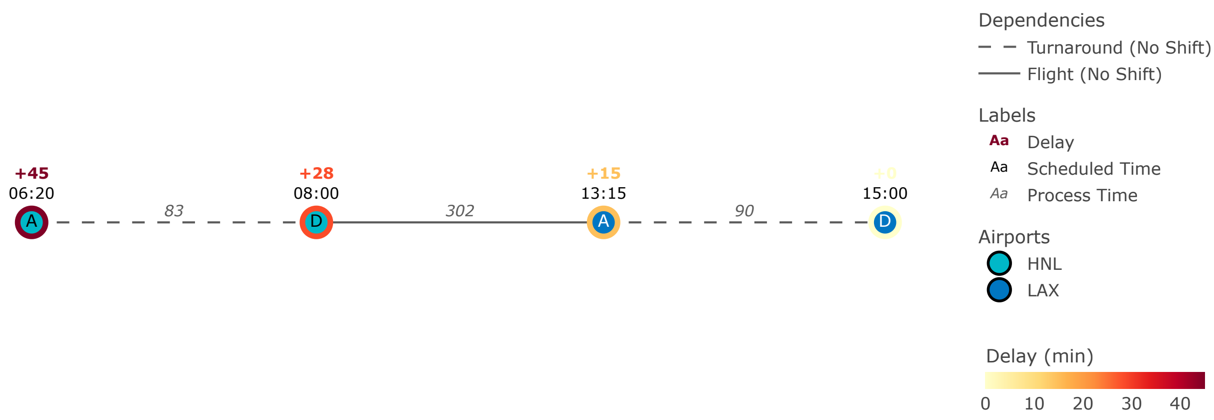

Figure 6. Delay propagation for scenario V1 with only flight and turnaround precedence constraints (no passenger transfers).

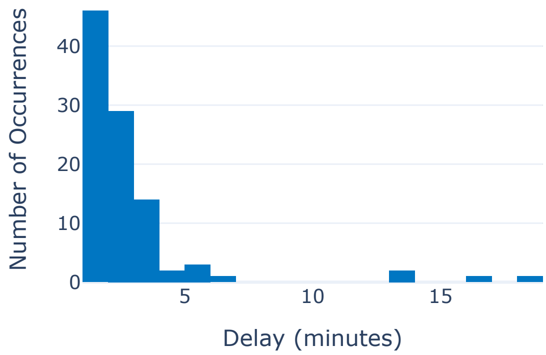

Figure 7. Delay distribution of the structural delay.

delay. The computed delay including structural delay is calculated by subtracting the scheduled timetable vector $d(k)$ from the computed event time vector $x(k)$. The non-structural computed delay is calculated by calculating the adjusted timetable vector $d_{struc}(k) = d(k) + z_{struc}$ and then subtracting $d_{struc}(k)$ from $x(k)$. z_{struc} is the structural delay vector.

Table 3 illustrates the differences between structural and non-structural delay by summarising model outputs for the three verification scenarios (V1–V3) and a baseline case with no initial delay. The table also reports delay multipliers (DM), as introduced in Section II.A, and settling times. The settling time is the scheduled time between the first and last delayed events. For both indicators, the non-structural initial and propagated delays are used.

Table 3. Delay statistics for verification scenarios, with and without adjustment for structural delay. Times indicated in hh:mm format.

Scenario:	None	V1	V2	V3
Iteration count	1	2	3	2
Total delay	03:56	09:15	14:26	13:17
Delay including structural delay				
Initial	00:00	00:45	00:45	00:45
Propagated	03:56	08:30	13:41	12:32
Structural	03:56	07:53	11:48	07:53
Delay excluding structural delay				
Total	00:00	01:22	02:36	05:24
Initial	00:00	00:39	00:45	00:45
Propagated	00:00	00:43	01:51	04:39
Delay indicators excluding structural delay				
Delay multipl.	0.00	2.11	3.48	7.20
Settling time	00:00	06:55	21:15	04:25

It is notable that for scenario V1, the initial delay decreases after adjustment for structural delay from 45 minutes to 39 minutes. This indicates that the event which is initially delayed has a structural delay of 6 minutes, which is removed when removing the structural delay.

3. Evaluating the Accuracy of the Model's Predictions To evaluate the accuracy of the model's predictions, scenario V4 inputs the actual initial delays observed on February 21, 2017, excluding delay caused by late aircraft arrival. Because BTS data only provides delay causes for arrival events and the model matches real-world flight sequences only for inter-island flights, this comparison is limited to inter-island arrivals. 59 arrivals were actually delayed on this date, of which 9 due to propagated delay.

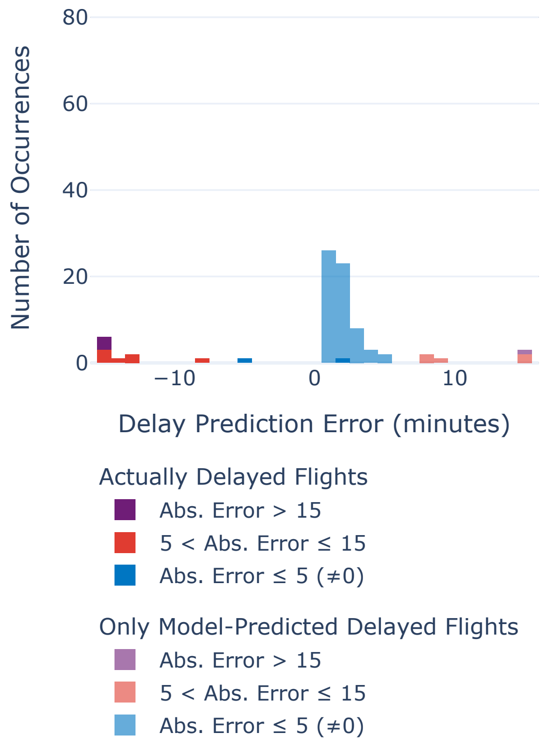

Figure 8. Distribution of model prediction error (before structural adjustment). Negative values indicate underprediction. Positive values indicate overprediction.

Figure 8 shows the delay prediction error, which is calculated as the computed delay minus the actual delay. Only flights with non-zero computed or actual delay are included. Events are grouped by:

- *Actual*: Flights that were delayed in reality.
- *Model-predicted only*: Flights that the model predicted would be delayed, but were not delayed in reality.

The figure shows that 3 of the 9 actually delayed flights were predicted within 10 minutes. The model overpredicts delay for 67 flights that were not delayed in reality.

Much of this overprediction is due to structural delay, as evident when comparing with Fig. 7. Removing structural delay from the computed values results in the distribution shown in Fig. 9. The number of overpredicted flights then drops from 67 to 13.

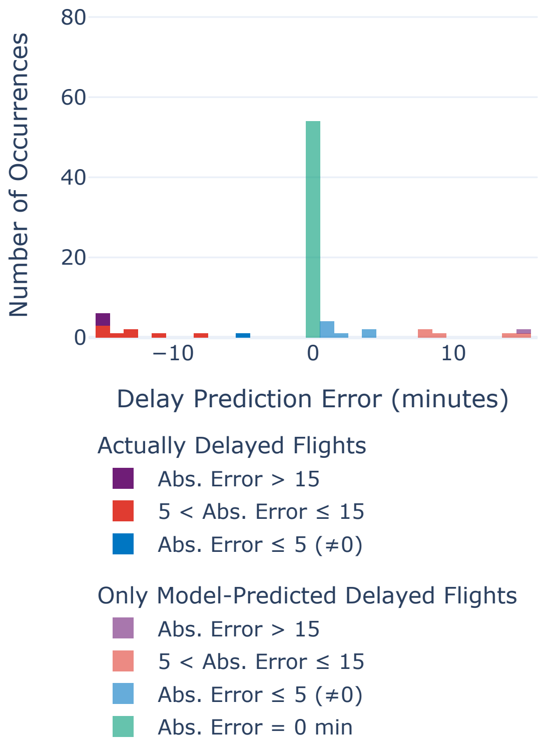

Figure 9. Distribution of model prediction error after subtracting structural delay.

To explore this further, Fig. 10 plots the prediction error (after structural delay adjustment) against the observed delay, for all flights where either the model or data show propagated non-structural delay. In total, 21 flights fall into this category. The model evidently underpredicts the flights that were actually delayed and overpredicts flights that were not actually delayed.

Table 4 summarises the key error statistics. Mean errors of 4–10 minutes suggest that while the model approximates overall delay magnitudes reasonably, individual prediction accuracy remains limited. Combining it with the insights from Figs. 8 to 10, one can see that the model is poorly suited for computing actual delays precisely given a specific initial input.

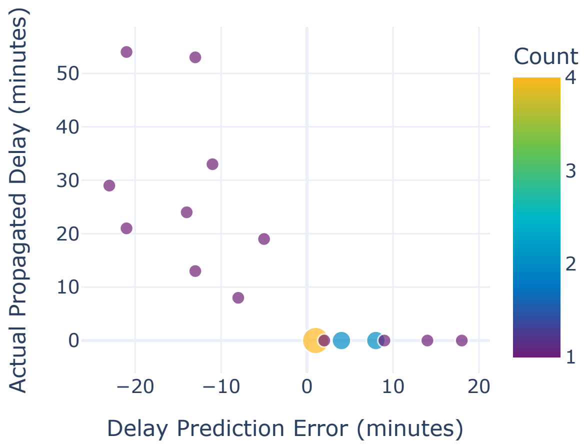

Figure 10. Prediction error versus observed delay (after structural delay adjustment).

Table 4. Model prediction accuracy for scenario V4, before and after structural delay adjustment.

Indicator	Including structural	Excluding structural
No. Flights	75	21
Actual Delay	04:14	04:14
Pred. Delay	05:13	03:16
Abs. Error	04:53	03:20
MAE	00:04	00:10
RMSE	00:06	00:12

D. Sensitivity to Process Time Assumptions

The accuracy of the delay propagation model is highly dependent on the assumed process times. These times are estimated from historical data, as discussed in Sections III.G and IV.E. This section investigates how the model's predictive performance changes as the used percentiles are varied.

A two-dimensional grid search was performed over combinations of TAT and block time percentiles, ranging from the 0^{th} to the 90^{th} percentile in steps of 10 percentage points. The 100^{th} percentile values skewed results as the system became unstable, and are excluded in this analysis. For each configuration, scenario V4 was simulated, which uses the actual initial arrival delays from February 21, 2017 as input. The model's predictions were then compared to the observed propagated delays for inter-island arrivals on that day. Passenger transfer precedence constraints were excluded in this analysis.

Figure 11 shows the mean absolute error (MAE) of the model's computed arrival delays across the sensitivity grid. The lowest MAE occurs for TATs between the 10^{th} and 30^{th} percentiles and block times between the 40^{th} and 60^{th} percentiles. This is in line with the values assumed in Sections III.G and IV.E.

However, a different picture emerges when only actual propagated delays are considered. Figure 12 shows that the model best detects these delays on the diagonal with increasing block time percentiles but decreasing TAT percentiles. In these configurations, the model can correctly compute delay all 9 flights that experienced actual propagated delay.

This improved prediction, however, comes at a cost. Figure 13 shows that as higher percentiles are used, the number of flights predicted to be delayed — but which were not actually delayed — increases sharply. This overprediction drives up the total error and diminishes the practical value of the predictions.

The combined effect is visible in Figure 14, which shows the total absolute error. While high process time assumptions seemingly increase the model's ability to correctly compute actual delay propagation, they also produce a large number of false positives.

E. Impact of Passenger Transfer Precedence Constraints

Previous results only considered the network without passenger transfer precedence constraints. While this allows for verification and validation of the model, the true power of max-plus linear systems lie in their efficient calculations for more connected networks. This section will add the passenger transfer precedence constraints and analyse the impact on the network stability and delay propagation behaviour. It will do so by studying the addition of passenger transfer constraints for those transfers which have more transferring passengers than the thresholds indicated in Table 5. The selected thresholds do not encompass a full linear scale, but were selected to showcase specific changes in the network around those threshold values.

First, the impact of the passenger transfers on the system stability characteristics will be discussed. Afterwards, the impact on the delay propagation effects under scenarios S1-S3 will be discussed.

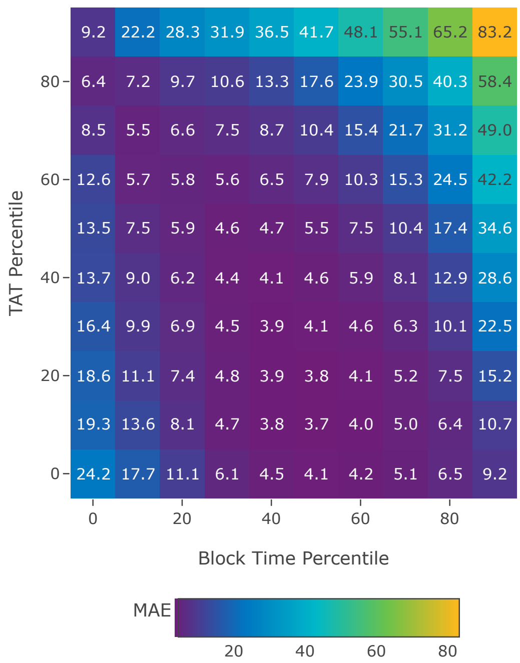

Figure 11. Mean absolute error (MAE) of model predictions under different combinations of TAT and block time percentiles.

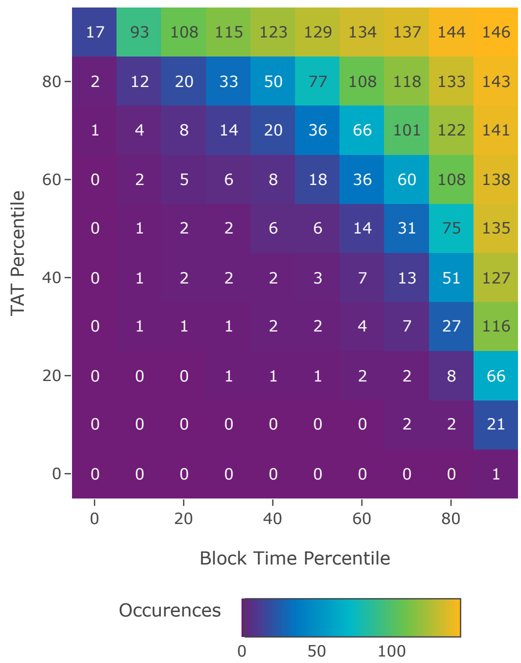

Figure 13. Number of flights overpredicted by more than 10 minutes. Overprediction increases with more conservative process time assumptions.

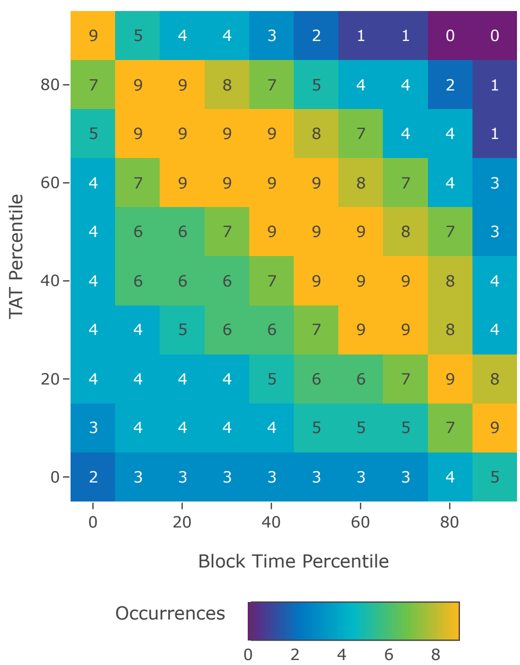

Figure 12. Number of correctly computed propagated delays (within 10 minutes of actual delay), as a function of TAT and block time percentiles.

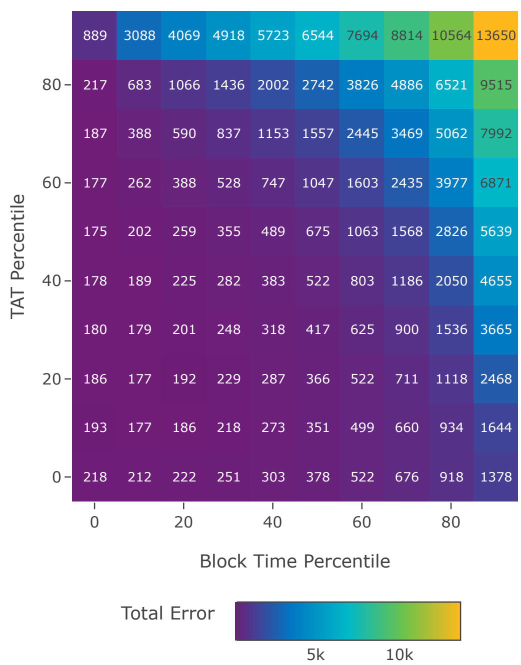

Figure 14. Total absolute prediction error across the sensitivity grid.

Table 5. Passenger transfer thresholds, including the number of passengers each threshold captures.

Threshold	Passengers captured
35.0	0.0%
18.0	10.4%
14.0	19.0%
11.5	30.0%
9.0	41.6%
5.5	59.7%
2.5	82.0%
0.0	100.0%

1. Impact on System Stability Characteristics

Table 6 shows how the inclusion of passenger transfer precedence constraints impacts the structural and stability characteristics of the max-plus system. As the threshold for passenger transfer inclusion is lowered, more passenger transfer arcs are added to the network, increasing the total number of precedence constraints from 404 (flight and turnaround only) to 1122 (flight, turnaround and all passenger transfers).

A notable structural shift occurs early: with the inclusion of only ~ 10% of passenger connections, the network transitions from three strongly connected components (SCCs) to a single fully connected component. This indicates that even a small number of key passenger connections is sufficient to bridge previously disconnected aircraft rotations into one interdependent structure.

However, this added connectivity reduces system robustness. The stability indicators Δ_1 and Δ_2 decrease as more passenger transfer precedence constraints are included. However, both values are relatively stable until ~ 40% of passenger connections are captured. It is likely that at this point, one of the night-time flights becomes part of the critical circuit, vastly increasing the maximum cycle mean. For this system, however, this suggests that it could be interesting to have aircraft wait if 12 or more passengers are expected for a transfer.

The final two columns confirm this increasing fragility, but do not exhibit a similarly significant tipping point. The number of recovery matrix entries with negative values rises fivefold, and the number

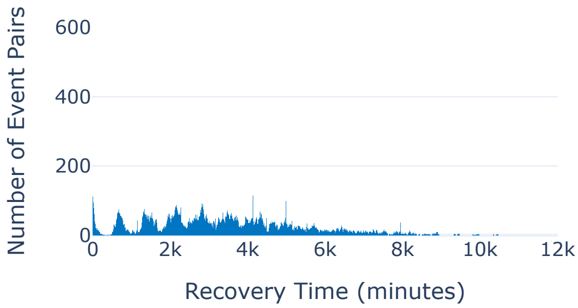

(a) Threshold = 14, ~ 20% of passengers captured.

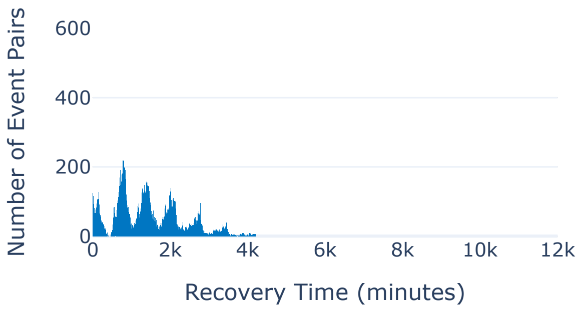

(b) Threshold = 9, ~ 40% of passengers captured.

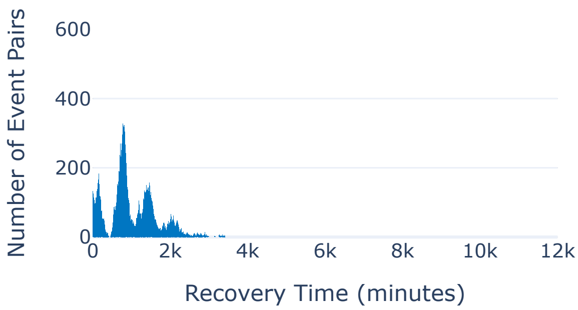

(c) Threshold = 5.5, ~ 60% of passengers captured.

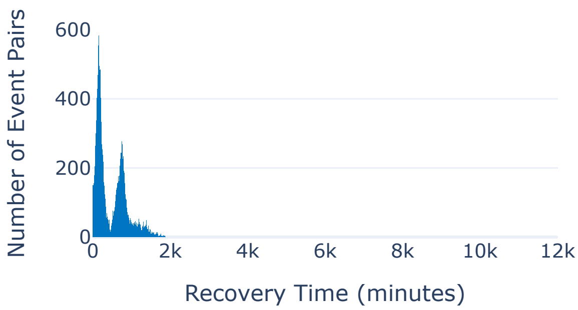

(d) Threshold = 2.5, ~ 80% of passengers captured.

Figure 15. Histogram of recovery times matrix including passenger transfer precedence constraints at various thresholds.

Table 6. Impact of varying the number of passenger transfer precedence constraints included in the network on the system stability characteristics.

Threshold	Captured	Arcs	SCC	Δ_1	Δ_2	$r_{ij} < 0$	$r_{ij} < 15$	Structural delay
35.0	0%	404	3	06:10	00:35	133	1266	03:55
18.0	~ 10%	415	1	06:10	00:35	152	1312	05:08
14.0	~ 20%	428	1	06:03	00:28	185	1388	06:50
11.5	~ 30%	448	1	06:03	00:28	218	1491	08:12
9.0	~ 40%	475	1	01:14	00:04	253	1627	10:03
5.5	~ 60%	528	1	01:14	00:04	333	1826	13:13
2.5	~ 80%	657	1	01:12	00:04	567	2340	23:52
0.0	100%	1122	1	01:12	00:04	635	2631	25:38

of short-slack event pairs ($r_{ij} < 15$ minutes) more than doubles. Note that the negative r_{ij} values correspond with the observations of structural delay in Section V.C. The findings in the table show that for those first ~ 30% of passengers captured, the increase of short-slack event pairs is limited (9.6%), although it does increase the no-slack event pairs by 64%.

Figure 15 shows the impact of increasing the number of passenger transfer precedence constraints on the full r_{ij} distribution. Clearly, adding more transfer constraints decreases r_{ij} for many event pairs. However, it is notable that even with 60% of passengers captured, a majority of event pairs retains a recovery time greater than the critical slack $\Delta_1 = 06{:}10$ at a threshold of 35.

2. *Impact on Delay Propagation*

Figure 16 shows the effect of including passenger transfer precedence constraints on delay propagation under scenario V1, where a 45-minute delay is applied to an early morning inbound mainland flight. The delay cascades across multiple aircraft rotations and propagates widely, far beyond the initially affected flight. This example clearly illustrates how passenger connections can magnify and extend the reach of delay propagation throughout the network.

Table 7 summarises the non-structural delay propagation for scenario V1 across various thresholds of included passenger transfer precedence relations. As more passenger connections are incorporated, the number of affected events, airports, and total propagated delay all increase. The delay multiplier (DM) rises steadily from 2.11 to 27.11 as the threshold is lowered. This highlights how seemingly minor delays in the morning can disrupt downstream flights throughout the day when the network is tightly coupled via passenger flows.

Table 8 shows similar results for scenario S2, which delays all morning inbound mainland flights. Again, a continuous increase in all propagation indicators is observed as more passenger transfers are added. The number of affected events increases from 27 to 188, and the delay multiplier rises from 1.28 to 8.84. These results are expected: the delayed mainland arrivals serve as feeders into nearly all inter-island aircraft rotations, making the entire schedule vulnerable once these connections are represented in the model. The slight decrease in propagated delay between the 80% and 100% thresholds is due to a rise in structural delay, which displaces part of the propagation effect.

Scenarios S1 and S3 behave differently. As shown in Tables 9 and 10, both scenarios exhibit relatively limited propagation when fewer than 30% of passenger connections are included. The delay multiplier and propagated delay only begin to rise meaningfully below the 10-passenger threshold. This is consistent with the earlier stability analysis: a small set of passenger transfers is not sufficient to significantly increase systemic vulnerability. It is likely that scenarios S1 and S3 do not interact strongly with the passenger transfer precedence constraints added first. After all, Table 6 shows that incorporating constraints which capture ~ 30% of passengers only adds 44 arcs to the

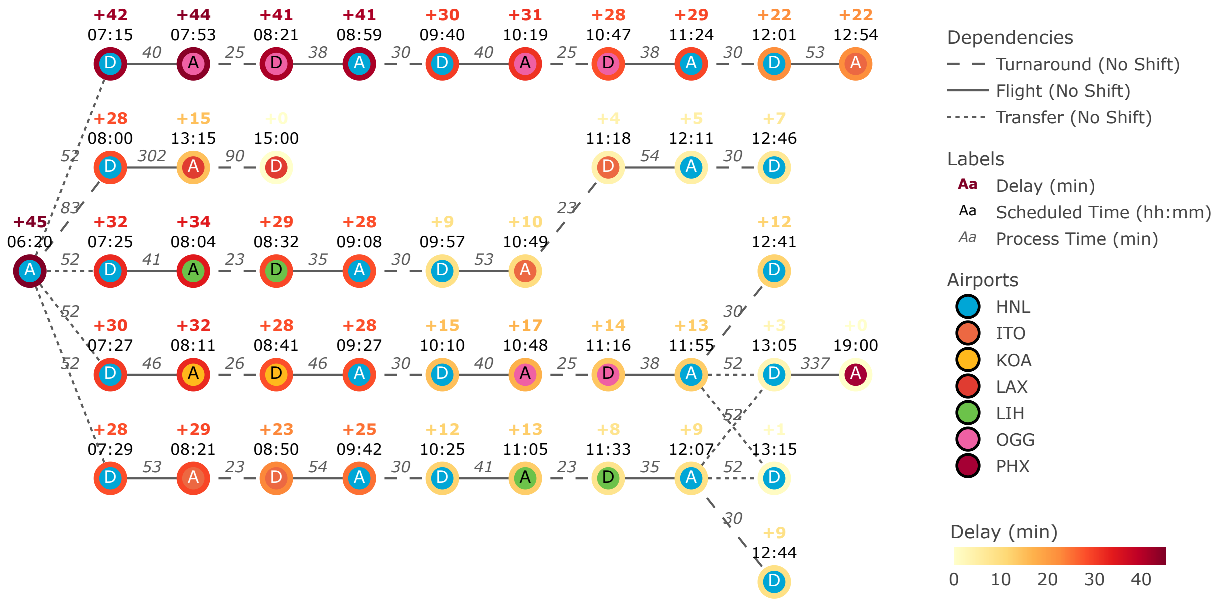

Figure 16. Delay propagation for scenario V1 with passenger transfer precedence constraints for all transfers involving more than 10 passengers. The delay branch ending at 12:54 continues to propagate delay until approximately 19:00, beyond what is shown here.

Table 7. Delay propagation behaviour for scenario V1 including passenger transfer precedence constraints at various thresholds. Times indicated in hh:mm format.

Threshold	Captured	Initial	Propagated	Events	Airports	DM	Mean	Settling Time
35.0	0%	00:39	00:43	3	2	2.11	14.4	06:55
18.0	~ 10%	00:39	00:43	3	2	2.11	14.4	06:55
14.0	~ 20%	00:39	07:19	21	4	12.26	20.9	11:34
11.5	~ 30%	00:39	12:11	38	5	19.75	19.2	11:34
9.0	~ 40%	00:39	15:57	53	6	25.54	18.1	11:34
5.5	~ 60%	00:39	15:04	53	6	24.18	17.1	11:34
2.5	~ 80%	00:39	15:48	58	8	25.31	16.3	13:35
0.0	100%	00:39	16:58	65	9	27.11	15.7	13:35

Table 8. Delay propagation behaviour for scenario S2 (morning inbound delay) including passenger transfer precedence constraints at various thresholds. Times indicated in hh:mm format.

Threshold	Captured	Initial	Propagated	Events	Airports	DM	Mean	Settling Time
35.0	0%	05:01	01:25	27	5	1.28	3.1	10:20
18.0	~ 10%	05:01	04:18	43	7	1.86	6.0	10:20
14.0	~ 20%	05:01	11:23	79	8	3.27	8.7	10:51
11.5	~ 30%	05:01	15:37	96	10	4.11	9.8	16:35
9.0	~ 40%	05:01	20:00	120	11	4.99	10.0	16:35
5.5	~ 60%	05:01	28:48	163	12	6.74	10.6	16:35
2.5	~ 80%	05:01	39:49	188	17	8.94	12.7	32:35
0.0	100%	05:01	39:19	188	17	8.84	12.6	32:35

Table 9. Delay propagation behaviour for scenario S1 (top 10 flights delayed) including passenger transfer precedence constraints at various thresholds. Times indicated in hh:mm format.

Threshold	Captured	Initial	Propagated	Events	Airports	DM	Mean	Settling Time
35.0	0%	03:41	02:45	19	7	1.75	8.7	17:10
18.0	~ 10%	03:41	05:47	27	9	2.57	12.9	17:10
14.0	~ 20%	03:41	06:24	35	9	2.73	11.0	17:10
11.5	~ 30%	03:41	06:14	37	9	2.69	10.1	17:10
9.0	~ 40%	03:41	11:17	53	10	4.05	12.8	17:10
5.5	~ 60%	03:41	17:49	71	10	5.82	15.1	17:10
2.5	~ 80%	03:41	28:43	150	21	11.48	15.5	32:10
0.0	100%	03:41	48:43	184	21	14.18	15.9	32:10

Table 10. Delay propagation behaviour for scenario S3 (KOA airport closure) including passenger transfer precedence constraints at various thresholds. Times indicated in hh:mm format.

Threshold	Captured	Initial	Propagated	Events	Airports	DM	Mean	Settling Time
35.0	0%	07:18	32:31	52	5	5.45	37.5	10:13
18.0	~ 10%	07:18	34:59	54	6	5.79	38.9	10:13
14.0	~ 20%	07:01	34:33	54	6	5.93	38.4	10:13
11.5	~ 30%	07:01	38:01	58	8	6.42	39.3	10:13
9.0	~ 40%	07:01	37:16	60	9	6.31	37.3	10:13
5.5	~ 60%	06:36	38:47	63	10	6.88	36.9	10:13
2.5	~ 80%	06:36	53:56	105	21	9.17	30.8	33:54
0.0	100%	06:36	30:06	186	23	17.08	34.2	33:54

network. It is plausible the transfers with the highest number of passengers are those from mainland flights to inter-island flights. These flights will then be added to the network first.

VI. Discussion

This section will discuss the results presented in Section V. First, Sections VI.A and VI.B will discuss the model's lack of prediction accuracy and the impact of respectively the process time selection and other assumptions on it. Section VI.C will then outline some further reflections on the delay propagation algorithm. Afterwards, Section VI.D outlines the effectiveness of the stability characteristics calculations. Section VI.E will discuss the impact of adding passenger transfers. Finally, Section VI.F presents the broader implications and gives suggestions for future work.

A. Impact of Process Time Selection on Prediction Accuracy

The scheduled max-plus linear system developed in this study offers a novel approach to modelling delay propagation in aviation networks. However, its predictive accuracy at the level of individual flights is limited. In scenario V4, where real-world initial delays were used as inputs, the predicted propagated delays corresponded only weakly to the actual values. This is likely due to inaccuracies in the selection of process times. Model performance is highly dependent on this selection, as shown by the sensitivity analysis in Section V.D.

In this study, process times were determined using a single percentile threshold for all turnaround times and another for all block times. In reality, the appropriate percentile can vary substantially by route or airport. For example, the 60^{th} percentile may best represent block times between HNL and OGG, while the 20^{th} percentile might better match HNL–JFK. Similarly, the optimal turnaround percentile could be much lower at HNL than at KOA.

Two possible strategies could improve tuning. First, selecting process times individually for each route or airport could better capture operational variability. However, this would tend towards solving an optimisation problem and would place much higher demands on the underlying data. Not only would the model require more data to robustly estimate per-route process time distributions, but it would also raise questions of how to validate the tuning process without overfitting to noise or temporary patterns.

Another strategy would be to use a longer time span (e.g. multiple months). This would increase the number of observations per origin-destination pair and airport, particularly for low-frequency routes. This could improve percentile estimates, but both turnaround and block times can vary seasonally and the available data quality is uncertain. For example, poor weather may increase block times for mainland flights, while holiday congestion can affect ground operations at popular tourist destinations. Selecting which dates to include in the analysis thus becomes an important choice. Moreover, preliminary inspection of current process times reveals no significant outliers, and it is unknown whether more data points would significantly change the process times.

These challenges suggest that fine-tuning process times using only publicly available data may be inherently limited by data quality. Access to operational data from airlines or airports, containing actual turnaround and block times, would enable more accurate process time estimation. Future work should explore the effectiveness of the model using these actual process time values.

B. Assumptions Impacting Prediction Accuracy

Besides the issues with the selection of process times discussed in Section VI.A, there are a number of assumptions which move the model away from reality, impacting prediction accuracy.

First and foremost is the lack of full knowledge about real-world flight sequences and precedence relations. Actual aircraft rotations were only known for the inter-island flights. While the greedy algorithm used to set up turnaround precedence constraints generates a roughly feasible set of aircraft rotations, it does not capture actual rotations. Moreover, crew rotation information was not available. Note that all the data above is known to airlines, and that including it would likely improve delay prediction accuracy.

Secondly, it was assumed that the network structure is fixed from day to day. In reality, precedence relations between events, especially for aircraft rotations, vary significantly between days due to operational considerations such as aircraft swaps, delays, maintenance, and ad hoc reassignments. The model assumes these relations are static and periodic. The timetable of February 21st also exhibits signs that there were aircraft reassignments, but this is difficult to know without data from the airline.

Even under perfect knowledge of precedence relations and process times, the assumption of deterministic process times imposes a strong limitation. Flight durations, turnaround times, and passenger connection processes are all inherently stochastic. The model uses percentile-based approximations for these process times, but still cannot capture day-to-day operational variability due to gate assignments, airport congestion or weather conditions. This means that while the model can estimate structural vulnerability to delay propagation, it is poorly suited for precise prediction of delay magnitudes in specific instances.

C. Further Reflections on the Delay Propagation Algorithm

While the predictions of the delay propagation algorithm are not highly accurate, the results indicate that it does function as intended. For the Hawaiian Airlines case, one useful insight is the identification of structural delay in the model, suggesting that certain scheduled flight pairings may be overly optimistic regarding required turnaround or block times. Minor adjustments to these schedules could therefore improve delay robustness without adding capacity.

From a methodological perspective, the delay propagation algorithm derives its main computational advantage from the max-plus algebraic formulation when delays propagate over many periods ($k \gg 1$). In such cases, the matrix-based max-plus approach efficiently captures the cumulative effects of delays. However, in scenarios with limited propagation depth ($k = 1$ or 2), a direct recursive implementation using conventional max and + operations may offer greater transparency without loss of efficiency.

D. Effectiveness of Stability Characteristics Calculations

One of the key advantages of scheduled max-plus linear systems lies in their ability to efficiently compute stability characteristics of a network, as demonstrated in Section V.B. Railway timetables are designed around short, typically hourly, repeating cycles, and the relevant stability characteristics are useful for these cases. However, aviation timetables often follow a 24-hour cycle. As a result, key indicators such as the maximum cycle mean and network throughput are dominated by structural features like overnight buffers. For example, in the Hawaiian Airlines network modelled in this study, the largest stability margins arose from aircraft parked overnight, which are functionally irrelevant to intra-day delay propagation. Consequently, in aviation networks, system-wide metrics such as the stability margin Δ_2 or the difference between the cycle time and maximum cycle mean Δ_1 may overstate the actual robustness of the schedule during core operating hours.

Nevertheless, these global metrics can still be useful in comparative or diagnostic contexts. For instance, sharp drops in stability indicators when passenger transfer precedence constraints are added, as observed in Section V.E, signal that specific connections introduce structural vulnerability. Similarly, large gaps between the maximum cycle mean and subsequent cycle means may suggest that the system contains bottlenecks or subnetworks that are disproportionately fragile.

Where the scheduled max-plus linear system does provide some value in the aviation context is in its ability to perform localised stability analysis. Tools such as the recovery matrix enable targeted diagnosis of which event pairs are vulnerable to propagation and can quantify the available slack between them. This is particularly valuable in settings with tight connections, such as hub-and-spoke operations or short-turn aircraft sequences, where a small amount of slack can make the difference between delay containment and wide-scale propagation.

Overall, the strengths of max-plus stability analysis lie in its efficiency and local interpretability. However, its global indicators show limited use when applied to aviation networks with long cycle times and big nighttime buffers. It could be useful for networks which do not exhibit these large nighttime buffers, such as networks including intercontinental flights. In general, however, the method is best used not to generate a general robustness score but as a tool for understanding structural delay propagation effects.

E. Modelling Passenger Transfer Precedence Relations

An important extension of the scheduled max-plus linear system in this study is the incorporation of passenger transfer precedence constraints. Holding an aircraft for delayed transferring passengers can prevent the operational and financial consequences of stranding passengers, but may increase delays elsewhere in the network. As shown in Section V.E, scheduled max-plus linear systems provide a quantitative framework for evaluating this trade-off.

Including only the top 10% of passenger transfer arcs by passenger volume transformed the network from multiple weakly connected components into a single strongly connected structure, greatly increasing its potential for delay propagation. This indicates that a small set of transfer constraints can create critical coupling between otherwise independent flight sequences. Recovery time analysis offers a practical tool to identify such precedence constraints. In the Hawaiian Airlines network, these are likely concentrated among morning mainland arrivals feeding subsequent inter-island departures.

The model also enables systematic testing of whether to include passenger transfer precedence constraints. By progressively lowering the inclusion threshold, it becomes possible to quantify how each additional passenger transfer constraint affects system robustness. For example, the sharp rise in delay multipliers when reducing the threshold from 18% to 14% suggests that certain high-volume transfers are structurally fragile. Conversely, other thresholds produce little change in propagated delay, indicating more resilient connections.

These findings demonstrate that scheduled max-plus linear systems can help determine whether aircraft should wait for specific inbound flights in the structural schedule design. Future work could investigate at a more granular level which passenger transfer precedence constraints can be included without significantly increasing delay propagation risk.

F. Broader Implications and Future Work

The application of scheduled max-plus linear systems to aviation networks, as explored in this study, shows that there is some potential to the technique, but that there are also many limitations. While the method is formally elegant and computationally efficient, its practical success hinges on the nature of the aviation network being modelled and the quality of the available data.

The most promising application of max-plus linear systems in aviation lies in their ability to characterise the structural properties of delay propagation, especially in networks where precedence relations are predictable and repeated over time. For airlines that operate largely periodic or cyclic schedules, max-plus linear systems offer a valuable tool for identifying fragile sequences, understanding propagation paths, and quantifying the slack between events. The recovery matrix, in particular, serves as a useful diagnostic for assessing the impact of local delays on the broader network.

At the same time, the results of this study show that actual predictive accuracy remains limited. When precedence relations are reconstructed using a greedy algorithm and process times are estimated globally, the model struggles to replicate observed delay propagation on a flight-by-flight basis. This shows that scheduled max-plus linear systems, in their current form, are better suited for strategic analysis and robustness evaluation than for real-time delay prediction or tactical decision-making.

Several avenues exist for future work. First, improvements in data quality and coverage could significantly enhance model calibration. Ideally, a future study would build the system based on actual data from an airline. Lacking this, one could explore more accurate estimation of process times on an airport or flight-by-flight basis.

Second, while this study focused on a single airline and one day of operations, the framework could be scaled to include larger, more complex networks. Doing so would test the computational scalability of the methods and allow investigation into multi-airline or multi-hub dynamics. The application of these models to low-cost carriers with greater daily periodicity or intercontinental networks without nighttime buffer could prove the stability characteristics discussed in this work to be useful for aviation networks after all.

Third, the max-plus framework could be extended to accommodate probabilistic or stochastic elements. Although traditional max-plus systems are deterministic, there has been research in stochastic max-plus algebra [16] that could, in principle, be applied to model uncertainty in process times or precedence relations. This could bridge the gap between structural modelling and real-world variability.

In summary, this study demonstrates that max-plus linear systems can be a useful tool for understanding specific characteristics of delay propagation in aviation networks. However, the model can still be refined with more accurate data, more airlines, stochastics max-plus algebra or hybrid approaches.

VII. Conclusion

This study set out to investigate to what extent scheduled max-plus linear systems can be applied to model, analyse, and simulate delay propagation in aviation networks. The results demonstrate that it is indeed possible to represent an aviation network as a scheduled max-plus linear system and to derive meaningful insights from such a representation. In particular, the framework proved effective in computing the recovery matrix and using it to identify structural weak points in the network, as well as in evaluating the impact of including passenger transfer precedence relations. These findings highlight the potential of max-plus linear systems to provide interpretable, system-level assessments of delay propagation vulnerabilities.

However, the research also revealed that not all analyses established in railway network applications translate directly to aviation. Some limitations stem from the specifics of the Hawaiian Airlines case study. The delay propagation algorithm showed limited predictive accuracy, largely due to uncertainty in process time selection and incomplete knowledge of actual precedence relations. The calculation of stability characteristics was also constrained by the large nightly buffer in the network, which reduces the relevance of cross-period stability metrics. Furthermore, this network, as most aviation networks, only exhibits cyclic periodicity only over a 24-hour cycle, meaning that propagation to subsequent periods is rare and operational interventions have ample time to mitigate it. This limits the power of the scheduled max-plus linear system's delay propagation.

Other limitations are inherent to the use of scheduled max-plus linear systems for aviation. Because network structure varies from day to day in airline operations, the model is generally unsuitable for realistically predicting propagation across multiple days. Moreover, aviation process times are inherently more stochastic than those in railway networks, which constrains the accuracy of real-time delay propagation estimation in a deterministic max-plus framework. While stochastic max-plus formulations could alleviate some of these issues, they would require more complex modelling and richer data inputs.

Future work should therefore prioritise improving data quality, ideally by incorporating operational data from airlines to obtain accurate process times and precise precedence relations. Applying the model to networks with shorter periodicities or without extensive overnight buffers, such as certain intercontinental operations, could also reveal additional insights into its capabilities and limitations. Finally, extending the framework to incorporate stochastic max-plus algebra offers a promising avenue to better capture the variability inherent in aviation processes, thereby enhancing its predictive power for real-world operations.

References

[1] Walker, C., "CODA Digest: All-Causes Delays to Air Transport in Europe - Annual 2023," , Dec. 2024. URL `https://www.eurocontrol.int/publication/all-causes-delays-air-transport-europe-annual-2023`.

[2] Bureau of Transportation Statistics, "Airline On-Time Statistics and Delay Causes," , Dec. 2024. URL `https://www.transtats.bts.gov/ot_delay/ot_delaycause1.asp?6B2r=FE&20=E`, dataset.

[3] OAG, "What are Minimum Connection Times (MCTs)?" , 2025. URL `https://www.oag.com/blog/minimum-connection-times-insiders-guide`.

[4] Bureau of Infrastructure and Transport Research Economics, "Domestic on time performance," , Sep. 2025. URL `https://www.bitre.gov.au/statistics/aviation/otphome`.

[5] Ball, M., Barnhart, C., Dresner, M., Hansen, M., Neels, K., Odoni, A., Peterson, E., Sherry, L., Trani, A., Zou, B., Britto, R., Fearing, D., Swaroop, P., Uman, N., Vaze, V., and Voltes, A., "Total Delay Impact Study: A Comprehensive Assessment of the Costs and Impacts of Flight Delay in the United States," , Oct. 2010. URL `https://rosap.ntl.bts.gov/view/dot/6234`.

[6] Scheelhaase, J., Braun, M., Maertens, S., and Grimme, W., "Costs for passengers and airlines due to the significant delays and other irregularities at European airports in the 2022 summer season," *Trans-*

portation Research Procedia, Vol. 75, 2023, pp. 96–105. doi: 10.1016/j.trpro.2023.12.012.

[7] Value Group, "The Economic Impact of Journey Disruptions: Analyzing Airlines and Railways," , Oct. 2023. URL https://valueg.com/the-economic-impact-of-disruptions/.

[8] AirHelp, "Cost of disrupted flights to the economy," , Oct. 2023. URL https://www.airhelp.com/en-int/blog/in-numbers-the-economic-impact-of-flight-disruptions/.

[9] Carvalho, L., Sternberg, A., Maia Gonçalves, L., Beatriz Cruz, A., Soares, J. A., Brandão, D., Carvalho, D., and Ogasawara, E., "On the relevance of data science for flight delay research: a systematic review," *Transport Reviews*, Vol. 41, No. 4, 2021, pp. 499–528. doi: 10.1080/01441647.2020.1861123.

[10] Wang, T., Zheng, Y., and Xu, H., "A Review of Flight Delay Prediction Methods," *2nd International Conference on Big Data Engineering and Education (BDEE)*, 2022, pp. 135–141. doi: 10.1109/BDEE55929.2022.00029.

[11] Huynh, T. K., Cheung, T., and Chua, C., "A Systematic Review of Flight Delay Forecasting Models," *2024 7th International Conference on Green Technology and Sustainable Development (GTSD)*, 2024, pp. 533–540. doi: 10.1109/GTSD62346.2024.10675123.

[12] Li, C., Mao, J., Li, L., Wu, J., Zhang, L., Zhu, J., and Pan, Z., "Flight delay propagation modeling: Data, Methods, and Future opportunities," *Transportation Research Part E: Logistics and Transportation Review*, Vol. 185, 2024, p. 103525. doi: 10.1016/j.tre.2024.103525.

[13] Goverde, R. M. P., "Railway timetable stability analysis using max-plus system theory," *Transportation Research Part B: Methodological*, Vol. 41, No. 2, 2007, pp. 179–201. doi: 10.1016/j.trb.2006.02.003.

[14] Goverde, R. M. P., "A delay propagation algorithm for large-scale railway traffic networks," *Transportation Research Part C: Emerging Technologies*, Vol. 18, No. 3, 2010, pp. 269–287. doi: 10.1016/j.trc.2010.01.002.

[15] Goverde, R. M. P., Heidergott, B., and Merlet, G., "A coupling approach to estimating the Lyapunov exponent of stochastic max-plus linear systems," *European Journal of Operational Research*, Vol. 210, No. 2, 2011, pp. 249–257. doi: 10.1016/j.ejor.2010.09.035.

[16] Goverde, R. M. P., "Punctuality of Railway Operations and Timetable Stability Analysis," PhD Dissertation, Delft University of Technology, Delft, Mar. 2005. URL https://resolver.tudelft.nl/uuid:a40ae4f1-1732-4bf3-bbf5-fdb8dfd635e7.

[17] Rengers, R., "Max-plus algebra and an application in aviation," BSc Thesis, Delft University of Technology, Delft, Jun. 2017. URL https://resolver.tudelft.nl/uuid:27c47f85-18e2-4789-9919-604fe1f7621e.

[18] Beatty, R., Hsu, R., Berry, L., and Rome, J., "Preliminary Evaluation of Flight Delay Propagation through an Airline Schedule," *Air Traffic Control Quarterly*, Vol. 7, No. 4, 1999, pp. 259–270. doi: 10.2514/atcq.7.4.259.

[19] Zhang, M., Zhou, X., Zhang, Y., Sun, L., Dun, M., Du, W., and Cao, X., "Propagation Index on Airport Delays," *Transportation Research Record*, Vol. 2673, No. 8, 2019, pp. 536–543. doi: 10.1177/0361198119844240.

[20] Wu, C.-L., and Law, K., "Modelling the delay propagation effects of multiple resource connections in an airline network using a Bayesian network model," *Transportation Research Part E: Logistics and Transportation Review*, Vol. 122, 2019, pp. 62–77. doi: 10.1016/j.tre.2018.11.004.

[21] Wang, P., Schaefer, L., and Wojcik, L., "Flight connections and their impacts on delay propagation," *The 22nd Digital Avionics Systems Conference, 2003. DASC '03.*, Vol. 1, 2003, pp. 5.B.4–5.1–9. doi: 10.1109/DASC.2003.1245858.

[22] Hao, L., Hansen, M., Zhang, Y., and Post, J., "New York, New York: Two ways of estimating the delay impact of New York airports," *Transportation Research Part E: Logistics and Transportation Review*, Vol. 70, 2014, pp. 245–260. doi: 10.1016/j.tre.2014.07.004.

[23] Kafle, N., and Zou, B., "Modeling flight delay propagation: A new analytical-econometric approach," *Transportation Research Part B: Methodological*, Vol. 93, 2016, pp. 520–542. doi: 10.1016/j.trb.2016.08.012.

[24] Baspinar, B., and Koyuncu, E., "A Data-Driven Air Transportation Delay Propagation Model Using Epidemic Process Models," *International Journal of Aerospace Engineering*, Vol. 2016, No. 1, 2016, p. 4836260. doi: 10.1155/2016/4836260.

[25] Wu, W., Zhang, H., Feng, T., and Witlox, F., "A Network Modelling Approach to Flight Delay Propagation: Some Empirical Evidence from China," *Sustainability*, Vol. 11, No. 16, 2019, p. 4408. doi: 10.3390/su11164408.

[26] Chen, J., and Li, M., "Chained Predictions of Flight Delay Using Machine Learning," *AIAA Scitech 2019 Forum*, American Institute of Aeronautics and Astronautics, San Diego, California, 2019. doi: 10.2514/6.2019-1661.

[27] Sun, J., Dijkstra, T., Aristodemou, K., Buzeţelu, V., Falat, T., Hogenelst, T., Prins, N., and Slijper, B., "Designing Recurrent and Graph Neural Networks to Predict Airport and Air Traffic

Network Delays," *10th International Conference for Research in Air Transportation*, edited by D. Lovell, FAA & EUROCONTROL, Tampa, FL, 2022. URL https://www.icrat.org/previous-conferences/10th-international-conference/papers/.

[28] Fleurquin, P., Ramasco, J. J., and Eguiluz, V. M., "Systemic delay propagation in the US airport network," *Scientific Reports*, Vol. 3, No. 1, 2013, p. 1159. doi: 10.1038/srep01159.

[29] Qin, S., Mou, J., Chen, S., and Lu, X., "Modeling and optimizing the delay propagation in Chinese aviation networks," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, Vol. 29, No. 8, 2019, p. 081101. doi: 10.1063/1.5111995.

[30] Wang, C., Hu, M., Yang, L., and Zhao, Z., "Prediction of air traffic delays: An agent-based model introducing refined parameter estimation methods," *PLoS ONE*, Vol. 16, No. 4, 2021. doi: 10.1371/journal.pone.0249754.

[31] Wu, C.-L., "Inherent delays and operational reliability of airline schedules," *Journal of Air Transport Management*, Vol. 11, No. 4, 2005, pp. 273–282. doi: 10.1016/j.jairtraman.2005.01.005.

[32] Baden, W., DeArmon, J., Kee, J., and Smith, L., "Assessing Schedule Delay Propagation in the National Airspace System," *47th Annual Transportation Research Forum*, New York, NY, 2006. doi: 10.22004/ag.econ.208044.

[33] Hansen, M., "Micro-level analysis of airport delay externalities using deterministic queuing models: a case study," *Journal of Air Transport Management*, Vol. 8, No. 2, 2002, pp. 73–87. doi: 10.1016/S0969-6997(01)00045-X.

[34] Wu, Q., Hu, M., Ma, X., Wang, Y., Cong, W., and Delahaye, D., "Modeling Flight Delay Propagation in Airport and Airspace Network," *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, 2018, pp. 3556–3561. doi: 10.1109/ITSC.2018.8569657.

[35] Long, D., and Hasan, S., "Improved Predictions of Flight Delays Using LMINET2 System-Wide Simulation Model," *9th AIAA Aviation Technology, Integration, and Operations Conference (ATIO)*, American Institute of Aeronautics and Astronautics, 2009. doi: 10.2514/6.2009-6961.

[36] Lin, Y., Li, L., Ren, P., Wang, Y., and Szeto, W. Y., "From aircraft tracking data to network delay model: A data-driven approach considering en-route congestion," *Transportation Research Part C: Emerging Technologies*, Vol. 131, 2021, p. 103329. doi: 10.1016/j.trc.2021.103329.

[37] Du, W.-B., Zhang, M.-Y., Zhang, Y., Cao, X.-B., and Zhang, J., "Delay causality network in air transport systems," *Transportation Research Part E: Logistics and Transportation Review*, Vol. 118, 2018, pp. 466–476. doi: 10.1016/j.tre.2018.08.014.

[38] Xiao, Y., Zhao, Y., Wu, G., and Jing, Y., "Study on Delay Propagation Relations Among Airports Based on Transfer Entropy," *IEEE Access*, Vol. 8, 2020, pp. 97103–97113. doi: 10.1109/ACCESS.2020.2996301.

[39] Wang, Y., Zheng, H., Wu, F., Chen, J., and Hansen, M., "A Comparative Study on Flight Delay Networks of the USA and China," *Journal of Advanced Transportation*, Vol. 2020, No. 1, 2020, p. 1369591. doi: 10.1155/2020/1369591.

[40] Xu, N., Donohue, G., Laskey, K. B., and Chen, C.-H., "Estimation of Delay Propagation in the National Aviation System Using Bayesian Networks," George Mason University, Fairfax, VA, 2005. URL https://www.semanticscholar.org/paper/Estimation-of-Delay-Propagation-in-the-National-Xu-Donohue/9c5a7c726315387acb5ff5cf8e849524046cd207.

[41] Liu, Y., and Wu, H., "A remixed bayesian network based algorithm for flight delay estimating," *International Journal of Pure and Applied Mathematics*, Vol. 85, No. 3, 2013, pp. 465–475. doi: 10.12732/ijpam.v85i3.3.

[42] Wu, W., and Wu, C.-L., "Enhanced delay propagation tree model with Bayesian Network for modelling flight delay propagation," *Transportation Planning and Technology*, Vol. 41, No. 3, 2018, pp. 319–335. doi: 10.1080/03081060.2018.1435453.

[43] Abdelghany, K. F., S. Shah, S., Raina, S., and Abdelghany, A. F., "A model for projecting flight delays during irregular operation conditions," *Journal of Air Transport Management*, Vol. 10, No. 6, 2004, pp. 385–394. doi: 10.1016/j.jairtraman.2004.06.008.

[44] AhmadBeygi, S., Cohn, A., Guan, Y., and Belobaba, P., "Analysis of the potential for delay propagation in passenger airline networks," *Journal of Air Transport Management*, Vol. 14, No. 5, 2008, pp. 221–236. doi: 10.1016/j.jairtraman.2008.04.010.

[45] AhmadBeygi, S., Cohn, A., and Lapp, M., "Decreasing airline delay propagation by re-allocating scheduled slack," *IIE Transactions*, Vol. 42, No. 7, 2010, pp. 478–489. doi: 10.1080/07408170903468605.

[46] Giannikas, V., Ledwoch, A., Stojković, G., Costas, P., Brintrup, A., Al-Ali, A., Chauhan, V., and McFarlane, D., "A data-driven method to assess the causes and impact of delay propagation in air transportation systems," *Transportation Research Part C: Emerging Technologies*, Vol. 143, 2022. doi: 10.1016/j.trc.2022.103862.

[47] Zhang, X., and Zhu, X., "Modeling and Delay Propagation Analysis for Flight Operation Based on Time Interval Petri Net," *The Open Automation and Control Systems Journal*, Vol. 6, 2015, pp. 433–438. doi: 10.2174/1874444301406010433.

[48] Shao, Q., and Xu, C., "Air transportation delay propagation analysis with uncertainty in coloured–timed Petri nets," *Proceedings of the Institution of Civil Engineers - Transport*, Vol. 173, No. 6, 2020, pp. 380–395. doi: 10.1680/jtran.17.00159.

[49] Brueckner, J. K., Czerny, A. I., and Gaggero, A. A., "Airline delay propagation: A simple method for measuring its extent and determinants," *Transportation Research Part B: Methodological*, Vol. 162, 2022, pp. 55–71. doi: 10.1016/j.trb.2022.05.003.

[50] Baspinar, B., Ure, N. K., Koyuncu, E., and Inalhan, G., "Analysis of Delay Characteristics of European Air Traffic through a Data-Driven Airport-Centric Queuing Network Model," *IFAC-PapersOnLine*, Vol. 49, No. 3, 2016, pp. 359–364. doi: 10.1016/j.ifacol.2016.07.060.

[51] Han, Y.-X., "Stability analysis for scheduled air traffic flow," *Aircraft Engineering and Aerospace Technology*, Vol. 95, No. 8, 2023, pp. 1201–1208. doi: 10.1108/AEAT-03-2022-0071.

[52] Baccelli, F., Cohen, G., Olsder, G. J., and Quadrat, J. P., *Synchronization and Linearity*, Wiley, 1992. doi: 10.2307/2583959.

[53] Heidergott, B., Olsder, G. J., and Woude, J. W. v. d., *Max Plus at work: modeling and analysis of synchronized systems: a course on Max-Plus algebra and its applications*, Princeton series in applied mathematics, Princeton University Press, Princeton, New Jersey, 2006. doi: 10.1515/9781400865239.

[54] De Schutter, B., and van den Boom, T., "Max-plus algebra and max-plus linear discrete event systems: An introduction," *Proceedings of the 9th International Workshop on Discrete Event Systems (WODES'08)*, Delft Center for Systems and Control, Göteborg, Sweden, 2008, pp. 36–42. doi: 10.1109/WODES.2008.4605919.

[55] Komenda, J., Lahaye, S., Boimond, J.-L., and van den Boom, T., "Max-plus algebra in the history of discrete event systems," *Annual Reviews in Control*, Vol. 45, 2018, pp. 240–249. doi: 10.1016/j.arcontrol.2018.04.004.

[56] Johnson, D. B., "Efficient Algorithms for Shortest Paths in Sparse Networks," *J. ACM*, Vol. 24, No. 1, 1977, pp. 1–13. doi: 10.1145/321992.321993.

[57] Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C., *Introduction to Algorithms*, 3rd ed., The MIT Press, Cambridge, MA, 2009. URL `https://mitpress.mit.edu/9780262046305/introduction-to-algorithms/`.

[58] Ahuja, R. K., Magnanti, T. L., and Orlin, J. B., *Network Flows: Theory, Algorithms, and Applications*, Prentice Hall, 1993.

[59] Cochet-Terrasson, J., Cohen, G., Gaubert, S., McGettrick, M., and Quadrat, J.-P., "Numerical Computation of Spectral Elements in Max-Plus Algebra," *IFAC Proceedings Volumes*, Vol. 31, No. 18, 1998, pp. 667–674. doi: 10.1016/S1474-6670(17)42067-2.

[60] Cook, G., and Goodwin, J., "Airline Networks: A Comparison of Hub-and-Spoke and Point-to-Point Systems," *Journal of Aviation/Aerospace Education & Research*, Vol. 17, No. 2, 2008. doi: https://doi.org/10.15394/jaaer.2008.1443.

[61] Herring, J., Lurkin, D. V., Garrow, D. L. A., John-Paul Clarke, D., and Bierlaire, D. M., "Airline customers' connection time preferences in domestic U.S. markets," *Journal of Air Transport Management*, Vol. 79, 2019, p. 101688. doi: 10.1016/j.jairtraman.2019.101688.

[62] EUROCONTROL, "Aviation Data Repository for Research," , 2025. URL `https://www.eurocontrol.int/dashboard/aviation-data-research`.

[63] Bureau of Transportation Statistics, "Airline On-Time Performance Data," , 2025. URL `https://transtats.bts.gov/DataIndex.asp`.

[64] Shepardson, D., "Alaska Airlines completes $1.9 billion acquisition of Hawaiian," , Sep. 2024. URL `https://www.reuters.com/markets/deals/alaska-airlines-completes-acquisition-hawaiian-2024-09-18/`, hOWPUBLISHED: Reuters.

[65] Villegas, M., "What It Takes to Become the Most Punctual U.S. Airline, Year after Year," , Jul. 2019. URL `https://newsroom.hawaiianairlines.com/blog/what-it-takes-to-become-the-most-punctual-u-s-airline-year-after-year`.

[66] Hawaiian Airlines, "Minimum Connection Times," , 2025. URL `https://www.hawaiianairlines.com/Minimum-Connection-Times`.

Part II

Literature Review and Research Proposal

1 Introduction

Note that the following literature review and research proposal were written 7 months before the completion of the thesis. As such, not all contents may represent the latest knowledge or work of the author. The research questions presented in this work also differ slightly from the research questions used in the final work.

In 2023, 71% of civil aviation flights in the European Civil Aviation Conference (ECAC) member states and 78% of United States domestic flights arrived on time (Bureau of Transportation Statistics, 2024; Walker, 2024). As the impact of these delays can be quantified in the order of tens of billions of dollars globally (AirHelp, 2023; Ball et al., 2010; Scheelhaase et al., 2023; Value Group, 2023), it is important to understand their cause. Multiple causes for delay exist, such as weather events, maintenance or security issues or a congested airspace. However, for both Europe and the US, delay propagation can be identified as the main contributor to delay (46% of delay minutes in Europe and 34% of delayed flights in the US[1]). Propagated or secondary delay is defined as any delay caused by the late arrival of aircraft, crew, passengers or loads from a previous flight. This in contrast to primary or root delay, which represent the root cause of a chain of previous flights being delayed. Propagated delay is inherently tied to scheduling, as the amount of slack in a schedule determines the disruption absorption capability of a schedule. For Europe, the breakdown by cause of the average delay minutes is shown in figure 1.1.

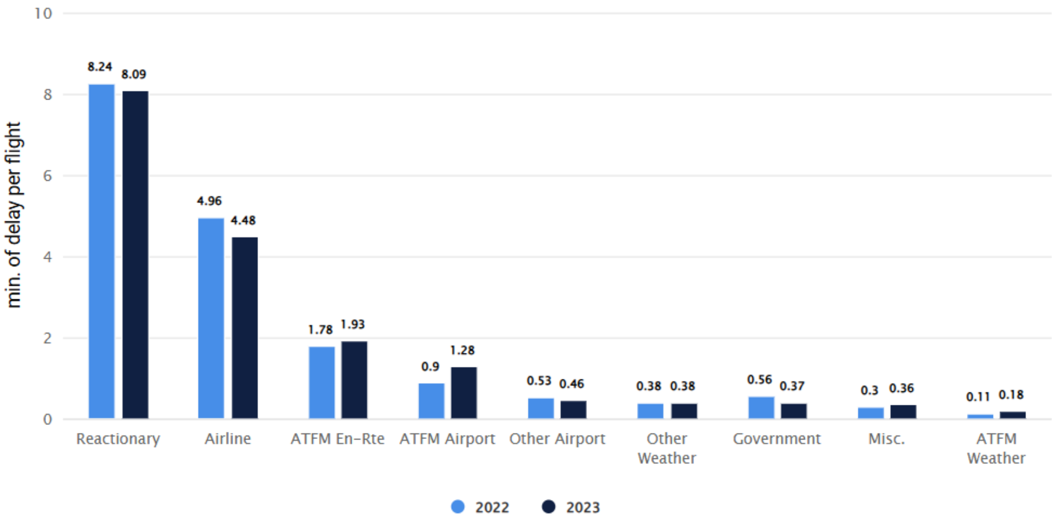

Figure 1.1: Breakdown by cause of the average delay minutes per flight 2022 vs. 2023. Note that reactionary delay is equivalent to propagated delay. ATFM refers to Air Traffic Flow Management. (Walker, 2024)

[1]The US Bureau of Transportation Statistics only reports propagated delay caused by late arriving aircraft, not by late arriving crew, passengers or loads from a previous flight.

Evidently, propagated delay is significant in both Europe and the US. Therefore, the project proposed in this literature review and research proposal will aim to introduce a new technique for modelling delay propagation in aviation networks. It will draw inspiration from railway network modelling to apply a mathematical technique known as max-plus algebra (Goverde, 2007). Max-plus algebra is especially powerful for discrete event systems in which there is some form of cyclic timetable and there is interdependency between certain events. While the cyclicity is less the case for aviation networks than railway networks, a cyclicity can be recognised. For instance, nearly all aviation timetables operate on a daily or weekly basis. Additionally, many carriers operate multiple short-haul flights a day, in which the timetable also presents a clear cyclicity. Aviation networks do present a lot of interdependency on their resources, with all of aircraft, cabin crew, cockpit crew and passenger connections causing interdependencies.

An application of max-plus algebra to a flight network was performed by Rengers (2017). However, the model was only applied to a very limited artificial hub-and-spoke network with two hubs and four spokes. It also does not extract any stability parameters from the model, which max-plus algebra is especially suitable for, nor does it calculate delay propagation. The project discussed in the research proposal thus aims to expand on this work and tackle current gaps in the knowledge of delay propagation modelling.

The research proposal itself will be discussed in chapter 3 of this report. First, however, chapter 2 contains a comprehensive literature review to understand the state of the art. Some conclusions will be drawn in chapter 4.

2
Literature Review

This chapter aims to give an overview of the state of the art of delay propagation modelling in aviation networks and extract a gap in the research. It will lay out a variety of sources which should collectively present the current knowledge of the field.

First, a short introduction to aviation networks will be given in section 2.1. Afterwards, how these networks are scheduled will be discussed in section 2.2. Section 2.3 will discuss the role of delay propagation in delay estimation, after which section 2.4 will discuss a variety of delay propagation modelling methods. An alternative modelling method called max-plus algebra, which is based on Petri nets, will be introduced in section 2.5. Afterwards, sections 2.6 and 2.7 will present applications of Petri nets and max-plus algebra in respectively railway and aviation network modelling. Finally, section 2.8 will present the gap in the research.

2.1. An Introduction to Aviation Networks

The project discussed in this research proposal aims to improve the modelling of delay propagation in aviation networks. Before diving into technical specifics, it makes sense to first understand what aviation networks look like. This will be introduced in this section.

An aviation network consists of flights connecting a variety of airports. These flights can either be considered part of commercial aviation or general aviation (International Civil Aviation Organization, 2009). Commercial aviation typically concerns all flights which transport goods and people under a commercial structure. This includes ticketed passenger and cargo flights, but also chartering entire aircraft. General aviation concerns all other types of flights, for research, instruction, pleasure or other purposes. As the most public information is available on commercial aviation and as commercial aviation typically concerns the most complex networks, research on aviation networks will typically focus on commercial aviation networks. For this reason, this research and the subsequent discussions will also focus on commercial aviation.

Different commercial airline setups exist. Some airlines focus exclusively on passenger or cargo transport, while others operate both. From a network perspective, the most typical distinction is between a hub-and-spoke network and a point-to-point network, as shown in figure 2.1. Note that the hub-and-spoke network requires much fewer flights to connect all 9 airports. However, tickets are typically offered from one point to another including a connection, forcing aircraft to wait for incoming aircraft's passengers in case of delay.

Older, traditional airlines such as American Airlines, KLM or Lufthansa typically operated hub-and-spoke networks. Newer so-called Low-cost carriers, such as Ryanair, EasyJet or JetBlue, typically operated point-to-point networks. However, modern airlines will usually operate a blended model, rather than operating exclusively via a hub or exclusively point-to-point (Cook & Goodwin, 2008). For example, United Airlines now also operates point-to-point flights and Ryanair also offers connecting flights via hubs (airliners.de, 2017; Slotnick, 2021). The frequency of flights between certain airports varies

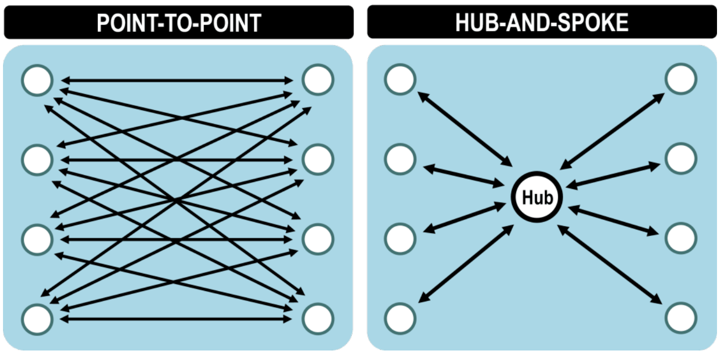

Figure 2.1: The difference between a hub-and-spoke and a point-to-point network. (Rodrigue, 2024)

between airlines and airports. Some airlines only operate 1 flight between two airports daily, while others operate many. For example, Air France operates between 8 and 10 flights between Amsterdam and Paris on a daily, fixed schedule.

To maximise the efficiency of possible transfers, hub-and-spoke carriers will often let all inbound flights at their hub arrive in the same time frame, followed by a large number of outbound flights. Hub airports will often adjust their runway use to accommodate for these inbound and outbound peaks. A combination of an inbound and an outbound peak is referred to as a block. For example, KLM operates a seven-block system at its hub airport (Gordijn, 2020).

As there are a vast number of flights (over 27,000 daily in ECAC member states (EUROCONTROL, 2024)) sharing limited resources and under external effects, delay is common and can occur in any of the different flight operation stages. Some of these stages are schematically represented in figure 2.2. To distinguish between different delay causes, the International Air Transport Association (2016) (IATA) specified different delay codes. For example, flight time is highly dependent on local weather conditions (73), and the time to approach can be dependent on local congestion in the airspace (81/83) or at the airport itself (83/88). Taxi time is dependent on the distance between runways and gates (87), and the time spent at a gate is highly dependent on the availability and speed of ground operations (09-39; 61-69). Note that the delay causes shown in figure 1.1 are based on the IATA delay codes.

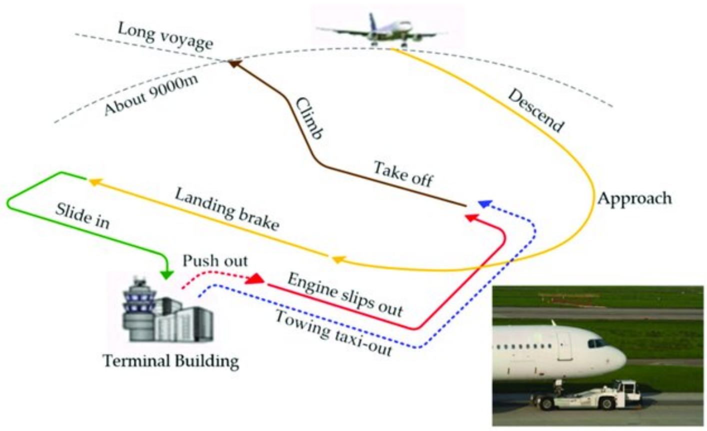

Figure 2.2: Schematic diagram of commercial aircraft arrival, departure and taxiing process. (J. Qin et al., 2023)

If a flight is delayed, this can have delay impacts beyond the delayed flight itself. These ”reactionary” or propagated delays are caused by the interdependency between flights. In the IATA delay codes, these are labelled as 91-96. While other exist, the most common interdependencies between flights are the transfer of the aircraft, the cockpit crew, the cabin crew, the passengers and the cargo. Note that if a flight is delayed on the ground, it can also take up resources which other flights might need, such as a gate or a de-icing station. These are also interdependencies. A flight chain representation of interdependencies is shown in figure 2.3. In a flight chain, flights are modelled as nodes and connections between flights are shown as edges. This is in contrast to an airport network, where airports are nodes and flights are edges.

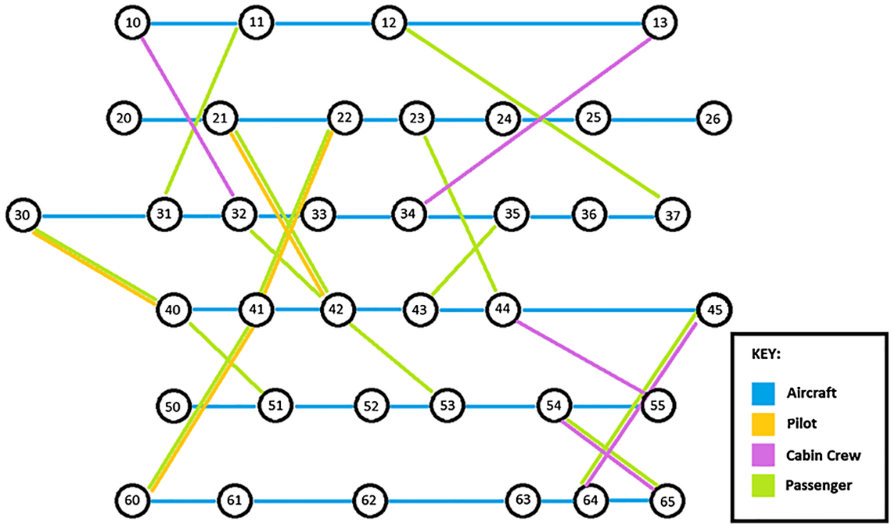

Figure 2.3: The flight network under study in C.-L. Wu and Law (2019), represented as a flight chain. Every number represents a flight. The interdependencies between flights are shown. (C.-L. Wu & Law, 2019)

As delay is random and in some cases, unavoidable, airlines use buffer time and mitigation measures to manage delay. Buffer time refers to adding extra time to the minimum required time for a specific part of the flight process. Airlines will often make a distinction between in-flight buffer time and ground or turnaround buffer time (C. Li et al., 2024). Buffer time is sometimes also referred to as slack in a schedule (AhmadBeygi et al., 2010). If buffer time is insufficient to absorb delay, airlines can deploy operational mitigation measures. This could be swapping out an aircraft or crew, or reallocating passengers or cargo to a different flight. With these measures, the interdependency changes. However, this requires manual intervention and will typically only occur once delay starts to become so significant that buffer time will not be sufficient to absorb the impact over multiple cycles. If the initial schedule is disrupted in this way, airlines must also work to recover from the disruption and get aircraft, crew and passengers back to their scheduled places. This problem is called the airline recovery problem (S. Wu et al., 2025).

Over 2023 in ECAC member states, 70.6% of flights were considered to arrive on time (Walker, 2024). Note that it is industry practice to consider an aircraft which is delayed less than fifteen minutes to be considered ”On Time”, while any aircraft which is delayed more than fifteen minutes to be considered ”Delayed” (Huynh et al., 2024). This is also evident from figure 2.4. Note that only 6.4% of flights were delayed more than 1 hour. Typically, this would be a threshold where mitigation measures would start to be relevant. Note that in the European Union, airlines are required to compensate passengers when their flight is delayed more than 2 hours (Kouris et al., 2020).

2.2. The Scheduling Problem in Aviation Networks

Airlines are well aware of the risks of delay and its financial and reputational impacts. However, properly preparing for it through buffer time in scheduling can be difficult, as the airline scheduling problem (ASP) is already very complex. Additionally, adding buffer time to a schedule reduces the utilisation rate of resources, increasing cost. To properly understand the difficulties in scheduling sufficient buffer time at the right moments, this section will give a short overview of the airline scheduling problem. Generally, the ASP is split into four sequential problems to reduce the computational complexity (Zhou et al., 2020). Most of these are formulated as a mixed integer programming problem. Each of the problems will be discussed below:

1. The schedule design problem: The SDP concerns setting up the timetable of flights. Typically, this problem is mainly solved based on marketing and sales decisions. If mathematical optimisation is involved, it typically uses passenger demand per flight leg and the market share of the airline.

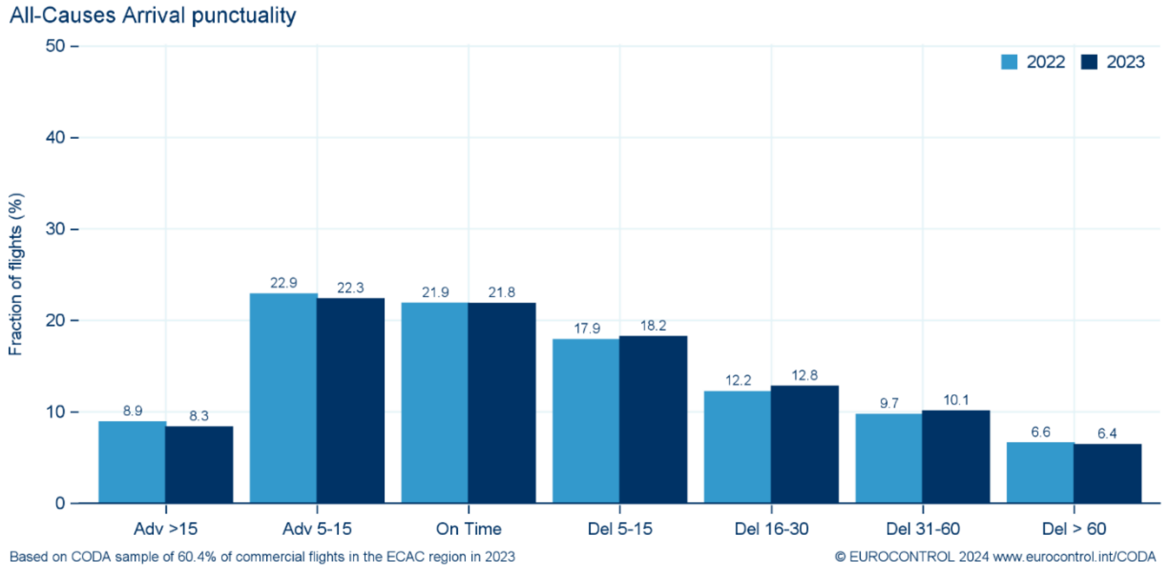

Figure 2.4: Arrival punctuality in ECAC member states in 2023 compared to 2022. Note that Adv indicates flights arriving before their scheduled time and that Del indicates flights arriving after their scheduled time. Also note that against its own 15-minute convention, EUROCONTROL here denotes an aircraft to arrive on time if it arrives 5 minutes before or after its scheduled time. (Walker, 2024)

The SDP is usually solved at least 6 months to a year in advance. (Eltoukhy et al., 2017)

2. The fleet assignment problem: In the FAP, each flight is assigned an aircraft type. Different variations exist, some also including passenger demand and aircraft type capacity. The FAP is usually solved not long after the SDP.
3. The aircraft routing problem: The ARP determines which specific aircraft is assigned to which flight. An important limitation in the ARP is maintenance, which an aircraft is typically required to undergo maintenance after a number of flights, flight hours or days/months. The ARP is typically solved months in advance.
4. The crew scheduling problem: The CSP is the final problem in the ASP, and concerns the assignment of cabin and cockpit crew to specific flights. It is usually split into two subproblems: The crew pairing problem (CPP) and the crew rostering problem (CRP). A crew pairing is a series of flights spanning one or more days which can be serviced consecutively by a crew member. The CPP selects the combination of crew pairings which minimises the total cost. The CRP then allocates specific crew members to certain pairings based on people-related parameters such as vacation, mandatory rest, fairness or seniority. The CPP is typically solved months in advance, while the CRP can only be solved one or two months in advance.

Together, the four sub-problems above generate a complete schedule for an airline network. Be aware that none of these problems are trivial, and extensive literature has been written on each of them. New technologies have allowed some of the sub-problems to be combined. Most notably, Y. Xu et al. (2021) integrated the SDP, the FAP and the ARP and expanded the traditional modelling approach with expected propagated delay. However, the estimation of the propagated delay is still rudimentary.

A final note is on the time frame in which each sub-problem is solved. The timetable, including the buffer time, is determined very far in advance. However, as the CSP is only solved in the last few months before a flight, the interdependencies between flights are only clear late in the process.

2.3. Delay Estimation in Aviation Networks

Before diving into delay propagation specifically, it makes sense to first understand the place of delay propagation modelling in the broader scope of delay estimation in aviation. A number of authors have published reviews of delay estimation methods over the past years. Three of these will be discussed in this section. First, an overview of each review will be given. Afterwards, two relevant conclusions will be extracted and presented.

The most expansive recent attempt to categorise different delay research approaches was performed by Carvalho et al. (2021), which included research up to March 2020. They categorised different approaches based on a number of categorisations, most notably their perspective (arrival, departure, propagation, airline, airport, or air system delay), the data sources and data pre-processing techniques used, and the fundamental method (classification, clustering, machine learning, network analysis, pattern mining, regression, statistical analysis) used. One particular conclusion of interest was that only 6.5% of analysed articles could be considered reproducible. This should thus be a particular focus for the work in this project.

Huynh et al. (2024) performed a similar exercise including work up to April 2024, but focused their work on short time frame delay forecasting models which had undergone comprehensive review. Due to this, only 33 sources were considered. Most of these sources used machine learning techniques, which are only a subset in the categorisation of Carvalho et al. (2021). This latter work included 123 sources. Huynh et al. (2024) thus exclude some of the more unconventional techniques, and do not segregate based on which perspective each paper takes. As such, it is not immediately clear what the place of delay propagation modelling is in the papers under study. Still, they do organise the models under study along different time frames and horizons. Through this, they demonstrate that the main interest of current study is in (very) short-term forecasting. This is in line with the needs of the industry, as aircraft delay rarely exceeds one hour (see figure 2.4).

T. Wang et al. (2022) performed a literature review including sources up to 2021 which was more in the style of Carvalho et al. (2021), albeit less expansive. They also categorised flight delay prediction methods into three main categories: traditional statistical analysis, modelling simulation and queuing theory, and machine learning methods.

Two conclusions on the place of delay propagation modelling in delay research can be drawn from the three literature reviews. Figure 2.5 shows the prevalence of certain perspectives in delay propagation modelling. As evident from the figure, delay propagation accounts for a similar amount of research as research from the perspective of airports, airlines, arrivals or departures. This contrasts, however, with the significance of propagated delay in the total resulting delay as presented in chapter 1 and figure 1.1 specifically. T. Wang et al. (2022) also explicitly mention delay propagation as an area for future research, stating there is necessity to ”increase the quantitative assessment and efficient delineation of delays caused by the propagation of flight delays.” There thus seems to be a need for future work in the area of delay propagation in aviation. A deep-dive into this topic will be presented in section 2.4.

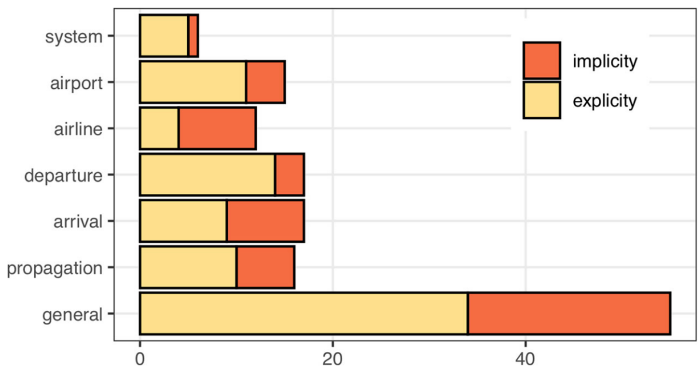

Figure 2.5: The prevalence of the perspective of delay research in the review of Carvalho et al. (2021). (Carvalho et al., 2021)

Another conclusion that can be drawn relates to the techniques used to estimate delay. Figure 2.6 shows the number of publications per year categorised by their main underlying method used. Evidently, machine learning methodologies are increasingly popular. Note that classification and cluster analysis algorithms can, in some cases, also be considered machine learning algorithms. The findings of Huynh et al. (2024) show this trend as well. This increase is in line with the recent global trend of research into machine learning technologies, as they are powerful and versatile tools. However, Carvalho et al. (2021) explicitly mentions that network analysis techniques are underexplored. The max-plus algebra techniques introduced in chapter 1 could thus be an interesting new analysis technique to explore. A more in-depth analysis of max-plus algebra and its potential applications will be presented in sections 2.5 to 2.7.

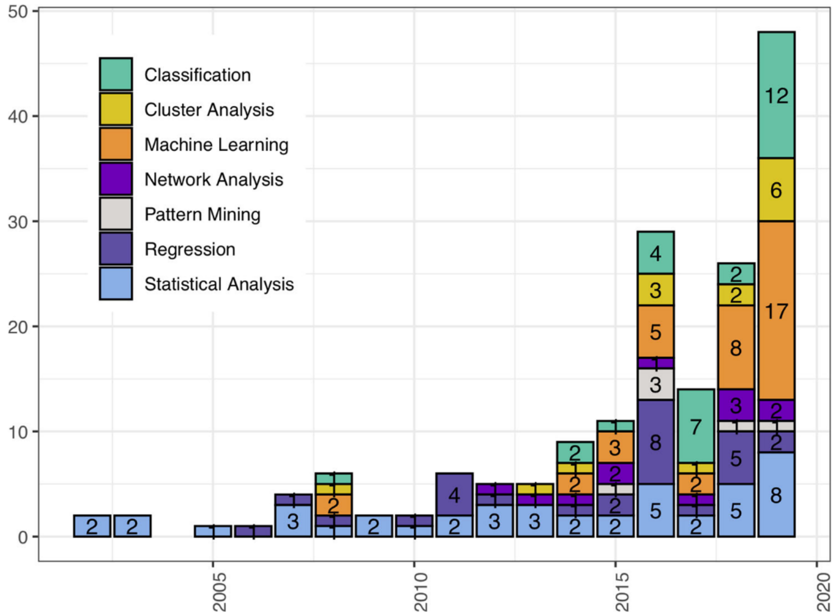

Figure 2.6: The number of publications on delay estimation in the review of Carvalho et al. (2021) per year, categorised by their main method used. (Carvalho et al., 2021)

2.4. Delay Propagation Modelling in Aviation Networks

Although delay propagation is less explored than delay research related to primary delay according to T. Wang et al. (2022), there is still a lot of work that has been performed in the area of delay propagation modelling in aviation. Most recently, C. Li et al. (2024) performed an extensive literature review of delay propagation modelling in aviation networks. They included 120 articles, published up to September 2022. This section serves as a summary of their work, including certain highlighted articles and additional literature from more recent years.

C. Li et al. (2024) categorised the papers under research into seven distinct categories based on the modelling method. While in no way the only possible categorisation, it does provide a clear overview of the state of the art. As the research proposed in this paper focuses on an uncommon modelling method, it is relevant to understand what other methods exist. The seven categories are:

1. Economic models
2. Statistical models
3. Machine learning models
4. Epidemic spreading models
5. Simulation-based models
6. Queueing models
7. Network representation models

Be aware that certain works can be combinations of multiple methods and can thus be mentioned in multiple categories. Each category will be discussed below.

Economic models are models which focus on describing the impact of delay of delay propagation. For example, the delay multiplier index ($DM = \frac{\text{initial delay+propagated delay}}{\text{initial delay}}$) as first presented by Beatty et al. (1999) falls under this category. This index has since been expanded on, most recently by C.-L. Wu and Law (2019). In their work, they proposed the expected delay multiplier and the expected propagated delay multiplier and applied these to a delay propagation tree model with a Bayesian network. Kafle and Zou (2016) developed an econometric modelling framework that jointly quantified the effect of various factors on the initiation and progression of propagated delay.

Statistical models use data analysis to aid in quantifying the significance of different factors influencing delay propagation. They also help in recognising certain patterns in the data. Examples of methods used in statistical models are survival analysis, comparative analysis, or regression. One limitation of statistical models indicated by C. Li et al. (2024) is that they can struggle with complex, non-linear or causal relationships, which are prevalent in aviation networks.

Machine learning methods focus on processing large amounts of data to forecast delay values. Besides information on delay, they are often also fed a number of other useful parameters such as weather or active runways. These methods are very powerful, and as shown in section 2.3, they are growing in popularity. However, machine learning methods often require expert knowledge and significant computation power to apply them successfully.

Epidemic spreading models are models which base their mathematical modelling on the propagation behaviour of infections in a population. In this application, the population will typically be either airports or individual flights. Particularly airport-based models bear similarity with network representations, and complex network theory will also be applied in some cases. Generally, epidemic spreading models can help illustrate the spread of delay through a network. However, they can tend to oversimplify the complex causal relationships behind an aviation network.

Simulation-based models are models in which the complex aviation network is simulated in a computer. The simulation takes some simplifying assumptions to simplify the otherwise computationally intensive modelling of the real system. Typically, two types can be recognised: agent-based models and general simulation-based models.

Agent-based models are highly versatile modelling techniques. They will typically model either a flight or an airport as an agent. These then interact with each other and the environment to simulate the real aviation system. A notable work in this category is that of Fleurquin et al. (2013). For the US aviation network, they identified passenger and crew connections as the most relevant internal factor contributing to delay propagation. This could be explained by the fact that existing research does not include these connections in the analysis. Typically, this is as the information is not publicly available, while aircraft connections typically are. However, S. Qin et al. (2019) performed a similar analysis for the Chinese aviation network, and concluded that aircraft connections were the most relevant internal factor contributing to delay propagation. This could thus be network-specific as well. Fleurquin et al. (2013) also concluded that US aviation network can be unstable, even under normal operating conditions. Generally speaking, agent-based models can generate precise micro-level insights into the complex aviation networks. However, they can be computationally very intensive.

General simulation-based models include all other simulation-based models which do not use agent-based modelling. They utilise either flight chains or airport network-based approaches. A notable work in this category is that of Baden et al. (2006), who applied a backtracking algorithm to trace cumulative delays back to their source. They then used the insights gained to understand what sources caused delay propagation. However, the approach could not be used for prediction purposes. Interestingly, C.-L. Wu (2006) utilise a simulation model to evaluate the effectiveness of implemented buffer time, in an approach also utilising discrete event systems. Discrete event systems include dynamic models formulated in max-plus algebra. Generally speaking, general simulation-based algorithms are less granular than their agent-based counterparts, but are highly adaptable to a variety of different scenarios in a macro-level context.

A relatively unique approach is those of the queueing models, which treat aircraft as ’customers’ and runways as ’servers’ where flights are processed. Aircraft then have to queue before they can get ’served’ by a runway. The approach is powerful for modelling the complex relationships in the use of the airspace and airport resources, where flight order and capacity can significantly impact the delay propagation. It can also be used in a network approach, in which different queues are connected to each other. In general, the queueing models are adept at identifying operational capacity restrictions, but may not encompass all complexities of aviation systems. In particular, they can ignore external factors such as weather or seasonal impacts.

The final and most expansive category of modelling techniques concerns network representation models. Generally, four types of network representation models can be recognised: Bayesian networks, complex networks, graph neural networks and alternative spatio-temporal models. Each will be discussed here.

Bayesian networks (BN) are probabilistic graphical models that represents a set of variables and their conditional dependencies via a directed acyclic graph. Due to the stochastic nature of aviation networks and the extensive interdependency of parameters, Bayesian networks are used extensively in a wide array of application in aviation networks. For instance, N. Xu et al. (2005) used BN to model the delay propagation between three airports. Their model is shown in figure 2.7. Note that it does not show a full aviation network but only a single flight chain. However, it does take into account external factors such as weather and the time of day. In a more recent application, C.-L. Wu and Law (2019) used Bayesian networks in combination with a delay propagation tree model. In this model, the departure delay of a flight was stochastically determined through a Bayesian network. Their work notably included crew and passenger connections as an input, as well as an approach to estimate this information for an actual network. They were also able to improve the schedule using their model. However, their approach still faced the difficulty which is typical of BN, mainly related to the complexity and technical understanding required to set up the correct network topology and estimate the various parameters.

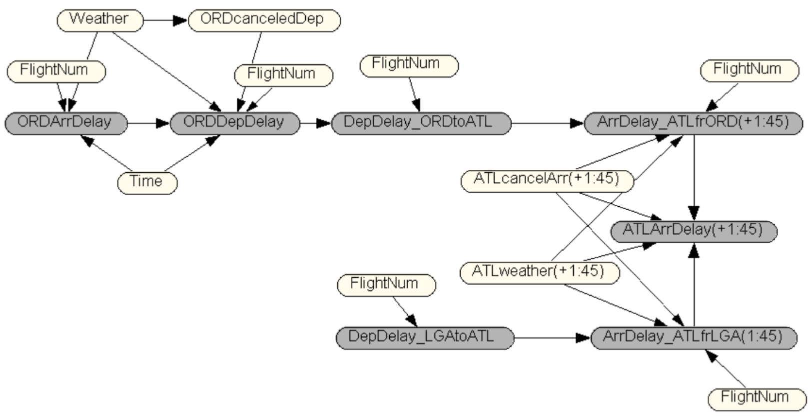

Figure 2.7: The Bayesian network used in the work of N. Xu et al. (2005). (N. Xu et al., 2005)

Complex network theory can be used to extract specific metrics or characteristics of aviation networks. One approach is to set the network up with all airports represented as nodes and direct flights as edges. Another approach is to let nodes represent airports, and the presence of edges between nodes show causal relationships. In this approach, whether an edge is included is determined by a statistical significance test. The dominant type is the Granger causality test. In both approaches, metrics such as betweenness centrality, degree or reciprocity can be extracted. A recent work is that of Y. Li et al. (2024), who developed a fully-connected delay propagation network based on transfer entropy and compared it to other existing methodologies. Generally, complex network theory metrics help to analyse how strongly connected the network is, and thus how fast delay could spread across the network at a macro level. However, they can ignore individual micro-level interactions or external factors.

Graph neural networks are a category of deep learning modelling methods. A recent contribution in this category is that of Y. Wu et al. (2024), who proposed a space-time separable graph convolutional network to model delay propagation. Graph neural networks are still a new area of research for aviation delay propagation modelling. Still, it is already proving promising in analysing the complex behaviours of delay propagation in large networks with a high level of detail. However, it does require significant computational resources.

Finally, alternative spatio-temporal models include all network representation modelling techniques which do not fall into any of the earlier categories. A notable contribution in this category is the work of AhmadBeygi et al. (2008). They proposed to use a delay propagation tree (DPT), an example of which is shown in figure 2.8, to estimate the delay of subsequent flights in a chain. AhmadBeygi et al. (2010) built on this work, utilising a DPT model to reallocate scheduled slack to improve the robustness of the schedule without changing initial scheduling problem. C.-L. Wu and Law (2019) and W. Wu and Wu (2018) expanded on this approach to combine it with the Bayesian networks presented earlier. Giannikas et al. (2022) constructed a multi-layer network model including aircraft, passenger and crew connections to model delay propagation. They then argue that delay networks are better represented as an acyclic graph than in a tree structure, as an acyclic graph can model a single aircraft being delayed due to multiple other delayed flights. They also show that 64.5% of delay propagation connections were only passenger-related and that delay propagation is underestimated without a multi-layer approach.

Timed Petri Nets also fall under the category of alternative spatio-temporal models, and are a key topic of this paper. The state of the art of their application in aviation networks will be discussed in section 2.7.

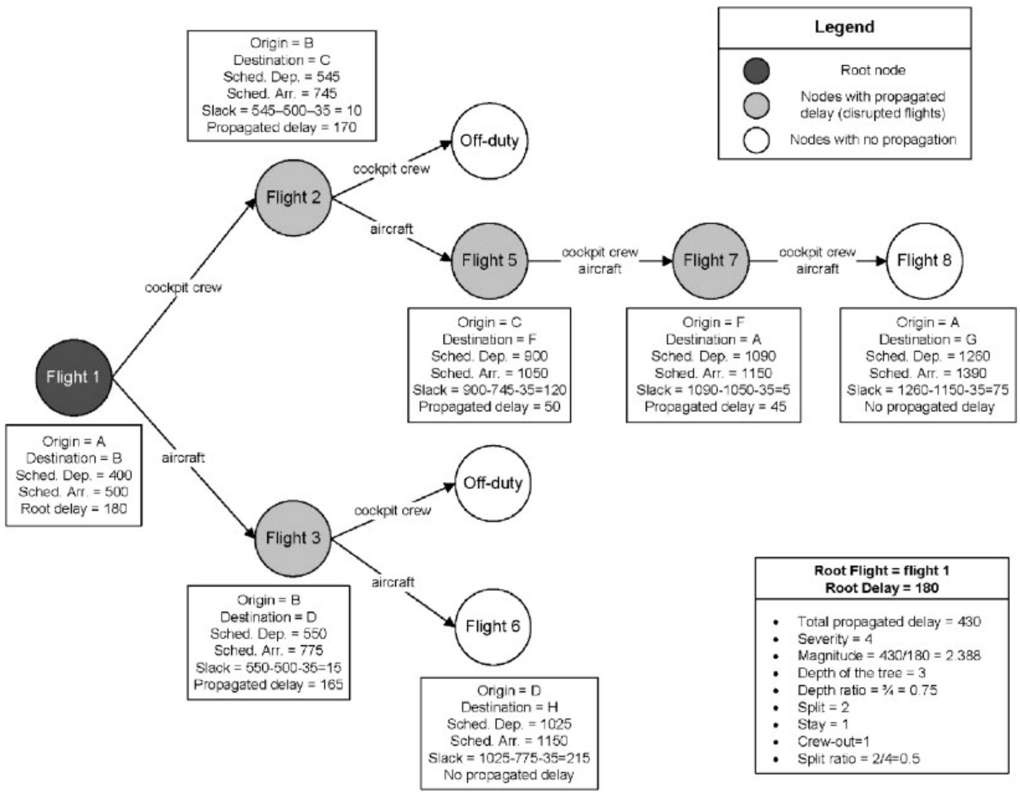

Figure 2.8: Example of a delay propagation tree. (AhmadBeygi et al., 2008)

This section has aimed to give an overview of the state of the art of delay propagation modelling in aviation. Specifically, it showcased seven main categories of methods used to model delay propagation. C. Li et al. (2024) also quantified the prevalence of certain modelling methods in their review, as shown in figure 2.9. In interesting contrast to the conclusions by Carvalho et al. (2021), network representation models are widely used in delay propagation modelling. Looking further at figure 2.10 shows that the majority of those models use Bayesian networks.

2.5. Introduction to Petri Nets and Max-Plus Algebra

Discrete event systems (DES) are a class of dynamic systems in which the system state changes through the occurrence of discrete, asynchronous events (Cassandras, 2005). They are sometimes also referred to as discrete event dynamic systems (DEDS). An aviation network can be modelled as a form of a DES. For example, an aircraft in the system can change in state from ”on the ground” to ”in-flight” through an event ”take-off”.

DES are often presented as so-called Petri nets. A Petri net shows the transitions (events) which change a part of the system from one place (state) to another. In a Petri net, one or more tokens are

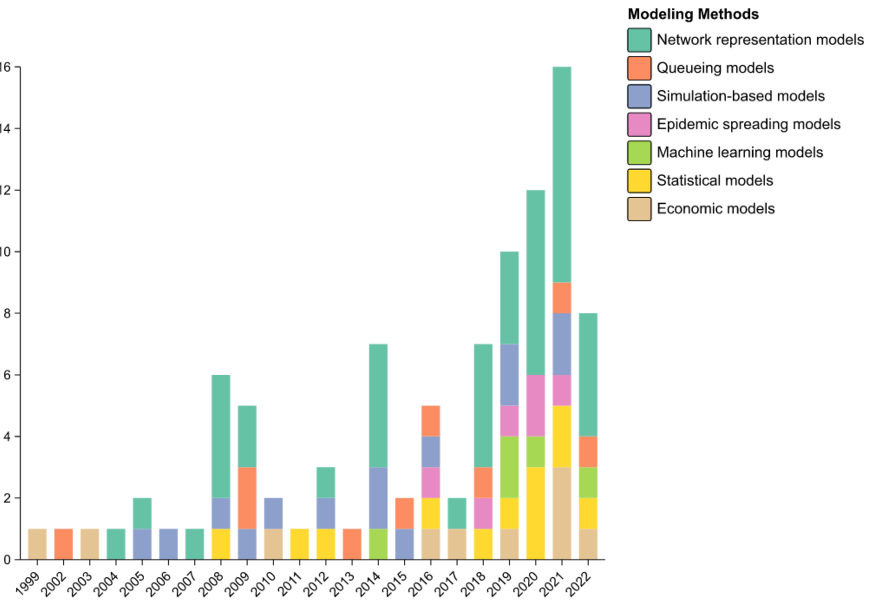

Figure 2.9: Number of publications by year with respect to different modelling methods in the review of C. Li et al. (2024). (C. Li et al., 2024)

assigned to a place if the condition described by that place is satisfied. The combination of all assigned tokens is called a *marking*. The marking thus defines the state of a Petri net at a specific moment. A Petri net shifts between states by moving a token from one place to another. Tokens can only move when a transition is enabled, a so-called *firing* of the transition. This occurs when there are sufficient tokens to traverse the weighted arcs leading into that specific transition (Cassandras, 2005). Different *firing rules* exist, each of which denote the circumstances under which a transition can fire (Komenda et al., 2018).

An example Petri net is shown in figure 2.11. By convention, in graph notation places are denoted with a circle and transitions are denoted with a bar. A token is denoted with a dark dot.

Places, transitions and arcs in Petri nets can be assigned a weight. Ordinary Petri nets are defined as Petri nets in which all arc weights are equal to 1 (Komenda et al., 2018). A specific class of ordinary Petri nets are timed event graphs (TEG), also known as decision-free timed Petri nets. TEGs are timed Petri nets in which each place has a single input transition and a single output transition. Note that transitions can still have multiple inputs or outputs. Weights of places and transitions denote the time that this place or transition takes. As TEGs avoid conflict in where tokens are consumed or where they are supplied to, they can be mathematically modelled using max-plus algebra (Komenda et al., 2018).

In max-plus algebra, the standard algebraic addition and multiplication operations are replaced by a maximum and addition operation. Max-plus algebra is somewhat of a misnomer, as it is not formally an algebra in the strictly mathematical sense. Rather, max-plus algebra refers to the idempotent semiring (dioid) of extended real numbers $(\mathbb{R} \cup \{-\infty\}, \max, +)$ (Komenda et al., 2018). Max-plus algebra is therefore sometimes also referred to as dioid algebra (Cassandras, 2005). In mathematical notation, the following operations are defined:

$$a \oplus b = \max\{a, b\} \quad (2.1)$$

$$a \otimes b = a + b \quad (2.2)$$

In conventional addition, the zero element 0 is absorbing: Adding 0 to a value returns the initial value. In max-plus addition, $\varepsilon = -\infty$ is defined as the zero element and is similarly absorbing: $a \oplus \varepsilon = a = \varepsilon \oplus a$ and $a \otimes \varepsilon = \varepsilon = \varepsilon \otimes a$ for all $a \in \mathbb{R}$ (De Schutter & van den Boom, 2008; Goverde, 2007). Note that as ε is the zero element, a non-zero element automatically refers to a finite element. (Goverde, 2007)

De Schutter and van den Boom (2008) use ε to define the number space of max-plus algebra as

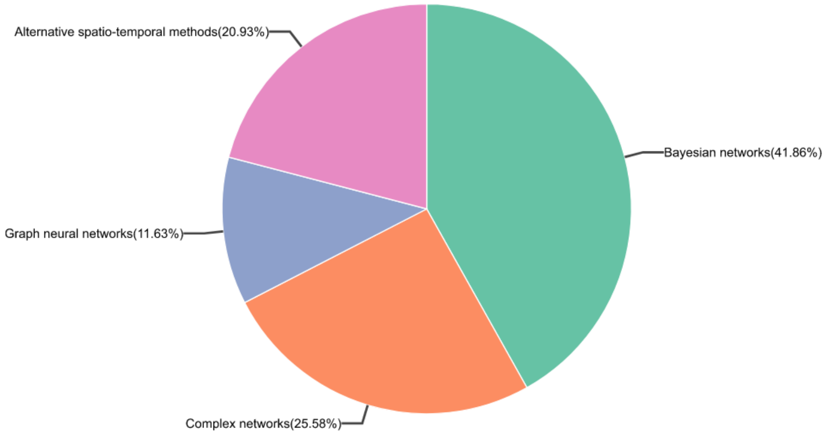

Figure 2.10: Proportions of different types of network representation models in the review of C. Li et al. (2024). (C. Li et al., 2024)

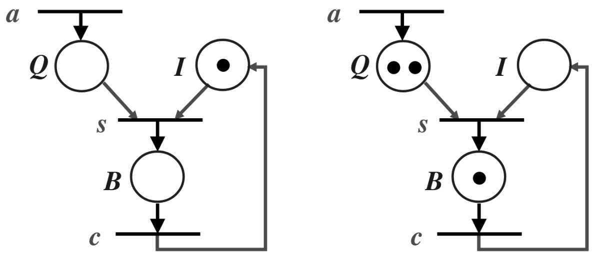

Figure 2.11: An example Petri net. Two states of the system are shown. (Cassandras, 2005)

$\mathbb{R}_\varepsilon \equiv \mathbb{R} \cup \{-\infty\}$. The structure $(\mathbb{R} \cup \{-\infty\}, \max, +)$ (Komenda et al., 2018) can then be replaced by the equivalent structure $\{\mathbb{R}_\varepsilon, \oplus, \otimes\}$.

Max-plus addition is commutative, associative and idempotent ($a \oplus a = a$). Due to the idempotency, max-plus addition has no reverse operation. Max-plus multiplication is associative and has a unit element $e = 0$. Multiplication is distributive over addition, meaning that $\otimes$ has priority over $\oplus$. $a \otimes b$ can therefore also be written as ab, as is common in conventional algebra (Goverde, 2007).

The max-plus-algebraic power operation for an $x \in \mathbb{R}$ and $r \in \mathbb{R}$ is defined as $x^{\otimes r}$ and is equivalent to rx in conventional algebra. Therefore, $x^{\otimes 0} = 0$ and $x^{\otimes -1} = -x$. For the special case $\varepsilon^{\otimes r}$, $\varepsilon^{\otimes r} = \varepsilon$ if $r > 0$, $\varepsilon^{\otimes r}$ is undefined if $r < 0$, and $\varepsilon^{\otimes 0} = 0$ by definition (De Schutter & van den Boom, 2008).

Max-plus algebra is a powerful tool, as it can be used to translate non-linear evolution equations in conventional algebra describing a TEG to linear equations in max-plus algebra. This allows for the use of linear algebra techniques and for representing the system in a state-space notation (Komenda et al., 2018). As such, it makes sense to also define matrix operations for max-plus algebra. If $A, B \in \mathbb{R}_\varepsilon^{m\times n}$ and $C \in \mathbb{R}_\varepsilon^{n\times p}$ then for all i, j: (De Schutter & van den Boom, 2008)

$$(A \oplus B)_{ij} = a_{ij} \oplus b_{ij} = \max(a_{ij}, b_{ij}) \tag{2.3}$$

$$(A \otimes C)_{ij} = \bigoplus_{k=1}^{n} a_{ik} \otimes c_{kj} = \max_k (a_{ik} + c_{kj}) \tag{2.4}$$

The max-plus-algebraic zero and identity matrices, $\mathcal{E}_{m\times n}$ and E_n, are defined as follows. $(\mathcal{E}_{m\times n})_{ij} = \varepsilon$ for all i, j; $(E_n)_{ii} = 0$ for all i and $(E_n)_{ij} = \varepsilon$ for all i, j with $i \neq j$. Note that $\mathcal{E}_{m\times n}$ has dimensions $m \times n$ and E_n has dimensions $n \times n$. Max-plus algebraic matrix power is defined as $A^{\otimes k} = A \otimes A^{\otimes k-1}$ for $k = 1, 2, \ldots$ and $A \in \mathbb{R}_\varepsilon^{n\times n}$. Additionally, $A^{\otimes 0} = E_n$ (De Schutter & van den Boom, 2008). Note that sometimes, the denotion $\otimes$ is not included in the power and only A^k is shown (Goverde, 2007).

Max-plus algebra is closely related to graph theory, and specifically precedence graphs. A precedence graph $\mathcal{G}(A)$ of a matrix $A \in \mathbb{R}_{\varepsilon}^{n \times n}$ is a weighted directed graph with vertices $1, 2, \ldots, n$ and an arc (j, i) with weight a_{ij} for each $a_{ij} \neq \varepsilon$ (De Schutter & van den Boom, 2008). An example precedence graph, taken from De Schutter and van den Boom (2008), is shown in figure 2.12. It corresponds to the matrix:

$$A = \begin{bmatrix} 0 & \varepsilon & 2 \\ 2 & 0 & 4 \\ 1 & 2 & 3 \end{bmatrix}$$

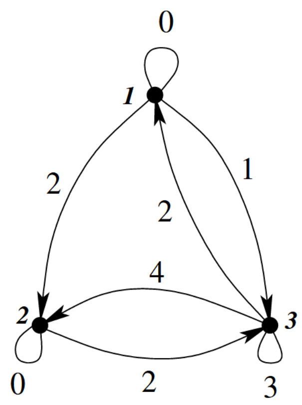

Figure 2.12: An example precedence graph. (De Schutter & van den Boom, 2008)

The max-plus-algebraic matrix power operation can be used to extract interesting properties from a precedence graph. For $k \in \mathbb{N}_0$, $A^{\otimes k}$ gives the weight of the path from j to i with length k which has the greatest weight. If a path does not exist, then by definition that weight is equal to ε. Using max-plus algebra to calculate $A^{\otimes k}$ for all $k \in 1, 2, \ldots, n-1$, one can thus determine which vertices in the precedence graph $\mathcal{G}(A)$ are connected with each other. If all vertices are connected, the precedence graph $\mathcal{G}(A)$ is considered to be *strongly connected*, and the corresponding matrix A is considered *irreducible*. In a more mathematical definition, matrix $A \in \mathbb{R}_{\varepsilon}^{n \times n}$ is *irreducible* if: (De Schutter & van den Boom, 2008)

$$(A \oplus A^{\otimes 2} \oplus \cdots \oplus A^{\otimes n-1})_{ij} \neq \varepsilon \text{ for all } i, j \text{ with } i \neq j \tag{2.5}$$

Goverde (2007) denotes the matrix $A \oplus A^{\otimes 2} \oplus \cdots \oplus A^{\otimes n-1}$ as the *critical path matrix* A^{+}.

For the example in figure 2.12, taken from De Schutter and van den Boom (2008), the precedence graph is strongly connected by inspection. However, it can also be shown mathematically. As equation (2.5) holds, matrix A is irreducible and precedence graph $\mathcal{G}(A)$ is strongly connected:

$$A \oplus A^{\otimes 2} = \begin{bmatrix} 0 & \varepsilon & 2 \\ 2 & 0 & 4 \\ 1 & 2 & 3 \end{bmatrix} \oplus \begin{bmatrix} 0 & \varepsilon & 2 \\ 2 & 0 & 4 \\ 1 & 2 & 3 \end{bmatrix}^{\otimes 2} = \begin{bmatrix} 0 & \varepsilon & 2 \\ 2 & 0 & 4 \\ 1 & 2 & 3 \end{bmatrix} \oplus \begin{bmatrix} 3 & 4 & 5 \\ 5 & 6 & 7 \\ 4 & 5 & 6 \end{bmatrix} = \begin{bmatrix} 3 & 4 & 5 \\ 5 & 6 & 7 \\ 4 & 5 & 6 \end{bmatrix}$$

An advantage of the matrix representation is that a max-plus-algebraic eigenvalue problem can be set up: For an $A \in \mathbb{R}_{\varepsilon}^{n \times n}$, if there exists $\lambda \in \mathbb{R}_{\varepsilon}$ and $v \in \mathbb{R}_{\varepsilon}^{n}$ with $v \neq \mathcal{E}_{n \times 1}$ such that $A \otimes v = \lambda \otimes v$, then λ is a max-plus-algebraic eigenvalue of A and v is the corresponding eigenvector. It has been shown that every square matrix with entries in $\mathbb{R}_{\varepsilon}$ has at least one eigenvalue. Additionally, an irreducible matrix has only one eigenvalue (De Schutter & van den Boom, 2008).

Sections 2.6 and 2.7 will outline the usage of Petri nets, max-plus algebra and their properties in respectively rail and aviation network delay propagation modelling.

2.6. Max-Plus Algebra in Rail Network Delay Propagation Modelling

Max-plus algebra has already been extensively applied to rail network delay propagation modelling. A lot of notable work has been performed by Goverde, which will be discussed in this section.

Goverde (2007) focused on modelling train departures as discrete events, under the assumption of a periodical (typically hourly) timetable in which each event occurs once every period. He proposed a model in which the kth departure time $x_i(k)$ of train i is dependent on the scheduled departure time $d_i(k)$ and a variety of precedence constraints related to the departure time of preceding trains $x_j(k-\mu_{ij})$. Here, μ_{ij} represents how many periods ago the preceding departure event j occurred. This model is represented through the following max-plus algebra recursive equations:

$$x_i(k) = \bigoplus_{j=1}^{n} (a_{ij}(\mu_{ij}) \otimes x_j(k-\mu_{ij})) \oplus d_i(k) \text{ for all } i = 1, \dots, n \tag{2.6}$$

where $a_{ij}(\mu_{ij})$ represents a scheduled process duration from departure j of μ_{ij} periods before to i. n is the total number of departure events in the timetable and $a_{ij} = \varepsilon$ for any unconstrained train pair (j, i). For the full derivation of equation (2.6) and subsequent equations, one is referred to Goverde (2007) or the subsequent work under discussion.

Setting $l = \mu_{ij}$ and collecting the process times a_{ij} into a matrix $A_l = (a_{ij}(l))$ with $[A_l]_{ij} = \varepsilon$ if no such term exists, equation (2.6) can be written in vector notation as follows:

$$x(k) = A_0 x(k) \oplus \cdots \oplus A_p x(k-p) \oplus d(k) = \bigoplus_{l=0}^{p} A_l x(k-l) \oplus d(k) \tag{2.7}$$

for some $p \in \mathbb{N}$, which Goverde (2007) denotes as the *order* of the system. Furthermore, in a scheduled max-plus linear system, which this is, the departure vector is defined as $d(k) = d_0 \otimes T^{\otimes k}$, where $d_0 = d(0)$.

Goverde (2007) goes on to define the backward-shift operator γ as $\gamma x(k) = x(k-1)$. They then use this to modify equation (2.7) to form the following:

$$x(k) = \bigoplus_{l=0}^{p} A_l \gamma^{\otimes l} x(k) \oplus d(k) = A(\gamma) x(k) \oplus d(k) \tag{2.8}$$

where $\bigoplus_{l=0}^{p} A_l \gamma^{\otimes l}$ is defined as a polynomial matrix in the shift operator γ.

Goverde (2007) applied equation (2.8) to the system shown in figure 2.13. The same system is shown as a timed event graph in figure 2.14. When applied in this way, equation (2.8) gives the following:

$$A(\gamma) = A_0 \oplus A_1\gamma = \begin{bmatrix} \varepsilon & 1 & 28 & \varepsilon \\ \varepsilon & \varepsilon & 28 & \varepsilon \\ \varepsilon & \varepsilon & \varepsilon & \varepsilon \\ \varepsilon & \varepsilon & 1 & \varepsilon \end{bmatrix} \oplus \begin{bmatrix} 52 & \varepsilon & \varepsilon & \varepsilon \\ 52 & \varepsilon & \varepsilon & \varepsilon \\ \varepsilon & 28 & \varepsilon & 57 \\ \varepsilon & 28 & \varepsilon & 57 \end{bmatrix} \gamma = \begin{bmatrix} 52\gamma & 1 & 28 & \varepsilon \\ 52\gamma & \varepsilon & 28 & \varepsilon \\ \varepsilon & 28\gamma & \varepsilon & 57\gamma \\ \varepsilon & 28\gamma & 1 & 57\gamma \end{bmatrix}.$$

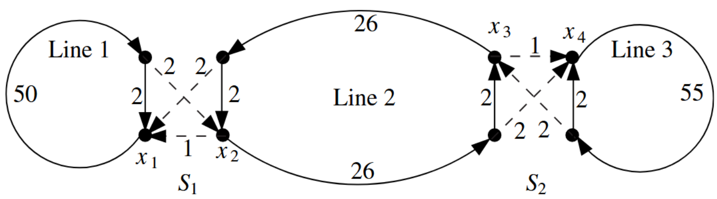

Figure 2.13: Example network in the work of Goverde (2007). (Goverde, 2007)

Using the formal definitions above, Goverde (2007) extracts a variety of useful insights. For instance, he shows that if $A(\gamma)$ is an irreducible polynomial matrix with acyclic subgraph $G(A_0)$, then $A(\gamma)$ has a unique eigenvalue $\lambda > \varepsilon$ and finite eigenvectors $v > \varepsilon$. This eigenvalue represents the maximum cycle mean of the system under study. The maximum cycle mean is the shortest possible time in which the system can complete one cycle. In other words, it is the cycle time of the critical circuit in the system. The eigenvectors represent initial timetable vectors $d_0 = v$ with which the system exhibits periodic behaviour with minimal cycle time $T = \lambda$.

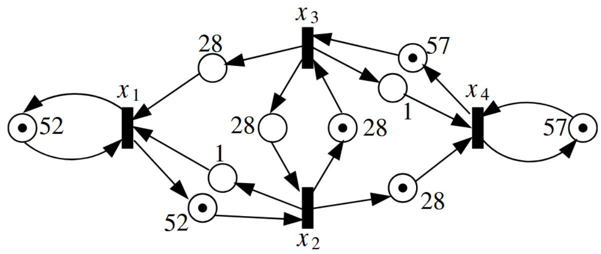

Figure 2.14: Timed event graph of the example network in the work of Goverde (2007). (Goverde, 2007)

Obviously, if the cycle time is less than the maximum cycle mean, the timetable will always generate and propagate delay. In other words, a timetable is unstable if $\lambda > T$ and, conversely, a timetable is stable if $\lambda < T$. If $\lambda = T$, the system is critical.

Note that the general statements above hold if the polynomial matrix $A(\gamma)$ is irreducible and the corresponding marked precedence graph (or timed event graph) is strongly connected. If this is not the case, then the eigenvalue λ is not necessarily unique. Rather, each subnetwork will have a supereigenvalue and subeigenvector corresponding to it. The supereigenvalue which corresponds to the most critical strongly-connected component of the system is denoted as λ_0 and represents the minimal attainable cycle time. Substituting λ with λ_0, a timetable is then stable if $\lambda_0 < T$.

Another indicator of stability is the delta $\Delta_1 = T - \lambda$, which gives the average amount of slack on critical circuits. A higher slack means that the system will be faster to recover from a disruption. Similarly, the *stability margin* Δ_2 specifies the maximum simultaneous increase of all process times such that the train network can still be operated with cycle time T. More formally, the stability margin is given by $\Delta_2 = \nu^{-1}$, where ν is the solution of the eigenvalue problem $A_T \otimes v = \nu \otimes v$ with $A_T = A(T)$.

If eigenvectors represent initial timetable vectors $d_0 = v$ with which the system exhibits periodic behaviour with minimal cycle time $T = \lambda$, it makes sense that an inverted formulation is also true: A periodic timetable is realisable only if d_0 is a generalised subeigenvector of $A(\gamma)$ associated to T. Formally, a periodic timetable $d(k) = d_0 \otimes T^{\otimes k}$ is realisable only if $d_0 \geq A(T^{\otimes -1}) \otimes d_0$.

For a stable system, there is slack which can absorb any disruptions in the system. The *recovery matrix* $R = r_{ij}$ is the $n \times n$ matrix of recovery times, where r_{ij} is the minimum total slack time over all paths from event j to i. Using max-plus algebra, the recovery matrix can be calculated using $r_{ij} = d_i^0 - d_j^0 - [A_T^+]_{ij}$, where $r_{ij} = \infty$ if there is no path from j to i. Note that in many cases, it is not necessary to calculate the full recovery matrix, reducing computation time.

Besides the stability insights above, Goverde (2007) also presents an algorithm to model delay propagation using max-plus algebra. He assumes $d_0 > \varepsilon$ and expands the scheduled max-plus linear system with an output vector $z(k)$ denoting the delays in period k:

$$\begin{cases} x(k) = A(\gamma)x(k) \oplus d(k), & k \in \mathbb{N}, \\ d(k) = d_0 \otimes T^{\otimes k}, & k \in \mathbb{N}, \\ z(k) = D^{-1}(k) \otimes x(k), & k \in \mathbb{N}, \\ x(1^-) = x_1, x(l) = x_l, & 1 - p \leq l \leq 0, \end{cases} \tag{2.9}$$

where $D^{-1} = diag(d_1^{\otimes -1}(k), \ldots, d_n^{\otimes -1}(k))$. This leads to $z_i(k) = d_i^{\otimes -1}(k) \otimes x_i(k) = x_i(k) - d_i(k)$.

As equation (2.9) presents a recursive system, the solution is simple given an initial delay vector $z(k)$. The algorithm proves to be effective and efficient, and that it can be used to extract the settling period k_s, average secondary delay, number of delayed trains and number of stations with delays. Goverde (2007) applies the work to a real Dutch railway timetable in the program PETER.

Goverde (2010) expands on the work of Goverde (2007) to include also arrival and passage events, besides only departure events. Additionally, he proposes a bucket-based delay propagation algorithm. In this algorithm, the delay propagation algorithm used in Goverde (2007) is improved upon by only computing the event time for delayed events, ordered by scheduled event time. After all, computing the event time for non-delayed events is nonsensical.

As the full set of matrices $A_0, \dots, A_p$ associated with a (non-scheduled) max-plus linear system $x(k) = \bigoplus_{l=0}^{p} A_l \otimes x(k-l)$ can be very large and sparsely populated, Goverde (2010) proposes a sparse representation based on *marked precedence graphs* to efficiently visualise these matrices. A marked precedence graph generalises a precedence graph of a single matrix $A \in \mathbb{R}_{\varepsilon}^{n \times n}$ to a *marked precedence graph* of the full set of matrices $A_0, \dots, A_p$. In the marked precedence graph, dependency on an event in a previous period is shown with a *marking* $\mu_{ij} = l \in \mathbb{N}_0$ on an arc. If the dependency is on an event within the current period, $\mu_{ij} = 0$. If the dependency is on an event from l periods back, $\mu_{ij} = l$. In a graph, the marking will be represented with l dots on the arc. Note that while there are similarities between timed event graphs and marked precedence graphs, they are different visual representations of discrete event systems. An example of a marked precedence graph is given in figure 2.15. A sparse representation of the same system is given in table 2.1.

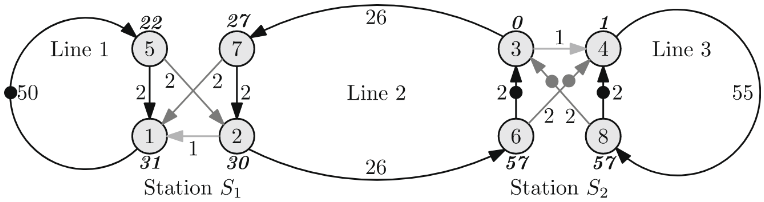

Figure 2.15: A marked precedence graph of the network in figure 2.13 (Goverde, 2010). Note the markings in line 1 and station S_2.

Table 2.1: The sparse notation of the marked precedence graph shown in figure 2.15.

Arc	1	2	3	4	5	6	7	8	9	10	11	12	13	14
j	1	3	3	4	5	5	7	7	2	2	6	6	8	8
i	5	4	7	8	1	2	1	2	1	6	3	4	3	4
μ_{ij}	1	0	0	0	0	0	0	0	0	0	1	1	1	1
a_{ij}	50	1	26	55	2	2	2	2	1	26	2	2	2	2

The asymptotic growth rate of $\{x(k)\}$ is defined as $\lim_{k\to\infty} \frac{x_k(k)}{k}$ for $1 \leq j \leq n$. If these limits exist and all have the same value (which is the case an irreducible system), this is known as the Lyapunov exponent of the sequence $\{A(k)\}$. The Lyapunov exponent is an important stability indicator. In stochastic max-plus linear systems, the Lyapunov exponent can be difficult to calculate. Goverde et al. (2011) proposes a coupling approach to estimate the Lyapunov exponent of stochastic max-plus linear systems. The result is an interval in which the Lyapunov exponent lies with a predetermined degree of confidence.

2.7. Petri Nets and Max-Plus Algebra in Aviation Network Delay Propagation Modelling

Petri nets and max-plus algebra have also been used in aviation networks to aid in modelling delay. This section aims to give an overview of various applications, focusing first on Petri nets and shifting then to max-plus algebra methods.

Ding et al. (2009) used Petri nets to model delay between arriving aircraft. They split aircraft arriving at the same runway into two groups: priority and non-priority. Aircraft in the priority group can have different levels of priority. They then set up a Petri net model for each group, which models the relationship between two consecutive flights based on different possible states each flight can be in. These states are related to the priority of each flight over the other or to the delay time of each flight in comparison to each flight's scheduled time properties. The model is then expanded to include other flights in a flight chain, and used to draw some general conclusions on flight delay propagation metrics. However, Ding et al. (2009) do not model a full network using Petri nets.

Zhang and Zhu (2015) propose Time Interval Petri Nets, which are then used to calculate delay propagation in a schedule. They use the max-plus operators to linearise the problem, but do not go so far as to present a matrix notation. To ensure reasonable complexity, they do not calculate the entire network after a delay, but generate a sub-network, which their algorithm then simplifies. This is possible, as they only take into account interdependency between aircraft, not between crew, passengers or cargo.

Shao and Xu (2020) proposed a coloured-timed Petri net approach for a cyclical multi-airport network. In their work, they included uncertain factors such as weather and ATC in the Petri net. This was achieved by extracting probability functions of airport or airspace capacity based on weather or ATC conditions from historical data. Afterwards, this probability was translated into expected capacity by feeding in real conditions of a specific day. Using this, the Petri net simulation could then estimate delay propagation for each flight. Based on the algorithm, an improved schedule was also proposed which reduced delay by over 50%. Shao and Xu (2020) again only included interdependency between aircraft.

Petri nets have also been used to develop support tools for air traffic controllers. For example, Kammoun et al. (2012) used Petri nets to reschedule flight plans in case certain sectors of airspace are capacity-limited due to adverse conditions. They did so by combining Petri nets with binary decision diagrams. Huang and Chung (2011) modelled the full air traffic control procedure in a Petri net. While delay was included to some extent, delay propagation was not modelled here.

On max-plus algebra, the application is relatively limited. P. Wang et al. (2003) was one of the first to propose utilising a max function in delay propagation estimation, albeit without max-plus notation. Additionally, they argued that the buffer time is ideally greater than the maximum typical deviation from the default flight and turnaround times. However, this is not always the case, and the deviation is always stochastic.

Abdelghany et al. (2004) presented a delay propagation model based on a directed acyclic graph, which can also be partially useful in max-plus algebra. However, they used the classical shortest path algorithm to solve the model, and concluded that aircraft interdependencies are the predominant cause of propagated delay.

Giannikas et al. (2022) and Kafle and Zou (2016) both recognised the need of the $\max$ operation, but did not linearise their equations using the $\otimes$ operator. Han (2023) uses max-plus algebra to set up a limited queueing model for an airport approach network. They then state they use this to draw conclusions about the stability and robustness of the network. However, the results are unclear and they do not yet expand this into a multi-airport network.

One of the most notable works in this category is that of Rengers (2017), who set up a scheduled max-plus linear system for aviation networks analogous to that of Goverde (2007). He modelled a small aviation network consisting of two hubs, each with two spokes, for a total of six airports. The network can be seen in figure 2.16. He then calculated the eigenvalue and established a realisable timetable. However, he only modelled a first-order system, while in reality an aviation network is often of a higher order. Additionally, he did not extract any stability parameters of actual timetables. Finally, he did not set up a delay propagation algorithm.

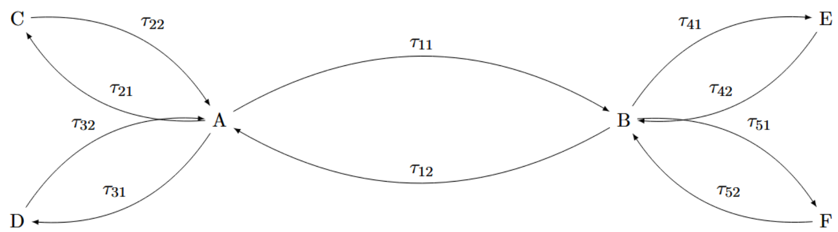

Figure 2.16: The network used in the work of Rengers (2017). (Rengers, 2017)

2.8. Gap in the Research

The past sections have described the current state-of-the-art of delay propagation modelling in aviation networks, particularly using Petri nets and max-plus algebra. The research showed that a wide array of modelling techniques is used to model delay propagation in aviation networks.

Economic, statistical, epidemic spreading, complex network theory will often generate reasonable results on a macro level, but will not provide a high level of granularity. On the other hand, agent-based models, graph neural networks and machine learning techniques can provide precise micro-level insights into complex aviation networks, they can be computationally very intensive. Queueing models are a powerful tool to identify operational capacity restrictions, but tend to only solve a subset of the greater airline network. Bayesian network models are varied in their use, and can be granular for a subset of the network or at lower levels of detail for the entire network. In summary, there does not yet seem to be an efficient modelling technique to model delay propagation at a high level of granularity for an entire network.

A potential solution for this problem could be the modelling of an aviation network as a scheduled max-plus linear system with a corresponding marked precedence graph. Applications of max-plus algebra to rail networks have shown to be highly effective at extracting interesting stability parameters and modelling delay propagation over multiple periods of a timetable. Additionally, C. Li et al. (2024) specifically mention that future research should explore methods to consider delay propagation between different aircraft, including their crew, passenger or cargo interdependencies. Max-plus algebra is particularly suited for modelling these interdependencies.

Applying max-plus algebra to aviation networks could thus be an interesting contribution to the existing work. To the author's knowledge, Rengers (2017) performed the most extensive application of max-plus algebra to an aviation network to date. However, this work only considered a limited network, a first-order max-plus system and he did not calculate stability parameters or delay propagation. As such, there is still a lot of potential to improve on and expand the modelling of delay propagation in an aviation network as a scheduled max-plus linear system. A research proposal based on this gap in the research will be presented in chapter 3.

3 Research Proposal

As introduced in chapter 1, the on-time performance of flights remains low, with only 71% of flights considered to be on time in the ECAC member states. It was shown that delay propagation was a major cause of this delay. Based on the literature review presented in chapter 2 and specifically the gap in the research discussed in section 2.8, this chapter will propose a research project to improve on the estimation of delay propagation. The research objective is to improve the estimation of delay propagation and related schedule robustness by modelling aviation networks as timed event graphs and analysing them using max-plus algebra.

This research proposal is structured as follows: First, section 3.1 will present the research question and its subquestions. Afterwards, section 3.2 will outline the research methodology and the planning. Finally, section 3.3 will discuss the data and tools which will be used in the project.

3.1. Research Questions

To guide the research, the research objective has been translated into a general research question and five subquestions. These are:

How can scheduled max-plus linear systems be leveraged to model, analyse, and predict delay propagation in aviation networks?

1. What data is available and how must it be processed to prepare it to model delay propagation in aviation networks?
2. How can scheduled max-plus linear systems be adapted and applied to model delay propagation in aviation networks?
3. How can the interdependencies between aircraft, crew, passengers, and cargo be integrated into a scheduled max-plus linear system, and what different roles do they play in the propagation of delays according to this system?
4. Can stability parameters commonly used in railway delay modelling (e.g., cycle time, stability margin, and recovery matrix) be effectively calculated for aviation networks using max-plus algebra, and what insights do they provide about network robustness?
5. How can an efficient delay propagation algorithm based on scheduled max-plus plus linear systems be developed for aviation networks?
6. How scalable is a max-plus algebra model for large-scale aviation networks, and how well does it capture the real-world complexities of delay propagation compared to smaller test networks?

3.2. Research Methodology

This section will discuss the research methodology to answer the research questions described in section 3.1. First, the general approach will be outlined. Afterwards, this approach will be translated into discrete work packages and a schedule outlined for all work packages.

The research proposed in this paper consists of two phases:

1. **Test Case:** First, the research will focus on a simple, theoretical test case. For this case, the fundamental scheduled max-plus linear system of equations for an aviation network will be established. Afterwards, these will be implemented in a computer program. Once the implementation is complete, it makes sense to extract relevant stability parameters and set up a delay propagation algorithm. The simple test case should make it easy to test the working of the computer program implementation. However, the program should be designed with a larger case in mind.
2. **Real-world case:** Once the program is proven to work for the test case, it makes sense to then test a real-world case. For the data used, one is referred to section 3.3. It seems likely, however, the data will first have to be cleaned, corrected and completed before it can be implemented in the model. Once that is complete, relevant stability parameters can be extracted and the delay propagation algorithm can be executed. The estimated propagated delay can then be compared to the actual delay experienced by an aircraft.

The two phases of the project have each been split into four work packages (WP 2-5 and WP 6-9). These can be found in table 3.1. The planned start and end dates for each work package can also be found there. The table also includes three additional work packages, which concern data availability research (WP 1) and reporting work (WP A, B and C). The timeline has been fully presented in a Gantt chart, which is included in appendix A. There are also some dates which are not part of the research work itself, but are relevant for this project as it concerns a Master thesis. These are shown in table 3.2.

Table 3.1: The work packages which together form the research.

WP	Content	Start	End
WP 1	Research what data is available and how it must be processed to model delay propagation in aviation networks	13/01	31/01
WP 2	Set up precedence graphs and corresponding scheduled max-plus linear systems of equations for different types of interdependencies	13/01	24/01
WP 3	Implement the scheduled max-plus linear systems into a computer model	27/01	14/02
WP 4	Compute relevant stability parameters from the scheduled max-plus linear systems for a test case	17/02	28/02
WP 5	Develop a delay propagation algorithm for a test case	03/03	21/03
WP A	Write the scientific paper introduction, methodology and the preliminary results of the test case	24/03	28/03
WP 6	Collect, clean and complete data for a real-world case	31/03	11/04
WP 7	Implement the real-world case into the computer modelþ	14/04	25/04
WP 8	Extract relevant stability parameters for a real-world case	28/04	16/05
WP 9	Calculate the delay propagation for a real-world case	28/04	16/05
WP B	Write and visualise the results of the real-world case, the discussion, the conclusions and any other sections of the thesis.	04/05	06/06
WP C	Further improve the thesis and incorporate feedback from the greenlight review.	09/06	11/07

Table 3.2: Relevant dates for the thesis process

Date	Event
03/04	Midterm review
05/06	Deadline first thesis draft
19/06	Greenlight review
11/07	Expected hand in date thesis
31/07	Expected defence date

There is a risk that certain work packages take longer than foreseen. Particularly understanding how to implement max-plus algebra in Matlab and setting up the real-world case could be unexpectedly challenging tasks. If this proves to be the case, the research can be scaled down: For instance, the choice could be made to only focus on the delay propagation algorithm or only focus on the relevant stability criteria. There is also a risk that the computational complexity of the real-world case is too large, making the model unsolvable in a reasonable timeframe. In this case, it could make sense to look at a smaller subproblem of the real-world case. Future work could then improve the computational efficiency.

3.3. Data and Tools

In modern research, data and modelling tools are a significant part of the work. This is also the case for the research proposed here. Therefore, both the data and modelling tools to be used will be discussed in this section.

C. Li et al. (2024) showed that the data used in delay propagation modelling research is mainly supplied by the US Bureau of Transportation Statistics (BTS), followed by data from airlines' aviation operations departments and data from aviation authorities. They argue this is mainly due to the BTS data's public accessibility. Similarly, they show most research is focused on the USA or China. As Carvalho et al. (2021) specifically mentioned only 6.5% of work could be classified as easily reproducible, reproducibility should be a focus of this report. As such, it makes sense to focus the research on BTS data. Specifically, the BTS provides an extensive on-time performance database[1], which can be used to set up a timetable schedule for an airline. It also includes the scheduled and actual arrival and departure times of flights.

A difficulty with the BTS data is that it does not contain information on aircraft interdependencies. This information is typically not publicly available, as crew, passenger or cargo connections are not shared by airlines. While this could be requested from an airline, C.-L. Wu and Law (2019) have also an approach to reconstruct these interdependencies. For sake of reproducibility, this approach will also be applied in this research.

For the same reason of reproducibility, it makes sense to use a widely adapted programming method. Python quickly comes to mind, as it is free and open source. However, there are no available Python packages which enable complete max-plus algebra support. These have been developed already for Matlab. Matlab is not free and open source, but is widely used in the scientific community. As such, Matlab is selected as programming and modelling tool. However, it could make sense to write the data preprocessing in Python. In that way, any external user can take the BTS dataset and perform the same data preprocessing without the need for a Matlab license.

[1] https://transtats.bts.gov/ONTIME/

4 Conclusion

The research proposed in this report aims to improve the estimation of delay propagation and related schedule robustness by modelling aviation networks as timed event graphs and analysing them using max-plus algebra. Delay propagation remains a significant contributor to the total amount of delay minutes in aviation. A better modelling of it could aid in building better schedules, which absorb disruptions more easily.

In this paper, first a literature review was conducted to examine the state of the art of delay propagation modelling in aviation networks. The literature review also introduced max-plus algebra as an alternative modelling technique and showcased its efficacy in railway network delay propagation modelling. Afterwards, current applications of max-plus algebra in aviation network delay propagation modelling were discussed, and the gap in the research was presented.

With the gap in the research clear, a research proposal was presented. This included the research questions, methodologies and the data and tools to be used. The max-plus algebra model will be implemented in Matlab, and tested using data from the US Bureau of Transportation Statistics.

References

Abdelghany, K. F., S. Shah, S., Raina, S., & Abdelghany, A. F. (2004). A model for projecting flight delays during irregular operation conditions. *Journal of Air Transport Management*, *10*(6), 385–394. https://doi.org/10.1016/j.jairtraman.2004.06.008

AhmadBeygi, S., Cohn, A., Guan, Y., & Belobaba, P. (2008). Analysis of the potential for delay propagation in passenger airline networks. *Journal of Air Transport Management*, *14*(5), 221–236. https://doi.org/10.1016/j.jairtraman.2008.04.010

AhmadBeygi, S., Cohn, A., & Lapp, M. (2010). Decreasing airline delay propagation by re-allocating scheduled slack. *IIE Transactions*, *42*(7), 478–489. https://doi.org/10.1080/07408170903468605

AirHelp. (2023, October). *Cost of disrupted flights to the economy*. Retrieved January 15, 2025, from https://www.airhelp.com/en-int/blog/in-numbers-the-economic-impact-of-flight-disruptions/

airliners.de. (2017, May 17). *Ryanair hat jetzt auch Umsteigeflüge im Programm* [airliners.de]. Retrieved January 4, 2025, from https://www.airliners.de/ryanair-umsteigefluege/41146

Baden, W., DeArmon, J., Kee, J., & Smith, L. (2006). Assessing schedule delay propagation in the national airspace system. *47th Annual Transportation Research Forum*. https://doi.org/10.22004/ag.econ.208044

Ball, M., Barnhart, C., Dresner, M., Hansen, M., Neels, K., Odoni, A., Peterson, E., Sherry, L., Trani, A., Zou, B., Britto, R., Fearing, D., Swaroop, P., Uman, N., Vaze, V., & Voltes, A. (2010, October 1). *Total delay impact study: A comprehensive assessment of the costs and impacts of flight delay in the united states* (Technical Report No. 55). NEXTOR. Washington, DC. https://rosap.ntl.bts.gov/view/dot/6234

Beatty, R., Hsu, R., Berry, L., & Rome, J. (1999). Preliminary evaluation of flight delay propagation through an airline schedule. *Air Traffic Control Quarterly*, *7*(4), 259–270. https://doi.org/10.2514/atcq.7.4.259

Bureau of Transportation Statistics. (2024, December 22). *Airline on-time statistics and delay causes* [Dataset]. Retrieved December 22, 2024, from https://www.transtats.bts.gov/ot_delay/ot_delaycause1.asp?6B2r=FE&20=E

Carvalho, L., Sternberg, A., Maia Gonçalves, L., Beatriz Cruz, A., Soares, J. A., Brandão, D., Carvalho, D., & Ogasawara, E. (2021). On the relevance of data science for flight delay research: A systematic review. *Transport Reviews*, *41*(4), 499–528. https://doi.org/10.1080/01441647.2020.1861123

Cassandras, C. G. (2005). Discrete-event systems. In D. Hristu-Varsakelis & W. S. Levine (Eds.), *Handbook of networked and embedded control systems* (pp. 71–89). Birkhäuser. https://doi.org/10.1007/0-8176-4404-0_4

Cook, G., & Goodwin, J. (2008). Airline networks: A comparison of hub-and-spoke and point-to-point systems. *Journal of Aviation/Aerospace Education & Research*, *17*(2). https://doi.org/https://doi.org/10.15394/jaaer.2008.1443

De Schutter, B., & van den Boom, T. (2008). Max-plus algebra and max-plus linear discrete event systems: An introduction. *Proceedings of the 9th International Workshop on Discrete Event Systems (WODES'08)*, 36–42. https://doi.org/10.1109/WODES.2008.4605919

Ding, J., Chen, T., & Xu, T. (2009). Propagated analysis of airport delays based on timed petri nets. *2009 International Conference on Computational Intelligence and Security*, *1*, 608–614. https://doi.org/10.1109/CIS.2009.256

Eltoukhy, A., Chan, F., & Chung, S. (2017). Airline schedule planning: A review and future directions. *Industrial Management and Data Systems*, *117*(6), 1201–1243. https://doi.org/10.1108/IMDS-09-2016-0358

EUROCONTROL. (2024, December 19). *Aviation long-term outlook*. Retrieved January 4, 2025, from https://www.eurocontrol.int/publication/eurocontrol-forecast-2024-2050

Fleurquin, P., Ramasco, J. J., & Eguiluz, V. M. (2013). Systemic delay propagation in the US airport network. *Scientific Reports*, *3*(1), 1159. https://doi.org/10.1038/srep01159

Giannikas, V., Ledwoch, A., Stojković, G., Costas, P., Brintrup, A., Al-Ali, A., Chauhan, V., & McFarlane, D. (2022). A data-driven method to assess the causes and impact of delay propagation in air transportation systems. *Transportation Research Part C: Emerging Technologies*, *143*. https://doi.org/10.1016/j.trc.2022.103862

Gordijn, E. (2020, November). *Nieuw Normen- en Handhavingstelsel Schiphol*. Advanced Decision Systems Airinfra & To70 BV. Schiphol. Retrieved January 16, 2025, from https://www.rijksoverheid.nl/documenten/rapporten/2021/02/16/bijlage-3-mer-nnhs-2020-deel-1-hoofdrapport

Goverde, R. M. P. (2007). Railway timetable stability analysis using max-plus system theory. *Transportation Research Part B: Methodological*, *41*(2), 179–201. https://doi.org/10.1016/j.trb.2006.02.003

Goverde, R. M. P. (2010). A delay propagation algorithm for large-scale railway traffic networks. *Transportation Research Part C: Emerging Technologies*, *18*(3), 269–287. https://doi.org/10.1016/j.trc.2010.01.002

Goverde, R. M. P., Heidergott, B., & Merlet, G. (2011). A coupling approach to estimating the lyapunov exponent of stochastic max-plus linear systems. *European Journal of Operational Research*, *210*(2), 249–257. https://doi.org/10.1016/j.ejor.2010.09.035

Han, Y.-X. (2023). Stability analysis for scheduled air traffic flow. *Aircraft Engineering and Aerospace Technology*, *95*(8), 1201–1208. https://doi.org/10.1108/AEAT-03-2022-0071

Huang, Y.-S., & Chung, T.-H. (2011). Modelling and analysis of air traffic control systems using hierarchical timed coloured petri nets. *Transactions of the Institute of Measurement and Control*, *33*(1), 30–49. https://doi.org/10.1177/0142331208095623

Huynh, T. K., Cheung, T., & Chua, C. (2024). A systematic review of flight delay forecasting models. *2024 7th International Conference on Green Technology and Sustainable Development (GTSD)*, 533–540. https://doi.org/10.1109/GTSD62346.2024.10675123

International Air Transport Association. (2016, January). Codes to be used in aircraft movement and diversion messages. In *Airport handling manual* (36th ed., pp. 692–696). https://xwiki.avinor.no/download/attachments/4030721/36th%20AHM%20730%20CodesUsedinAircraftMovementDiversionMessages.pdf?api=v2

International Civil Aviation Organization. (2009). Review of the classification and definitions used for civil aviation activities. https://skybrary.aero/sites/default/files/bookshelf/4416.pdf

Kafle, N., & Zou, B. (2016). Modeling flight delay propagation: A new analytical-econometric approach. *Transportation Research Part B: Methodological*, *93*, 520–542. https://doi.org/10.1016/j.trb.2016.08.012

Kammoun, M. A., Sava, A., & Rezg, N. (2012). A flight plan rescheduling in air traffic management problem: A time discret event system approach. *45*, 285–290. https://doi.org/10.3182/20120523-3-RO-2023.00077

Komenda, J., Lahaye, S., Boimond, J.-L., & van den Boom, T. (2018). Max-plus algebra in the history of discrete event systems. *Annual Reviews in Control*, *45*, 240–249. https://doi.org/10.1016/j.arcontrol.2018.04.004

Kouris, S., Steer, & Directorate-General for Mobility and Transport (European Commission). (2020). *Study on the current level of protection of air passenger rights in the EU*. Publications Office of the European Union. https://data.europa.eu/doi/10.2832/529370

Li, C., Mao, J., Li, L., Wu, J., Zhang, L., Zhu, J., & Pan, Z. (2024). Flight delay propagation modeling: Data, methods, and future opportunities. *Transportation Research Part E: Logistics and Transportation Review*, *185*, 103525. https://doi.org/10.1016/j.tre.2024.103525

Li, Y., Cai, K., Zhu, Y., & Yang, Y. (2024). Modeling delay propagation in airport networks via causal biased random walk. *IEEE Transactions on Intelligent Transportation Systems*, *25*(5), 4692–4703. https://doi.org/10.1109/TITS.2023.3321398

Qin, J., Wu, H., Lin, Q., Shen, J., & Zhang, W. (2023). The recovering stability of a towing taxi-out system from a lateral instability with differential braking perspective: Modeling and simulation. *Electronics*, *12*, 2170. https://doi.org/10.3390/electronics12102170

Qin, S., Mou, J., Chen, S., & Lu, X. (2019). Modeling and optimizing the delay propagation in chinese aviation networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, *29*(8), 081101. https://doi.org/10.1063/1.5111995

Rengers, R. (2017, June). *Max-plus algebra and an application in aviation* [BSc Thesis]. Delft University of Technology. https://resolver.tudelft.nl/uuid:27c47f85-18e2-4789-9919-604fe1f7621e

Rodrigue, J.-P. (2024, April 30). *The geography of transport systems* (6th ed.). Routledge. https://doi.org/10.4324/9781003343196

Scheelhaase, J., Braun, M., Maertens, S., & Grimme, W. (2023). Costs for passengers and airlines due to the significant delays and other irregularities at european airports in the 2022 summer season. *Transportation Research Procedia*, *75*, 96–105. https://doi.org/10.1016/j.trpro.2023.12.012

Shao, Q., & Xu, C. (2020). Air transportation delay propagation analysis with uncertainty in coloured–timed petri nets. *Proceedings of the Institution of Civil Engineers - Transport*, *173*(6), 380–395. https://doi.org/10.1680/jtran.17.00159

Slotnick, D. (2021, March 25). *United airlines unveils 26 new point-to-point routes from midwest* [The points guy]. Retrieved January 4, 2025, from https://thepointsguy.com/news/united-new-routes-may-2021-midwest-east-coast-hawaii/

Value Group. (2023, October 13). *The Economic Impact of Journey Disruptions: Analyzing Airlines and Railways*. Retrieved January 15, 2025, from https://valueg.com/the-economic-impact-of-disruptions/

Walker, C. (2024, December 16). *CODA digest: All-causes delays to air transport in europe - annual 2023*. EUROCONTROL. Retrieved December 22, 2024, from https://www.eurocontrol.int/publication/all-causes-delays-air-transport-europe-annual-2023

Wang, P., Schaefer, L., & Wojcik, L. (2003). Flight connections and their impacts on delay propagation. *The 22nd Digital Avionics Systems Conference, 2003. DASC '03.*, *1*, 5.B.4–5.1–9. https://doi.org/10.1109/DASC.2003.1245858

Wang, T., Zheng, Y., & Xu, H. (2022). A review of flight delay prediction methods. *2nd International Conference on Big Data Engineering and Education (BDEE)*, 135–141. https://doi.org/10.1109/BDEE55929.2022.00029

Wu, C.-L. (2006). Improving airline network robustness and operational reliability by sequential optimisation algorithms. *Networks and Spatial Economics*, *6*(3), 235–251. https://doi.org/10.1007/s11067-006-9282-y

Wu, C.-L., & Law, K. (2019). Modelling the delay propagation effects of multiple resource connections in an airline network using a bayesian network model. *Transportation Research Part E: Logistics and Transportation Review*, *122*, 62–77. https://doi.org/10.1016/j.tre.2018.11.004

Wu, S., Liu, E., Cao, R., & Bai, Q. (2025). Airline recovery problem under disruptions: A review. *Computers & Operations Research*, *175*, 106915. https://doi.org/10.1016/j.cor.2024.106915

Wu, W., & Wu, C.-L. (2018). Enhanced delay propagation tree model with bayesian network for modelling flight delay propagation. *Transportation Planning and Technology*, *41*(3), 319–335. https://doi.org/10.1080/03081060.2018.1435453

Wu, Y., Yang, H., Lin, Y., & Liu, H. (2024). Spatiotemporal propagation learning for network-wide flight delay prediction. *IEEE Transactions on Knowledge and Data Engineering*, *36*(1), 386–400. https://doi.org/10.1109/TKDE.2023.3286690

Xu, N., Donohue, G., Laskey, K. B., & Chen, C.-H. (2005). Estimation of delay propagation in the national aviation system using bayesian networks. https://www.semanticscholar.org/paper/Estimation-of-Delay-Propagation-in-the-National-Xu-Donohue/9c5a7c726315387acb5ff5cf8e849524046cd207

Xu, Y., Wandelt, S., & Sun, X. (2021). Airline integrated robust scheduling with a variable neighborhood search based heuristic. *Transportation Research Part B: Methodological*, *149*, 181–203. https://doi.org/10.1016/j.trb.2021.05.005

Zhang, X., & Zhu, X. (2015). Modeling and delay propagation analysis for flight operation based on time interval petri net. *The Open Automation and Control Systems Journal*, *6*, 433–438. https://doi.org/10.2174/1874444301406010433

Zhou, L., Liang, Z., Chou, C.-A., & Chaovalitwongse, W. A. (2020). Airline planning and scheduling: Models and solution methodologies. *Frontiers of Engineering Management*, *7*(1), 1–26. https://doi.org/10.1007/s42524-020-0093-5

A

Gantt Chart Research Proposal

On the following page, a Gantt chart of the research plan can be found.

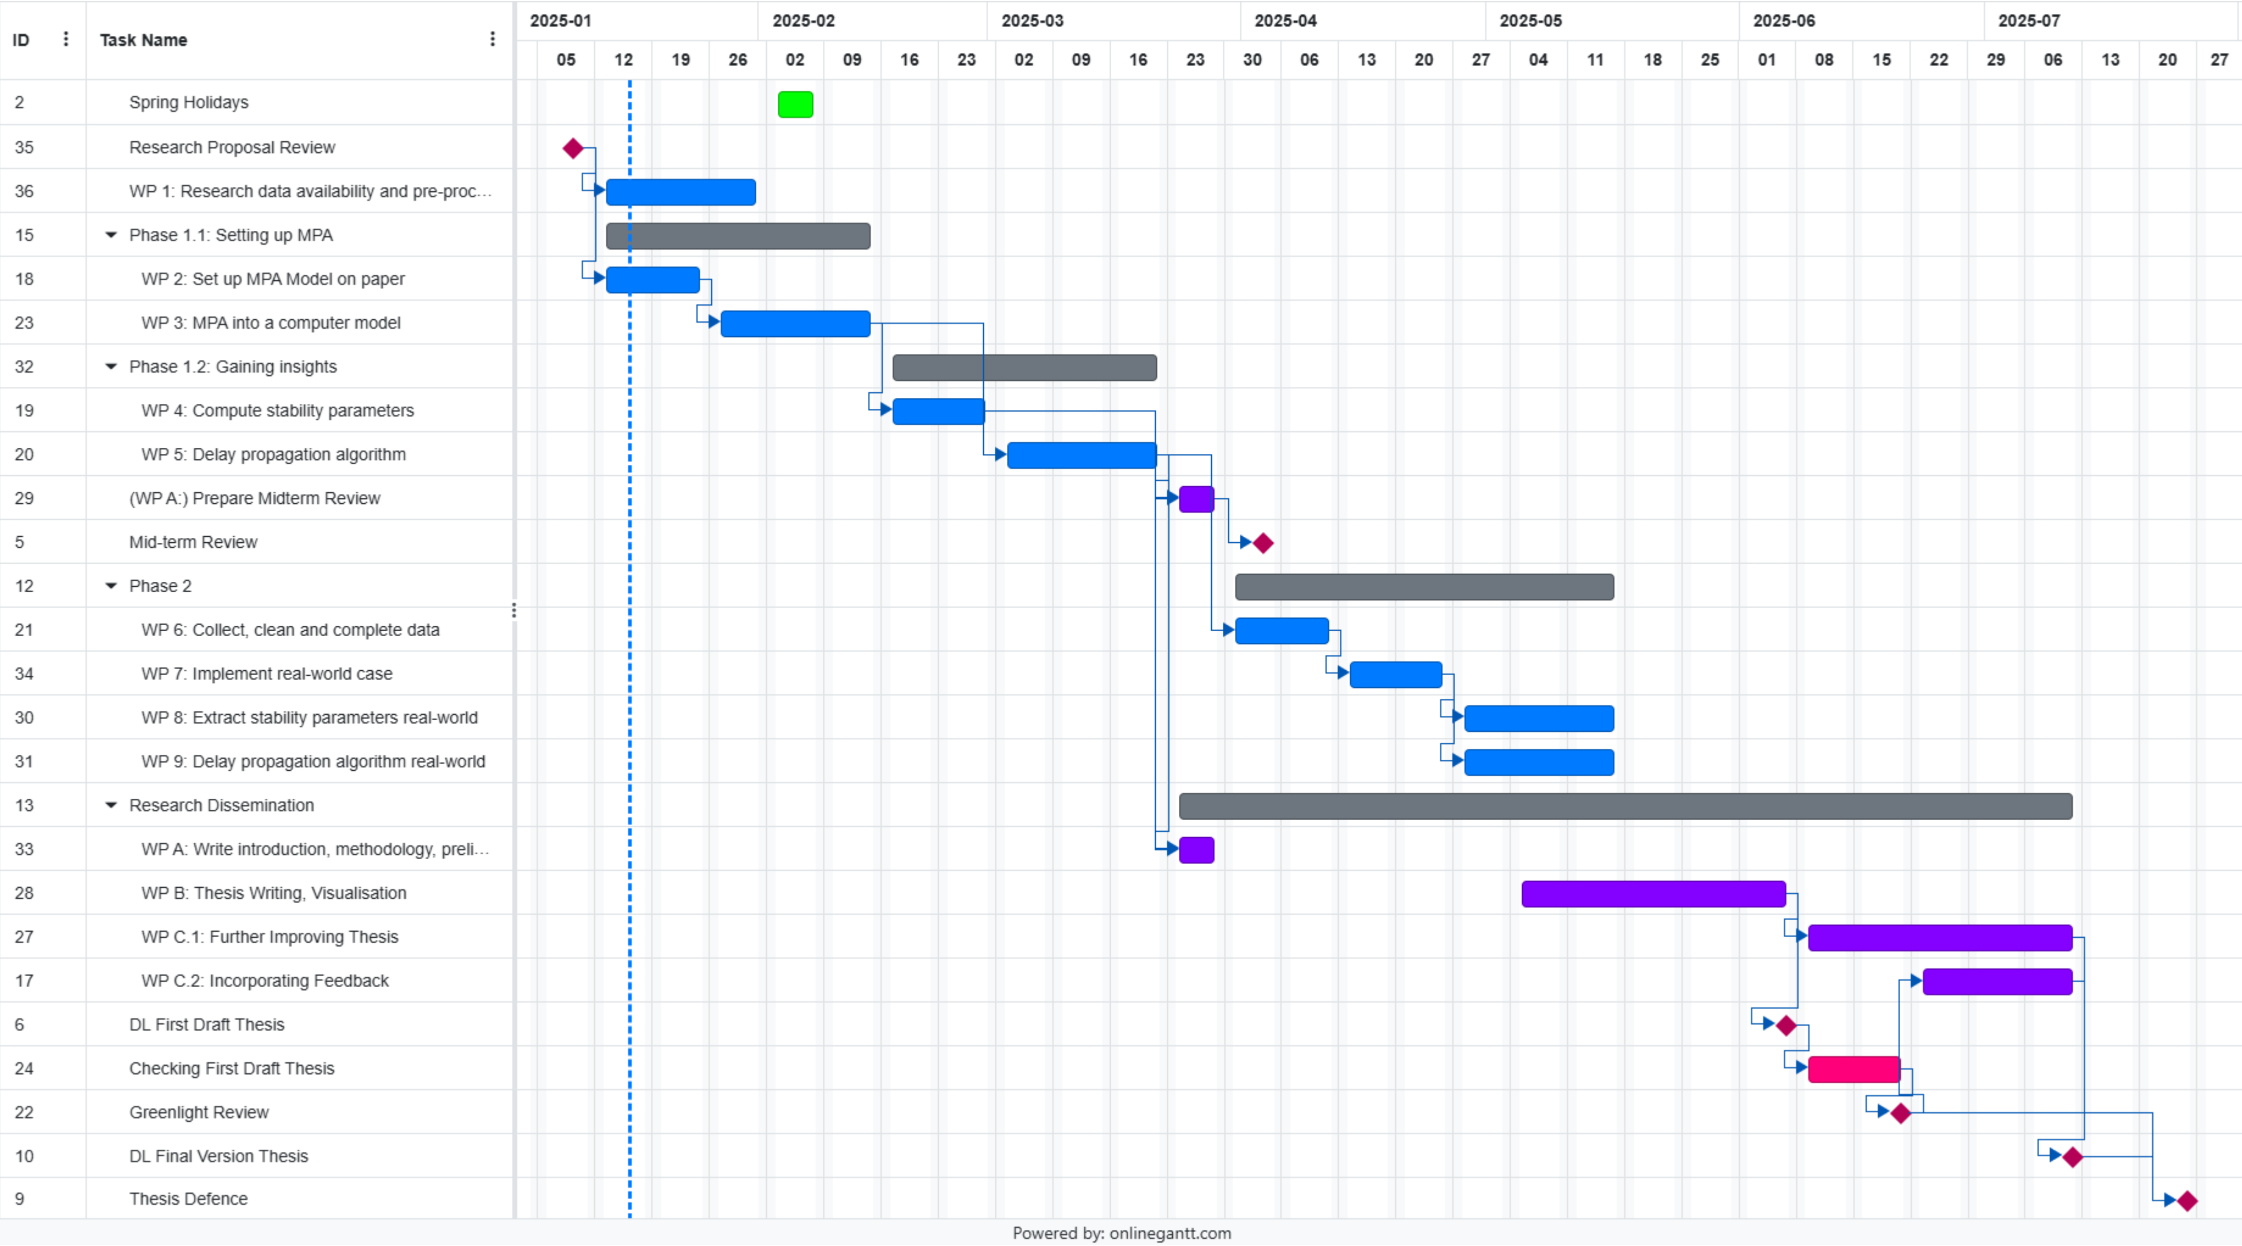
ID
Task Name
2025-01
2025-02
2025-03
2025-04
2025-05
2025-06
2025-07
05 12 19 26 02 09 16 23 02 09 16 23 30 06 13 20 27 04 11 18 25 01 08 15 22 29 06 13 20 27
2 Spring Holidays
35 Research Proposal Review
36 WP 1: Research data availability and pre-proc...
15 Phase 1.1: Setting up MPA
18 WP 2: Set up MPA Model on paper
23 WP 3: MPA into a computer model
32 Phase 1.2: Gaining insights
19 WP 4: Compute stability parameters
20 WP 5: Delay propagation algorithm
29 (WP A:) Prepare Midterm Review
5 Mid-term Review
12 Phase 2
21 WP 6: Collect, clean and complete data
34 WP 7: Implement real-world case
30 WP 8: Extract stability parameters real-world
31 WP 9: Delay propagation algorithm real-world
13 Research Dissemination
33 WP A: Write introduction, methodology, preli...
28 WP B: Thesis Writing, Visualisation
27 WP C.1: Further Improving Thesis
17 WP C.2: Incorporating Feedback
6 DL First Draft Thesis
24 Checking First Draft Thesis
22 Greenlight Review
10 DL Final Version Thesis
9 Thesis Defence
Powered by: onlinegantt.com

B

Notes on the use of generative AI

Generative AI is a powerful tool for evaluating and creating code and text. As per the TU Delft guidelines for generative AI chatbots in unsupervised assessment[1], AI chatbots can be used for unsupervised assessment. However, it is best practice to outline how they were used to avoid confusion over created content, fraud and achievement of learning objectives. This will be done in this section.

To aid in understanding literature for this report and in drafting research questions, ChatGPT was used. Some example queries will be included in this chapter and shortly discussed.

The primary use of ChatGPT was to aid in understanding complex scientific literature. For this, it proved an extremely valuable tool. Being able to ask questions related to specific parts of articles allowed for a deeper understanding. For instance, the author initially struggled with certain mathematical formulations, but an additional explanation from ChatGPT, often with an example, helped a lot. It remained important to be prudent, however, as ChatGPT could sometimes also give inaccurate or even blatantly incorrect responses. In a clear example of this, when asked ”Could you show and explain a timed petri net applied to a rail timetable network?”, ChatGPT gave a correct Petri net formulation in words, but generated the obviously incorrect figure B.1 as an illustration.

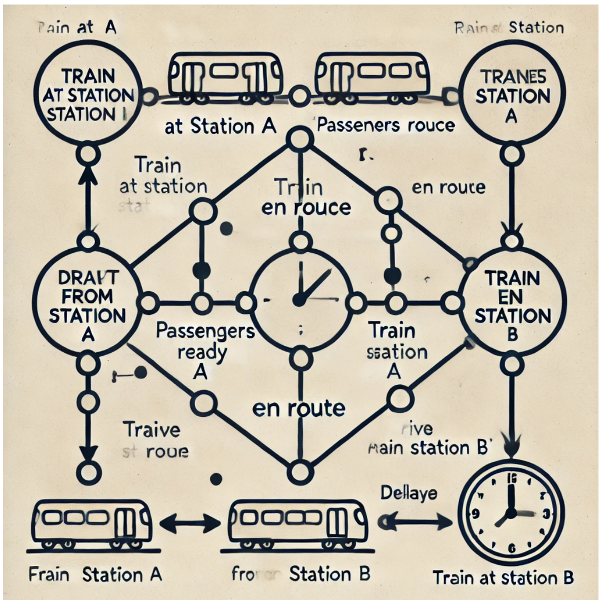

Figure B.1: An image generated by ChatGPT, supposed to represent a Petri net.

[1] https://www.tudelft.nl/teaching-support/didactics/assess/guidelines/ai-chatbots-in-unsupervised-assessment

Another limitation of ChatGPT became apparent when asking questions about specific sections of reports in a more conversational style. This is best demonstrated through the transcript of a conversation with ChatGPT on the work of Goverde (2007). The complete transcript of queries is given below:

1. Could you explain section 2 of this paper? (with the work of Goverde (2007) as attachment)
2. Could you explain the properties (Idempotency, Commutativity and associativity of $\oplus$, Distributivity of $\otimes$ over $\oplus$.) as well as the other technical terms in this section (such as semiring, dioid, etc)
3. Could you further explain the matrix operations? Please provide a small example as well.
4. Could you explain how matrix power works in max-plus algebra?
5. Related to this, what is a precedence graph? Is the critical path matrix a precedence graph G?
6. Could you explain the period delay described in section 3.1 of the paper?
7. What are the main take-aways from section 3.2?
8. Why is the gamma-representation presented in section 3.3 relevant for the paper?
9. How would one then understand equation 8? Let's say gamma is equal to 2. How would you then approach calculating this?
10. The last two paragraphs of this section go into a lot of mathematical detail. Could you explain this?
11. Could you explain the last paragraph of section 3.4 in detail? Sentence by sentence?
12. This is not the last paragraph of section 3.4. The last paragraph of section 3.4 starts with "As previously defined a homogeneous higher-order..." and ends with "...corresponds to a timed event graph that has no incoming paths or trees." Please explain this section of the text sentence by sentence.
13. Please explain the following section sentence by sentence: <quote from Goverde (2007)>
14. I would like to ask you some questions about chapter 4 of the paper. Could you first summarise the general aim and contents of this section?
15. You started summarising section 5, on timetable stability. I would like you to summarise the section on Max-Plus Spectral Analysis.
16. Could you explain the following statement? <quote from Goverde (2007)>
17. Could you explain what the acyclic subgraph A_0 is in more detail? The total graph A(c) is cyclic right? Why is this subgraph always acyclic?
18. Could you explain the following section? Please include an example on the calculation <quote from Goverde (2007)>
19. Could you explain Lemma 1? Please include why it is relevant as well.
20. This is not lemma 1. Lemma 1 discusses the subeigenvector. A rough copy from the text is: <quote from Goverde (2007)>
21. I still don't fully understand the subeigenvectors (or supereigenvalues). Could you explain more about what these are?

Although at the beginning, ChatGPT was able to recognise what section 2 was, it later lost track when asked about specific sections ("This is not the last paragraph of section 3.4."; "I would like to ask you some questions about chapter 4 of the paper. [...] You started summarising section 5..."). An interesting note is that the author was forced to switch (due to the limited use of GPT-4o on free models) from GPT-4o to GPT-4o mini at question 10. This could explain the drop in quality at this point.

One final point where generative AI proved to be very useful was in the drafting of the research questions. When given the literature review and asked "Based on this literature review, what research questions would you suggest?", ChatGPT was able to come up with 10 interesting research questions, 5 of were used as a basis or directly copied to be used in this research.

Generally, the use of generative AI tools in the preparation of this work demonstrated that they can be useful tools to aid in understanding and distilling insights from existing text. They can also be helpful to brainstorm. However, any results they produce should never be taken for the truth and always be validated with primary sources.