Transportation Research Part E 161 (2022) 102687

Contents lists available at ScienceDirect

Transportation Research Part E

journal homepage: www.elsevier.com/locate/tre

Timescales of delay propagation in airport networks

Yanjun Wang [a,c,*], Max Z. Li [b], Karthik Gopalakrishnan [c], Tongdan Liu [d,a]

[a] College of Civil Aviation, Nanjing University of Aeronautics and Astronautics, Nanjing 210016, China
[b] Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI, USA
[c] Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA, USA
[d] SIPPR Engineering Group Co., Ltd, Zhengzhou 450007, China

ARTICLE INFO

Keywords:
Delay propagation
Air transportation
Network science
Statistical analysis

ABSTRACT

Flight delays persist and spread in airport networks due to high interconnectivity in the air transportation infrastructure. How quickly delay propagates between two airports is determined by factors such as the number of flights between airports, the duration of the flight, presence of disruptions, and schedule buffers. Accurate estimation of the time for delay propagation can improve system predictability and reliability. However, noisy airport delay data, along with a lack of visibility into airline scheduling and disruption management strategies, result in a challenging estimation problem for such propagation timescales. We present an algorithm to estimate statistically significant time lags between airport delays from noisy, aggregate operational data. The algorithm uses sliding correlation windows to extract the airport pairs with stable delay lags. We apply our method to identify different timescales of interactions for US airport delays in 2017. Our analysis yields two main results: (1) The most stable lags between airport delays involve the Northeast airports; (2) The stable lags between two airports are negatively correlated to the scheduled flight times between the same two airports. These results regarding delay propagation speeds have potential implications for delay prediction models and airline schedule design.

1. Introduction

A flight is considered to be delayed if its actual takeoff or landing time is later than its scheduled time. Flight delays result in significant costs to the airlines, lost time for passengers, and additional environmental emissions. In the US, nearly 2 million, or 20% of all arriving and departing commercial flights were considered delayed in 2018. While some level of flight delays is indicative of the system operating efficiently, i.e., close to its capacity limits, excess delay remains problematic. In the US, the largest cause of flight delays is due to late-arriving aircraft (see Fig. 1) since the same aircraft typically operates more than one flight per day. Thus, when an aircraft operating one flight is delayed, this aircraft may continue to be delayed for all subsequent legs it operates. This indicates that during times when there is a demand-capacity imbalance, there is the potential for a significant amount of delay propagation within the air transportation system.

The time it takes for delays from one airport to manifest their effects at another airport is what we refer to as the *timescale* of delay propagation. For example, the timescale of delays at two airports connected by a direct flight may be equal to that direct flight's scheduled flight time. Other factors may result in faster or slower timescales for delays to be coupled across different airports. As an example, consider a day with poor weather across a wide geographical area, such as a snowstorm in the northeast US. This snowstorm may result in multiple airports in major US East Coast cities (e.g., New York City, Boston, Philadelphia) simultaneously experiencing

* Corresponding author.
E-mail addresses: ywang@nuaa.edu.cn (Y. Wang), maxzli@umich.edu (M.Z. Li), karthikg@mit.edu (K. Gopalakrishnan), ltd@nuaa.edu.cn (T. Liu).

https://doi.org/10.1016/j.tre.2022.102687
Received 23 May 2021; Received in revised form 9 March 2022; Accepted 23 March 2022
Available online 13 April 2022

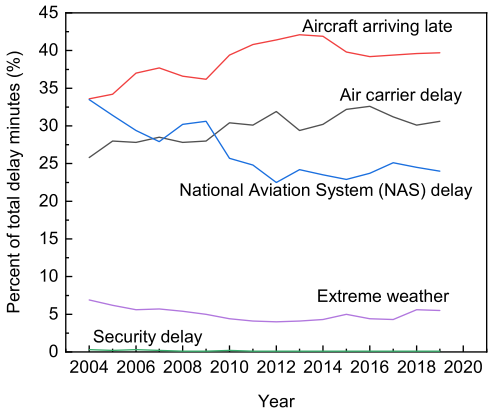

Fig. 1. Cause of flight delays in the US over the 2004–2020 time period (Bureau of Transportation Statistics, 2021).

flight delays. Furthermore, due to resultant traffic management initiatives such as ground stops or airspace flow programs, these delays may begin and end at approximately the same time. In this scenario, it would appear as if the delays propagated instantly and the delay propagation timescale was zero. On the other hand, the timescale may also be slower if there are no direct flights between two airports: Consider two airports with no direct flights between them, but there is a flight that connects the two airports via a third airport. Thus, delays would take a longer time to propagate between the original two airports, relative to the flight time if there had been a direct flight.

The above examples indicate that several mechanisms may result in the observed delays at different airports being coupled. In addition, each coupling could have a variable timescale as well. However, in reality, not all airports influence each other to the same extent, and these timescales of delay propagation may not be statistically significant due to the high degree of variability in the system. Therefore, we are only interested in identifying airport pairs whose timescale is both *statistically significant* (in terms of the coupling between two airports) as well as *stable* (in terms of the timescale magnitudes being consistent). We note that we use *stable* and *consistent* interchangeably in this paper.

Estimating and predicting the timescale of delay propagation from historical data is useful for delay prediction models and schedule design for airlines. When we know these delay propagation timescales, we can anticipate the next airports that will be delay-impacted, as well as when such impacts will occur, as soon as there is a disruption at one airport. Airlines can use these estimated timescales to calibrate schedule buffers dynamically to obtain a more robust schedule.

We now provide some motivations regarding the technical gap that our work addresses when it comes to estimating delay propagation timescales. We note that analytically extracting such timescales is theoretically possible: This requires several high-fidelity, potentially proprietary data sources such as the flight schedules, airline schedule buffers, statistics on external (i.e., weather-induced) as well as internal (e.g., maintenance-related) disruptions, and most importantly, a model for airlines' response to these disruptions. However, airlines generally do not share sensitive operational information, nor do we have perfect information regarding the impact of disruptions in an air transportation network. Hence, our work seeks to rigorously estimate the timescales of delay propagation from historical data.

We encounter several challenges in estimating the timescales of delay propagation from historical data. First, airport delay data reflect complex and nonlinear interactions in the system, which results in time-varying correlations. Second, analysis of the data can be challenging: In particular, the occurrence of delays can be considered as a poorly-characterized stochastic process, and the characteristics of delay propagation also depend on the time-of-day. This makes it challenging to search for statistically significant and stable temporal patterns (e.g., typical timescales for delays to propagate) using real-world data. There is a need to balance between requiring a high degree of statistical significance and stability while accommodating for temporal variability without capturing spurious trends. Finally, the data from a network of airports can be high-dimensional, leading to challenges in computational tractability.

Given the challenges in dissecting and analyzing flight delays across a network of airports, we first scope our goals to the following: Given an aggregate data set of airport delay measurements, how do we extract *statistically significant* estimates of the time taken for delay signals to propagate between two airports? In other words, how do we use airport delay data to identify the typical timescales for delay propagation between airports? To address this question, we first review the significant body of previous models and perspectives on flight delay propagation for a majority of the remainder of this section (Section 1.1), then conclude this section by explicating our contributions in the context of these previous work (Section 2). Next, we describe our technical approach in Section 3, where we formalize our aforementioned desired qualities of statistical significance and measurement stability. We then detail in Section 4 the US airport delay data used in our case study, from which we summarize and discuss the results in Sections 5 and 6, respectively. We conclude with some directions for future work in Section 7.

1.1. Prior work

Modeling and prediction of flight delay propagation have been studied for several decades. Popular techniques to study how delays propagate including queuing models (Peterson et al., 1995a; Pyrgiotis et al., 2013), Monte Carlo simulations and data-driven learning (Peterson et al., 1995b; Khan et al., 2021; Wang et al., 2021), intra-airline delay propagation (Beatty et al., 1999; Kafle and Zou, 2016; Wu and Law, 2019) and causal inference (Du et al., 2018). While these methods have significantly improved our ability to predict delays, especially for an airline who has access to many of the proprietary operational parameters that are required by the models, they are less helpful in identifying statistically significant patterns of delay propagation from historical data. In other words, these models, while calibrated to improve prediction accuracy, do not answer the higher-level question that motivates our work. Specifically, if an airport is currently delayed, when should we expect to see its effects on other airports throughout the system, and can we do so without proprietary, extensive airline-specific data. This is primarily because prior works tend to use historical data to calibrate and fine-tune data-driven models for delay propagation, instead of directly seeking to answer this question without making any assumptions or assuming a model for the propagation of delays.

An alternative approach is to purely view this as a spatio-temporal data analysis problem, where we seek to identify statistically significant delay propagation patterns. The multi-dimensional, large-scale, correlated data across airports seem well suited for network analysis tools. In fact, several researchers have used network models for studying delay propagation and dynamics the air transportation system (Guimerà et al., 2005; Li and Jing, 2021; Bao et al., 2021; Guo et al., 2022; Kim and Park, 2021). Most prior works focused on analyzing the structure and connectivity patterns of these networks using historical data. More recently, researchers have also incorporated these static measures of network connectivity to develop dynamic models for delay evolution (Fleurquin et al., 2013; Baspinar et al., 2021; Guo et al., 2022). However, the network analysis tools that have been employed to study airport networks primarily focus on flight connectivity patterns, and not the dynamics and evolution of delays as standalone signals. Furthermore, these works do not consider the system as a dynamic network with time-evolving signals that demand further analysis (see Cong et al., 2016; Li et al., 2021 for an overview of how airport delays can be represented as time-varying signals on graphs). Our aim is to view historical data as realizations of node-signals in a functional, dynamically evolving network and perform further statistical analysis.

We note that statistical analysis of time series embedded within multidimensional data is not a new concept. A common approach, used across multiple domains, for analyzing multidimensional data is to construct a graph by using the Pearson correlations of node signals as edge weights. These graphs have been referred to as *functional networks*, subsequently studied in several contexts. For instance, brain function networks have been used for causal analysis (van den Heuvel and Hulshoff Pol, 2010; Jiang et al., 2019) and to reveal the role of specific regions in neurological diseases (Jiang et al., 2019). Other applications of such networks include exploring synchronization patterns in rainfall events (Boers et al., 2019) and studying correlated physiological states during sleep (Bashan et al., 2012). Our work extends the functional network theory for studying "synchronizations" and temporal patterns in infrastructure system performance metrics — specifically, we focus on flight delays in the air transportation network. To accomplish this, we build upon a class of *time delay stability* (TDS) algorithms first proposed by Bashan et al. (2012) to understand sleep patterns. In particular, we introduce our algorithm to study the temporal couplings of delays between groups of airports, and investigate the theoretical properties of such an algorithm (Appendix B).

2. Our contributions

In this paper, we develop a time delay stability (TDS) algorithm to estimate statistically significant timescales for airport delay propagation from noisy measurements. The main contributions of our work are as follows:

1. We develop a TDS algorithm that allows for the extraction of statistically significant lags between airport delay time series based on piece-wise cross-correlation values. In addition, the intuitive properties of our algorithm are demonstrated using an idealized setting with sinusoidal delay signals.
2. We construction a pair of weighted directed graphs — a Delay Stability Network (DSN) and a Delay Lag Network (DLN) — to easily interpret and analyze the results from the TDS algorithm. The DSN encodes the number of stable lags per origin–destination pair, whereas the DLN encodes the average time duration of the lags on an origin–destination pair.
3. We demonstrate the utility of our algorithm by constructing DSNs and DLNs from 2017 US airport delay data, and analyze the DSNs and DLNs to uncover the timescales of delay propagation in the US airport network. From this analysis, we find that distances between the airports appear to be negatively correlated with delay propagation timescales.

As we have detailed in Section 1.1, previous research provides important insights into flight delay propagation. In addition to the contributions we list above, our work approaches delay propagation from a perspective of signals coupled by some underlying network. This statistical signal processing perspective of flight delay propagation has some precedence in studies such as Diana (2009), but has received relatively less attention than other models we survey in Section 1.1. The network characteristics we uncover in the form of stable time lags between airport delay signals could validate the effectiveness of initiatives taken to mitigate or reduce the impact of flight delays.

3. Methodology

In this section, we detail our method for analyzing timescales of airport delay propagation. In our representation of the air transportation system, airports are given as nodes, with edges between nodes weighted by the correlation between airport delays. Recall that the input for our algorithm is the historical delay time series at each airport. We develop a TDS algorithm in order to estimate statistically significant airport delay propagation timescales. This algorithm ensures that the estimated timescales between these interactions are a stable, persistent characteristic. We visualize these results using two graph representations: The Delay Stability Network (DSN) and the Delay Lag Network (DLN). DSNs represent airport pairs with stable timescales of delay propagation, i.e., the presence of an edge indicates that delays at one airport consistently affects the other in a statistically significant way. The DLN represents the magnitude of the delay propagation timescales for stable sequences. In a DLN, an edge is present between two nodes only for stable timescales, and the edge weight represents the magnitude of the delay propagation timescale. Thus, both the DSN and DLN help us identify statistically significant airport delay interactions as well as the timescale of such interactions, respectively.

3.1. Setup and notation

We use $\mathbb{R}$ to denote the set of reals, $\mathbb{R}_{\geq 0}$ for the set of non-negative reals, $\mathbb{R}^{m\times n}$ for the space of $m \times n$ matrices with real-valued entries, and $\mathbb{N}$ for the set of natural numbers. We also denote by $\wedge$ and $\vee$ the logical-AND and logical-OR operators, respectively. Let N denote the number of airports, and let Δt denote a time discretization, with time $t_i = i \times \Delta t$ for $i \in \mathbb{N}$. The observed total delay (arrival plus departure delay) at airport $a \in \{1, \ldots, N\}$ at time t_i is given by $x_i^a \in \mathbb{R}_{\geq 0}$. The time series of delays at airport a is denoted by $X^a = \left\{x_1^a, x_2^a, \ldots, x_n^a\right\}$, where $|X^a| = n$ is the length of time series X^a. Time series are denoted in capital letters, e.g., $X^a \in \mathbb{R}^{|X^a|\times 1}$, and the j^{th} element of the time series is indexed using square brackets. For example, $X^a[j] \in \mathbb{R}$ is the sum of the arrival and departure delays for all flights scheduled at airport a within the time interval of length Δt beginning at time t_j. Let $d_{\text{thr}}(X^a, \eta) \in \mathbb{R}_{\geq 0}$ with $\eta \in [0, 1]$ denote the $(\eta \times 100)$th percentile delay threshold for X^a, i.e., $(\eta \times 100)$ percent of delay observations in X^a are below $d_{\text{thr}}(X^a, \eta)$. Let $r : \mathbb{R}^{M\times 1} \times \mathbb{R}^{M\times 1} \to [-1, 1]$ be the standard sample Pearson correlation. For two time series X_1 and X_2 of length M, it is given by

$$r(X_1, X_2) = \frac{\sum_{\ell=0}^{M-1}\left(X_1[\ell] - \overline{X}_1\right)\left(X_2[\ell] - \overline{X}_2\right)}{\sqrt{\sum_{\ell=0}^{M-1}\left(X_1[\ell] - \overline{X}_1\right)^2}\sqrt{\sum_{\ell=0}^{M-1}\left(X_2[\ell] - \overline{X}_2\right)^2}}, \tag{1}$$

where $\overline{X}_i = \frac{1}{M}\sum_{j=0}^{M-1} X_i[j]$. We define $X_{t_i}^a = \left\{X^a[t_i], \ldots, X^a[t_i + w]\right\}$ to be a sub-series of X^a with length w starting at t_i. Next, we define a *directional sliding correlation* $r_{t_i,k}^{a\to b}$ given by

$$r_{t_i,k}^{a\to b} := r\left(X_{t_i}^a, X_{t_i+k\Delta t}^b\right), \qquad k = 0, 1, \ldots, k_{\max}. \tag{2}$$

Intuitively, *directional* refers to the fact that we are "sliding" the sub-series in one direction only, i.e., in the direction of increasing $k = 0, 1, \ldots, k_{\max}$. The p-value of the Pearson coefficient defined in (2) is denoted by $p\left(r_{t_i,k}^{a\to b}\right) \in [0, 1]$, and calculated as

$$p\left(r_{t_i,k}^{a\to b}\right) = 2 \times \mathbb{P}\left(T > \frac{r_{t_i,k}^{a\to b}\sqrt{\left|X_{t_i}^a\right| - 2}}{\sqrt{1 - \left(r_{t_i,k}^{a\to b}\right)^2}}\right), \tag{3}$$

with $T \sim \text{StudentT}\left(\left|X_{t_i}^a\right| - 2\right)$.

We now introduce some notation and variables that we will use throughout the remainder of this section. We collect the two aforementioned parameters w and $k_{\max}$, along with w_t, ε, and p_{thr} into a parameter tuple $\mathcal{P} := \left\{w, k_{\max}, w_t, \varepsilon, p_{\text{thr}}, \eta\right\} \in \mathbb{N}_{\geq 1}^2 \times \mathbb{R}_{\geq 0}^2 \times [0, 1]^2$. The parameters w_t, ε, and p_{thr} are the minimum stable lag time duration, maximum allowable lag perturbation, and p-value threshold, respectively; we will use them later in the TDS algorithm and the construction of the DSN and DLN.

Let $G = \left(V, A = \left[a_{ij}\right]\right)$ be a weighted, directed graph with $|V| = N$ nodes and adjacency matrix A, where $a_{ij} \neq a_{ji} \in \mathbb{R}$ and $a_{ij} = 0$ indicates no directed edge from node i to j. Denote by $G_{t_i} = \left(V, A_{t_i}\right)$ the time-varying graph at t_i, and a sequence of such graphs by $\mathcal{G}_{t_i,t_j} = \left\{G_{t_1}, \ldots, G_{t_j}\right\}$.

3.2. Time delay stability algorithm (TDS)

For any pair of airports (a, b), the TDS algorithm calculates the lag between the two time-series X^a and X^b based on the Pearson correlation coefficient between sub-series samples. The algorithms ComputeG and ComputeDN map the computed lag from TDS to directed graphs representing lag magnitudes and stability. The magnitude of the lag, along with its stability in time, provides a

Algorithm 1: Time Delay Stability (TDS); Compute Lag Graph (ComputeG); Compute Delay Networks (ComputeDN). Note that the output produces the DSN and DLN.

$\text{TDS}\left(X^a, X^b, t_i, \mathcal{P}\right)$
  Retrieve $k^\star_{t_i}$ from Eq. (4)
Result: $k^\star_{t_i}$

$\text{ComputeG}\left(\left\{X^1, \dots, X^N\right\}, t_i, \mathcal{P}\right)$
  $a_{uv} = \text{TDS}\left(X^u, X^v, t_i, \mathcal{P}\right), \qquad \forall u, v \in \{1, \dots, N\}$
Result: $G_{t_i} = \left(V, A_{t_i} = \left[a_{uv}\right]\right)$

$\text{ComputeDN}\left(\mathcal{G}_{t_i,t_{i'}} = \left\{G_{t_i}, \dots, G_{t_{i'}}\right\}, \mathcal{P}\right)$
check $i \in [0, n-2] \wedge i' \in [i+1, n]$
  for $u \in \{1, \dots, N\}$ **do**
    for $v \in \{1, \dots, N\}$ **do**
      Retrieve all valid $k^{uv}_{t_j}$ from Eq. (5)
      Construct $\mathcal{K}^{u\to v}$ using $k^{uv}_{t_j}$ and Eq. (6)
      $a^{\text{DSN}}_{uv} = \left|\mathcal{K}^{u\to v}\right| / \left(i' - i\right)$
      $a^{\text{DLN}}_{uv} = \left|\mathcal{K}^{u\to v}\right|^{-1} \sum_{k^{uv}_{t_j} \in \mathcal{K}^{u\to v}} k^{uv}_{t_j}, \qquad \forall j \in \left[i, i'\right]$
Result: $G^{\text{DSN}}_{t_i,t_{i'}} = \left(V, A^{\text{DSN}} = \left[a^{\text{DSN}}_{uv}\right]\right)$ and $G^{\text{DLN}}_{t_i,t_{i'}} = \left(V, A^{\text{DLN}} = \left[a^{\text{DLN}}_{uv}\right]\right)$

measure of timescale between two airport delay time series. We outline TDS, ComputeG, and ComputeDN in Algorithm 1, and provide a schematic of the workflow in Fig. 2.

TDS takes a sub-series $X^a_{t_i} \subset X^a$ of length w beginning at time t_i, and $X^a_{t_i}$ will be the *fixed reference sub-series* against which sub-series of X^b will be compared. TDS then takes multiple sub-series from X^b, all of length w, starting at the same time index t_i, and shifting forward by some *stride length* $s \le w$ a total of $k_{\max}$ number of times. We denote each X^b sub-series by $X^b_{t_i+k\Delta t}$ for $k = 0, \dots, k_{\max}$. This selection of a fixed reference sub-series from X^a and $k_{\max}$ number of comparison sub-series from X^b is depicted in Fig. 2(*a*). Specifically, Fig. 2(*a*) depicts two delay time series (BOS and ORD), where a sub-series of delays at BOS of length w is taken to be the fixed reference sub-series. Note that due to the asymmetry between how sub-series is selected from X^a versus X^b, TDS is an asymmetric measure of time delay stability between airports a and b, i.e., the resulting sequence of correlation coefficients when X^a is the fixed reference will be different than if X^b is the fixed reference. This asymmetry property is important to keep in mind because this results in the DSN and DLN being directed graphs.

Proposition 1. *Given X^a, X^b, t_i, and $\mathcal{P}$, the* TDS *algorithm is in general asymmetric, i.e.,*

$$\text{TDS}\left(X^a, X^b, t_i, \mathcal{P}\right) \neq \text{TDS}\left(X^b, X^a, t_i, \mathcal{P}\right).$$

Proof. See Appendix A. □

Without loss of generality, we describe the remainder of the TDS algorithm assuming that X^a is the reference time series. For each fixed reference sub-series $X^a_{t_i}$, TDS computes a corresponding *lag* based on the sequence of directional sliding correlation and p-value pairs $\left\{\left(r^{a\to b}_{t_i,k}, p\left(r^{a\to b}_{t_i,k}\right)\right)\right\}_{k=0}^{k=k_{\max}}$. Specifically, for a pre-selected p-value threshold (in this paper, we set $p_{\text{thr}} = 0.05$), TDS retrieves a lag value $k^\star_{t_i}$ that is maximized over the set of statistically significant directional sliding correlation values, formalized in (4). This procedure of selecting $k^\star_{t_i}$ is depicted in Fig. 2(*b*). The fixed reference sub-series for the BOS example in Fig. 2(*b*) is used to compute correlations between the various sub-series from ORD, and this correlation (which is only kept if it is statistically significant) is plotted with respect to k.

$$k^\star_{t_i} = \underset{k=0,\dots,k_{\max}}{\text{argmax}} \left\{ r^{a\to b}_{t_i,k} \,\middle|\, p\left(r^{a\to b}_{t_i,k}\right) \le p_{\text{thr}} \right\} \tag{4}$$

As a technical remark, we characterize the behavior of TDS when applied to *idealized sinusoidal signals*; this analysis can be found in Appendix B.

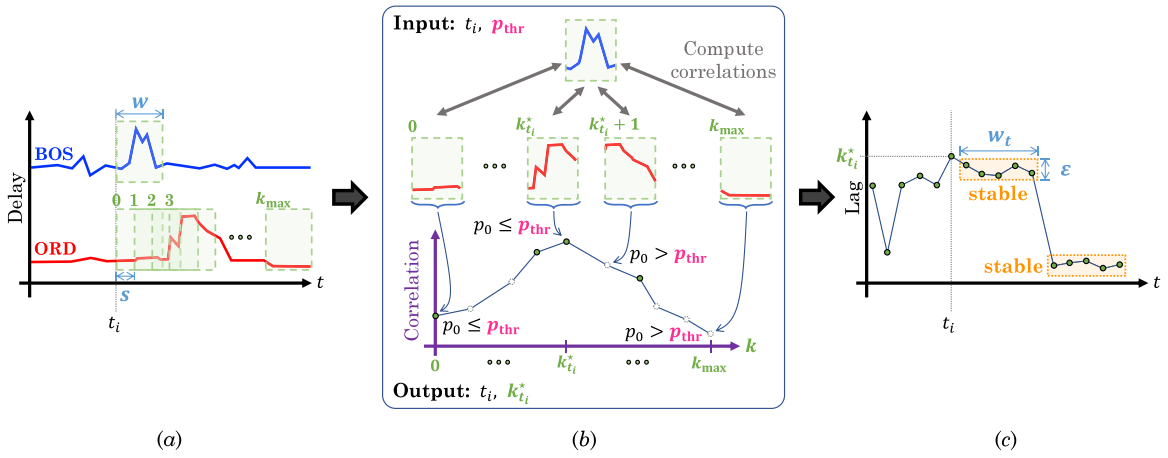

Fig. 2. Schematic diagram of the TDS algorithm workflow, specifically the (*a*) sub-series selection; (*b*) computation of correlations and selection based on statistical significance (solid green dots indicate retained correlations). The DSN and DLN are then produced after (*c*) extraction of stable sequence of lags.

3.3. Delay lag network (DLN) and delay stability network (DSN) construction

The TDS algorithm that we describe in the previous subsection extracts statistically significant, temporally-stable (i.e., by way of the w_t minimum stable lag time duration parameter), and consistent (i.e., by way of the ε maximum allowable lag perturbation parameter) lags. We now describe the process to encode these extracted lag values into a graph representation, amenable to various network analysis tools and more easily comparable with previous air transportation network-wide results. Through ComputeG, the lag value $k^\star_{t_i}$ is computed for every pair of airports (u, v), and is set as the weight of the edge going from airport u to airport v in graph G_{t_i}. We now will use ComputeDN to locate *stable* sequences of lags across time index interval $[t_i, t_i']$. This step in the workflow is depicted in Fig. 2(*c*). In particular, Fig. 2(*c*) plots the lag values extracted from Fig. 2(*b*) against time, and identifies stable sequences of lags based on parameters w_t (minimum duration) and ε (maximum perturbation). We first collect the sequence of graphs together corresponding to $[t_i, t_i']$ in $\mathcal{G}_{t_i,t_i'} = \left\{G_{t_i}, \ldots, G_{t_i'}\right\}$, and give the graph sequence $\mathcal{G}_{t_i,t_i'}$ to ComputeDN along with the parameter tuple $\mathcal{P}$.

We now motivate the need for a delay threshold η included within the parameter set $\mathcal{P}$: When we apply ComputeDN to a data set of airport delays, the algorithm may return lags $k^\star_{t_i}$ that correspond to times with negligible airport delays. For example, suppose that both airports a and b experience a nominal operational day, with possibly occasional accumulations of trivial delays on the order of minutes due to inconsequential issues that have nothing to do with network-wide disruptions. We consider this to be undesirable, as these delays should *not* cause airports a and b to be linked by some lag $k^\star_{t_i}$; thus, we use the η from $\mathcal{P}$ to threshold delays at a specific airport. The delays at that airport must surpass this threshold at time t_i in order for the corresponding lag value $k^\star_{t_i}$ to be considered *valid*. Formally, this thresholding can be written as

$$k^{uv}_{t_j} = \left\{ \left[a_{uv}\right] \in A_{t_j} \,\middle|\, j \in \left[i, i'\right] \wedge \underbrace{x^u_j > d_{\text{thr}}\left(X^u, \eta\right) \wedge x^v_j > d_{\text{thr}}\left(X^v, \eta\right)}_{\text{Delay threshold condition at airports } u \text{ and } v} \right\}. \tag{5}$$

Within ComputeDN, stable sequence of lags is defined by two parameters from $\mathcal{P}$: w_t and ε, where w_t is the *minimum length* of a time interval that contains consecutive lags within *an acceptable deviation* ε. Specifically, we have that the set of stable lag sequences $\mathcal{K}^{u\to v}$ between airport pairs (u, v) is:

$$\mathcal{K}^{u\to v} = \left\{ k^{uv}_{t_j} \,\middle|\, \underbrace{\delta^+ - \delta^- \ge w_t}_{\substack{\text{min. stable lag} \\ \text{time duration}}} \wedge \underbrace{\max_{\xi\in[\delta^-,\delta^+]} k^{uv}_{t_\xi} - \min_{\xi\in[\delta^-,\delta^+]} k^{uv}_{t_\xi} \le \varepsilon}_{\text{max. allowable lag perturbation}} \right\}, \tag{6}$$

where $j \in [i, i']$ in (6) and there exists valid δ^+, δ^- such that $\delta^- \le j \le \delta^+$. Note that δ^+ and δ^- describe time intervals wherein the minimum stable lag time duration is satisfied, as well as the maximum allowable lag perturbation.

The last two lines in Algorithm 1 for ComputeDN prior to returning the DSN and DLN inform the interpretation of the DSN and DLN. The DSN is a weighted graph $G^{\text{DSN}}_{t_i,t_i'}$ that encodes the *stability* of the lag between airport delays on its directed edges. Regardless of the lag value, the total time duration during which that lag is stable is given by $|\mathcal{K}^{u\to v}|$, and the edge weight a^{DSN}_{uv} stores the value $|\mathcal{K}^{u\to v}|$ normalized by the total time duration of interest $i' - i$. On the other hand, the DLN encodes the actual lag value between airport delays, averaged across all stable lag durations. We emphasize that these graph representations can be interpreted directly

with respect to the airport origin–destination pairs that appear in the DSN and the DLN. For a concrete example of these graph-based representations, we direct the reader to Fig. 5, where we display the actual DSN and DLN for the US airport delay case study.

Finally, we note that the outputs of TDS and ComputeDN depend on input parameter choices $\mathcal{P}$. Thus, to ensure that our results are robust, we would like the results to not be overly sensitive to our choice of w (minimum sub-series length), w_t (minimum stable lag time duration), ϵ (maximum allowable lag perturbation), and η (airport delay threshold). Since the airport delay threshold naturally depends on the input airport delay data set, we discuss setting η in Section 4. For the parameters w, w_t, and ϵ that directly affects ComputeDN, we conduct a sensitivity analysis on these parameters in Section 6.1, and discuss their influence on our results for the US airport network. Note that we fix $p_{\text{thr}} = 0.05$ as a standard significance threshold, and we fix $k_{\max}$ to be 12 units (6 h) since that is the longest flight within the contiguous US.

4. US airport delay data

We use US airport delay data retrieved from the Aviation System Performance Metrics (ASPM) maintained by the Federal Aviation Administration (FAA) (Anon., 2020). We restrict our analysis to 77 airports tracked by ASPM, known as the "ASPM 77". The data spans a year, starting at $t_0 \overset{\Delta}{=}$ 00:00 UTC on January 2, 2017, and ending at 23:59 UTC on December 31, 2017. We use a time discretization of 30 min. The departure or arrival delay of individual flights that were scheduled to operate within each 30 min time windows were added up and considered as the airport delay. Note that flights with negative delays (i.e., early arrivals) are not reported in the ASPM database; these flights are treated as having 0 delays. However, we note that our method can work for negative-valued signals, as noted in Appendix B. Between the ASPM 77 airports, there were 5.6 million individual flights operated within the time frame of interest. The number of airport pairs within the ASPM 77 connected by direct flights is 5175. The aggregate total airport delay (i.e., departure plus arrival delays) across the ASPM 77 for the time frame of interest is 177 million minutes, with average hourly delays of slightly over 2 h across the 77 airports.

We give an example of running $\text{TDS}\left(X^a, X^b, t_i, \mathcal{P}\right)$ using 2017 US airport delay data in Fig. 3. We take the fixed reference to be a = BOS, the sliding time series to be b = DCA, and the time intervals of interest t_i to be between 00:00 UTC on March 5, 2017 through 00:00 UTC on March 7, 2017. Fig. 3*(a)* plots the BOS and DCA airport delay time series, along with their delay thresholds obtained with $\eta = 0.5$ (in the succeeding section, we will discuss how we selected the 50th percentile of total airport delays for η). Each delay data point is the sum of arrival and departure delay of all flights that were scheduled to arrive at or depart from BOS/DCA within a 30 minute time-frame. The set of maximal directional sliding correlation values (i.e., the set of $k^\star_{t_i}$ obtained from (4)) $r^{\text{BOS}\to\text{DCA}}_{t_i,k}$, along with the accompanying p-value, is shown in Fig. 3*(b)*.

Lastly, the set of lag values $k^\star_{t_i}$ corresponding to the maximal $r^{\text{BOS}\to\text{DCA}}_{t_i,k}$ values are given in Fig. 3*(c)*. Valid lags are denoted in red, indicating lag values at time indices that not only meet the delay threshold criterion, but also pass the p-value threshold $p_{\text{thr}} = 0.05$. Note that a *stable* (or consistent) sequence of lag values would consist of temporally continuous valid lags (i.e., red dots in Fig. 3*(c)*) that meet a minimum time length given by w_t. Furthermore, this stable sequence of lags cannot vary too much in magnitude, i.e., the variation with respect to the vertical axis is constrained by the maximum perturbation parameter ϵ. Thus, a sequence of valid lags can fail to be stable for two reasons: There may be invalid lags interspersed, or the lag magnitude may vary excessively. Finally, the lag values in Fig. 3*(c)* are passed to ComputeG and ComputeDN to be encoded into a DSN and DLN. Intuitively, the edges of the DSN is weighted according to the total number of valid lags. Similarly, the edges of the DLN are weighted according to the average lag magnitude, given a stable sequence of lags.

4.1. Selecting a delay threshold η

We now discuss selecting the delay threshold η, as it is an important design criterion to ensure interpretable and operationally relevant DSN (and correspondingly, the DLN). In particular, if the threshold is too high, the number of edges in the DSN become sparse, resulting in a large number of isolated nodes. Contrarily, if the threshold is too low, spurious edges representing stable lags between inconsequential airport delays arise. With this in mind, we set a reasonable threshold η by observing the change in network connectivity (i.e., number of directed edges) and number of isolated nodes (i.e., number of nodes with no adjacent nodes) within the DSN. For the 2017 US airport network delay data, we plot the network connectivity and number of isolated nodes as a function of the delay threshold in Fig. 4.

In Fig. 4, we sweep η from $\eta = 0.1$ (i.e., the 10th percentile of delays) to $\eta = 0.9$ (i.e., the 90th percentile) in 0.1 increments. We observe that prior to the 50th percentile, the number of isolated nodes (airports) is essentially constant, and consists of less than 15% of all airports. If a delay threshold higher than the 50th percentile is selected, more lag values are invalidated, resulting in the loss of edges and an exponential increase in the number of isolated airports. Thus, we select a threshold such that we achieve a balance between the loss of important edges versus retaining spurious ones. We fix the 50th percentile, i.e., $\eta = 0.5$ as the delay threshold to use for the analysis and results in this paper.

5. Results

In this section, we present our main results pertaining to the timescales of delay propagation within the US airport network, using the TDS algorithm detailed in Section 3, as well as the data set and parameter settings discussed in Section 4.

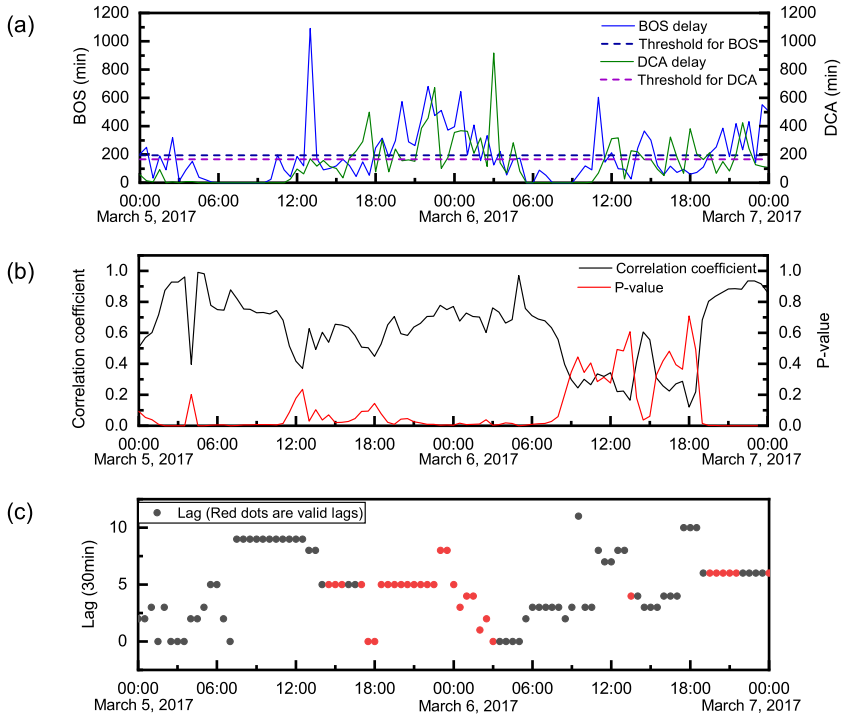

Fig. 3. Example of time delay stability between BOS and DCA. Note that the horizontal time axis is in UTC. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

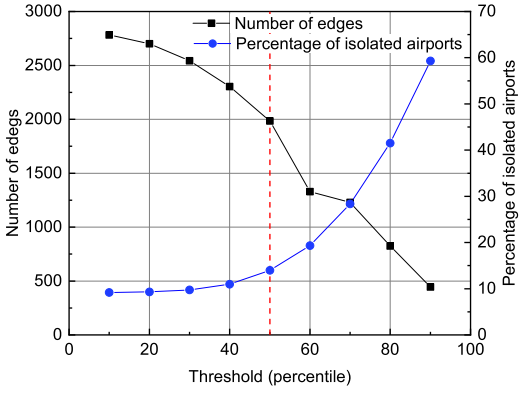

Fig. 4. Number of edges and isolated nodes (airports) in the DSN as a function of η.

5.1. Airport pairs with stable delay lags

We first analyze the DSNs that are identified from our analysis. Recall that a "stable" lag means that the time scale of delay interaction between the two airports is consistent (i.e., are stable and the lag magnitude does not change with time). In Fig. 5 (left), we show the top 5 percentile of OD pairs based on their edge weight $[a_{uv}^{DSN}]$ (refer to Algorithm 1 for the exact definition of the edge weight). We note a distinct geographical pattern: The most stable lags involve airports in the Northeast part of the US. This means that many of the east coast airports end up with delays that are locked "in-phase" with those of other airports. Of course, the lag could be different for different OD pairs, as shown in Fig. 5 (right). However, the delays at these airport pairs still remain locked in phase longer than other pairs of airports.

A potential explanation for the high stability of airport lags involving the east coast airports is their geographical proximity. This can lead to correlated weather impacts, as well as coordinated traffic management initiatives across these airports, both of which

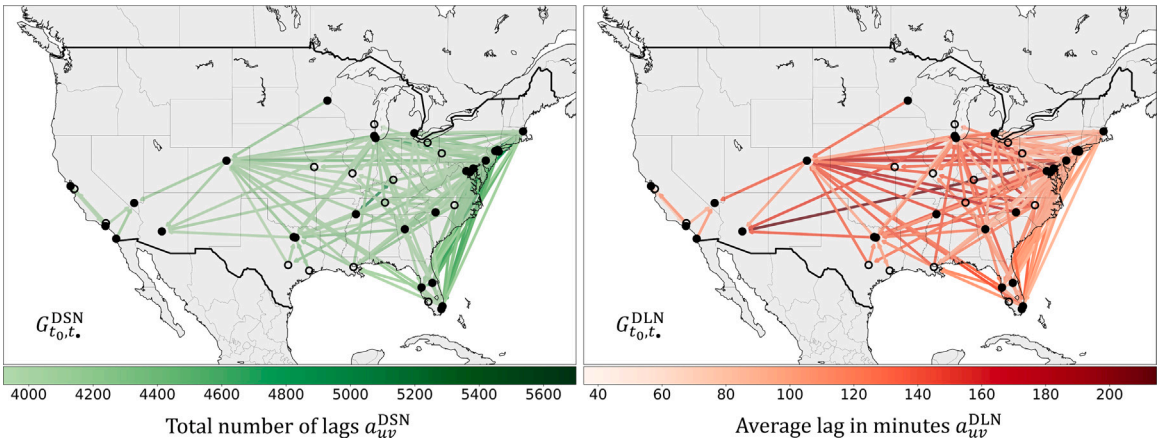

Fig. 5. Visualization of the DSN (*left*) and DLN (*right*) corresponding to the 2017 US airport data. Core airports are shown as filled circles, and non-core airports are hollowed circles.

can result in high magnitude of delays that occur simultaneously at these airports. Such a scenario would result in stable delay lags between the affected airports.

5.2. Relation between delay lag and flight duration

Our data-driven stability and lag analysis to identify the timescale at which delay propagates between airports helps to identify the dominant modes of delay propagation. But first, we will revisit the two common mechanisms through which delay propagates. The most intuitive mode is due to a delayed aircraft making multiple hops. When the schedule buffer at an airport is not sufficient to absorb the delay of an inbound aircraft, the subsequent departure of that flight will add to the delay metric of the airport. Note that the time-scale for the propagation of delays through this mechanism is roughly the flight duration between the two airports. For example, if the flight duration from New York to San Francisco is 6 h, then the time-scale of delay propagation is also 6 h. We refer to this mode of delay propagation as *tail-propagated delay*.

Another mode of delay propagation occurs when the FAA issues certain traffic management policies. In particular, a Ground Delay Program (GDP) can lead to an instantaneous propagation of delays to an airport because its departures are headed to an airport which is congested. To further elaborate, a GDP is a strategic measure taken by the FAA to prevent excessive airborne delays at an airport with reduced capacity (e.g., due to poor weather). Thus, flights that are arriving at reduced-capacity airports are delayed on the ground at their departure points itself. This results in instantaneous delay propagation across airports which may not even have any instantaneous traffic between them. We refer to this mode of delay propagation as *GDP-propagated delay*.

These different delay propagation modes all have different delay propagation timescales. The observed delay lag is a result of three factors: The first factor is the typical flight time between the two airports, also known as the schedule block time (SBT); lower the SBT, lower the lag. The second factor is the extent to which GDPs in one airport affects the other. The third factor that affects the lag in route $a \to b$ is the nature of inbound traffic at airport b; if there is significant inbound traffic at b from other airports with different time scales of interaction, then the overall delay at b will be more strongly governed by those interactions than from a.

Figs. 6*(a)–(c)* show the delay lag in relation to the SBT. We also produced the empirical cumulative distribution functions (CDFs), as well as pertinent empirical CDF thresholds, corresponding to the difference between SBT and delay lag. For brevity, we have relegated these to Appendix D, along with additional auxiliary discussions. These lags (and their corresponding airport pairs) are grouped into three categories depending on whether the lag is higher than, lower than, or nearly equal to the SBT. For brevity, we will define $\Delta(\mathrm{SBT}, \mathrm{Lag})$ to denote this difference:

$$\Delta(\mathrm{SBT}, \mathrm{Lag}) = \mathrm{SBT} - \mathrm{Lag}.$$

Fig. 6*(c)* shows airport pairs where the delay lag is significantly smaller than the SBT for that route, i.e., $\Delta(\mathrm{SBT}, \mathrm{Lag}) \gg 0$. This means that the delay actually propagates much faster than the time it takes for a flight to travel from one airport to the other. In particular, we note that most of these airport pairs are correspond to transcontinental routes. This suggests that the delay interactions between these airports are strongly affected by local traffic (with a smaller SBT) and instantaneous propagation due to GDPs.

Fig. 6*(a)* shows airport pairs where the delay lag was significantly larger than the SBT, i.e., $\Delta(\mathrm{SBT}, \mathrm{Lag}) \ll 0$. This means that very little tail delay propagation occurred between these two airports. This is also consistent with the edges highlighted in Fig. 6*(a)*, as they are often small regional airports which typically only have flights to one central hub. Finally, Fig. 6*(b)* shows airports pairs where the lag is comparable to the SBT, i.e., $\Delta(\mathrm{SBT}, \mathrm{Lag}) \approx 0$.

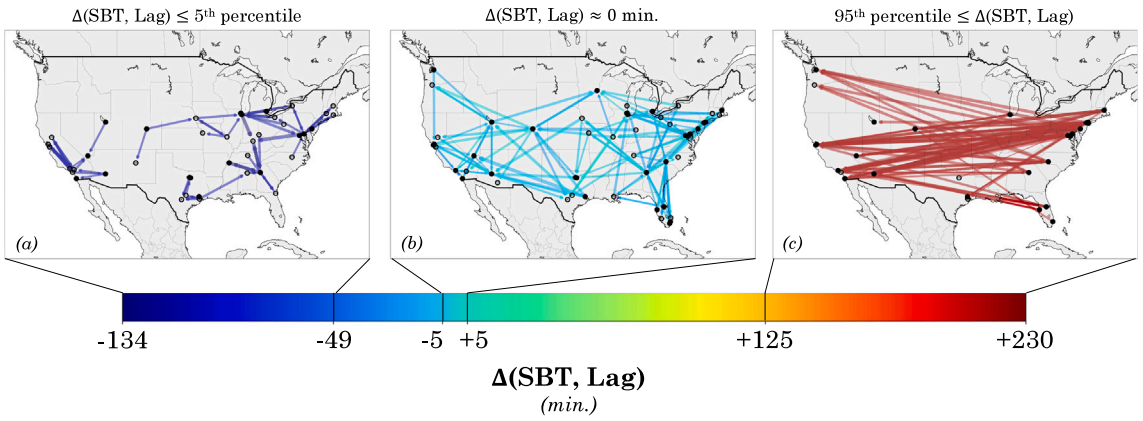

Fig. 6. OD pairs with *(a)* SBT significantly smaller than lag values, $\Delta(\text{SBT, Lag}) \ll 0$; *(b)* SBT approximately equal to lag values, $\Delta(\text{SBT, Lag}) \approx 0$; *(c)* SBT significantly larger than lag values, $\Delta(\text{SBT, Lag}) \gg 0$.

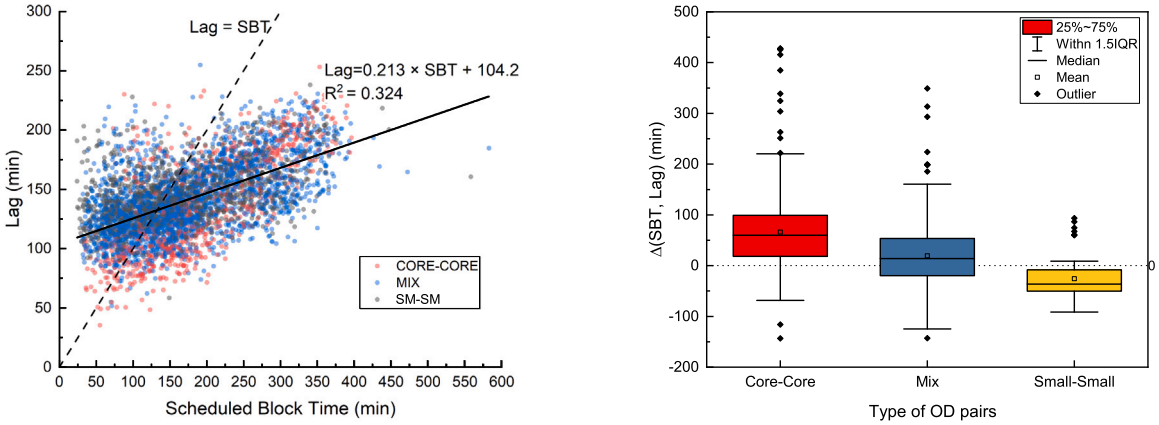

Fig. 7. (*left*) Comparison between valid lag values of delays between airport pairs and the actual block time of flights between airport pairs; (*right*) Distribution of Δ(SBT, Lag), split by airport pair type.

Examining Figs. 6*(a)–(c)*, a natural question to ask is the following: What factors does the speed of propagation of delays between two airports depend on? We investigate two of such factors: The size of the airports (in terms of the number of flight operations) involved in propagating delays, and the geographical distance between these airports. The reason why we investigate the role of these two factors is two-fold: (1) We would expect stronger delay influences between larger airports that have more direct flights between them, and (2) shorter geographical distance dictate that the airports tend to experience correlated weather impacts, resulting in near-instantaneous delay propagation. These results are presented in Fig. 7.

In Fig. 7 we plot the lag as a function of the typical SBT for each airport pair. In the figure, we also distinctly mark the three types of airport pairs based on the size: Airport pairs where both are large airports (core-core), airport pairs where one is large and another is smaller (denoted as “mix”), and airport pairs where both are small airports (small-small). Across all types of airport pairs, we see a weak positive correlation between the lag and the SBT in the left panel of Fig. 7. We also plot the diagonal line indicating SBT equal to lag, and the linear fit for the observed data points for comparison. We observe that when the SBT is high, the lag is typically lower than the SBT, whereas when the SBT is low, the lag is higher than the SBT. While the left panel of Fig. 7 analyzed the SBT-Lag curve as a function of the SBT, we perform a more detailed analysis on the role of airport size on Δ(SBT, Lag), and plot the results in the right panel of Fig. 7. Here, we draw standard box-plots for Δ(SBT, Lag) for the three different types of OD pairs. When both airports are small, the lag is typically larger than the SBT (i.e., $\Delta(\text{SBT, Lag}) < 0$). On the other hand, when both airports are large, then the lag is typically smaller than the SBT (i.e., $\Delta(\text{SBT, Lag}) > 0$). This may indicate that airport size also influences the average timescale of delay propagation, although we caveat that more analysis is needed to separate the two factors (e.g., core-core airports may also tend to be farther apart from each other, resulting in a higher SBT).

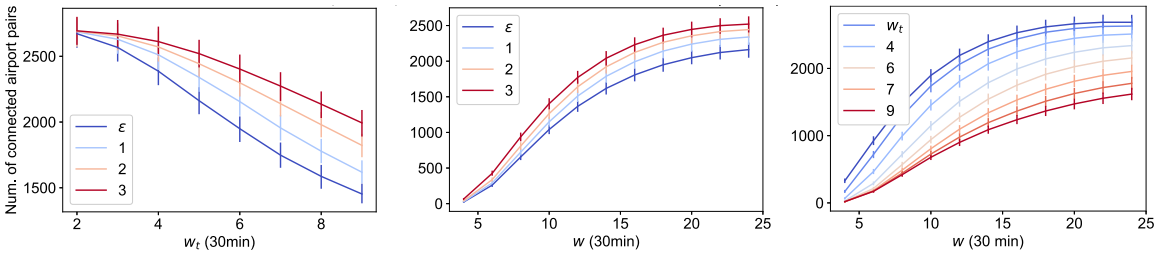

Fig. 8. Sensitivity of number of edges in DSN (averaged across a day, for all days in 2017) given our parameter choices $w = 24$ (12 h), $w_t = 5$ (2.5 h), and $\epsilon = 1$ (30 min). The error bars denote the 95% confidence interval around the average number of edges.

6. Discussion

We now focus in on the case study results from the previous section with discussions regarding four aspects: The sensitivity of our method and results to specific members of the parameter set $\mathcal{P}$ (Section 6.1), the asymmetries in delay propagation timescales between two airports (Section 6.2), the relative importance between various factors when it comes to delay lags (Section 6.3), and the advantages of graphical representations such as the DSN and DLN (Section 6.4). We close this discussion section with a summary of potential applications of our delay propagation timescales work in Section 6.5.

6.1. Sensitivity to parameters

We will now discuss the sensitivity of the resultant DSN and DLN with respect to the input parameters for TDS, ComputeG, and ComputeDN. Recall that the η parameter dictates an airport delay threshold level, wherein a higher η results in a loss of edges in the DSN and DLN, and a lower η produces spurious edges. Since the selection of η is highly dependent on the raw delay values themselves, we already discussed selecting η in Section 4.1 when giving an overview of the US airport delay data. We use 30 minute time discretizations, and we set the value of $k_{\max}$ to be 12 time units (6 h) corresponding to approximately the longest contiguous flight duration. For statistical significance, we choose a p-value threshold of $p_{\text{thr}} = 0.05$.

Our sensitivity analysis centers on w, w_t, and ϵ. Recall that w is the sub-series length, w_t is the minimum stable lag time duration, and ϵ is the maximum allowable lag perturbation. These three parameters are directly related to the directional sliding correlation (by way of w) and the extraction of stable lags (by way of w_t and ϵ). If w is too small, then the directional sliding correlations may become statistically insignificant, since the p-value is directly proportional to $\sqrt{w}$. We also want to avoid values of w that are too large, since delay signal aliasing may occur. We want to balance between identifying too many lags as stable, versus rejecting significant stable lags: We do this by adjusting w_t and ϵ. For our results, we set w to be 24 units (12 h), w_t to be 5 units (2.5 h), and ϵ to be 1 unit (30 min). We provide the sensitivity results for our selection of w, w_t, and ϵ in Fig. 8, where one parameter within w, w_t, and ϵ is fixed, the remaining parameters are varied, and sensitivity is measured with respect to the number of connected airport pairs, i.e., edges in the DLN and DSN. Additional sensitivity plots are given in Appendix C. The error bars on all sensitivity plots represent the 95% confidence interval of the mean number of edges, averaged across a day.

Our selection of w, w_t, and ϵ takes into account interpretability as well as robustness to changes in the other parameters. We set w to be approximately the length of an operational day in the US NAS (12 h), and consider lag values within ±15 min of each other (i.e., $\epsilon = 1$ unit) as indistinguishable from an airport delay perspective. Our choice of ϵ was motivated by the fact that a flight is considered to be on-time if it arrives less than 15 min after its published arrival time (Anon., 2016). With this selection of w and ϵ, we see from Fig. 8 that the number of edges does not vary significantly around $w = 24$ units; the same is true for fixing ϵ to be 1 unit. Finally, since the number of edges seems to decrease linearly with w_t, we choose a value for w_t in the middle (i.e., $w_t = 5$ units) where edges are not rejected too frequently nor spuriously included.

6.2. Asymmetry in airport delay interactions

In Section 3, we noted (and proved rigorously) that the lag between two airports is not symmetric. This means that estimating the lag of one airport with respect to the other is not the same as compared to if we estimate the lag of the second airport's delays with respect to the first. Here, we will present some empirical results to highlight this further and discuss its implications.

We list in Table 1 the 10 airport pairs with the most stable delay lags (as measured by the total length of the stable lag periods). For these airport pairs, we report the mean SBT, as well as the mean delay lags obtained from our algorithm, for each airport pair, distinguishing between the order. We note that the SBT for the airport pairs differs significantly based on the direction of travel. This may be a consequence of efficient and robust scheduling by airlines that takes into account (1) typical wind patterns that influences travel times in one direction versus another, and (2) significant delays when attempting to land at a congested airport due to traffic management initiatives such as a GDP, while departures tend to be less impacted. For example, the SBT for flights going from Boston (BOS), a less congested airport to the New York City Airports (LGA and EWR), which are more congested is higher than for flights in the other direction.

Table 1
Comparison on the lags differences (in minutes)

Airport a	Airport b	Total length	SBT		Lag		Lag difference
			$a \to b$	$b \to a$	$a \to b$	$b \to a$	
BOS	EWR	5701	84.7	71.0	49.2	50.6	−1.4
LGA	BOS	4887	78.9	83.4	101.5	66.9	34.5
PHL	BOS	4863	84.4	95.7	85.3	46.9	38.4
PHL	ORD	4695	141.8	123.6	111.7	81.2	30.5
DCA	BOS	4679	93.5	103.6	97.4	72.8	24.5
MCO	EWR	4610	163.7	172.8	90.6	87.3	3.2
MCO	JFK	4598	161.0	180.0	95.7	92.2	3.4
ATL	EWR	4560	139.2	149.6	79.4	83.7	−4.3
LGA	DCA	4546	87.9	80.3	89.0	77.3	11.7
FLL	EWR	4540	175.3	182.3	80.5	89.4	−8.9

Table 2
The importance of features in predicting the delay lag, ranked via random forest.

Feature	Importance
Distance between two airports	0.885
Number of direct flights operated between two airports	0.056
Ratio of number of direct flights to the total number of flights between two airports	0.023
Time difference between two airports	0.017
Airport size (Core-Core)	0.008
Airport size (Small-Core)	0.004
Airport size (Core-Small)	0.003
Airport size (Small-Small)	0.003

Empirically, the lag between two airports is also asymmetric, which matches our theory as stated and proved in Proposition 1. This difference can be small, on the order of a few minutes (e.g., BOS-EWR versus EWR-BOS), or large (30 min or more, such as in the case of PHL-BOS versus BOS-PHL). The routes with high lag differences highlight the significant asymmetry in the US airport network, even between two relatively major airports such as Philadelphia and Boston. We believe that further investigation is warranted to better understand both the cause of such a large difference in the delay propagation timescale as well as its potential implications.

6.3. Stable lag factors: Relative importance via random forest

Thus far, we have identified that airport delay lags depend on airport size (e.g., core-core, mix, or small-small) as well as the SBT between airports pairs (i.e., see Fig. 7). However, we did not quantitatively consider the *relative importance* of these factors, even though a visual inspection of the left panel in Fig. 7 would suggest that the dominant factor was the SBT. In this subsection, we will quantify the relative importance of various factors in determining the delay lags (i.e., the timescale of delay propagation).

For this analysis, we set up a random forest classifier (Breiman, 2001) with the delay lag set as the predicted (dependent) variable, along with the following features: Distance between the airports, number of direct flights between the airports, fraction of direct flights in comparison to the total number of flights, time zone difference, and the size of the airports (core-core, mix, small-small). The random forest outputs are presented in Table 2. We note that even with a completely different class of methods (i.e., random forest), the results indicate that the distance between the airports – a quantity directly related to the SBT – is the most important factor in determining the delay lag. Furthermore, we can quantitatively see that although the airport size is important, the distance between airports play a much more dominant role as a factor in predicting the delay lags.

6.4. Using the DSN and DLN representations

We choose to represent our results from TDS in the form of two graphs, namely, the Delay Stability Network (DSN) and the Delay Lag Network (DLN), in order to open up future analyses that leverage graph-theoretic and network analysis tools. Although the focus in this paper is on presenting the methodology to construct such graphs, we believe that this representation is extremely powerful. Specifically, the DSN and DLN graphs are amenable to the application of several analytical tools, examples of which include node or graph clustering, community detection, and computation of node-based centrality measures. These graph-theoretic tools help to quickly identify OD pairs that consistently propagate delays. This can play a major role in the development of airport delay cascade and diffusion models (David and Jon, 2010), as well as the exploration of optimal network structures that minimize delays (Crainic et al., 2021). The former models are examples of *tactical* disruption management strategies, whereas the latter is an example of *strategic* network design optimization.

6.5. Potential delay propagation timescale applications

Our analysis of the timescales of delay propagation can have a direct impact on improving airline schedules; examples include decisions regarding the routes to fly, schedule buffer design, and standby crew positioning. Our analysis of historical airport delay data is designed to capture statistically significant and consistent delay propagation effects. Such a method could serve as a valuable, computationally tractable tool for airlines to evaluate the performance of new scheduling practices.

We present an illustrative example of how a result such as lag asymmetry (e.g., Table 1) could be considered as a factor during airline scheduling. Specifically, we discuss this example in the context of adding *buffers* to airline schedules (AhmadBeygi et al., 2010; Arıkan et al., 2013; Brueckner et al., 2021).

Example (*Airline Schedule Buffers*)**.** Suppose that an airline conducts a delay analysis via TDS (i.e., Algorithm 1) using historical delay signals from its own network, and observes Table 1. In particular, PHL-BOS has a stable delay lag that is highly asymmetric. Historical delays observed at PHL lead delays at BOS by 85 min; however, when delays are observed at BOS, it tends to lead PHL delays by only 46 min, or about twice as fast as the "PHL leads BOS" case. The airline notes that the corresponding SBTs are around 84 and 95 min, respectively. The significant difference between the BOS-PHL SBT (95.7 min) and the BOS-PHL lag (46.9 min) could be flagged for further investigation by the airline as a candidate OD pair for schedule compression. Such schedule compressions, particularly if the airline operates a shuttle-like schedule (e.g., the Delta Shuttle service offered by Delta Air Lines Lines, 2021), could allow for additional scheduled flights, thereby increasing service frequency.

Furthermore, with DSNs and DLNs, airlines can visualize clearly the airports which have a strong coupling in terms of their delays. While some of these delay coupling effects may be desirable, as they allow delay to diffuse throughout the system rather than being concentrated at a hub airport, other coupling effects may be undesirable. Our method and workflow provide a viable way to identify these OD pairs, and is an important first step towards a comprehensive analysis of desirable and undesirable airport delay propagation effects.

Our work in this paper was motivated by the airport delay propagation problem; however, the underlying phenomena that we study – signals indicative of how the underlying networked system is performing – is more general. In particular, our setup in essence comprises of a networked system (i.e., high-dimensional system with correlated data from its sub-components) generating a time series of scalar signals (e.g., delays) at its nodes. Other systems amenable to this type of abstraction include other transportation infrastructure such as bikeshare networks and road traffic networks, as well as examples from other domains such as power grids and energy networks, biological systems (e.g., ECG signals from the heart, or EEG signals from the brain), and epidemiological models. A common theme among all of these examples is that the observed data is potentially noisy, and that there are multiple pathways through which one node might affect another, resulting in multiple timescales of interaction between the different system sub-components. Thus, finding dominant and consistent timescales of interaction from data help to elucidate the underlying dynamics of these systems, paving the way towards designing controllers for these systems.

7. Conclusion and future work

The delays at different airports within an air transportation system interact through traffic (e.g., direct flights) and operational (e.g., traffic management initiatives) interconnectivity. In this work, we provide a rigorous method through which the timescales of these airport delay interactions can be identified and studied. Specifically, we propose TDS which retrieves a *lag* between two airport delay time series. This lag value represents a statistically significant temporal offset between time periods of increased airport delays across a pair of airports. Then, these delay lag values can be retrieved for an entire network of airports, and encoded as a *Delay Stability Network* (DSN) or *Delay Lag Network* (DLN). The former focuses on persistent, or stable lags, whereas the latter focuses on the actual magnitude of the lags. By studying the resultant DSN and DLN, we can quantitatively analyze airport delay interaction timescales, as well as visualize them as weighted, directed graphs.

We applied our method to 2017 US airport delay data focused on the ASPM 77 set of airports, and observed that the most stable lags tend to be incident on Northeast airports. We also compared the difference between the scheduled block time, a proxy for flight duration, and the lag magnitude. This difference quantifies a delay interaction timescale relative to the timescale of a direct flight between the two airports. In addition, we discuss many characteristics of our method as it relates to the data-driven case study, e.g., the sensitivity of our results to input parameters, and ranking (using random forest) the relative importance of operational factors for determining the delay lag. The results from the application of our methods have potential implications for building network-wide airport delay propagation models, as well as for airline schedule design.

In terms of future work, computing delay interaction timescales for each individual airline would be an important extension. In particular, it may be interesting to examine if different operational practices, such as having a hub-and-spoke network versus a more point-to-point one would result in differing DSN and DLN characteristics. Another direction for future work would be to consider externalities such as convective weather explicitly, and investigate its impacts on delay propagation timescales. For this, we would have to create DLNs and DSNs for time periods that only include days with a particular type of weather-induced impact.

Finally, we could consider the dispersion of the data points in the left panel of Fig. 7 as a measure of the *predictability* (or lack thereof) within the system. In other words, the variance along the y-axis (the lag) could be considered as a measure of the inherent variability in delay propagation timescales for a fixed OD pair. Thus, another interesting future research direction would be to develop techniques that reduce the variability in propagation dynamics.

CRediT authorship contribution statement

Yanjun Wang: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Funding acquisition. **Max Z. Li:** Conceptualization, Methodology, Formal analysis, Investigation, Writing – original draft, Writing – review & editing. **Karthik Gopalakrishnan:** Conceptualization, Methodology, Formal analysis, Investigation, Writing – original draft, Writing – review & editing. **Tongdan Liu:** Conceptualization, Methodology, Software, Formal analysis, Investigation.

Acknowledgments

This research was supported by the National Natural Science Foundation of China (Grant Nos. U2033203, U1833126, 61773203, 61304190). The authors thank Professors Hamsa Balakrishnan and Amedeo Odoni for their helpful comments and suggestions.

Appendix A. Proof for Proposition 1

Proof. Suppose that X^a is selected to be the reference time series. Then, the fixed reference sub-series is $X^a_{t_i}$ with mean value $\mu^a_{t_i} = \frac{1}{\left|X^a_{t_i}\right|}\sum_{x^a_i \in X^a_{t_i}} x^a_i$, and for $k = 0, \dots, k_{\max}$ we have a collection of comparison sub-series $X^b_{t_i+k\Delta t}$ each with mean $\mu^b_{t_i,k} = \frac{1}{\left|X^b_{t_i+k\Delta t}\right|}\sum_{x^b_i \in X^b_{t_i+k\Delta t}} x^b_i$. We assume that the two time series are sufficiently different, i.e., there exists at least one $k \in \{0, \dots, k_{\max}\}$ such that $\mu^a_{t_i} \neq \mu^b_{t_i,k}$ and $\left|X^a_{t_i} \cap X^b_{t_i+k\Delta t}\right| < N = \left|X^a_{t_i}\right| = \left|X^b_{t_i+k\Delta t}\right|$. Note that N is fixed, since the sub-series length is fixed at w for both $X^a_{t_i}$ and $X^b_{t_i+k\Delta t}$.

Then, for each $k = 0, \dots, k_{\max}$, the resultant directional sliding correlation and its associated p-value is:

$$\begin{aligned} r^{a\to b}_{t_i,k} &= \frac{\sum_{\ell=1}^{N}\left(X^a_{t_i}[\ell] - \mu^a_{t_i}\right)\left(X^b_{t_i+k\Delta t}[\ell] - \mu^b_{t_i,k}\right)}{\sqrt{\sum_{\ell=1}^{N}\left(X^a_{t_i}[\ell] - \mu^a_{t_i}\right)^2}\sqrt{\sum_{\ell=1}^{N}\left(X^b_{t_i+k\Delta t}[\ell] - \mu^b_{t_i,k}\right)^2}} \\ &= C_{X^a_{t_i}} \frac{\sum_{\ell=1}^{N}\left(X^b_{t_i+k\Delta t}[\ell] - \mu^b_{t_i,k}\right)}{\sqrt{\sum_{\ell=1}^{N}\left(X^b_{t_i+k\Delta t}[\ell] - \mu^b_{t_i,k}\right)^2}}, \end{aligned} \tag{7}$$

and

$$p\left(r^{a\to b}_{t_i,k}\right) = 2 \times \mathbb{P}\left(T > \frac{r^{a\to b}_{t_i,k}\sqrt{N-2}}{\sqrt{1-\left(r^{a\to b}_{t_i,k}\right)^2}}\right), \tag{8}$$

with $T \sim \text{StudentT}(N-2)$. Note that if X^a is the reference time series, then for all $k = 0, \dots, k_{\max}$, the term $\sum_{\ell=1}^{N}\left(X^a_{t_i}[\ell] - \mu^a_{t_i}\right) / \sqrt{\sum_{\ell=1}^{N}\left(X^a_{t_i}[\ell] - \mu^a_{t_i}\right)^2} = C_{X^a_{t_i}}$ is constant. On the other hand, if X^b is the reference time series, then at each $k = 0, \dots, k_{\max}$ the directional sliding correlation is:

$$r^{b\to a}_{t_i,k} = C_{X^b_{t_i}} \frac{\sum_{\ell=1}^{N}\left(X^a_{t_i+k\Delta t}[\ell] - \mu^a_{t_i,k}\right)}{\sqrt{\sum_{\ell=1}^{N}\left(X^a_{t_i+k\Delta t}[\ell] - \mu^a_{t_i,k}\right)^2}}, \tag{9}$$

now with $\sum_{\ell=1}^{N}\left(X^b_{t_i}[\ell] - \mu^b_{t_i}\right) / \sqrt{\sum_{\ell=1}^{N}\left(X^b_{t_i}[\ell] - \mu^b_{t_i}\right)^2} = C_{X^b_{t_i}}$ being constant for all k. Suppose first that $k = 0$ is not one of the indices where the two time series differ, which implies that $r^{a\to b}_{t_i,0} = r^{b\to a}_{t_i,0}$ and $C_{X^a_{t_i}} = C_{X^b_{t_i}}$. However, since X^a and X^b are sufficiently different on some k, there must be some $k' = 1, \dots, k_{\max}$ for which $r^{a\to b}_{t_i,k'} \neq r^{b\to a}_{t_i,k'}$. Suppose now that $k = 0$ is one of the indices where the two time series differ, then we are done since symmetry breaks at $k = 0$ with $r^{a\to b}_{t_i,0} \neq r^{b\to a}_{t_i,0}$. □

Appendix B. Idealized TDS behavior with sinusoidal signals

We focus on the behavior of the TDS algorithm applied to *idealized* periodic signals modeled as phase-shifted sinusoids. The motivation for us to do so is twofold: many real-life signals attain some aspect of periodicity (e.g., airport delays can be periodic with respect to the length of an operational day), so understanding and contrasting the behavior of TDS on ideal sinusoidal signals may be important. Secondly, no such theoretical benchmarks were evaluated in Bashan et al. (2012), where the original time delay stability algorithm was developed.

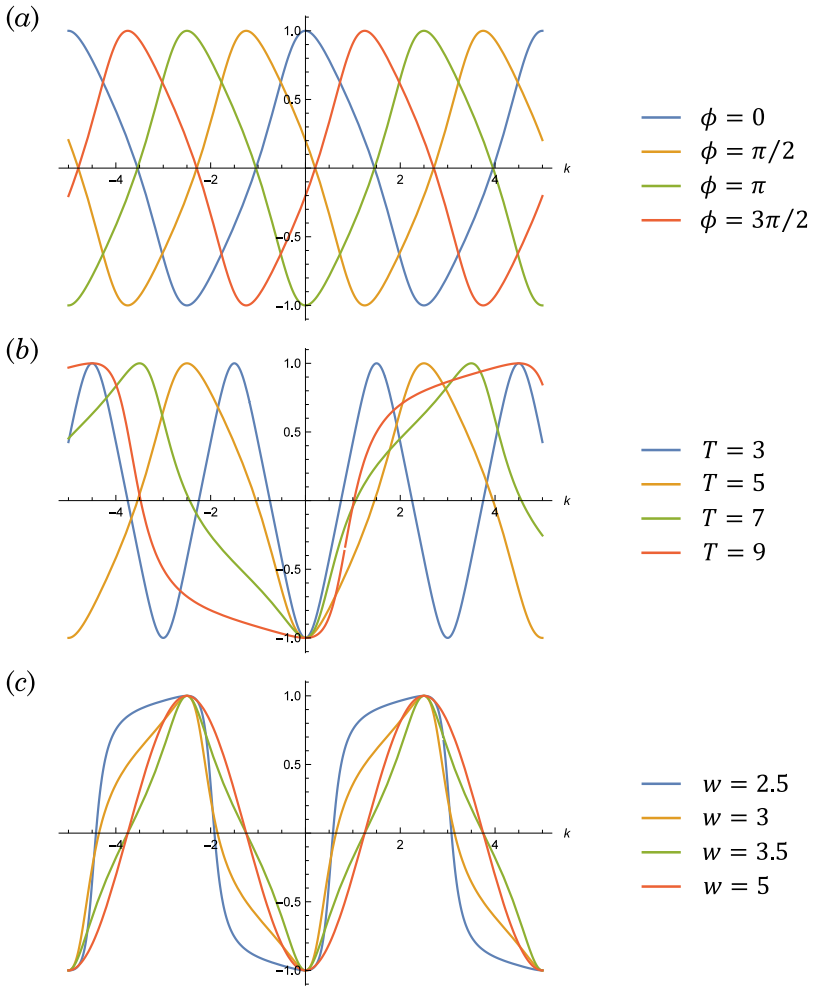

Fig. B.1. Plot of (11) with *(a)* $T = 5, w = 4$, and varying phase parameter ϕ; *(b)* $\phi = \pi, w = 4$, and varying period parameter T; and *(c)* $\phi = \pi, T = 5$, and varying sub-series length w.

Two questions that we formalize in this analysis are as follows: Given two discretized, sinusoidal, and possibly phase-shifted time series signals, when should we expect correlations to be *maximized*, and what are the effects of signal parameters (i.e., period, phase) or TDS parameters (i.e., length of sub-sequence w) on the relationship between the correlation versus lag (i.e., the sub-sequence index k). With respect to the workflow of TDS, we give an answer to the first question in Proposition 2 and provide some insights to the second in Remark 1.

To maintain consistency with the rest of the methodology section, we let time series X^a remain as the reference, and compare it with time series X^b. We assume that both time series are discrete sines with period $|T|$, sampled with respect to index ℓ, and with X^b having a phase shift parameterized by $\phi \in [0, 2\pi)$:

$$X^a[\ell] = \sin\left(\frac{2\pi\ell}{T}\right), \qquad X^b[\ell] = \sin\left(\frac{2\pi\ell}{T} + \phi\right). \tag{10}$$

Following the workflow for TDS described in Section 3, we take the fixed reference sub-series $X^a_{t_i}$ with length $\left|X^a_{t_i}\right| = w$, and without loss of generality set $t_i = 0$. We then examine the directional sliding correlations $r^{b\to a}_{t_i=0,k}$, given by:

$$r^{b\to a}_{t_i=0,k} = \frac{\sum_{\ell=0}^{w-1}\left(X^a_0[\ell] - \mu^a_0\right)\left(X^b_{k\Delta t}[\ell] - \mu^b_{0,k}\right)}{\sqrt{\sum_{\ell=0}^{w-1}\left(X^a_0[\ell] - \mu^a_0\right)^2}\sqrt{\sum_{\ell=0}^{w-1}\left(X^b_{k\Delta t}[\ell] - \mu^b_{0,k}\right)^2}} \tag{11}$$

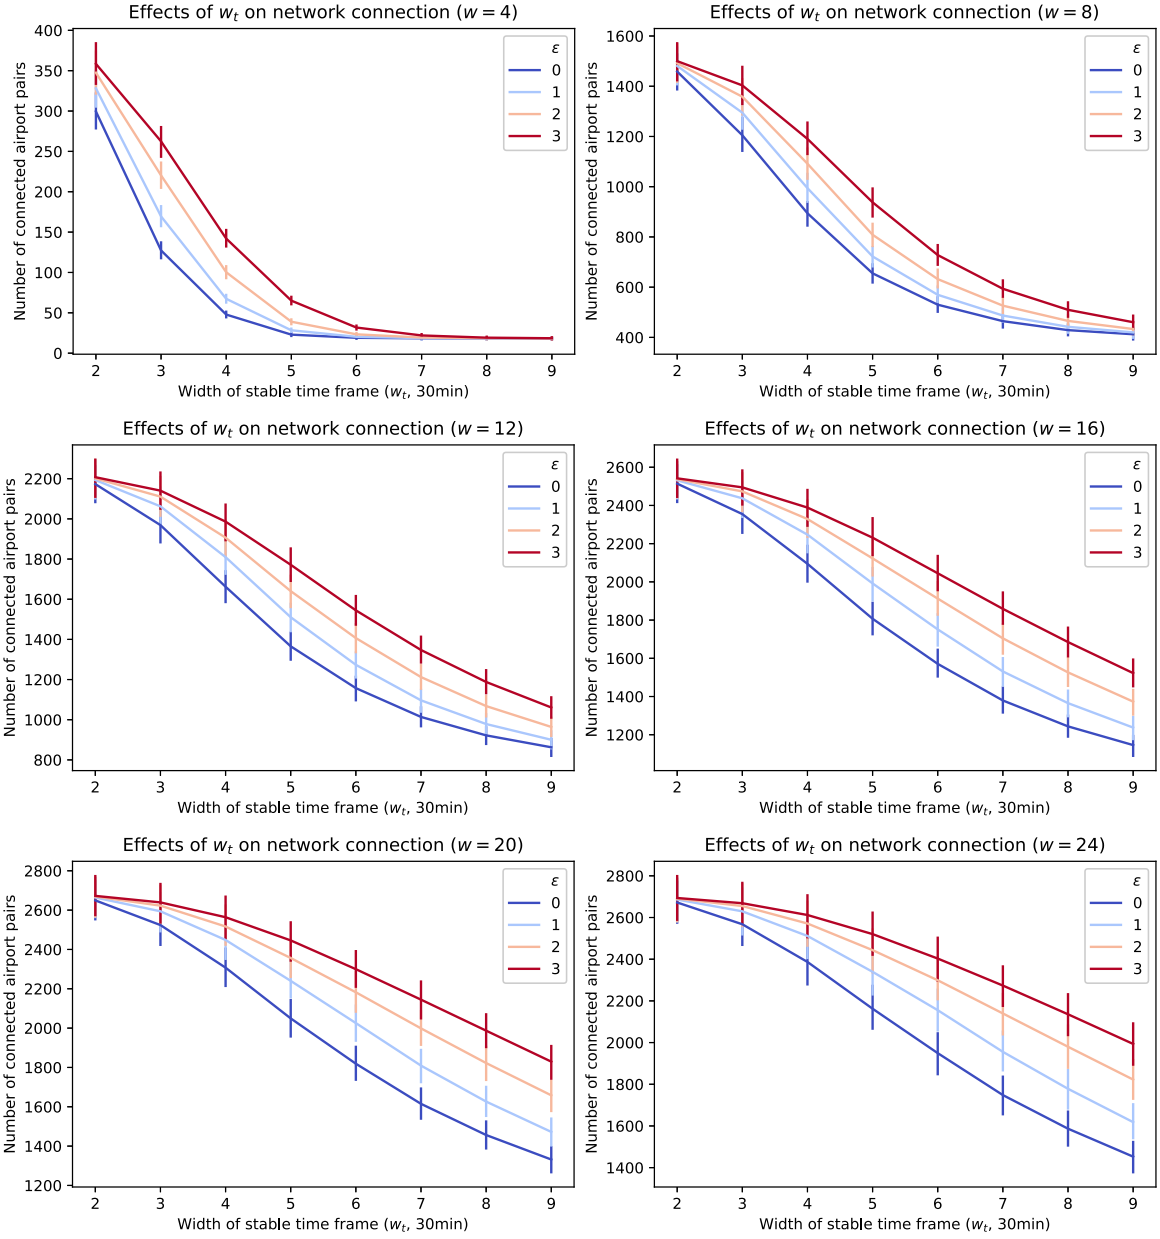

Fig. C.1. Sensitivity of number of edges in DSN (averaged across a day, for all days in 2017) with respect to the maximum allowable lag perturbation ϵ and the minimum stable lag time duration w_t. The sub-series length w is held constant. The error bars denote the 95% confidence interval around the average number of edges.

with sample means

$$\begin{aligned}
\mu_0^a &= \frac{1}{w}\sum_{\ell=0}^{w-1}\sin\left(\frac{2\pi\ell}{T}\right) = \frac{1}{2w}\csc\left(\frac{\pi}{T}\right)\left(\cos\left(\frac{\pi}{T}\right) - \cos\left(\frac{\pi(2w-1)}{T}\right)\right), \\
\mu_{0,k}^b &= \frac{1}{w}\sum_{\ell=0}^{w-1}\sin\left(\frac{2\pi(\ell+k)}{T}+\phi\right) = \frac{1}{w}\csc\left(\frac{\pi}{T}\right)\sin\left(\frac{\pi w}{T}\right)\sin\left(\frac{\pi(2k+w-1)}{T}+\phi\right).
\end{aligned} \tag{12}$$

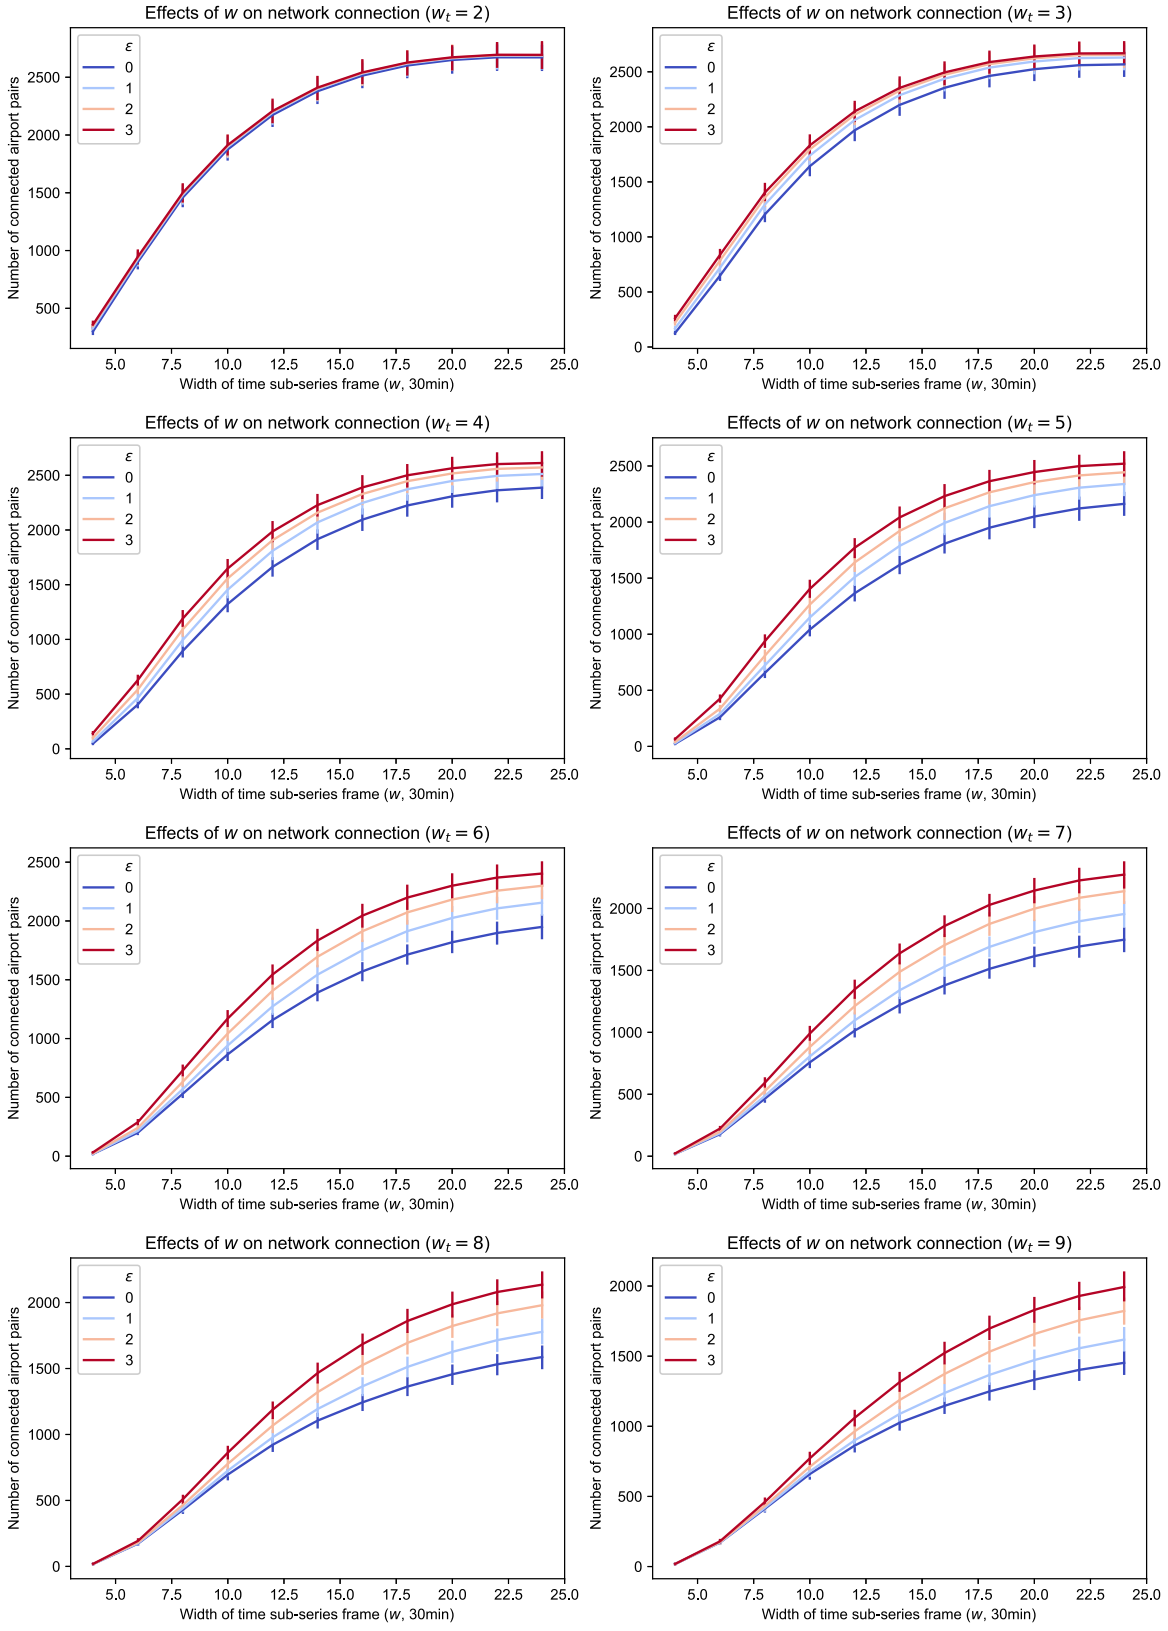

Fig. C.2. Sensitivity of number of edges in DSN (averaged across a day, for all days in 2017) with respect to the maximum allowable lag perturbation ϵ and sub-series length w. The minimum stable lag time duration w_t is held constant. The error bars denote the 95% confidence interval around the average number of edges.

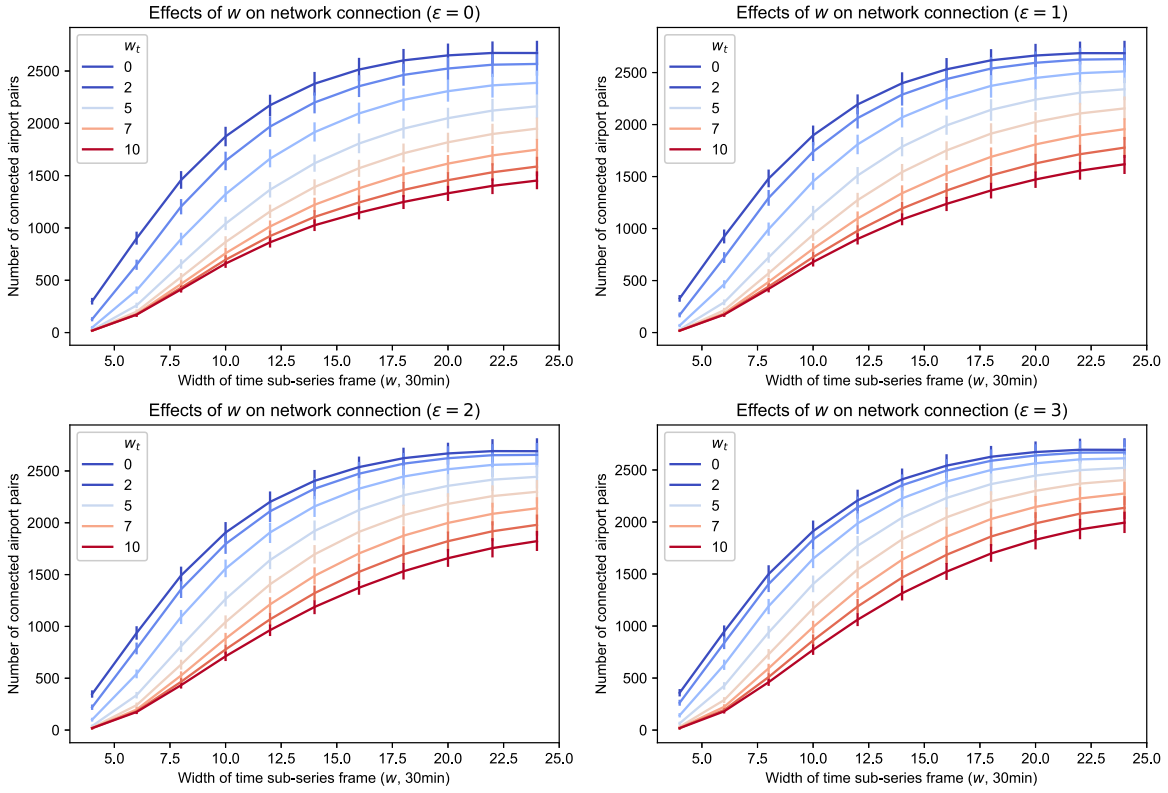

Fig. C.3. Sensitivity of number of edges in DSN (averaged across a day, for all days in 2017) with respect to the minimum stable lag time duration w_t and sub-series length w. The maximum allowable lag perturbation ϵ is held constant. The error bars denote the 95% confidence interval around the average number of edges.

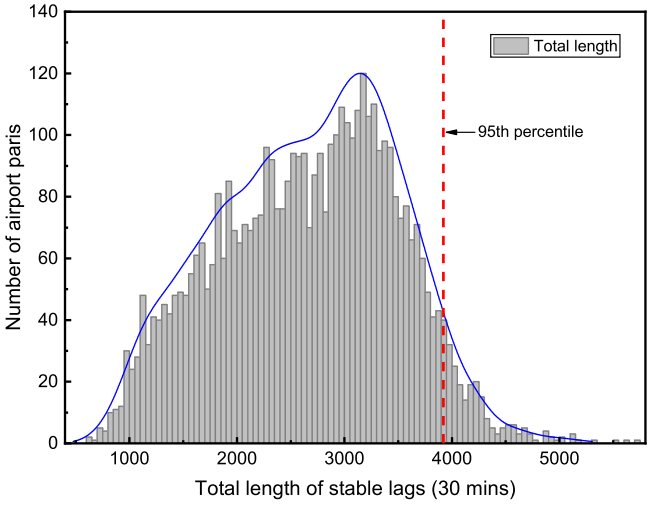

Fig. C.4. Empirical distribution of total length of stable lags (in 30 min) across all the OD pairs (bin size = 50). Blue line is the kernel smooth fitting to the data. Red dash line indicates the 95th percentile of the total length.

Proposition 2. *If the sub-series length is fixed to be $w = cT$ for any positive integer $c \geq 1$, then the directional sliding correlation $r^{b\to a}_{t_i=0,k}$ from (11) is given by*

$$r^{b\to a}_{t_i=0,k} = \cos\left(\frac{2\pi k}{T} + \phi\right), \tag{13}$$

and is maximized at

$$k = \frac{T}{2\pi}\left(2\pi\gamma_1 - \phi\right), \quad \gamma_1 \in \mathbb{Z}. \tag{14}$$

Proof. Assume without loss of generality that $\Delta t = 1$, so we can set $k \overset{\Delta}{=} k\Delta t$. Since w is equal to the period (or integral multiples of the period), the mean value of both sub-series is 0, i.e., $\mu_0^a = \mu_{0,k}^b = 0$. The directional sliding correlation Eq. (11) reduces to:

$$r_{t_i=0,k}^{b\to a} = \frac{\sum_{\ell=0}^{w-1} \sin\left(\frac{2\pi\ell}{T}\right)\sin\left(\frac{2\pi(\ell+k)}{T}+\phi\right)}{\sqrt{\sum_{\ell=0}^{w-1}\sin^2\left(\frac{2\pi\ell}{T}\right)}\sqrt{\sum_{\ell=0}^{w-1}\sin^2\left(\frac{2\pi(\ell+k)}{T}+\phi\right)}}. \tag{15}$$

Using the elementary trigonometric identity $\sin^2(x) = 1/2 - \cos(2x)/2$, and that the sub-series length is $w = cT$ for $c \in \mathbb{N}_{\geq 1}$, the denominator can be rewritten as follows:

$$\begin{aligned}&\sqrt{\sum_{\ell=0}^{w-1}\sin^2\left(\frac{2\pi\ell}{T}\right)\sum_{\ell=0}^{w-1}\sin^2\left(\frac{2\pi(\ell+k)}{T}+\phi\right)}\\ =&\sqrt{\left(\frac{cT}{2}-\frac{1}{2}\sum_{\ell=0}^{cT-1}\cos\frac{4\pi\ell}{T}\right)\left(\frac{cT}{2}-\frac{1}{2}\sum_{\ell=0}^{cT-1}\cos\left(\frac{4\pi(\ell+k)}{T}+2\phi\right)\right)} = \frac{cT}{2}.\end{aligned} \tag{16}$$

We now simplify the finite sum of products of sines in the numerator by using the identity $\sin(\alpha)\sin(\beta) = \cos(\alpha-\beta)/2 - \cos(\alpha+\beta)/2$,

$$\begin{aligned}r_{t_i=0,k}^{b\to a} &= \frac{2}{cT}\sum_{\ell=0}^{cT-1}\sin\left(\frac{2\pi\ell}{T}\right)\sin\left(\frac{2\pi(\ell+k)}{T}+\phi\right)\\ &= \frac{1}{cT}\left(\sum_{\ell=0}^{cT-1}\cos\left(-\frac{2\pi k}{T}-\phi\right) - \sum_{\ell=0}^{cT-1}\cos\left(\frac{2\pi k+4\pi\ell}{T}+\phi\right)\right)\\ &= \frac{1}{cT}\left(cT\cos\left(\frac{2\pi k}{T}+\phi\right) - 0\right) = \cos\left(\frac{2\pi k}{T}+\phi\right).\end{aligned} \tag{17}$$

Note that the second equality comes from cosine being an even function, and the second term vanishes since we are summing over complete periods. The set of k at which $r_{t_i=0,k}^{b\to a} = \cos(2\pi k/T + \phi)$ is maximized can then be found by first solving for k such that

$$\frac{\partial}{\partial k} r_{t_i=0,k}^{b\to a} = \frac{2\pi}{T}\sin\left(\frac{2\pi k}{T}+\phi\right) = 0. \tag{18}$$

Solving this first order condition gives

$$k = \begin{cases}\frac{T}{2\pi}\left(2\pi\gamma_1 - \phi\right), & \gamma_1 \in \mathbb{Z},\\ \frac{T}{2\pi}\left(2\pi\gamma_2 - \phi + \pi\right), & \gamma_2 \in \mathbb{Z}.\end{cases} \tag{19}$$

Using the second derivative test, we have that the second order condition for local maxima is

$$\frac{\partial^2}{\partial k^2} r_{t_i=0,k}^{b\to a} = -\frac{4\pi^2}{T^2}\cos\left(\frac{2\pi k}{T}+\phi\right) < 0. \tag{20}$$

For $k = \frac{T}{2\pi}\left(2\pi\gamma_2 - \phi + \pi\right)$ and $\gamma_2 \in \mathbb{Z}$, the second order condition is violated with $4\pi^2/T^2 > 0$, since T is a positive-valued period. For $k = \frac{T}{2\pi}\left(2\pi\gamma_1 - \phi\right)$ and $\gamma_1 \in \mathbb{Z}$, the second order condition is satisfied with $-4\pi^2/T^2 < 0$. □

Remark 1. Denote by $\mathcal{K}$ the set

$$\mathcal{K} = \underset{k\in\mathbb{R}}{\operatorname{argmax}}\left\{r_{t_i,k}^{b\to a}\right\}, \tag{21}$$

where $r_{t_i,k}^{b\to a}$ is computed from (11). The elements of this set, i.e., the lag k at which the directional sliding correlation is maximized, appear to be invariant under changes in the sub-series length w, as long as period and phase parameters T and ϕ remain constant. We plot realizations of (11) with fixed $\phi = \pi$, $T = 5$, and varying w in Fig. B.1(*c*). As expected, changing the phase parameter ϕ in Fig. B.1(*a*) does not change the shape of (11), and only affects its horizontal position. Changing the period parameter T in Fig. B.1(*b*) distorts the shape of (11). We note that, without assuming any relationships between ϕ, T, and w, the full form of (11) is extremely complicated. Our simple expression derived in Proposition 2 removes a degree of freedom by setting $w = cT$ for some positive integer $c \geq 1$.

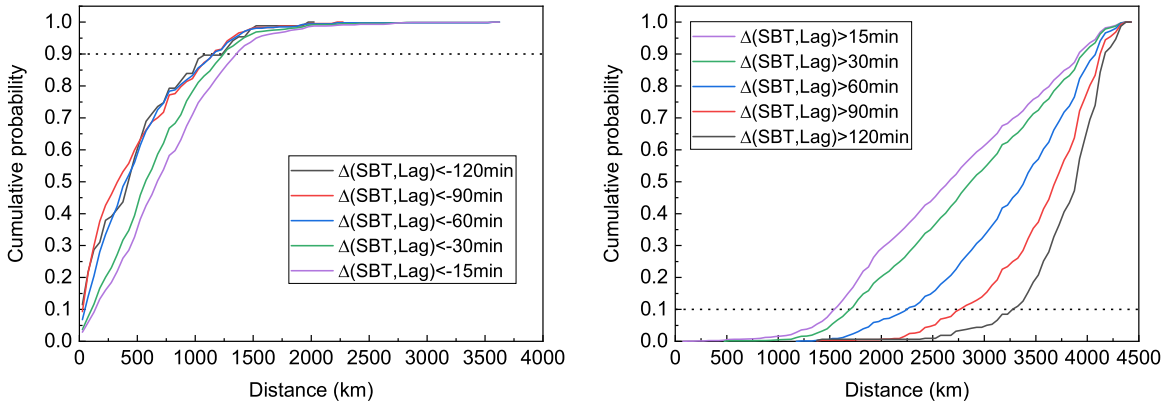

Fig. D.1. Empirical cumulative distribution functions (CDFs) of distance between OD pairs, given that the Δ(SBT, Lag) for that OD pair is greater than $\Delta_{\min}$ or lesser than $-\Delta_{\min}$, with $\Delta_{\min} \in \{15, 30, 60, 90, 120\}$ minutes. The 90th percentile of the empirical cdf for $\Delta(\text{SBT, Lag}) < -\Delta_{\min}$ demarcates an *upper cutoff distance* (UCD), whereas the 10th percentile of the empirical cdf for $\Delta(\text{SBT, Lag}) > \Delta_{\min}$ demarcates a *lower cutoff distance* (LCD).

Table D.1
Upper cutoff distance (UCD) and lower cutoff distance (LCD) corresponding to the 90th and 10th percentiles of the empirical cdfs for $\Delta(\text{SBT, Lag}) \in (-\infty, -\Delta_{\min}) \cup (\Delta_{\min}, \infty)$ in Fig. D.1, along with the number of OD pairs below the UCD or above the LCD.

$\Delta_{\min}$ (min)	$\Delta(\text{SBT, Lag}) < -\Delta_{\min}$		$\Delta(\text{SBT, Lag}) > \Delta_{\min}$	
	# ODs	UCD, 90th perc. (km)	# ODs	LCD, 10th perc. (km)
15	1302	1370	1906	1566
30	1006	1266	1603	1723
60	532	1177	1015	2280
90	237	1174	576	2789
120	87	1139	318	3307

Appendix C. Sensitivity of TDS parameters

See Figs. C.1–C.4.

Appendix D. The distance between airports as a function of Δ(SBT, lag)

To further investigate how distance between airports may affect Δ(SBT, lag), we plot the Cumulative Distribution Functions (CDFs) of the distance between OD pairs for ten different cases (i.e., ten partitions of Δ(SBT, Lag) values) in Fig. D.1. The left part of Fig. D.1 shows the scenarios where the speed of delay propagation is much slower than SBT, i.e., $\Delta(\text{SBT, lag}) \ll 0$ (here we refer as *slower cases*), while the right part shows the cases with $\Delta(\text{SBT, lag}) \gg 0$ (*faster cases*).

Table D.1 presents the detail of upper cutoff distance (UCD) and lower cutoff distance (LCD) as well as the number ODs for all the ten cases. We observe that the UCDs of the slower cases are within the range of $[1000\ \text{km}, 1500\ \text{km}]$. Over 90% OD pairs is less than 1300 km. In contrast, the LCDs for faster cases cover the range from 1500 km to 3300 km. When two airports are separated by more than 3300 km, delay propagates much faster ($\Delta(\text{SBT, lag}) > 120$ min).

References

AhmadBeygi, S., Cohn, A., Lapp, M., 2010. Decreasing airline delay propagation by re-allocating scheduled slack. IIE Trans. 42 (7), 478–489.
Anon., 2016. 14 CFR, Part 234 - Airline Service Quality Performance Reports, Code of Federal Regulations. United States.
Anon., 2020. Aviation system performance metrics (ASPM). https://aspm.faa.gov/apm/sys/main.asp. (Accessed: 30 September 2020).
Arıkan, M., Deshpande, V., Sohoni, M., 2013. Building reliable air-travel infrastructure using empirical data and stochastic models of airline networks. Oper. Res. 61 (1), 45–64.
Bao, J., Yang, Z., Zeng, W., 2021. Graph to sequence learning with attention mechanism for network-wide multi-step-ahead flight delay prediction. Transp. Res. C 130, 103323.
Bashan, A., Bartsch, R.P., Kantelhardt, J.W., Havlin, S., Ivanov, P.C., 2012. Network physiology reveals relations between network topology and physiological function. Nature Commun. 3, 702–709.
Baspinar, B., Tutku Altun, A., Koyuncu, E., 2021. Event-based air transport network resiliency management with meta-population epidemic model. J. Aerosp. Inf. Syst. 18 (9), 632–644.
Beatty, R., Hsu, R., Berry, L., Rome, J., 1999. Preliminary evaluation of flight delay propagation through an airline schedule. Air Traffic Control Q. 7 (4), 259–270.

Boers, N., Goswami, B., Rheinwalt, A., Bookhagen, B., Hoskins, B., Kurths, J., 2019. Complex networks reveal global pattern of extreme-rainfall teleconnections. Nature 566 (7744), 373–377.
Breiman, L., 2001. Random forests. Mach. Learn. 45 (1), 5–32.
Brueckner, J.K., Czerny, A.I., Gaggero, A.A., 2021. Airline mitigation of propagated delays via schedule buffers: Theory and empirics. Transp. Res. E 150, 102333.
Cong, W., Hu, M., Dong, B., Wang, Y., Feng, C., 2016. Empirical analysis of airport network and critical airports. Chin. J. Aeronaut. 29 (2), 512–519.
Crainic, T.G., Gendreau, M., Gendron, B., 2021. Network Design with Applications to Transportation and Logistics. Springer, Cham.
David, E., Jon, K., 2010. Networks, Crowds, and Markets: Reasoning About a Highly Connected World. Cambridge University Press, USA.
Diana, T., 2009. Do market-concentrated airports propagate more delays than less concentrated ones? A case study of selected U.S. airports. J. Air Transp. Manag. 15 (6), 280–286.
Du, W.-B., Zhang, M.-Y., Zhang, Y., Cao, X.-B., Zhang, J., 2018. Delay causality network in air transport systems. Transp. Res. E 118, 466–476.
Fleurquin, P., Ramasco, J.J., Eguiluz, V.M., 2013. Systemic delay propagation in the US airport network. Sci. Rep. 3.
Guimerà, R., Mossa, S., Turtschi, A., Amaral, L., 2005. The worldwide air transportation network: Anomalous centrality, community structure, and cities' global roles. Proc. Natl. Acad. Sci. USA 102 (22), 7794–7799.
Guo, Z., Hao, M., Yu, B., Yao, B., 2022. Detecting delay propagation in regional air transport systems using convergent cross mapping and complex network theory. Transp. Res. E 157, 102585.
Jiang, X., Shen, Y., Yao, J., Zhang, L., Xu, L., Feng, R., Cai, L., Liu, J., Chen, W., Wang, J., 2019. Connectome analysis of functional and structural hemispheric brain networks in major depressive disorder. Transl. Psychiatry 9 (1), 1–12.
Kafle, N., Zou, B., 2016. Modeling flight delay propagation: A new analytical-econometric approach. Transp. Res. B 93, 520–542. http://dx.doi.org/10.1016/j.trb.2016.08.012.
Khan, W.A., Ma, H.-L., Chung, S.-H., Wen, X., 2021. Hierarchical integrated machine learning model for predicting flight departure delays and duration in series. Transp. Res. C 129, 103225.
Kim, M., Park, S., 2021. Airport and route classification by modelling flight delay propagation. J. Air Transp. Manag. 93, 102045.
Li, M.Z., Gopalakrishnan, K., Pantoja, K., Balakrishnan, H., 2021. Graph signal processing techniques for analyzing aviation disruptions. Transp. Sci. 55 (3), 553–573.
Li, Q., Jing, R., 2021. Characterization of delay propagation in the air traffic network. J. Air Transp. Manag. 94, 102075.
Lines, D.A., 2021. Delta shuttle. https://www.delta.com/content/www/en_US/traveling-with-us/where-we-fly/flight-partners/delta-shuttle.html/.
Peterson, M.D., Bertsimas, D.J., Odoni, A.R., 1995a. Decomposition algorithms for analyzing transient phenomena in multiclass queueing networks in air transportation. Oper. Res. 43 (6), 995–1011.
Peterson, M.D., Bertsimas, D.J., Odoni, A.R., 1995b. Models and algorithms for transient queueing congestion at airports. Manage. Sci. 41 (8), 1279–1295.
Pyrgiotis, N., Malone, K.M., Odoni, A., 2013. Modelling delay propagation within an airport network. Transp. Res. C 27, 60–75.
Bureau of Transportation Statistics, 2021. Airline on-time statistics and delay causes. https://www.transtats.bts.gov/OT_Delay/OT_DelayCause1.asp.
van den Heuvel, M.P., Hulshoff Pol, H.E., 2010. Exploring the brain network: A review on resting-state fMRI functional connectivity. Eur. Neuropsychopharmacology 20 (8), 519–534.
Wang, C., Hu, M., Yang, L., Zhao, Z., 2021. Prediction of air traffic delays: An agent-based model introducing refined parameter estimation methods. PLoS One 16 (4), 1–22.
Wu, C.-L., Law, K., 2019. Modelling the delay propagation effects of multiple resource connections in an airline network using a Bayesian network model. Transp. Res. E 122, 62–77.