MIT Libraries | DSpace@MIT

MIT Open Access Articles

Estimation of Arrival-Departure Capacity Tradeoffs in Multi-Airport Systems

The MIT Faculty has made this article openly available. ***Please share*** how this access benefits you. Your story matters.

As Published: http://dx.doi.org/10.1109/CDC.2009.5400462

Publisher: Institute of Electrical and Electronics Engineers

Persistent URL: http://hdl.handle.net/1721.1/58731

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Attribution-Noncommercial-Share Alike 3.0 Unported

Massachusetts Institute of Technology

Estimation of Arrival-Departure Capacity Tradeoffs in Multi-Airport Systems

Varun Ramanujam and Hamsa Balakrishnan

***Abstract*— The accurate estimation of airport capacity is critical for the efficient planning of landing and takeoff operations, and the mitigation of congestion-induced delays. The analysis of tradeoffs between arrival and departure capacity at an airport, represented by the airport capacity envelope, has been the main focus of prior research. The increasing demand for air traffic operations has resulted in the growth of multi-airport systems, in which several major airports that are in close proximity of each other serve the same geographical region. The arrival and departure flows into these airports interact with each other, and it is necessary to consider inter-airport arrival-departure capacity tradeoffs while scheduling operations.**

This paper proposes a statistical technique based on quantile regression, for systematically analyzing arrival-departure capacity tradeoffs in multi-airport systems using observations of flight operations. The proposed technique enables the identification of key factors (such as, runway configuration geometry, weather conditions, etc.) that influence both the capacity envelopes of individual airports, and the capacity envelope of the multi-airport system as a whole. The approach is demonstrated through an analysis of the capacity envelopes of the New York area multi-airport system (comprising Newark (EWR), John F. Kennedy (JFK) and LaGuardia (LGA) airports).

I. Introduction

The safe and efficient planning of airport and terminal-area operations is essential for meeting the predicted increase in demand in the Next Generation Air Transportation System (NextGen) without incurring unacceptably large delays. The effects of congestion are beginning to be seen even in today's system: in the United States, between 2006 and 2007, there was a 30% rise in delays due to terminal-area volume, while there was only a 1% increase in traffic [4,6]. A critical step to meeting the expected increase in demand is through the improved utilization of airport capacity, especially in congested terminal-areas. With the growth of secondary and even tertiary airports in the most congested regions of the country (for example, the New York area, the San Francisco Bay area, the Los Angeles area, Boston and Chicago [3]), the problem of coordinating operations in multi-airport systems to use terminal-area capacity more efficiently has also become increasingly important.

The operating capacity of an airport is given by its arrival capacity (the number of aircraft landings per hour) and its departure capacity (the number of departures per hour). Due to the shared nature of ground resources such as runways and taxiways, there is a tradeoff between the arrival and departure capacity at an airport [2,14]. An airport capacity envelope is the boundary (generally approximated as a convex polygon on the plane with the arrival and departure rates as axes) that defines the envelope of the maximum capacities that can be achieved under specified operating conditions. Fig. 1 illustrates a representative capacity envelope for an airport, for a given runway configuration. The capacity envelope describes the capacity available for a runway configuration under a specific set of conditions, and captures the tradeoff between the maximum arrival and departure rates [5]. The extent of the tradeoff depends on factors such as the relative alignment of the active runways, meteorological factors like wind and visibility, as well as the aircraft fleet mix.

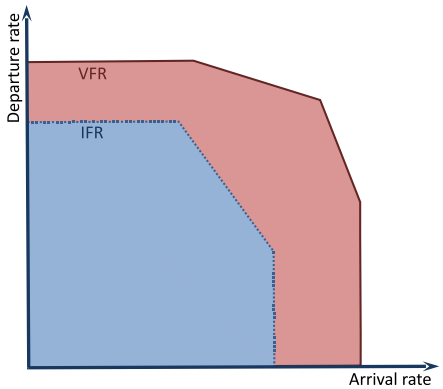

Fig. 1. Illustration of capacity envelope for an airport, under a particular runway configuration, for different meteorological conditions: (1) Visual Flight Rules (VFR), and (2) Instrument Flight Rules (IFR). The shaded regions represent the feasible operating points.

The emergence of several core, secondary and regional airports (known collectively as a *metroplex* [15]) in an already congested airspace results in the complex interaction of traffic flows. The interactions between arrival and departure operations at proximate airports could potentially make it infeasible to simultaneously operate all the airports at their individual optimal runway configurations. This implies that the resource allocation and schedule optimization at an airport would have to not only take into consideration the arrival-departure tradeoffs for that airport, but also the tradeoffs with arrival and departure capacities at the other airports within the multi-airport system. While there is anecdotal and descriptive evidence of such interactions between air traffic flows into and out of neighboring airports, there have been no attempts to quantify these interactions and their impact on capacity using operational data.

A detailed understanding of airport capacity tradeoffs and their dependence on external factors such as operating conditions and airport layout, both in the single- and multi-

Varun Ramanujam is a graduate student at the department of Civil and Environmental Engg. at Massachusetts Institute of Technology, Cambridge, MA, USA. varunr@mit.edu.

Hamsa Balakrishnan is an assistant professor in the department of Aeronautics and Astronautics at Massachusetts Institute of Technology. hamsa@mit.edu.

airport settings, is necessary for the efficient utilization of airport capacity. While these tradeoffs have not been studied so far in the multi-airport setting, the estimation of single-airport capacity envelopes has traditionally been restricted to theoretical approaches [2,11,12]. Empirical methods that use operational data have only been employed in recent times, but have generally been *ad hoc* approaches [10]. This paper proposes a systematic statistical inference technique, based on quantile regression, to estimate piecewise-linear, convex capacity envelopes using observed throughput data, for both individual airports as well as the corresponding inter-airport dependencies in a multi-airport system. The problem of estimating the capacity envelopes as quantiles (or percentiles) of the reported data is reduced to the solution of a linear program. The proposed technique is illustrated using reported data on arrivals and departures at the three major New York area airports, namely John F. Kennedy (JFK), Newark (EWR) and LaGuardia (LGA), for the years 2006 and 2007. The analysis identifies the key factors that influence the intra- and inter-airport capacity tradeoffs, and determines the associated capacity envelopes.

II. Related work

Airport capacity is affected by various external factors such as, pilot and air traffic controller procedures, approach and departure speeds, runway and taxiway occupancy times, weather, etc. Theoretical approaches to capacity estimation have traditionally modeled these factors through simplified stochastic models of aircraft behavior, and estimated the capacity using the expected separation time between successive aircraft operations [2,11,12,19,22,23]. These models construct the capacity envelope through linear interpolation between capacity estimates estimated at specific arrival/departure mix ratios. Newell [20] and Odoni et al. [21] provide comprehensive overviews of analytical and simulation methods that adopt the above approach.

Empirical estimation approaches have the potential to reflect the practical operating capacity envelopes at airports, as opposed to idealized models of capacity tradeoffs. Recognizing this, Gilbo [10] proposed a quasi-statistical procedure for estimating the capacity envelope of a single airport, for a given runway configuration. 15-min arrival and departure counts were used to estimate the capacity envelope as the convex hull (polytope) of the planar scatter of the counts. Frequency-based filtering was employed on the outer perimeter of the data to eliminate outliers, and alternative outlier rejection criteria such as proximity to neighboring observations and rank order statistics were also mentioned. In contrast, this paper proposes a systematic statistical approach to estimating intra- and inter-airport capacity envelopes from observed data.

On the statistical estimation front, *quantile regression* techniques attempt to determine statistics such as the median or a general percentile of the conditional distribution of a response variable which is a function of observed variables [17]. For example, Bernini et al. modeled the production frontier in classical economics as a higher-order quantile ($90-100$%ile), and studied the interaction between underlying determinants at intermediate quantiles [1]. We extend these techniques to the case of airport capacity envelope estimation; in particular, our approach is suitable for determining piecewise-linear, concave quantile functions.

III. Problem statement

Given the counts of arrival and departure throughput at an airport per time interval (say, 15 minutes), we would like to determine the capacity envelope by suitably excluding outliers; we would also like to identify key factors that affect the capacity tradeoffs; finally, we would like to develop a technique that is capable of analyzing both intra- and inter-airport capacity envelopes.

Our approach to solving this problem draws from the field of quantile regression. A data point is said to be at the τth quantile (or $(100-\tau)^{\text{th}}$ percentile) if it is larger than a proportion $\tau/100$ of the data points, and less than a proportion $(1-\tau/100)$ [17]. Similar to least-square regression techniques that estimate the mean of a response variable given values of the predictor variables, quantile regression techniques estimate other statistics, such as the median or a quantile [17]. Since the airport capacity envelopes represent the upper limits of operating capacity, quantile regression techniques, with τ sufficiently large, are suitable mechanisms for estimating them. In other words, if $\tau = 99.5$, we would like to draw the capacity envelope such that 99.5% of all reported operating points fall within the feasible region.

The quantile (value of τ) chosen to represent the airport capacity is conceptually similar to the frequency-based filter adopted by Gilbo [10] to discard spurious data (for example, reporting errors). While quantile regression has traditionally been used to determine linear quantile functions, in the case of airport capacity envelopes we would like to determine a piecewise-linear, concave, continuous function that represents the quantile. We show that the estimation of parameters of such a quantile function can be conducted by solving a linear program. In addition to airport-specific capacity envelopes, we use the proposed approach to study the presence and magnitude of tradeoffs between operations at neighboring airports.

IV. Model description and solution framework

In this section, we describe the modeling of capacity envelopes, and the formulation of a linear optimization problem to estimate them.

A. Data sources

The FAA's Aviation System Performance Metrics (ASPM) database provides records of flight activity at 77 of the major airports in the United States [7]. For each airport, for every 15-min interval, the database includes reports of the number of arrivals, the number of departures, prevailing weather conditions (Visual or Instrument Flight Rules, wind speed and direction), and the runway configuration that was used. This paper focuses on the three major New York area airports. Overnight periods of operation (midnight-6AM) and

15 minutes before and after runway configuration changes are filtered out from the analysis, due to the increased tendency for reporting errors during these periods.

B. Capacity envelopes

The capacity envelope representation is decomposed into two parts: an arrival rate threshold and a departure capacity function. The arrival rate is treated as the independent variable, and the departure capacity for any time interval is estimated as a function of the arrival rate using quantile regression. This decision is based on the observation that arrivals are generally given priority at airports over departures. It is therefore reasonable to assume that while the arrival rate depends on the demand, departure capacity is traded off to accommodate arrivals. In other words,

$$\text{Departure capacity} = f(\text{arrival rate}). \tag{1}$$

The arrival capacity is defined as the maximum number of arrivals that can be accommodated in a time period. It is also called the *unhindered arrival capacity* since it spans operating points at which the presence of departures do not reduce the arrival rate. It corresponds to the *arrivals-only point* on the capacity envelope [5], and is estimated independently of the departure capacity.

Since arrival and departure counts are reported for 15-minute intervals, the counts are used as measures of the arrival and departure rates. Since the capacity envelope forms a convex polygon, the objective of this paper is to determine a piecewise-linear, concave quantile function for the departure capacity that reflects the tradeoffs with arrival rate, for the range of observed arrival rates. Recent research on piecewise quantile regression models have assumed no knowledge of the potential locations of the breakpoints where the slope of the quantile function changes [9,16]. In contrast, we leverage knowledge of the structure of the capacity envelope to estimate the piecewise linear function. There is a finite number of potential breakpoints on the capacity envelope due to the fact that the arrival count (in a 15-minute interval) is a nonnegative integer, and is bounded (typically less that 20). This allows the estimation of the linear segments over all unit intervals of arrival count with minimal computational effort. The piecewise-linear capacity envelope then has the following form:

$$Q_\tau(y|x) = \sum_i \alpha_k^i \theta^i + (\sum_i \beta_k^i \theta^i)x, \text{ for } (k-1) \le x \le k, \ \forall k \tag{2}$$

where
y is the departure count,
$x \in \{0, 1, \cdots, x_{\max}\}$ is the arrival count,
$Q_\tau(y|x)$ is the τ-quantile function of y with respect to x, which represents the capacity envelope,
$k \in \{1, \ldots x_{max}\}$ denotes the kth interval of the arrival count,
$\theta^i \in \theta$ are the factors influencing the capacity envelope (for example, VFR/IFR conditions or runway alignment), and
α_k^i and β_k^i are the intercept and slope contributions of the factor θ^i, for the kth linear segment.

C. Formulation of estimation problem

Given a data set of N operating observations involving y, x and θ at a given airport, the process of estimating the piecewise coefficients (α_k, β_k) for a chosen quantile τ involves solving the following linear optimization problem:

$$\text{Minimize} \quad \sum_{n=1}^{N} Z_n \tag{3}$$

subject to:

$$Z_n \ge y_n - \left[\sum_i \alpha_k^i \theta_n^i + (\sum_i \beta_k^i \theta_n^i)x_n\right] \qquad \text{if } k-1 \le x_n \le k, \forall n \tag{4}$$

$$Z_n \ge \omega_\tau \left[\sum_i \alpha_k^i \theta_n^i + (\sum_i \beta_k^i \theta_n^i)x_n - y_n\right] \qquad \text{if } k-1 \le x_n \le k, \forall n \tag{5}$$

$$\sum_i \beta_{k+1}^i \theta^{i,m_i} \le \sum_i \beta_k^i \theta^{i,m_i}, \forall k \in \{1 \cdots x_{\max} - 1\}, \quad \forall m_i \in \{\min, \max\}, \ \forall i \tag{6}$$

$$\sum_i \beta_1^i \theta^{i,m_i} \le 0, \quad \forall m_i \in \{\min, \max\}, \ \forall i \tag{7}$$

$$\sum_i \alpha_k^i \theta^{i,m_i} + (\sum_i \beta_k^i \theta^{i,m_i})k = \sum_i \alpha_{k+1}^i \theta^{i,m_i} + (\sum_i \beta_{k+1}^i \theta^{i,m_i})(k) \quad \forall k \in \{1, 2, \ldots, x_{\max} - 1\}, \ m_i \in \{\min, \max\} \ \forall i \tag{8}$$

In the above formulation, constraints (4)-(5) define the absolute deviation of the estimated quantile function from the observed value for y. Constraints (6)-(7) help ensure concavity and non-positivity of the piecewise slope estimates for all values of $\theta^i \in \theta$ in the range $[\theta^{i,\min}, \theta^{i,\max}]$. This feature is motivated by the observation that the magnitude of tradeoffs monotonically increases with arrival rate. Constraints (8) ensure continuity of adjacent segments. $\omega_\tau = (100 - \tau)/\tau$ is the asymmetric weight applied upon the negative deviations, where τ denotes the order of the quantile estimated [18]. The size of the formulation depends on the number of observations in the data set (N), the maximum arrival count in a 15-minute interval ($x_{\max}$) and the number of factors considered (that is, the size of θ).

The unhindered arrival (or arrivals only) capacity can be estimated using a simplified version of the above formulation. The data set is restricted to observations where the arrival rate is not impeded by departures. The unhindered arrival capacity is estimated by solving the following linear program:

$$\text{Minimize} \quad \sum_{n=1}^{N} Z_n \tag{9}$$

subject to:

$$y_{n,\min}^{cap} = \sum_i \alpha_{x_{\max}}^i \theta_n^i + (\sum_i \beta_{x_{\max}}^i \theta_n^i)x_{\max} \tag{10}$$

$$\left.\begin{aligned} Z_n &\ge x_n - \textstyle\sum_i \gamma^i \phi_n^i \\ Z_n &\ge \omega_\tau[\textstyle\sum_i \gamma^i \phi_n^i - x_n] \end{aligned}\right\}, \text{ if } y_n \le y_{n,\min}^{cap} \tag{11}$$

$$Z_n \ge 0, \quad \text{otherwise} \quad \forall n \tag{12}$$

where γ_k^i and ϕ^i are the counterparts of α_k^i and θ^i in Equation (2). The lowest value of the departure capacity ($y_{\min}^{cap}$) is realized at the highest observed arrival count ($x_{\max}$) owing to the concavity assumption.

The choice of quantile τ for the unhindered arrival capacity estimate is independent of its counterpart for the departure capacity function. The principle governing this choice in both cases is effective outlier elimination [10]. The choice is determined by iterating over a progression of quantiles descending from 100%ile at chosen step sizes (set at 0.25%ile in this study), until stable functional parameter estimates are obtained. In other words, the choice of τ was such that the capacity envelope did not change significantly for a small change in τ. The statistical significance of each incremental vector of features (θ and ϕ) is ascertained through quantile likelihood ratio tests [18].

D. Comparison of proposed method with frequency-based filtering

As pointed out earlier, the chosen quantile (τ) for representing the capacity envelope in the proposed approach is notionally similar to frequency-based filtering. However, in contrast to frequency-based filtering, a regression-based approach has the ability to quantify underlying factors that influence the capacity curve, through hypothesis testing over a range of specifications. For example, in addition to estimating capacity envelopes for each runway configuration, the proposed approach can identify the characteristics of the configuration (such as, the angle between the active arrival runways) that impact the capacity envelope to a statistically significant extent. Also, the proposed LP-based estimation framework will consider factors such as the proximity of a point to the quantile function, in addition to its observation frequency, when discarding points, thereby simultaneously accounting for two of the potential outlier rejection criteria suggested by Gilbo [10]. This property is enabled by the concavity constraints (6)-(7) of the LP formulation for estimating the capacity quantle function, and is illustrated in Figure 2.

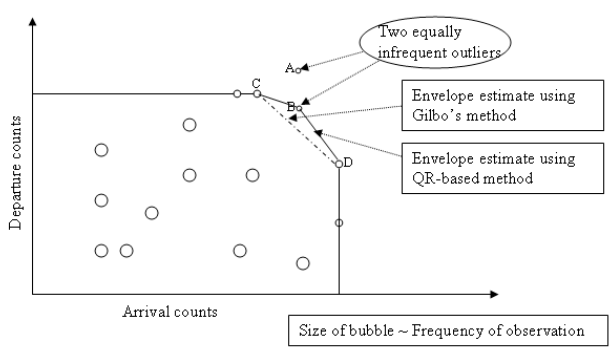

Fig. 2. Outlier elimination: Comparison of proposed approach with Gilbo's frequency-based filtering method

In Figure 2, frequency-based filtering would eliminate operating points A and B owing to their infrequency, while the quantile regression methodology could retain point B due to its conformance with the concave shape of the capacity envelope, as determined by the neighboring high-frequency observations (C and D).

V. Case study: The New York multi-airport system

The proposed formulations were applied to obtain capacity envelope estimates for the three major airports in the New York region: JFK, LGA and EWR. In addition, the interaction between operations at different airports within the system was also investigated using pairwise inter-airport capacity envelopes (for example, JFK arrivals vs. LGA departures, JFK departures vs. EWR departures, etc.). This decomposition of inter-airport relationships into pairs is motivated by the observation that at any given time, there is a dominant inter-airport capacity interaction. Under this framework, the relationships among operational capacities at the three NY airports can be represented through 15 capacity envelopes (3 intra-airport, and 3×4 inter-airport pairs). The capacity for an operation at any time would then be determined by the most restrictive of its pairwise envelopes for that time interval.

The choice of the dependent and independent variables in the pairwise capacity envelopes is determined by the relative congestion experienced at the airports, combined with the earlier observation that arrivals have greater priority than departures. In particular, from analyzing the data for the years 2005 and 2006, we arrive on the following precedence order: JFK arrivals $\succeq$ LGA arrivals $\succeq$ EWR arrivals $\succeq$ JFK departures $\succeq$ LGA departures $\succeq$ EWR departures.

The following sections describe the estimation results for the pairwise capacity profiles at the New York area airports, and discuss the implications.

VI. Results and discussion

As described in Section IV, quantile function specifications were developed and estimated for the 15 pairwise capacity relationships (3 intra- and 12 inter-airport pairs) involving the three New York airports (JFK, EWR and LGA). The unhindered arrival capacities of the three airports were also estimated. The three airports are equipped with 4, 3 and 2 runways respectively, and their relative alignments are shown in Figure 3. We note that the principal runways at the three airports are aligned with each other.

Data for capacity estimation was extracted from the ASPM archives, which were described in Section IV-A. Apart from arrival and departure counts for every 15-min interval, the database also provides the corresponding information on the prevailing weather conditions (wind and visibility), the active runway configuration and the operating conditions (VFR/IFR). The estimation data set covered the years 2005 and 2006 for JFK and EWR, but was restricted to the year 2006 for LGA due to inconsistencies in the throughput reports during 2005. The linear programs for estimating the capacity envelopes were coded in AMPL [8], and solved

Fig. 3. A map of the New York area, showing the approximate locations of the three core airports and their relative layouts. Note that the airport layouts are not to scale with the map.

using CPLEX [13] with the default primal-dual simplex method.

A. Intra-airport capacity tradeoffs

The influencing factors on the capacity tradeoffs that were considered included visibility (VFR or IFR), alignment of the arrival and departure runways (parallel or crossing), and the number of additional runways for arrival or departure operations (beyond the primary runway). These factors were statistically tested for their influence on the intra-airport capacity envelopes.

The orders of quantile (τ) that yield robust estimates of the capacity envelope were found to be 99%ile, 99%ile and 99.5%ile for JFK, EWR and LGA respectively. The τ values for the corresponding unhindered arrival capacities were found to be relatively higher, at 99.75%ile, 99.5%ile and 99.75%ile respectively. Figures 4-6 illustrate the effect of each of the factors found to have a statistically significant influence on the capacity envelopes, under VFR and IFR conditions, for the three airports.

Figures 4-6 show that visibility has a significant influence on the unhindered arrival capacities, but does not appear to have a tangible effect on departure capacities. The relative alignment of the primary runways used for arrivals and departures plays a critical role in determining the capacity envelopes. Figures 4-5 show that the area under the capacity envelope progressively increases as we go from mixed arrival/departure operations on a single runway, through separate arrival and departure runways that intersect or converge, to additional parallel runways, demonstrating the benefit of independent operations on runways. The use of an additional departure runway at JFK and EWR provides the most benefit in terms of departure capacity at low values of arrival rate. Figure 4 also shows that the use of an additional arrival runway at JFK flattens the slope of the tradeoff curve, indicating the effective redistribution of operations across the two runways, while also increasing the unhindered arrival capacity. It is also observed that the use of the parallel runway configurations (22R, 22L) or (4R, 4L) at JFK results in a lower capacity as compared to their perpendicular counterparts (31R, 31L) or (13R, 13L). This is possibly explained by the smaller separation between the

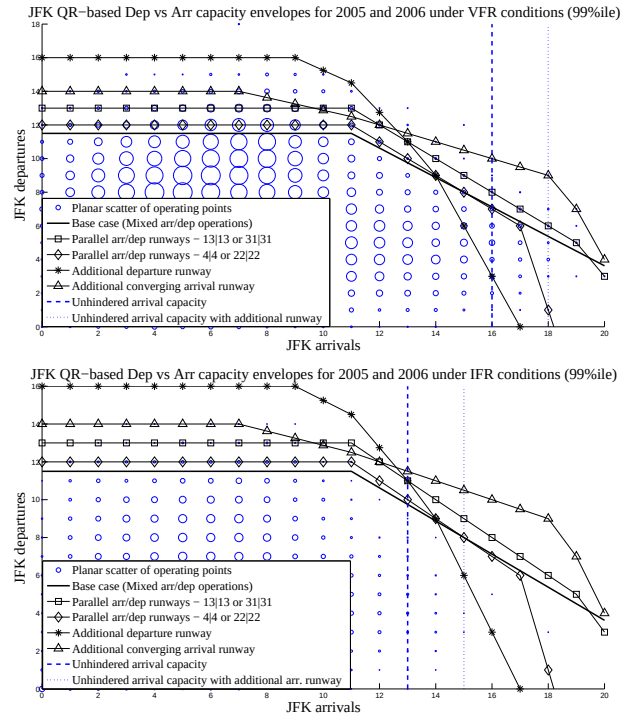

Fig. 4. JFK capacity envelopes for VFR (top) and IFR (bottom) conditions.

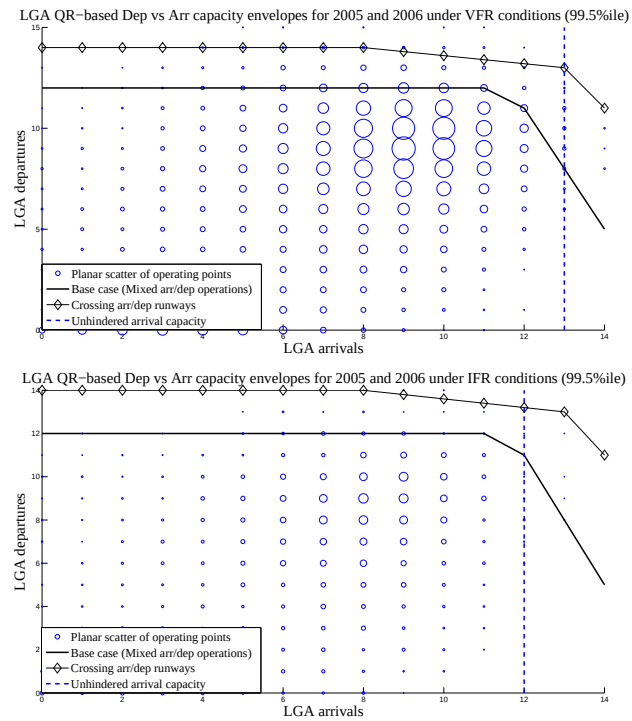

Fig. 5. LGA capacity envelopes for VFR (top) and IFR (bottom) conditions.

runway centerlines of the former pairs (resulting in a greater coupling of operations).

B. Inter-airport capacity tradeoffs

Since inter-airport interactions are expected to involve the airspace rather than the airport surface, the overlap between approach or departure paths is considered instead of the

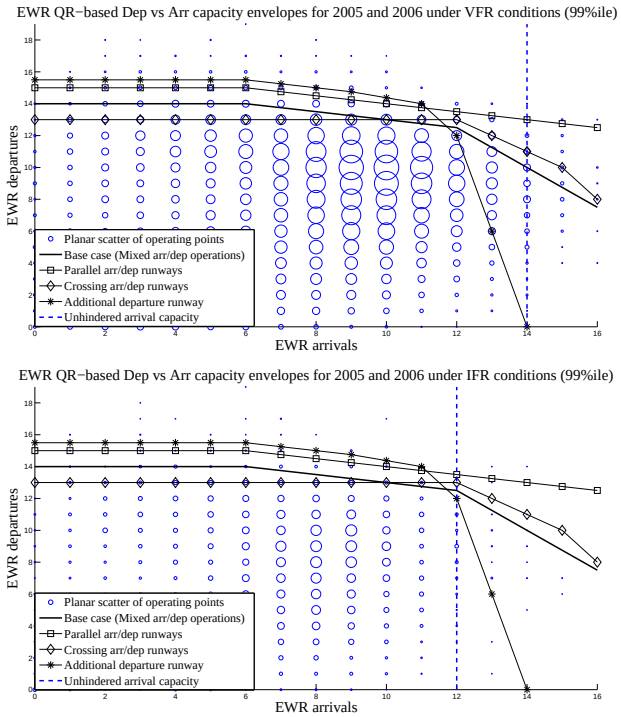

Fig. 6. EWR capacity envelopes for VFR (top) and IFR (bottom) conditions.

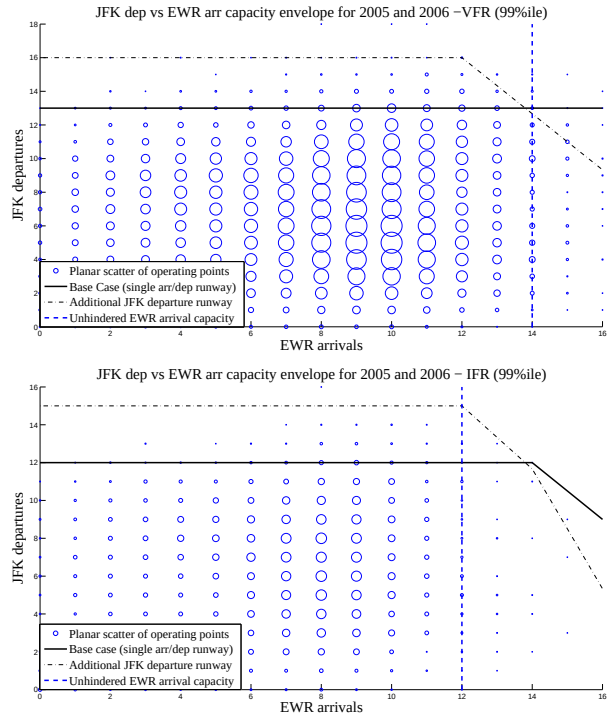

Fig. 7. Capacity envelopes for JFK departures vs EWR arrivals under VFR (top) and IFR (bottom) conditions (99%ile).

runway alignment attribute used for intra-airport capacity envelopes in Section VI-A. The approach and departure paths were approximated by two-dimensional conics extrapolated from the runway in the direction of operation, and binary terms were used to signify the intersection of these 2D conics. Representative inter-airport capacity envelopes for pairs of airports and arrival-departure operations, under different flight conditions are shown in Figures 7-9.

From the inter-airport capacity envelopes, it is observed that capacity tradeoffs are prominent at higher throughput values than those seen in the intra-airport envelopes. This observation suggests that airport (ground) capacity is a more dominant operational bottleneck than the capacity of the surrounding airspace. Another consequence of this observation is the impact of IFR conditions and single runway configurations on the inter-airport capacity curves, mainly due to the reduced throughput. Figures 7-9 illustrate this phenomenon for selected airport operational pairs that exhibited tradeoffs close to the limits of their respective operational capacities. The approach path overlap attribute was not found to be statistically significant for any of these pairwise envelopes, possibly due to the limited operational range over which the tradeoff effects were found to be prominent.

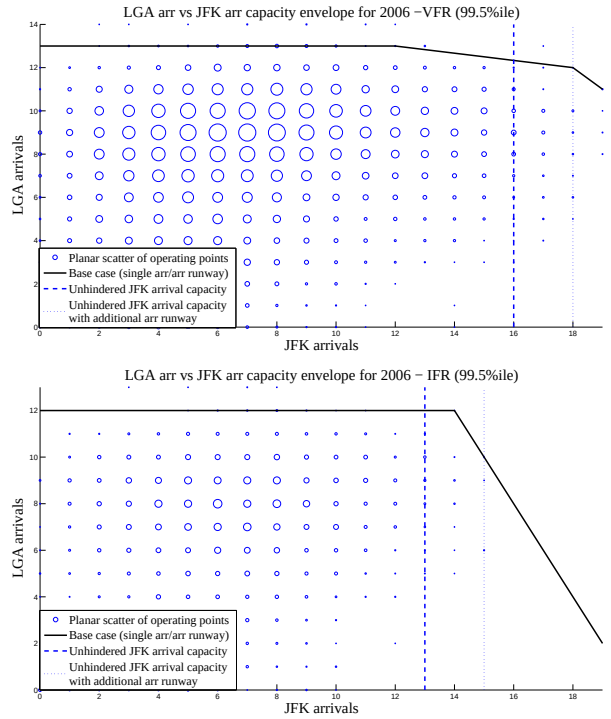

Fig. 8. Capacity envelopes for LGA arrivals vs JFK arrivals under VFR (top) and IFR (bottom) conditions (99.5%ile).

VII. Conclusion

This paper proposed a statistical framework for identifying and quantifying arrival-departure tradeoffs in a multi-airport system, and for estimating airport capacity envelopes using observations of flight operations. Quantile regression, the technique adopted within this framework, is shown to be well-suited for systematic outlier elimination. The proposed approach generates credible capacity estimates, and also assesses the influence of underlying factors that impact the capacity tradeoffs.

The estimation of capacity envelopes was conducted in two parts: the first part estimated a piecewise-linear function

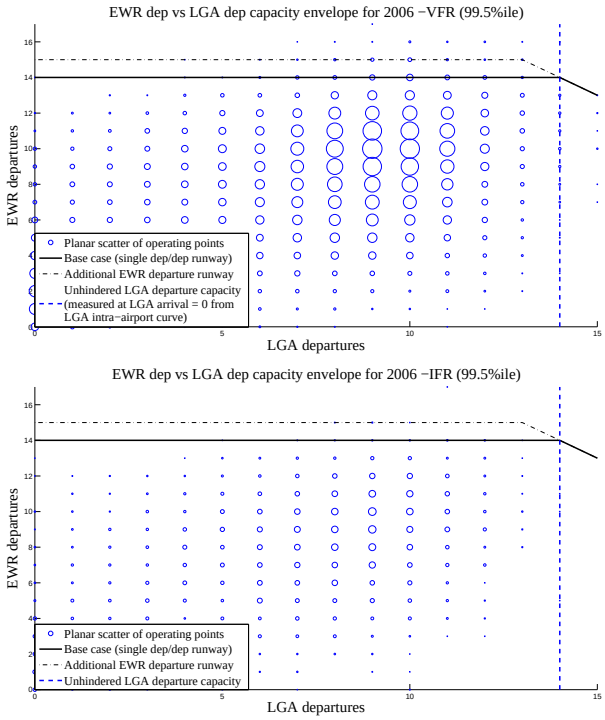

Fig. 9. Capacity envelopes for EWR departures vs LGA departures under VFR (top) and IFR (bottom) conditions (99.5%ile).

that related the departure capacity to the arrival rates, and the second part required the estimation of the unhindered arrival capacity. It was shown that both estimation problems can be solved using linear programming. The proposed approach was demonstrated through the estimation of capacity envelopes for the three New York area airports: JFK, EWR and LGA. This paper also extended the estimation methodology to analyze interactions between operations at different (and nearby) airports, by considering the New York multi-airport system.

References

[1] C. Bernini, M. Freo, and A. Gardini. Quantile estimation of frontier production function. *Empirical Economics*, 29:373–381, 2004.

[2] A. Blumstein. The landing capacity of a runway. *Operations Research*, 7(1):752–763, 1959.

[3] Philippe A. Bonnefoy. Emergence of Secondary Airports and Dynamics of Regional Airport Systems in the United States. Master's thesis, Massachusetts Institute of Technology, 2005.

[4] Bureau of Transportation Statistics. BTS website, 2007. `http://www.bts.gov`.

[5] Richard de Neufville and Amedeo Odoni. *Airport Systems: Planning, Design and Management*. McGraw-Hill, 2003.

[6] Federal Aviation Administration. OPSNET database, 2007. `http://aspm.faa.gov/main/opsnet.asp`.

[7] Federal Aviation Administration. Aviation System Performance Metrics, accessed September 2008. `http://aspm.faa.gov`.

[8] R. Fourer, D. M. Gay, and B. W. Kernighan. *AMPL: A Modeling Language for Mathematical Programming*. The Scientific Press (now part of Wadsworth Publishing), 1993.

[9] M. Furno. Parameter instability in quantile regression. *Statistical Modelling*, 7(4):345–362, 2007.

[10] E.P. Gilbo. Airport capacity: representation, estimation, optimization. *IEEE Transactions on Control System Technology*, 1(3):144–153, 1993.

[11] R.M. Harris. Models for runway capacity analysis. Technical report, The MITRE Corporation, MTR-4102, Rev. 2, Langley, Virginia, USA, 1972.

[12] S.L.M. Hockaday and A. Kanafani. Development in airport capacity analysis. *Transportation Research*, 8(3):171–180, 1974.

[13] ILOG, Inc. ILOG CPLEX 11.0.1, 2008. `http://www.ilog.com`.

[14] M. Janic and V. Tosic. Terminal airspace capacity model. *Transportation Research Part A*, 16(4):253–260, 1982.

[15] Joint Planning and Development Office. Next Generation Air Transportation System Integrated Plan, December 2004. `http://www.jpdo.aero/integrated_plan.html`.

[16] Mi-Ok Kim. Quantile regression with varying coefficients. *The Annals of Statistics*, 35(1):92–108, 2007.

[17] R. Koenker and K.F. Hallock. Quantile regression. *Journal of Economic Perspectives*, 15(4):143–156, 2001.

[18] Roger Koenker. *Quantile Regression*. Cambridge University Press, 2005.

[19] D.A. Lee, P.F. Kostiuk, R.V. Hemm, W.E. Wingrove, and G. Shapiro. Estimating the effects of the terminal area productivity program. Technical report, Logistics Management Institute, NS301R3, 1997.

[20] G.F. Newell. Airport capacity and delays. *Transportation Science*, 13(3):201–241, 1979.

[21] A. R. Odoni, J. Bowman, D. Delahaye, J.J. Deyst, E. Feron, R.J. Hansman, K. Khan, J.K. Kuchar, N. Pujet, and W.R. Simpson. Existing and required modeling capabilities for evaluating ATM systems and concepts. Technical report, International Center for Air Transportation, Massachusetts Institute of Technology, NAG2-997, 1997.

[22] W.J. Swedish. Upgraded airfield capacity model. Technical report, The MITRE Corporation, MTR-81W16, 1981.

[23] C. S. Venkatakrishnan, A. Barnett, and A. R. Odoni. Landings at Logan airport: Describing and increasing airport capacity. *Transportation Science*, 27(3):211–227, 1993.