Airport Congestion During Relief Operations

Michael Veatch
Department of Mathematics, Gordon College Wenham, MA 01984
mike.veatch@gordon.edu
Jarrod Goentzel
Center for Transportation & Logistics, MIT, Cambridge, MA 02139
goentzel@mit.edu

January 30, 2015

Abstract

Scheduling the airlift of relief supplies into a damaged or small airport during a crisis is a major challenge. The volume of cargo and flights can temporarily overwhelm the airport's capacity and the nature of the traffic is also complex, generated by international donors, the national government, foreign governments, and private parties (VIPs, press, etc.). We analyze new data from 2010 Haiti earthquake response to develop insights and parameters for planning. Using these data, we model various scenarios to explore alternative methods for scheduling flights, managing bottlenecks such as parking space and offloading resources, and prioritizing aircraft type.

1 Responding to disasters

Large, rapid onset humanitarian crises have become more frequent and humanitarian organizations have increased their preparation for responding to large crises. For disasters, there is typically a short window of seven to 10 days when relief supplies can have the most humanitarian impact, making airlift an important mode of delivery. In many parts of the world, the volume of flights needed to respond to a major crisis can temporarily exceed the capacity of the local airport, particularly the parking and unloading capacity.

While the U.N. World Food Program Logistics Cluster and military operations have effective airlift procedures, each humanitarian response presents unique challenges that need to be addressed. When these organizations support the government in coordinating airport operations, ad-hoc scheduling methods and operating procedures are often created. The airport operations for the 2010 Haiti earthquake response, which have been well documented, offer a particularly interesting case study. Our research explores the logistical questions of how to schedule an airport to deliver as much cargo as possible in a short window after the disaster using this case study.

First, we review the experience of the U.S. Air Force in scheduling the Port-au-Prince airport after the 2010 Haiti earthquake. Next we analyze a new dataset documenting operations on the ground in Port-au-Prince to develop insights and numerical parameters that can be referenced in future planning efforts. Finally, using these new parameters, we develop simulation models to quantify the effects of various airport scenarios and scheduling strategies. We measure effectiveness in terms of the tons of cargo delivered and the frequency with which arriving flights must be diverted or held. The analysis suggests that the mix of aircraft used in the Haiti response greatly impacted the volume of aid, that scheduling parking spaces beyond the number of aircraft that can be unloaded at once is crucial for buffering, and that the volume of flights scheduled should be increased beyond what is standard practice. The last finding is due to the variability inherent in the system, such as the tendency of flights to no show.

2 Context and Relevant Literature

2.1 Airport Operations in Humanitarian Response

Airports are critical hubs for the urgent supplies required in a humanitarian response. An airport near the affected community may be constrained due to damage by a natural disaster or by conflict. Small airports have further constraints such as limited paved parking space and limitations on the size of aircraft. When parking space is occupied, additional aircraft cannot land without blocking the runway or taxiway. The ability to move supplies away from the ramp and ground transportation may also be limiting factors, though the focus of this research is on parking and unloading. The rate at which cargo is unloaded is very sensitive to the type of cargo and whether it is loaded as rolling freight, palletized, or stacked or is passengers only. Rates also vary by aircraft type, with larger aircraft generally having much higher unloading rates, in tons/hour. Our research explores how aircraft type can be considered in planning operations.

For a capacity constrained airport, the goal is to regulate the arrival of aircraft to match the unloading capacity. Extra parking space creates a buffer between arriving aircraft and the unloading resource, insulating it from variation in the arrival times and unloading times. Limited parking space, and the cost and limitations on ground hold or air hold time, make scheduling crucially important. With no airflow management, arrivals may far exceed capacity, leading to air holds and massive numbers of diverted flights, as experienced immediately after the Haiti earthquake (Longmire, et al. 2012). Two airflow management strategies are a segmented slot reservation system, as done in Haiti, or an *air bridge*, where supplies are flown to a staging point by multiple governments and NGOs, then reloaded onto standardized aircraft with a more controlled schedule. For the 2004 tsunami response, air bridges were critical in utilizing the small airport in Banda Aceh. Jakarta functioned as a transfer point for most aid, the U.S. used the Utapao Naval Air Base in Thailand, and some aid was flown through Medan, Indonesia (Gempis 2005, Morris 2006).

Another fundamental tradeoff examined in this paper is how heavily to schedule the airport. Booking too many flights runs the danger of delayed or diverted flights and not being able to respond to last-minute high-priority flights. Scheduling too few will result in idle unloading resources because of no shows or variation in unloading times. Our analysis suggests that the load should be higher than standard practice.

2.2 Relevant Literature

Airport scheduling during a rapid-onset crisis is quite different than in the commercial context. Therefore, we did not attempt to review all of this literature. One area of some relevance is gate assignment, surveyed in Dorndorf et al. (2007). Some of this literature refers to cargo as well as passenger flights. Recent work on gate scheduling has included the treatment of uncertainty in arrival and departure times (Lim et al., 2005), which is important in the humanitarian context. Runway and taxiway scheduling, such as Anderson et al. (2000), makes use of queueing models and the physical layout of the airport.

The U.S. Air Force Air Mobility Center has detailed procedures and information systems for planning airlift operations. Extensive work has been done on optimization models (Baker et al., 2002, Granger et al., 2001). These models include ramp space constraints and ground times for each type of aircraft and airbase. With the Air Mobility Center's increasing role in humanitarian airlift operations, studies and policies have been written for them. Models of these humanitarian operations are developed by Mogilevsky (2013) and Penny (1996). The first paper considers transport between multiple locations, but does not model individual airports in detail. The second paper develops a queueing model of one small airport, and identifies variability of arrival times and ground times as a major concern. It differs from our model in that only military aircraft are considered and data from all military airlifts (not just humanitarian crises) is used. Literature analyzing the Haiti airport experience is discussed in the next section.

3 Haiti Airport Scheduling: Description and Data

On day 3 following the Haiti earthquake, the Haitian government gave control of the Toussaint Louverture Airport in Port-au-Prince to the U.S. government. The U.S. Air Force 601st Air and Space Operations Center/Air Mobility Division of the Florida Air National Guard at Tyndall Air Force Base quickly established the Haiti Flight Operations Coordination Center (HFOCC) (Davidson and Smith, 2011). Other Air Force units were already at the airport, dealing with flights arriving uncontrolled from all over the globe. With the U.S. military presence, the airport operated virtually continuously, averaging over 80 flights per day in the month following the earthquake.

HFOCC used a slot reservation system, adapted from plans for responding to domestic crises. Each parking space was reserved for the nominal time on ground (TOG) for that type of aircraft: one hour for small, two hours for narrowbody, and four hours for widebody aircraft. The ramp was divided into 11 parking spaces, one controlled by the United Nations because of their pre-existing stabilization mission. Light aircraft could also park on the grass ramp, controlled by the government of Haiti. The remaining 10 spaces were considered the parking maximum on ground (MOG). Time slots were only made available up to the working MOG, which is the maximum number of aircraft that can be simultaneously serviced (onload and offload of cargo and passengers). The working MOG was initially six (five narrowbody and one widebody), and later increased to nine (seven narrowbody and two widebody) (Longmire, et al. 2012, pp. 51 and A-12).

Of the working MOG, 15% of slots were allocated to the government of Haiti, which included some aid flights that they arranged directly with donors, and 4% were held back. Some of these allocations were for specific times of day. Not allocating all 10 spots in advance allowed 105 aircraft that arrived without slots (or outside of their slot) to be accommodated, including government of Haiti and some US military flights. Other contingencies during their 60 days of scheduling included VIP aircraft staying longer than the nominal two hours of ground time and four broken aircraft tying up spots.

In hindsight, were the correct number of slots allocated? The reports cited above and others all agree that there were wasted spots due to no shows. They recommend strategies for reducing no shows, but endorse the strategy of not scheduling all parking and unloading slots, particularly in the first week of a crisis. In general, the answer to this question requires balancing the costs of over-capacity, e.g., diverts, and under-utilization. For Haiti, diverts were less than 1% over the 60 day HFOCC operation, including arrivals without slots; two arrivals with slots were diverted in the first three days of slot scheduling (Jones 2011). Two other diverts were due to runway light failures, which fortunately occurred well after traffic had peaked, and one due to misunderstanding an aircraft size (Longmire, et al. 2012). Diverts were generally sent to Homestead, Florida, though some smaller aircraft went to Santo Domingo, Dominican Republic.

3.1 Analyzing Parking Spaces using Maximum on Ground (MOG) Data

Previous analyses of the Haiti airport in 2010 have not utilized data from the operations on the ground. For our analysis, we obtained new data in the form of a "flight log" documenting each flight in the first 31 days of the HFOCC operation. Major Matt Jones, who was Deputy Commander and Director of Operations for the Joint Task Force Port Opening (JTF-PO) at the Port-au-Prince airport in January 2010, provided these data in the form of a Microsoft Excel spreadsheet.

Aircraft generally departed as soon as they were unloaded (a few waited for VIPs or to load evacuees; air traffic delays were minimal), so the occupied spaces roughly measure the unloading work available. Figure 1 shows the distribution of occupied spaces, counted at arrival time events for 2052 flights. Days 1-5 have a mean of 5.8, while days 6-31 have a mean of 4.3. While it was not unusual to approach the parking MOG of 10, the average utilization of parking spots was fairly low. For days 1-5, if we assume that only 3 aircraft can be unloaded at a time (Jones 2011, p. 8) then there was virtually always enough work, but compared to the working MOG of 6 used by HFOCC, there were fewer than 6 occupied spaces for about 30% of the arrival times. There are some caveats to the data in Figure 1. It omits 67 flights with partial data missing. The arriving flight is counted, so there are never zero occupied spaces. Counting spaces after a departure event allows zeroes, and reduces the mean by 0.6. Measuring at uniformly spaced times is desirable but was not practical.

Figure 2 shows the results when parked aircraft are counted instead of spaces. The mean number of aircraft is 7.6 and 5.2 for days 1-5 and 6-31, respectively. Occasionally more than one aircraft is assigned to a parking space due to data errors or creative parking of small aircraft. Given these issues, the data in Figure 1 are preferable to Figure 2. The main insight from the data is that the number of parked aircraft fluctuates widely, often dropping below the unloading capacity, despite the scheduling efforts.

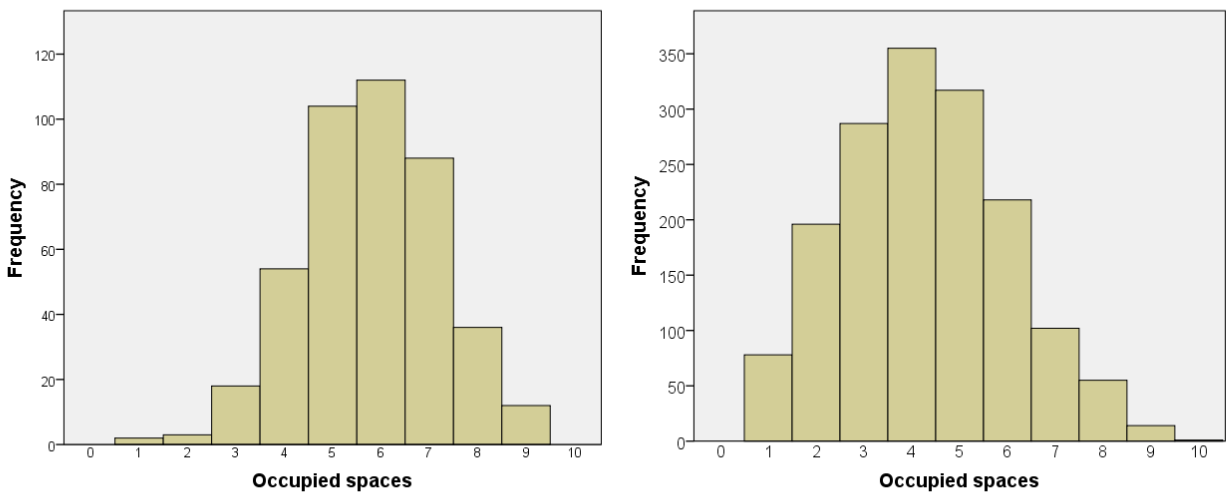

(a) (b)

Figure 1. Parking space utilization (a) Day 1-5. (b) Day 6-31.

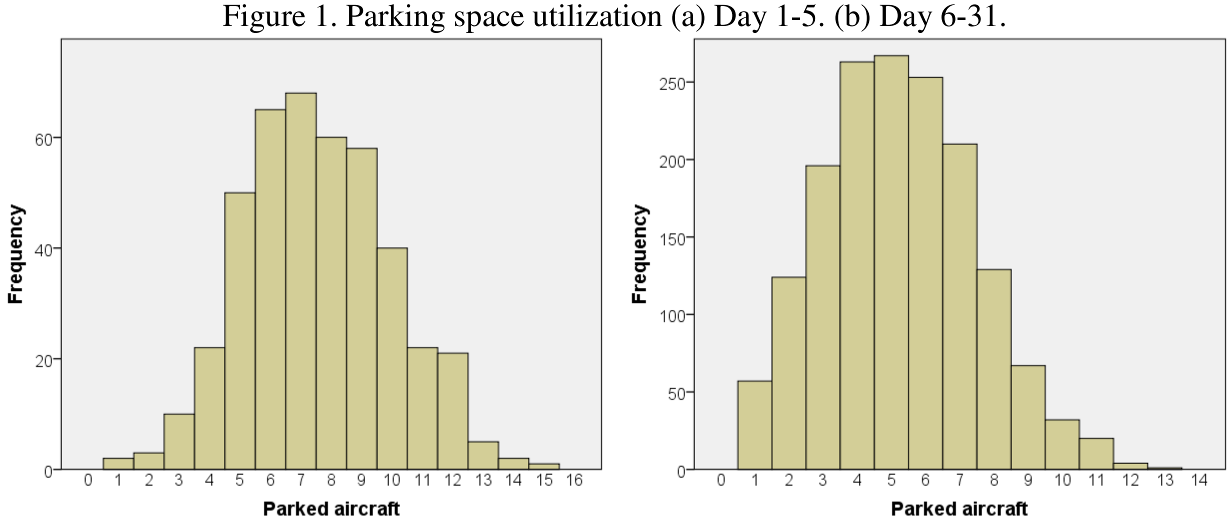

(a) (b)

Figure 2. Number of parked aircraft (a) Day 1-5. (b) Day 6-31.

3.2 Analyzing Unloading Operations with Time on Ground (TOG) Data

Slot scheduling depends on reasonable predictions for the time required to service the plane, which in this case is mostly cargo unloading. Hence, we used the flight log to also analyze unloading operations. Following the earlier analysis in Frank et al. (2011), we use TOG, which includes a short taxi time, as the time for which the parking space is occupied. Flights were excluded if they had (1) missing data, (2) parked on grass or Spot 11, (3) TOG < 15 minutes or (4) TOG > 7 hours. The reasons for this data cleaning are: only light aircraft parked on the grass and they are not relevant to the analysis of ramp parking; Spot 11 designates the U.N. ramp, which was managed separately; and the truncations eliminate likely data entry errors and screen out flights that did not deliver aid or were waiting for something other than unloading/loading. Figure 3 shows that TOG is highly variable, even when segregated by aircraft size. The large standard deviations in TOG are problematic in slot scheduling. A large proportion of narrowbody TOG is greater than the planned time of two hours. Table 1 summarizes the data and adds expected payloads, illustrating the unloading rates by size. Of the other variables recorded, only the aircraft type was helpful for explaining TOG, and then only for a few types such as C-130 and C-17; see Frank et al. (2011). Knowing the cargo type would obviously be helpful, but such data were not available at the time. A gamma distribution fit to this data is used in the analysis below. However, it should be noted that the standard deviation is sensitive to the upper truncation of 7 hours and data errors (we cleaned the data where we could), and that the data includes significant waiting times, not just unloading and loading time.

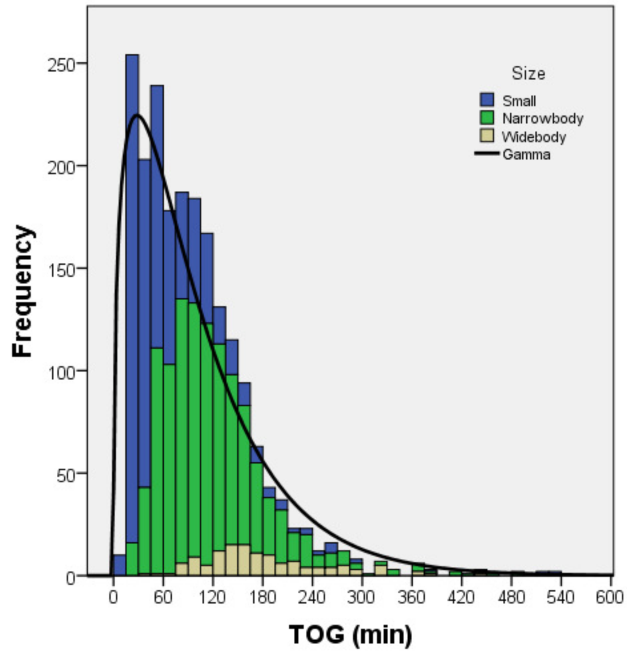

Figure 3. Time on ground, Haiti Day 1-31

Table 1. TOG statistics

Aircraft type	Number of sorties	TOG (min)		Unloading rate	
		Mean	Std. dev.	Tons/aircraft	Tons/hour/space
Small	840 (41.3%)	63	55	5	4.9
Narrowbody	1,065 (52.4%)	119	66	20	10.4
Widebody	129 (6.3%)	183	80	35	11.7
Combined	2034	99.8	71.5		

4 Models for Scheduling Airport Resources

We modeled airport operations in a humanitarian response considering the two potential bottlenecks we analyzed in the Haiti response: parking spaces and unloading capacity. Figure 4 provides an overview of the process where organizations request slots in a flight schedule that differentiates between three types of aircraft. The schedule is designed using nominal unloading times and we assume that all slots are filled by aircraft seeking space during the emergency. However, flights may not show up, or cancel in the last 24 hours, leaving an empty slot in the schedule. There may also be unscheduled high priority flights, but we assume that they do not use the same parking spaces. If the spaces are full, an arriving aircraft is diverted to another destination; there are no ground or air holds. Taxi time is assumed to be negligible and no air traffic delays are considered. Two simulation models were developed to understand different aspects of the problem. Long simulations were run to measure steady state behavior.

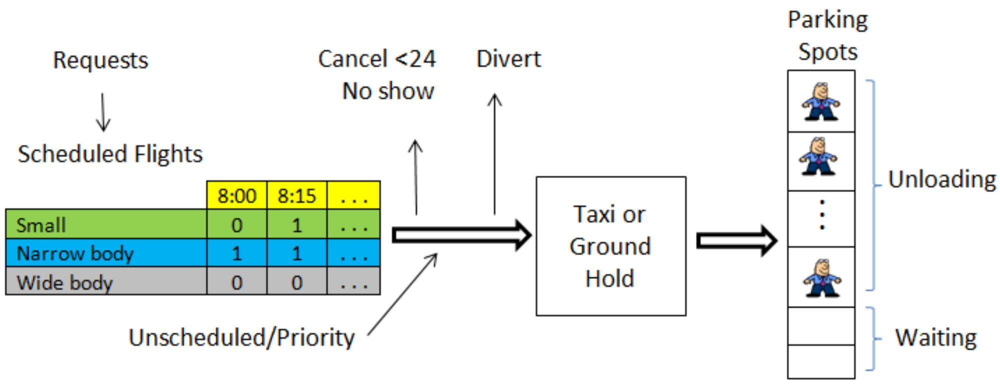

Figure 4. Airport parking and unloading bottleneck

4.1 *Scheduling Parking Slots*

The first model is a multi-server, finite buffer queue with balking. It does not distinguish between aircraft sizes in the arrival schedule. Aircraft are scheduled to arrive at equally spaced times. Each aircraft has the same no show probability. Unloading time has a gamma distribution with combined mean and standard deviation from Table 1. Three scenarios were analyzed (Table 2). We defined offered load as the fraction of time the unloading resource would be busy if there were no diverts. We chose an offered load higher than HFOCC actually scheduled to illustrate the potential throughput, and because the model does not considered additional high-priority flights.

Table 2. Airport scenarios

Scenario	Unloading spaces	Waiting spaces	Scheduled arrivals/hr	No show rate	Offered load
Small airport	1	1	2/3	25%	81%
Haiti Days 1-5	6	4	5	20%	108.5%
Haiti Days 6-31	8	2	5	20%	81%

For the small airport, limited waiting space has a large impact on the divert rate. Adding a second waiting space reduces the divert rate from 10.3% to 4.5%. Reducing the no show rate improves throughput somewhat (Figure 5). In this graph, the vertical distance between the requested and actual landings is the number of diverts per hour. The time between scheduled arrivals changes with the no show rate to keep the arrival rate (requested landings) constant.

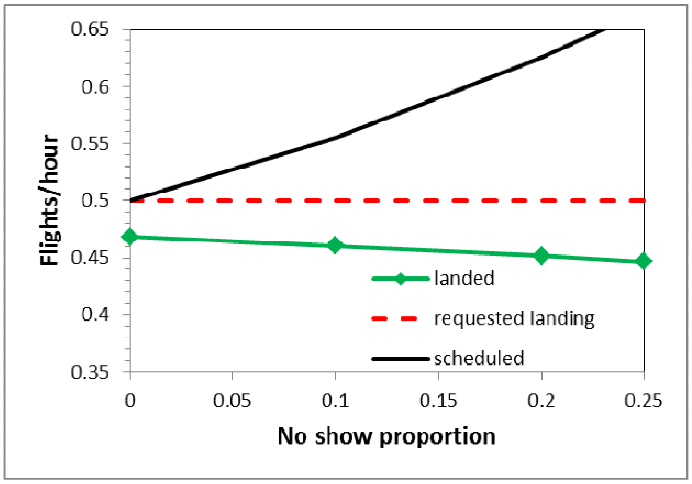

Figure 5. Impact of no show proportion, small airport

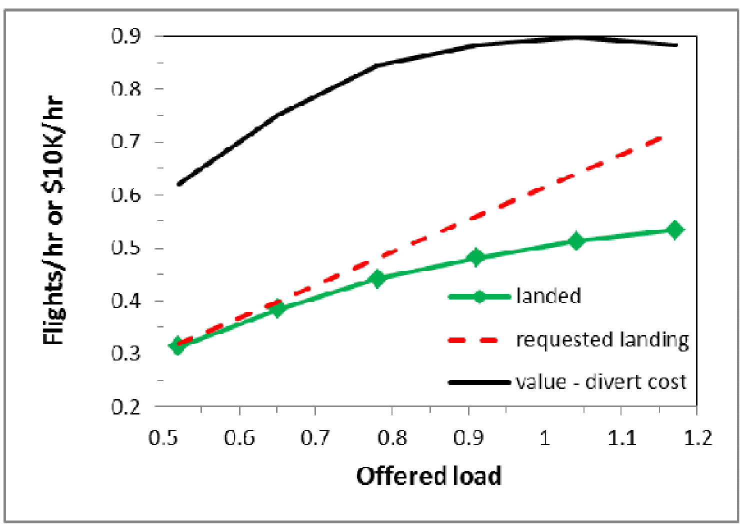

Figure 6. Benefits of overbooking, small airport

Schedules are typically built using a nominal TOG and MOG, with some under booking to allow for variability and contingencies. However, no shows and diverts lead to unused capacity. Overbooking compared to the nominal values is needed to improve throughput, although at the cost of more diverts. Figure 6 shows the tradeoff as the offered load increases. One way to make the tradeoff is by assigning a value per unloaded aircraft and a cost to divert. Using a value of $20,000 and cost of $10,000, Figure 6 shows that an offered load of 105% is optimal. This result uses a 25% no show rate, so the optimal overbooking is actually 31%.

In the Haiti days 6-31 scenario, a lower no show rate results in better throughput with only a modest increase in diverts (Figure 7). Figure 8 shows that diverts could be decreased (and throughput increased) by reducing TOG variability. Eliminating variability is more important when the system is heavily loaded, as in this day 6-31scenario. The baseline standard deviation of TOG from the Haiti data is 100 minutes. Completely eliminating unloading time variability decreased diverts from 11.1% to 8.5%, which is the lowest possible value, given the offered load of 108.5%. There is also variability in the system due to no shows, however, the four waiting spaces provide sufficient buffering so that no shows do not affect throughput (notice that the system is overloaded even after no shows are taken into account). We also tested a triangular distribution of unloading time and found that the results were very similar to those with a gamma distribution.

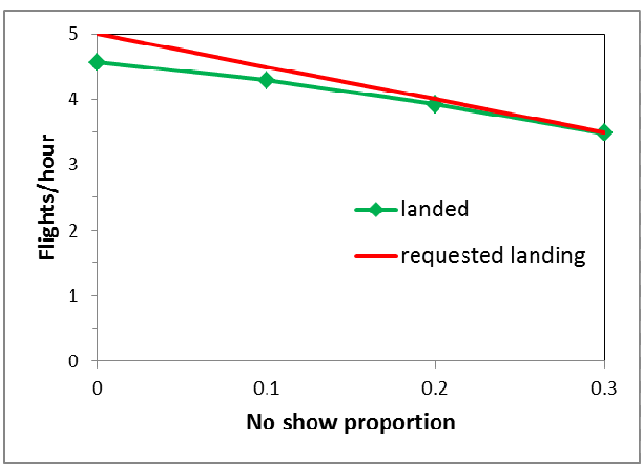

Figure 7. Impact of no show proportion on throughput, Haiti days 6-31

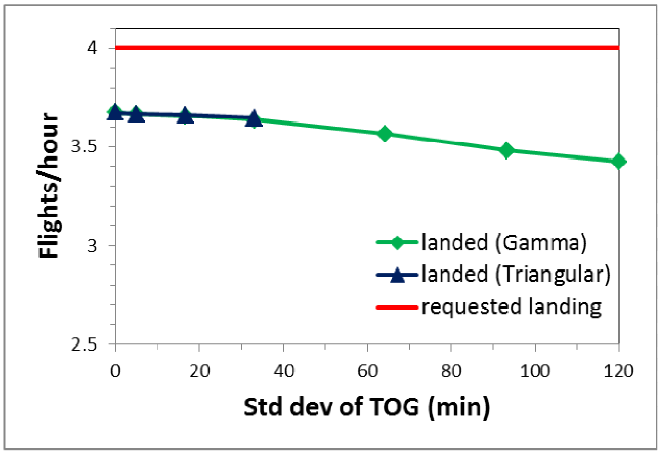

Figure 8. Impact of unload time variability on throughput, Haiti days 1-5

4.2 Scheduling by Aircraft Types

The second model allows for three aircraft types and prioritizing the types to reduce the diverting of high-priority flights. Prioritizing is done by reserving space on the ground for higher priority aircraft that will arrive within the reservation time. For example, suppose the reservation time is 60 minutes, a priority 2 flight is scheduled to arrive at 1:00, and no priority 1 flights are scheduled before 2:00. If a parking space is available at 12:00, it is reserved for the priority 2 flight. If a priority 3 flight arrives at 12:20 and all spaces are busy or reserved, it will be diverted. We assume that no shows occur before the reservation time, so that when a space is reserved for a flight it always arrives. No two aircraft arrive at exactly the same time and unloading is done first come first serve.

First, we explore the impact of plane mix on the overall throughput of the airport. We consider the Haiti days 6-31 scenario with no prioritization. Using the mix of aircraft, unloading times, and payloads from the Haiti operation, given in Table 1, results in the rightmost case in Figure 9. When the scheduled aircraft per hour is kept constant, the tons offloaded per hour increases as small aircraft are replaced with narrowbody. The use of smaller aircraft results in more flights but fewer tons delivered.

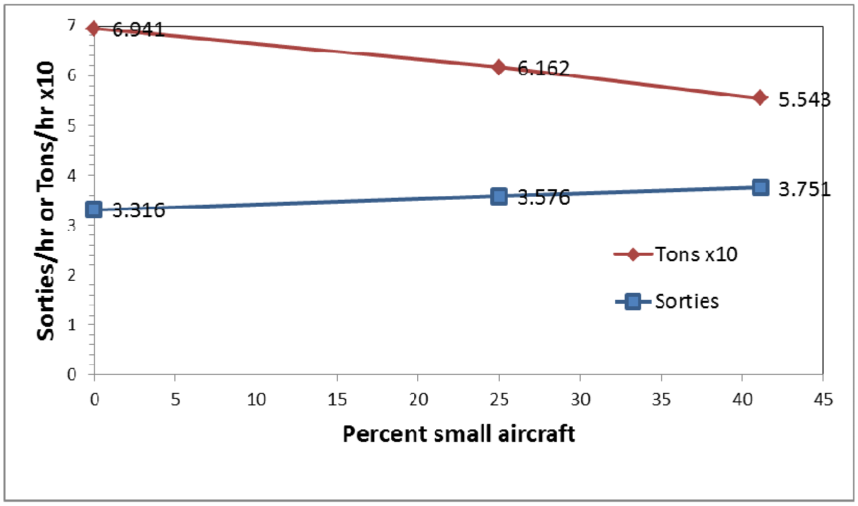

Figure 9. Replacing small aircraft with narrowbody, Haiti days 6-31

Next, we apply that insight to prioritize larger aircraft in the scheduling process. Priority 1 was given to widebody, priority 2 to narrowbody, and priority 3 to small aircraft. Figure 10 shows that a reservation time between 30 and 45 minutes maximizes throughput. When no prioritizing is done, the divert rate is 6% for all three types, and the mix of aircraft sizes is the same as in the schedule. With a 30 minute reservation time, the divert rates are 19.1% for small, 2.9% for narrowbody, and 0.5% for widebody. Longer reservation times push the large aircraft divert rates even lower but leave more spaces empty, leading to lower utilization of the unloading resource and less throughput.

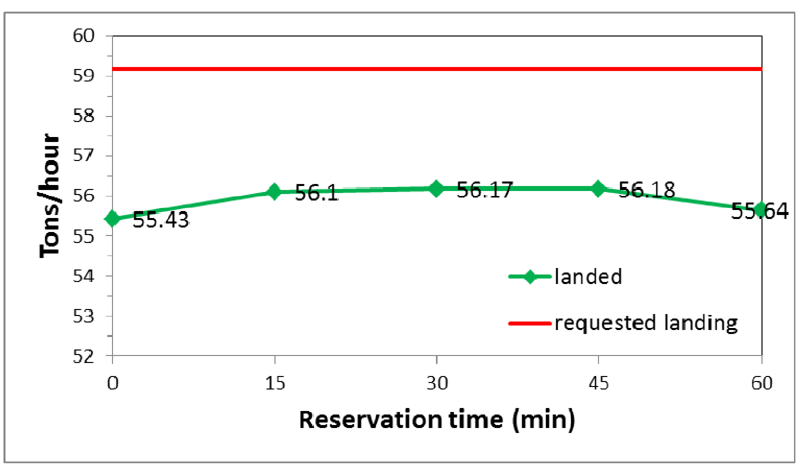

Figure 10. Reserving spaces for large aircraft, Haiti Day 6-31

Finally, we consider the impact of no shows on the operation, without prioritization. In the Haiti data, international and US civilian flights had higher no show rates than US military flights. Table 3 shows the impact of assuming the US military rate of 18% for all types. The Haiti days 6-31 scenario was used. The baseline offered load is 77%, making possible the higher throughput when there are fewer no shows. This result is similar to the single aircraft type model in Figure 7. For this analysis, flights are proportional to tons because the same mix of aircraft size was assumed for each sector.

Table 3. Impact of no shows by sector

	No show rate				Throughput (flights/day)	Diverts
	Int'l	US civil	US govt	Overall		
Haiti data	23%	29%	18%	23.8%	3.71	1.8%
Improved	18%	18%	18%	18%	3.95	3.0%

5 Conclusions and Recommendations

Our research utilizes new data from the Haiti earthquake response to provide insights and parameters to guide airport scheduling in future humanitarian crises. We conclude by offering some specific recommendations.

1. *Increased booking*. For a slot reservation system, we recommend that slots should be offered aggressively. For the critical first week of a response, the philosophy should not be that of a commercial airport, that essentially every scheduled flight is able to land, barring bad weather. Instead, the benefit of additional flights should be compared to the cost of a diverted flight. The amount of booking needed to keep the unloading resource busy could be estimated using a model such as those presented here. Potential reasons for more no shows and unloading time variability should be considered, since both require more booked flights to maximize throughput. For larger airports—those with

more parking slots—this heavier booking will have only a small effect on the divert rate. At smaller airports the effect on diverts will be larger.

2. *Air bridge.* The mix of aircraft and no show rate in the Haiti database demonstrates the potential of an air bridge to increase throughput. An air bridge is recommended when the destination airport is not able to handle widebody aircraft or when the desired volume of aid exceeds the unloading capacity. In these scenarios the increased throughput tends to outweigh the disadvantages of reduced agency autonomy and increased transport time. The air bridge will be effective after multiple agencies are delivering significant aid, and so might be implemented on day 3 or 4 of a crisis.
3. *Confirmed reservations.* No shows and late cancellations can lead to idle unloading resources at a small airport. If an air bridge is not used, a slot reservation system that discourages no shows should be used. Reservations should not be accepted far in advance and should be confirmed a day or so ahead. Priority or quota systems should be considered for allocating slots, rather than first-come first-served. Agencies should be incentivized to use the slots they reserve.
4. *Estimating unload time.* Our analysis of the Haiti ground operations data provides new estimates for unloading time of various aircraft sizes. Future operations should capture more data such as the type of cargo and the breakdown of idle and operational time on ground in order to refine estimates. More accurate estimates for unloading time would enable planners to more efficiently reserve time slots.
5. *Prioritize aircraft types.* The mix of aircraft type can have a major impact on throughput. Policies that promote the use of easily unloaded aircraft should be implemented when the desired volume of aid exceeds unloading capacity. These policies might include allocating the majority of landing slots to "efficient" aircraft types and only allowing

other aircraft in the remaining slots. Prioritizing large aircraft when making divert decisions can also increase throughput. The model results suggest reserving a ground spot for a larger aircraft that is due to arrive in the next 30 or 45 minutes, diverting smaller aircraft if necessary.

Recommendations 2-4, which seek to remove variability from the system, are consistent with recommendations in Longmire et al. (2012). However, the recommendation of increased booking is novel, based on insights from the parking space data and our model.

References

Anderson, K., F. Carr, E. Feron, and W. Hall. "Analysis and modeling of ground operations at hub airports." 3rd USA/Europe Air Traffic Management R&D Seminar, Napoli, 13-16 June 2000.

Baker, S.F., D.P. Morton, R.E. Rosenthal, and L.M. Williams. "Optimizing military airlift." Operations Research 50, no. 4 (2002): 582-602.

Davidson, S. and D. Smith. "Order from chaos: Haiti flight operations coordination center." Air Land Sea Bulletin, Air Land Sea Operations Center, January (2011): 4-9.

Dorndorf, U., A. Drexl, Y. Nikulin, and E. Pesch. "Flight gate scheduling: State-of-the-art and recent developments." Omega 35, no. 3 (2007): 326-334.

Frank, A., J. Pilch and R. Wilhelmy. "Airport operations in Haiti: Predicting estimated time on ground for post-2010 earthquake." MIT Humanitarian Response Lab report (2011).

Gempis, V. "Yokota C-130s continue to fly aid to tsunami victims." Air Force Print News, January 10, 2005.

Granger, J., A. Krishnamurthy, and S.M. Robinson. "Stochastic modeling of airlift operations." In Proceedings of the 33nd conference on Winter simulation. IEEE Computer Society (2001): 432-440

Jones, M. "CRG experience in Haiti." Air Land Sea Bulletin, Air Land Sea Operations Center, January (2011): 19–21.

Lim, A., B. Rodrigues, and Y. Zhu. "Airport gate scheduling with time windows." Artificial Intelligence Review 24, no. 1 (2005): 5-31.

Longmire, J., A. Rieger and M. Whiting. “Haiti flight operations coordination center: After action report.” U.S. Air Force 612 Air and Space Operations Center/Air Mobility Division, March 2010.

Mogilevsky, P. “Optimizing transportation of disaster relief material to support US Pacific Command foreign humanitarian assistance operations.” Naval Postgraduate School Dept. of Operations Research, 2013.

Morris, Richard R. “With compassion and hope: the story of operation UNIFIED ASSISTANCE, the Air Force support for tsunami relief operations in southeast Asia.” Office of History Headquarters, Pacific Air Forces, Hickam Air Force Base, Hawaii, January 2006.

Penny, D.C. “Crisis airlift management: effective scheduling.” Air Force Inst. of Tech. report AFIT/GTM/LAC/96S-10, 1996.

Wolfe, P., K. Nordstrom and G. Tiven. “Scheduling and prioritization of flights into Haiti after the January 2010 earthquake.” MIT Humanitarian Response Lab report, 2011.