Seeing the Trees for the Forest: Disentangling Delay Propagation in Airport Networks through Delay Propagation Trees

Bruno Campanelli, Riccardo Gallotti and José Ramasco

Instituto de Física Interdisciplinar y Sistemas Complejos IFISC (UIB-CSIC),
Campus UIB, 07122 Palma de Mallorca, Spain

During the last century, air transportation has grown steadily and faster than any other transportation mode, with only a minor setback in the wake of 9/11 attacks [1, 2]. The yearly cost of flight delays is measured in billions dollars/euros in the US and Europe [3, 4], and various socio-economical factors contribute to the difficulty of providing adequate infrastructural support, for example in the case of airports located in urban areas [5]. Having a thorough understanding of the processes regulating the spreading of disruption and delays in airport networks is therefore an important research goal, which is also interesting from a more theoretical point of view as these systems typically exhibit rich, complex behavior, and can be the substrate for phenomena not strictly related to air transportation, such as epidemic spreading. This type of study is made difficult not only by the domain-specific knowledge required, but also because of the scarce availability of high-quality data about flight performance outside of the US from institutional sources. In order to overcome this problem, we use data obtained from the Flightradar24 web service, which covers a period of almost two years.

In this work, we analyze daily traffic for both the US and European systems, with the goal of finding and classifying interesting, recurring patterns which characterize the days where significant perturbations took place. Of particular importance are the aircraft returning to a large hub multiple times during the day (especially in the European system, where airlines have a smaller number of operating bases), and/or depart earlier in the morning. Furthermore, the delay propagation trees can be reconstructed from the data, giving insights into the delay propagation patterns. In order to do this, we define a square lattice where each cell represents a single airport during a time window of fixed size. A directed link between two cells is established if there is at least one delayed flight departing from one cell and arriving in the other, and the weight is given by the total amount of delay. An example of this is shown in Fig. 1, where O'Hare International Airport (ORD) starts propagating delay early in the morning before receiving any significant input, from which we can infer that the delay was generated on-site. Fig. 2 shows the daily incoming and outgoing delays for the largest airports in the network during the same day. ORD is shown having both very high amounts of incoming and outgoing delay, as airports suffering from severe problems can request incoming flights to be delayed (for example to wait for better metheorological conditions). This exemplifies that the fact that delay is being trasmitted from one place to another is not enough to draw conclusions about causality, making an analysis of the topology of connections necessary.

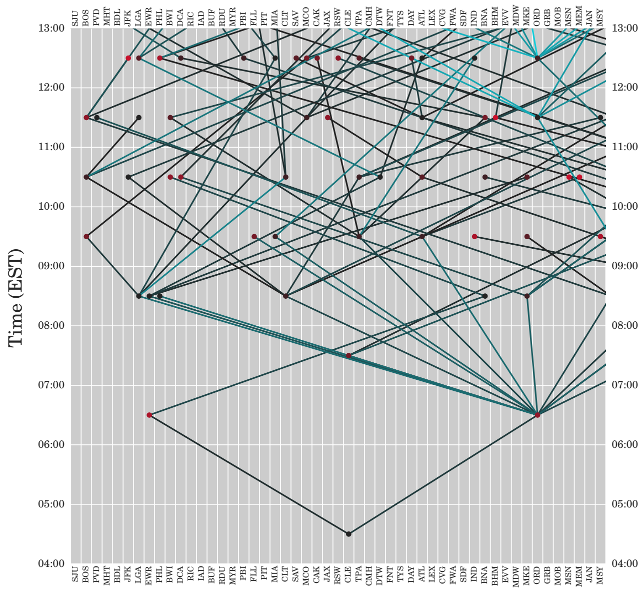

Figure 1: The beginning of a delay propagation tree from 23rd of March 2015. Brighter edges transfer higher delay, brighter nodes produce more delay than they receive.

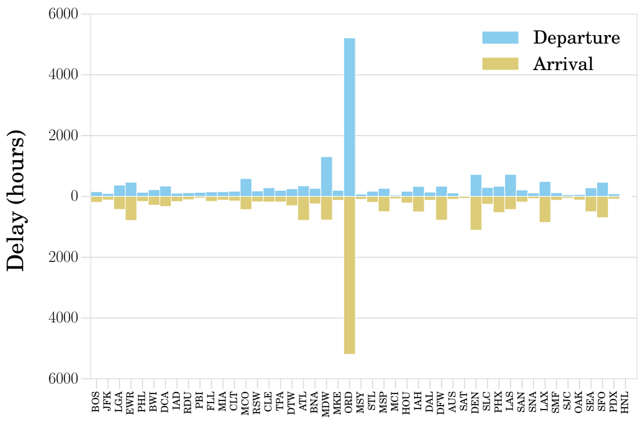

Figure 2: Daily input and output delay for the largest (in terms of movements) airports in the US in 23rd of March 2015.

[1] T.A. Heppenheimer, Turbulent Skies: The History of Commercial Avia- tion (Wiley, 1998).

[2] ICAO Air Navigation Report 2016 Edition

[3] Cook A. and Tanner G., European airline delay cost reference values, Performance Review Unit EUROCONTROL, 2011

[4] Join Economic Committee of US Congress, Your flight has been delayed Again: Flight delays cost passengers, airlines and the US economy billions.

[5] ITF (2014), Expanding Airport Capacity in Large Urban Areas, OECD Publishing, Paris. DOI: http://dx.doi.org/10.1787/9789282107393-en