Deconstructing Delay:

A Case Study of Demand and Throughput at the New York Airports

Amy Kim
Department of Civil Engineering
University of California, Berkeley
Berkeley, CA
amy_kim@berkeley.edu

Mark Hansen
Department of Civil Engineering
University of California, Berkeley
Berkeley, CA
mhansen@ce.berkeley.edu

Abstract - **This paper introduces an empirically driven, non-parametric method to isolate and estimate the effects of demand and throughput changes to observed changes in flight delay. Classical queuing model concepts were used to develop a method by which an intermediate queuing scenario could be constructed, in order to isolate the delay effects due to shifts in demand and throughput. This method includes the development of a stochastic throughput function that is based entirely on data and as a result has two advantages: it uses non-parametric, empirically-based probability distributions, and capacity need not be estimated explicitly. The method was applied to a case study of the three major New York airports of LaGuardia (LGA), John F. Kennedy (JFK), and Newark Liberty (EWR), for the peak summer travel seasons of 2006 and 2007, using data extracted from ASPM. This case study was of particular interest given that these airports experienced record levels of delay in 2007. The simulation results were consistent with both OPSNET and ASPM data, and were successful in quantifying the delay effects of demand and throughput changes from 2006 to 2007.**

Keywords - delay; demand; throughput; capacity; runway operations; New York airports; simulation; probability; ASPM; OPSNET.

I. INTRODUCTION

This paper introduces a method for estimating the effects of demand and throughput changes to observed changes in flight delay. As the delay observed over days, weeks or years changes from one time period to the next, we would like to know how much its evolvement can be attributed to demand and throughput changes. As a result, the motivation for this work is to address the following question: how can we isolate and measure shifts in delay caused by changes in demand and throughput when both are changing simultaneously?

There is an extensive body of literature and knowledge on methods to predict airport capacity and delay, both analytically [1] and through simulation. The purpose of this work is not to estimate the expected capacity outright [2], but to use empirical data that implicitly contains information about capacity to quantify how simultaneous changes in demand and throughput affect delay.

A new, empirically driven simulation procedure was developed from classical queuing concepts to address the question posed above. The main engine of this new procedure is a stochastic throughput function that was developed to have two key advantages. Firstly, this throughput function is driven by non-parametric probability distributions of throughput constructed from available data. Secondly, capacity need not be explicitly estimated, as the capacity of the operation under analysis is implicitly included in the probability distributions. This is advantageous because operational capacity is subject to a wide variety of factors and can be quite difficult to estimate well.

The simulation method is then applied to a case study of flight delay at the three major New York area airports: LaGuardia (LGA), John F. Kennedy (JFK), and Newark Liberty (EWR). Specifically, the arrival and departure operations at these airports were analyzed in order to determine how demand and throughput affected the delay changes observed between the 2006 and 2007 summer travel seasons.

The main goal in applying this new procedure is to provide information about the causes of delay shifts at one greater level of detail. The ability to isolate individual contributions of demand and throughput mechanisms to delay could be helpful in creating more focused, effective strategies and policies to address the delay problem.

II. BACKGROUND

During the summer of 2007, flight delays reached record high levels throughout the National Airspace System (NAS) and beyond. National and international headlines reported story after story describing the extreme wait times and missed connections that air travelers were subject to during this peak travel season. The three airports of the New York area experienced some of the highest delays within the NAS, with travelers spending 3.9 million more hours waiting for their aircraft to take off after leaving their gates in 2007 as compared to a decade earlier [3]. The increase in total operations from 2006 to 2007 at these airports was approximately 3-4%, but the increase in delay was in the order of about 28% [4]. In addition, in 2007 the New York airports accounted for about 40% of all delay in the NAS; in 2004 they accounted for only 15% [5].

Delay metrics can be found and/or calculated with relative ease from several data sources. One such source is OPSNET, which is the official source of historical NAS air traffic delays and operations. In OPSNET, an airport picks up a delay each

 ISBN: 978-0-615-20720-9

time a flight is held up 15 or more minutes due to runway congestion, weather, air holding, traffic flow restrictions, or other event that would cause a flight's realized schedule to deviate from its flight plan. Table 1 contains the results of OPSNET airport delay data extracted for LGA, EWR, and JFK for May through September of 2006 and 2007. The first half of the table indicates that the total number of operations have decreased at LGA and EWR, but have increased significantly at JFK. The second half of the table shows the number of flights that were delayed more than 15 minutes from their flight plans; it can be observed that the number of delayed flights has almost doubled at JFK from 2006 to 2007.

TABLE I. OPSNET DATA

Total Number of Arrival and Departure Operations			
	May-Sept 2006	**May-Sept 2007**	**% Change**
LGA	172,142	168,616	-2.0%
EWR	191,531	188,211	-1.7%
JFK	169,957	197,626	+16.3%
Total Number of Flights Delayed >15 Minutes			
	May-Sept 2006	**May-Sept 2007**	**% Change**
LGA	14,119	15,810	+12.0%
EWR	21,707	19,809	-8.7%
JFK	8,276	15,065	+82.0%

III. METHODOLOGY

Delay can be estimated using the traditional queuing model, where a queuing scenario is constructed from a cumulative demand curve and cumulative throughput curve [6]. An example of a simplified fictional queuing scenario is shown in Figure 1. The demand and throughput curves are actually step functions because customers (or vehicles, aircraft, etc.) are discrete entities. However, demand and throughput can be approximated as continuous functions (smoothed curves) over sufficiently long periods of time, which simplifies calculations. Under a classical deterministic approach, the throughput function at some time t can be determined as follows:

$$Q(t) = d(t) \quad \text{if } d(t) < c$$
$$= c \quad \text{if } d(t) \geq c \tag{1}$$

Where $Q(t)$ is the throughput at time t
$d(t)$ is the demand at time t
c is the fixed service capacity, constant over all t

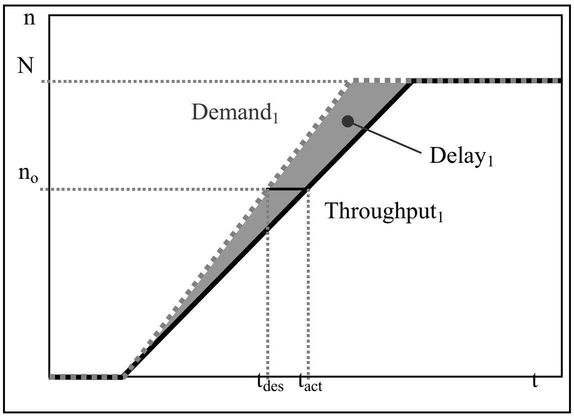

Figure 1. Queuing Scenario under Year 1 Demand and Year 1 Throughput

Note that in Figure 1 cumulative $Q(t)$ and $d(t)$ are shown to vary linearly with time. However, this is a simplification in that these quantities are most often non-linear, time-dependent functions.

Assuming first-in first-out (FIFO) conditions, the delay experienced by an arbitrary customer n is the difference between n's desired service time (t_{des}) and actual service time (t_{act}). This is also the horizontal distance between the two curves. The number of customers queued for service at time t is the vertical distance between the curves at t. Where the demand and throughput curves meet, customers are being served without any delay and as a result there are no standing queues for service; when the curves are apart, customers must queue for service. The throughput curve cannot cross the demand curve as per Equation (1) because customers cannot be served until they demand service. The area between the demand and throughput curves is the total delay experienced by customers over the total observation time T (we assume that our observations begin at time 0):

$$\tau = \int_0^T [Q(t) - d(t)]dt \approx dt * \sum_{j=1}^{j=J} [d(j) - Q(j)] \tag{2}$$

Where τ is total delay over time period (0,T)
$Q(t)$ is the throughput function at time t
$d(t)$ is demand at time t
T is total observation time
dt is the duration of a small time slice
j is the number of time slices over time T, from j=1 to j=J

An average delay per customer can then be determined by dividing this total delay by the total number of customers N that requested service over the observation time T. In these queuing diagrams, T could represent one day.

Figures 1 & 3 depict fictional queuing scenarios for an average day in an arbitrary year (Year 1) and the following year (Year 2), respectively. The areas between the demand and throughput curves represent the total delays in Year 1 and in Year 2. The change in total delay from Year 1 to Year 2 is the difference of the two areas; however, this difference could be

 ISBN: 978-0-615-20720-9

caused by changes in demand, changes in throughput, or both. In order to isolate the change in delay caused solely by a change in demand, we can construct a "counterfactual" scenario where the Year 2 demands are served using the Year 1 throughput function. The counterfactual scenario is represented in Figure 2, and the resulting delay is represented by the total area between the demand and throughput curves. The difference between the resulting counterfactual delay and the Year 1 delay (solid area in Figures 1 and 2) is the change in total delay due to the demand shift from Year 1 to Year 2 (depicted in cross-hatch). The Year 2 delay (total area between the curves in Figure 3) minus the counterfactual delay and Year 1 delay is the change in total delay due to the throughput shift from Year 1 to 2 (depicted by the unfilled area in Figure 3).

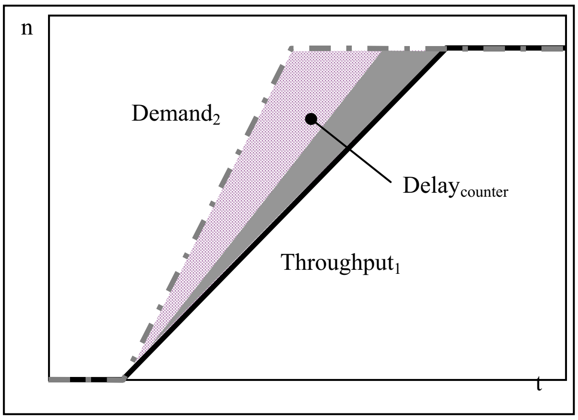

Figure 2. Queuing Scenario under Year 2 Demand and Year 1 Throughput (Counterfactual)

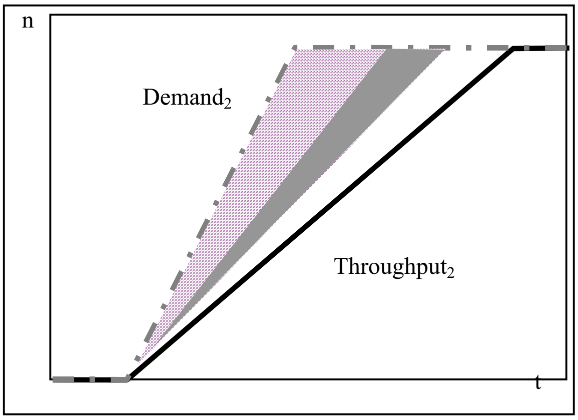

Figure 3. Queuing Scenario under Year 2 Demand and Year 2 Throughput

The figures show an increase in demand and a decrease in throughput from Year 1 to 2, but this trend was chosen for illustrative purposes only. The entire process is summarized in Table II.

TABLE II. DEMAND AND THROUGHPUT SCENARIOS

Demand	**Throughput**	**Total Delay**	**Δ in Total Delay**
for an average day in…			
Year 1	Year 1	(1) Year 1	n/a
Year 2	Year 1	(2) Counterfactual	(2)-(1); due to demand shift
Year 2	Year 2	(3) Year 2	(3)-(2); due to throughput shift

The Year 1 and Year 2 queuing scenarios can easily be constructed from available data (which will be discussed in detail later on), but the counterfactual scenario, because it does not actually exist, must be generated through simulation. The simulation is an iterative process that takes the demand in each time interval and, using a throughput function, assigns a throughput value. All aircraft not served in a time interval comprise the queue in that time interval, and from this a delay calculation can be made.

The classic definition of a deterministic throughput function was introduced in Equation (1). Based on available data sets that include arrival and departure demand, arrival and departure throughput counts, and weather information, we can construct a deterministic throughput function as follows:

$$q_o(t)=\min[d_o(t),c_o(w(t))] \tag{3}$$

Where $q_o(t)$ is the actual recorded throughput for operation type o in time interval t
$d_o(t)$ is the actual demand for operation o in time interval t
c_o is the fixed capacity for operation o
$w(t)$ is the weather condition at time t

Weather enters into the model as either visual or instrument flight rules (VFR or IFR), and is included as a factor in the model because of the significant impact it has on operational capacity. The operation types are either arrivals or departures.

The deterministic throughput function is an idealized situation and as such does not represent actual operations very well. c_o is a critical input to the function and several major assumptions are needed to determine its value(s). The alternative to the deterministic throughput function is a stochastic model that incorporates some levels of uncertainty. Based on the available data, a stochastic model that preserves the dependence of throughput on demand and weather can be constructed as follows:

$$P(Q_o(t)=q_o(t)|d_o(t),w(t)) = f_Q(q_o|d_o,w) \tag{4}$$

Where Q_o is a random variable representing throughput

 ISBN: 978-0-615-20720-9

f_Q is the conditional probability distribution function for throughput

The probability that Q_o takes some throughput value q_o, conditional on the demand and weather in time interval t, is taken from f_Q. It was necessary to include capacity as an explicit input to the deterministic model; however, in the stochastic throughput function it is implicitly captured in f_Q. f_Q can be constructed entirely from an appropriate data set without having to make assumptions about its shape and parameters. In fact, the non-parametric nature of f_Q is one of the main advantages of this model.

The counterfactual scenario discussed earlier was modeled using approximations of the stochastic throughput function, as described in Figure 4. The deterministic approximation uses mean throughput values conditional on demand, weather, and other known factors to simulate the counterfactual scenario. The stochastic approximation uses random number generation to simulate throughput values. For the purposes of the New York airports delay analysis, Year 1 will correspond to the time period of May through September 2006, while Year 2 represents that of the same months in 2007. Modeling the counterfactual scenario involves assigning simulated 2007 (Year 2) demand and 2006 (Year 1) throughput, calculating queue lengths in each time interval, and then calculating the average delay per flight over all time intervals from May through September. The following is the iterative procedure that was followed.

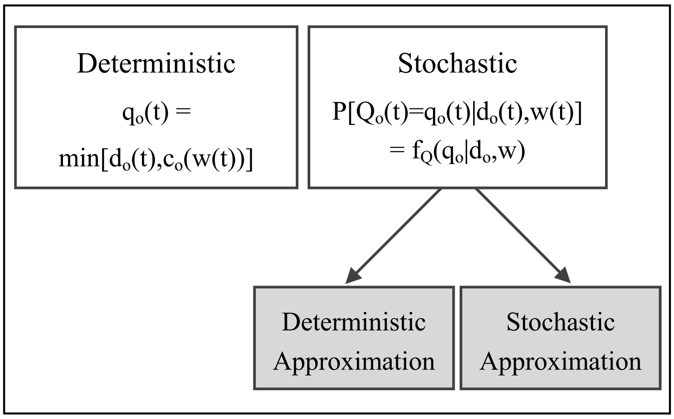

Figure 4. Specifications for the Throughput Function

1) At time interval t=1, initialize

$$\hat{D}_{o,07}(1) = D'_{o,07}(1) \tag{5}$$

Where $\hat{D}_{o,07}(t)$ is the simulated total (new & queued) 2007 demand for operation o in time interval t

$D'_{o,07}(t)$ is the "new" 2007 demand for operation type o in time interval t.

2) Find $\hat{Q}_{o,06}(t)$ conditional on $\hat{D}_{o,07}(t)$, $Q_{o,07}(t)$, & w, where $\hat{Q}_{o,06}(t)$ is the simulated 2006 throughput for operation type o in time interval t. $\hat{Q}_{o,06}(t)$ is determined using a stochastic throughput function.

3) If t=T, go to Step 4. Otherwise,

a) Set

$$\hat{D}_{o,07}(t) = D'_{o,07}(t) + [\hat{D}_{o,07}(t-1) - \hat{Q}_{o,06}(t-1)] \tag{6}$$

Where $\hat{D}_{o,07}(t)$ is comprised of the "new" demand of the current interval t in addition to the queued aircraft (those that are still waiting for service) from the previous time interval (t-1).

b) Update t=t+1.

c) Repeat Step 2.

4) Calculate the average delay per flight for operation o for the simulated counterfactual scenario.

$$\hat{\bar{\tau}}_o = \frac{\Delta t * \sum_{t=1}^{T} [\hat{D}_{o,07}(t) - \hat{Q}_{o,06}(t)]}{\sum_{t=1}^{T} \hat{Q}_{o,06}(t)} \tag{7}$$

Where $\hat{\bar{\tau}}_o$ is the simulated average delay per flight for operation type o, from t=1 to t=T, in minutes

Δt is the length of one time interval

The above procedure must be able to reproduce 2006 and 2007 operations as shown in the data such that when the counterfactual scenario is simulated using the same procedure, we can be confident of the results. In other words, the simulation method must produce good agreement between the actual and simulated baselines, which entirely depends on the specifications of the throughput function applied in Step 2. Deterministic approximations to the stochastic throughput function were first tested. These consisted of mean counts conditional on demand and weather, in addition to time of day effects and queue presence indicators, were first tested. Stochastic approximations of the throughput function, which involved randomly drawing from probability distributions of throughput conditional on demand and weather, were also tested. The methods above did not satisfactorily reproduce 2006 and 2007 operations, most likely due to underlying mechanisms not controlled for in the simulation. These phenomena might include serial correlation of demand and throughput between the quarter-hour intervals, arrival & departure interaction effects, and more. Finally, a stochastic approximation method that compares probability distributions of 2006 and 2007 counts, conditional on demand and weather, was tested. This approach, herein referred to as the "compared distribution" method, is able to, by design, identically replicates the baseline scenarios. As such the compared distribution method was chosen for use here.

In the compared distribution method, $\hat{Q}_{o,06}(t)$ is simulated in the following manner by starting with the 2007 (Year 2) data. All steps below are "substeps" of Step (2) from above.

 ISBN: 978-0-615-20720-9

The compared distribution method was used to generate the counterfactual scenarios for all New York airports under analysis. Note that the compared distribution method (as well the other stochastic approximation method) preserves time of day, day of week, and monthly effects from one year to the next, because the simulation is run time sequentially from t=1 to T in both 2006 and 2007.

1) Construct cumulative probability distributions (cdf) of counts conditional on demand and weather, $F(Q_{o,07}|D_{o,07},w_{07})$, for 2006 and 2007.

2) Find the cumulative probability of the empirical 2007 count for operation type o, conditional on 2007 demand and 2007 weather condition for some time interval t, $F(Q_{o,07}|D_{o,07},w_{07})$.

3) Based on the simulated 2007 demand, $\hat{D}_{o,07}$, find the interval in the 2006 count cdf that the 2007 probability found in the previous step falls into. From this, lower and upper bounds ($F_L(Q_{o,06}|\hat{D}_{o,07},w_{06})$ and $F_U(Q_{o,06}|\hat{D}_{o,07},w_{06})$, respectively) of the 2006 cdf and corresponding 2006 simulated count values ($Q_{o,06,L}$ and $Q_{o,06,U}$, respectively) are obtained.

4) Construct a probability value, $f(x)$, for the simulated 2006 count based on the 2007 count cdf's position between the lower and upper bounds of the 2006 cdf interval:

$$
\begin{aligned}
f(x) &= Pr(\hat{Q}_{o,06} = Q_{o,06,L}) \\
&= \frac{F_U(Q_{o,06}\mid\hat{D}_{o,07},w_{06}) - F(Q_{o,07}\mid D_{o,07},w_{07})}{F_U(Q_{o,06}\mid\hat{D}_{o,07},w_{06}) - F_L(Q_{o,06}\mid\hat{D}_{o,07},w_{06})}
\end{aligned}
\tag{8}
$$

$$
Pr(\hat{Q}_{o,06} = Q_{o,06,U}) = 1 - f(x) \tag{9}
$$

5) Generate random number n. If $n \le f(x)$, set count to lower bound 2006 count $Q_{o,06,L}$; otherwise set count to upper bound $Q_{o,06,U}$.

There are fewer count data recorded at very high demand values, and as a result the cumulative probability distributions of counts conditional on high demands are often based on small and incomplete data sets. To avoid reliance on probability distributions constructed using sparse data, all counts recorded with demands beyond the capacity threshold were combined into a single truncating probability distribution at the cut-off demand. For all simulated demands higher than that of the demand truncation point, this combined probability distribution is used for count simulation.

IV. Description of Data

The Aviation System Performance Metrics (ASPM) database is part of the Federal Aviation Administration's (FAA's) Operations and Performance Data system. Data from the "*Download/Airport*" section of the ASPM database was used for this analysis. The data includes hourly as well as quarter-hourly arrival and departure counts, demands, and visibility conditions (either visual (VFR) or instrument (IFR) flight rules). The data is available for 77 major airports in the United States.

ASPM count data are based on individual aircraft landing and take-off times as supplied through Airline Service Quality Performance (ASQP) data or Enhanced Traffic Management System (ETMS) messages.

ASPM provides the perfect data set to construct the counterfactual scenarios described in the previous section; however, some particular characteristics of the ASPM demand data selected for this analysis must be noted. Firstly, the demand data used here is based on the updated flight plan just before a flight is due to take off at the origin airport; it does not reflect demand as defined by airline schedules. As a result, for flights arriving at a given airport, the delay calculated in this analysis includes all delays that occur between the filed flight plan take-off time (demand) and actual landing time (count), but does not include the delays between scheduled and flight plan take-off times (although this information can also be found in the ASPM dataset). For flights departing the airport, the delay calculated in this analysis includes the delay incurred between the time that the flight was scheduled to depart according to the flight plan, and the time that it actually does depart. As a result, the calculated delay will not include the effects of ground delay programs (GDP), the effects of air traffic management (ATM), plus other mechanisms that would cause a flight to deviate from its schedule. Secondly, the reported demand represents the total number of aircraft that were available for operation o (arrival or departure) in time interval t. An aircraft will count towards demand in each and every time interval starting in the one when it was first available to land/depart until the time interval when it is actually able to do so. As a result, the demand $D_o(t)$, reported in t includes the "new" demand $D'_o(t)$, plus the queued (unserved) aircraft from the previous time interval $[D_o(t-1)-Q_o(t-1)]$. $D'_o(t)$ for each 15-minute interval is easily calculated from the ASPM dataset, and is used for input to the simulation.

$$
D'_O(t) = D_O(t) - [D_O(t-1) - Q_O(t-1)] \tag{10}
$$

Where $D'_o(t)$ is the "new" demand for operation type o in time interval t
Δt is the length of one time interval
$D_o(t)$ is the total demand for operation o in time interval t
$Q_o(t-1)$ is the throughput for operation o in time interval $t-1$.

If an aircraft's demand and service times fall within the same or adjacent intervals, its delay is recorded to be zero. For instance, if time intervals are 15 minutes in length, an aircraft will not be counted towards delay if its demand and actual service times are, for instance, 1 minute and 14 minutes into the interval respectively. Also, ASPM counts will never exceed the total demand in any given time interval, meaning that operations which occur earlier than scheduled are not counted as negative delay or a delay savings.

ISBN: 978-0-615-20720-9

Data from LGA, EWR, and JFK were obtained for May 1 through September 30 2006, and May 1 through September 30 2007. From the data, we can derive the following information about demands, throughput, and delay. Figure 5 displays cumulative arrival demands by hour averaged over all days from May 1 through September 30. It can be observed that the total daily demand (averaged over all days) at LGA and EWR has decreased (between 2% and 3%) from 2006 to 2007 while it has increased significantly (by approximately 18%) at JFK. As expected, departure demands exhibit very similar trends and as a result are not displayed here. Figure 6 displays the average arrival count recorded during VFR conditions plotted against demand. The average arrival count per demand was calculated by averaging all counts recorded at each demand level from 0 to 70+. Observe that the arrival counts match arrival demands up to a certain point, after which this trend stops as the facility cannot serve at the demanded rate any longer. After this peak count level, arrival counts remain steady or begin to decrease until the slope of the curve flattens out. Also beyond the peak, all demand cannot fully be served within the same time period any longer. The peak arrival count is the realized arrival capacity for a given airport [7]. Based on this simple yet reliable capacity estimation, Figure 6 suggests that the arrival capacities of all three airports have decreased from 2006 to 2007. One can also observe that higher arrival demands were reported at LGA and JFK in 2007, which suggests that there were longer queues, which in turn suggests that aircraft waited longer for service and therefore experienced greater delay in 2007. The same phenomenon, however, was not recorded at EWR. A similar analysis can be applied to the averaged departure counts in Figure 7, which implies that departure capacities have dropped at LGA and EWR but have increased at JFK from 2006 to 2007. However, much higher demands (and therefore queuing) were reported at JFK in 2007, which may be the result of increased demand and/or more severe demand peaking effects, as the data does not seem to suggest that capacity has decreased.

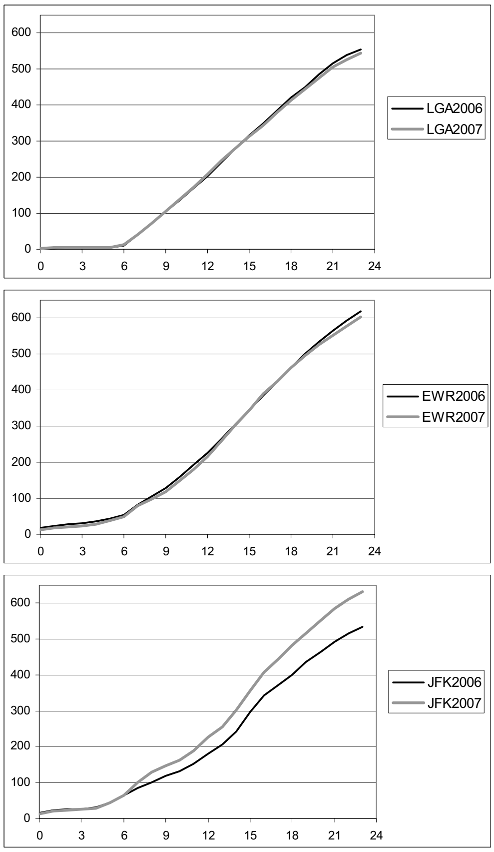

Figure 5. Cumulative Arrival Demands (Average Day of Ops), by Hour

ISBN: 978-0-615-20720-9

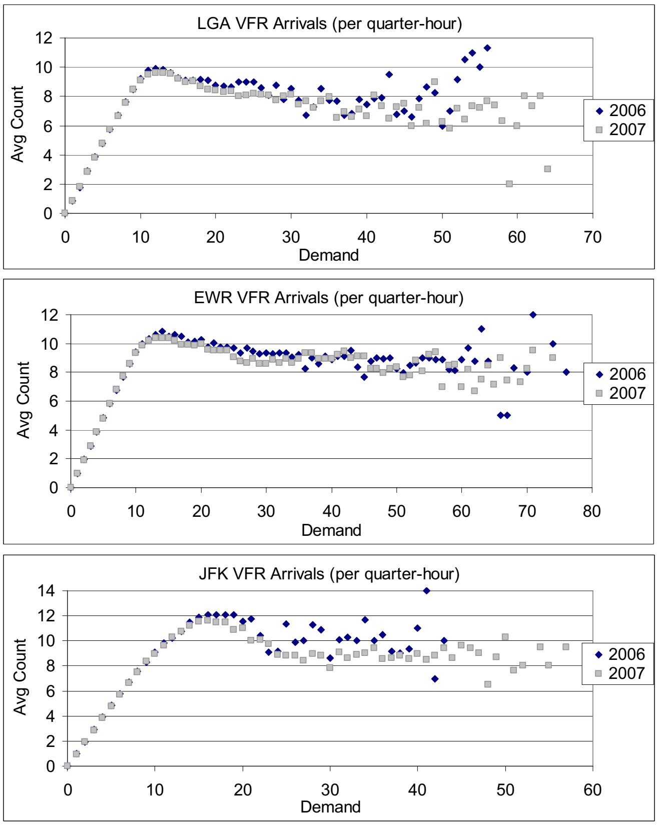

Figure 6. Average Arrival Counts vs. Demand, in VFR, by Quarter-Hour

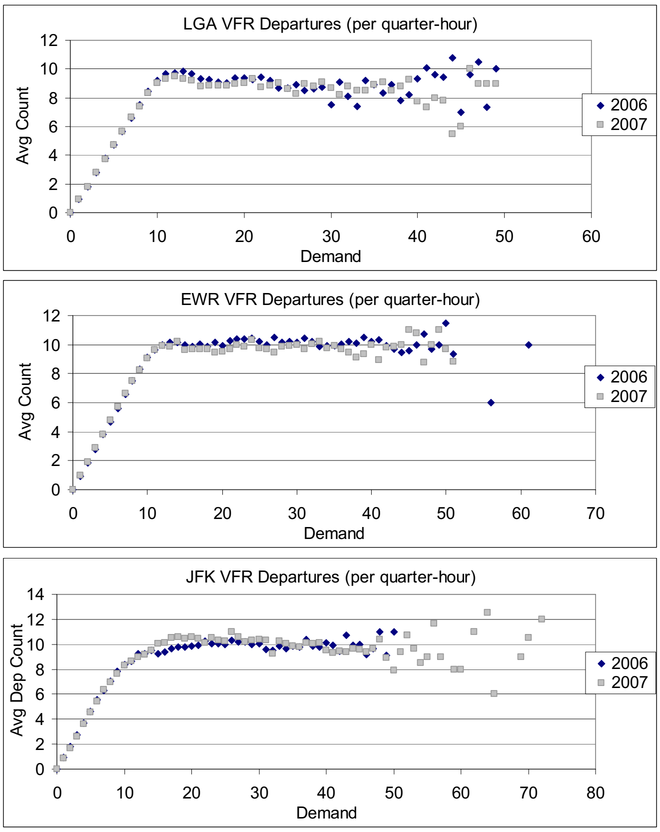

Figure 7. Average Departure Counts vs. Demand, in VFR, by Quarter-Hour

The average delay per flight was also calculated for arrival and departure operations at each airport from May through September of 2006 and 2007 as per Equation (7). Recall that delay is calculated against flight plan demand, and the data is tabulated in 15-minute intervals (such that Δt=15 min). The data set contains T=14,688 quarter-hour intervals.

The delay results are summarized in Table III.

TABLE III. AVERAGE DELAY PER FLIGHT, MAY-SEPT 2006 & 2007

		Average delay per flight (min)		**Change** *(from 2006 to 2007)*
		2006	***2007***	
LGA	Departure	8.56	10.72	+2.16
	Arrival	8.85	10.7	+1.85
EWR	Departure	11.53	9.95	-1.58
	Arrival	11.46	12.06	+0.60
JFK	Departure	12.06	14.38	+2.32
	Arrival	3.23	8.11	+4.88

The average delay per flight increased at both LGA and JFK between 2006 and 2007, and significantly so for JFK arrivals. Average delay has decreased by about 1.6 minutes per departing flight at EWR, and for arrival flights it has increased 0.6 minutes. These results from ASPM are consistent with the OPSNET data discussed previously.

Modeling the counterfactual scenarios involves recreating the structure of the ASPM demand and count data by assigning simulated 2007 demand and 2007 throughput values, and then calculating queue lengths and average delay in the same manner as was done for the data shown in Table III.

V. RESULTS

The simulation results are summarized in Table IV. The reported counterfactual delays are the average of 10 simulation runs for each scenario. The standard deviations of the 10 runs are also reported.

ISBN: 978-0-615-20720-9

TABLE IV. DELAY RESULTS

	Average Delay per Flight (min)			**Δ delay due to Δ demand *(implies)***	**Δ delay due to Δ throughput *(implies)***	**SD****
	2006	**CF***	**2007**			
LGA						
Departure	8.56	6.49	10.72	-2.08	4.24	0.115
				Demand has decreased	*Throughput has decreased*	
Arrival	8.85	6.03	10.70	-2.82	4.67	0.097
				Demand has decreased	*Throughput has decreased*	
EWR						
Departure	11.53	5.78	9.95	-5.76	4.18	0.064
				Demand has decreased	*Throughput has decreased*	
Arrival	11.46	6.65	12.06	-4.81	5.41	0.097
				Demand has decreased	*Throughput has decreased*	
JFK						
Departure	12.06	19.09	14.38	7.03	-4.71	0.138
				Demand has increased	*Throughput has increased*	
Arrival	3.23	4.91	8.11	1.68	3.20	0.084
				Demand has increased	*Throughput has decreased*	

* Counterfactual, referring to scenario with 2007 demand and 2006 throughput
** Standard deviation of counterfactual delay, for 10 simulation runs made

TABLE V. DELAY RESULTS (COUNTERFACTUAL SCENARIO II)

	Average Delay per Flight (min)			**Δ delay due to Δ demand *(implies)***	**Δ delay due to Δ throughput *(implies)***	**SD****
	2006	**CF***	**2007**			
LGA						
Departure	8.56	13.08	10.72	4.52	-2.36	0.093
				Throughput has decreased	*Demand has decreased*	
Arrival	8.85	13.36	10.70	4.51	-2.66	0.287
				Throughput has decreased	*Demand has decreased*	
EWR						
Departure	11.53	20.33	9.95	8.80	-10.38	0.088
				Throughput has decreased	*Demand has decreased*	
Arrival	11.46	16.40	12.06	4.94	-4.34	0.351
				Throughput has decreased	*Demand has decreased*	
JFK						
Departure	12.06	9.10	14.38	-2.97	5.29	0.060
				Throughput has increased	*Demand has increased*	
Arrival	3.23	6.52	8.11	3.29	1.59	0.084
				Throughput has decreased	*Demand has increased*	

* Counterfactual, referring to scenario with 2006 demand and 2007 throughput
** Standard deviation of counterfactual delay, for 10 simulation runs made

The results are consistent with the trends seen in the OPSNET data, as well as the ASPM data presented in the previous section and used for this simulation. At both LGA and EWR, arrival and departure demand changes have results in decreases in arrival and departure delay, implying that demand has declined. In addition, delays attributed to changes in throughput have increased, which would imply that throughput has dropped as well. At JFK, arrival and departure delays have increased due to changes in demand, suggesting that demands have gone up (with the departure demands having caused relatively significant increases in delay). However, increases in departure throughput have caused departure delays to drop while arrival throughput may have decreased and caused a subsequent increase in arrival delay.

The counterfactual scenario can also be constructed by swapping the demand and throughput years and simulating 2006 demand with 2007 throughput; in other words, using the same procedure described above but with the years switched. In this case, the difference between the counterfactual and 2006 base year delays can attributed solely to changes in throughput, and the difference between the 2007 base year and counterfactual scenario delays to changes in demand. Table V contains the results of this simulation.

The delay trends in Tables IV and V are consistent with one another. It also appears that the magnitudes of the changes in delay are consistent between the two analyses at LGA, EWR arrivals, and JFK, although there is greater discrepancy in the departure results for EWR. Because demand, throughput and delay are not necessarily related linearly, the "direction" in which the counterfactual scenario is simulated could have a significant effect on the delay results (of Tables IV and V). However, the choice regarding which way to simulate the counterfactual scenario is arbitrary, and consequently the two sets of delay results may serve to validate the simulation process. The differences between the two sets of results for EWR departures may be due to other dependent effects not accounted for or readily apparent in the simulation process. Also, as demands increase, delays also increase at much faster rates; conversely, when demands are lower an increase in throughput can result in a significantly greater delay reduction [2]. This may account for the fact that the Table V results for EWR show much larger changes in delay between the two years than Table IV.

We can make a few inferences based on the results in Tables IV and V above. Firstly, of the three airports JFK has experienced the largest overall increase in delay due to changes in throughput and demand. In particular, a substantial growth in arrival and departure demands has contributed to the large increase in delay at JFK. Departure throughputs have not similarly increased to offset this rise in demand, while the problem in the arrival operations is further exacerbated by a decrease in throughput. Decreased throughput does not necessarily mean a drop in airport capacity. In fact, sources at

 ISBN: 978-0-615-20720-9

the FAA believe that fleet mix changes in 2007 at JFK led to higher minimum in-trail separations, which would certainly reduce throughput. It has also been suggested that New York airspace controllers had grown more conservative about aircraft separations due to safety concerns. This would support the findings of Figures 6 and 7, which suggest that capacities have generally decreased (except for JFK departures) between 2006 and 2007. The drop in demand at LGA and EWR (Figure 5) occurred alongside a drop in throughput, and generally resulted in an overall increase in delay at these airports.

VI. CONCLUSIONS

The New York airports experienced a very significant rise in delays over the summer of 2007 compared to previous periods, most specifically that of summer 2006. The purpose of this work was to estimate how much of this change in delay was due to demand changes and how much was due to throughput changes. Because demand and throughput change simultaneously, the purpose of this work was to quantify how changes in each contribute to a change in delay, and ultimately provide information about the causes of delay at one greater level of detail. To do this, an empirically driven simulation procedure was developed from classical queuing concepts, and applied to a case study of the three major New York area airports in summer 2006 and 2007. This procedure consists of a stochastic throughput function whose main advantages are that it uses non-parametric, empirically-based probability distributions and that capacity need not be estimated explicitly. The throughput function was used to recreate the structure of the ASPM data and construct the intermediate “counterfactual” scenario, by which the delay changes from 2006 to 2007 could be attributed to either demand or throughput.

The simulation results confirmed the OPSNET and ASPM data results. The counterfactual scenario was first constructed with 2007 demand and 2006 throughput. At both LGA and EWR, arrival and departure demand changes have results in decreases in arrival and departure delay, implying that demand has declined. In addition, delays attributed to changes in throughput have increased, which would imply that throughput has dropped as well. At JFK, arrival and departure delays have increased due to changes in demand, suggesting that demands have gone up. However, increases in departure throughput have caused departure delays to drop while arrival throughput may have decreased and caused a subsequent increase in arrival delay. The counterfactual scenario was also constructed with 2006 demand and 2007 throughput, and the results of this simulation served to validate the previous simulation results.

VII. FURTHER WORK

This procedure is a starting point from which we can further analyze and deconstruct the causes of operational delay at airports in terms of demand and throughput. However, knowing only the demand and throughput effects on delay has limited importance; it would be beneficial to identify factors other than flight rule conditions that influence demand and throughput. This could, in turn, be used to re-specify the throughput function to control for additional factors not yet included in the model. Phenomenon yet uncontrolled for might include fleet mix changes, and arrival/departure interaction effects (the model as of yet assumes arrivals & departures to be independent of one another). Another direction for future work is to base delay calculations on a demand scenario other than that of the flight plan, such as demand recorded at the time flights are scheduled by the airlines to arrive or depart. Using this, the effects of GDP as well as all the effects of ATM at origin airports could be incorporated into the analysis.

ACKNOWLEDGMENT

The authors would like to thank the Air Traffic Organization at the FAA, particularly Joe Post, Dan Murphy and Michael Wells, for sponsoring this work as well as providing invaluable insights and suggestions. We would also like to thank Dave Knorr at the FAA and Ken Wright at MITRE for their assistance.

REFERENCES

[1] G.F. Newell, “Airport capacity and delays,” Transportation Science, Vol. 13, no. 3, 1979.

[2] Mark Hansen, "Post-deployment analysis of capacity and delay impacts of an airport enhancement: case of a new runway at Detroit," Air Traffic Control Quarterly, Vol. 12, pp. 339-365, 2004.

[3] Ken Belson, “Study puts price tag on delays at airports,” New York Times, December 2, 2007.

[4] Dave Knorr, “Delay Metrics: Why are Delays Increasing?”, Presentation given at the National Airspace System Performance Workshop, Asilomar, Pacific Grove, CA, 2007

[5] Ken Wright, “Pangs of New York: Reducing Delays at Kennedy, LaGuardia, and Newark Airports”, panel member, Transportation Research Board Annual Meeting, Washington, D.C., 2008.

[6] Carlos Daganzo, Fundamentals of transportation and traffic operations, 1st ed.. Oxford; New York: Pergamon, 1997.

[7] Yu Zhang, Jasenka Rakas, and Mark Hansen, “Methodology for Estimating Airport Capacity and Throughput Performance,” unpublished.

 ISBN: 978-0-615-20720-9