Optimizing Humanitarian Relief Operations with Transloads

Kevin Guadagno, Kevin Saval, Quinn Van Drew, Sophia Vasiliadis, Philip Cho, and Jesse Pietz

United States Air Force Academy, Colorado 80840

Corresponding author's Email: jesse.pietz@usafa.edu

Author Note: Kevin Guadagno, Kevin Saval, Quinn Van Drew and Sophia Vasiliadis are class of 2015 graduates of the United States Air Force Academy. Philip Cho and Jesse Pietz are operations research faculty members at the United States Air Force Academy. This work is a result of a year-long operations research capstone project partnering with Headquarters Air Mobility Command Directorate of Analyses, Assessments, and Lessons Learned.

Abstract: The 2010 earthquake in Haiti caused nearly 112,000 fatalities making it one of the deadliest natural disasters ever recorded in the western hemisphere. In the weeks following the disaster, the United States Air Force Air Mobility Command moved over 29,000 passengers and 18,000 tons of relief aid in support of the relief effort, Operation Unified Response. During the first 96 hours of the operation 59% of aircraft transporting relief aid to Haiti arrived late. In order to assist planners in responding to future disaster relief efforts, we introduce a mixed integer programming (MIP) model that reduces the time required to deliver available relief aid into Haiti. The aircraft routing schedule outputted by the model demonstrates that our optimized airlift network increases the amount of relief aid delivered in the first 96 hours of the operation. Due to issues with the tractability of the MIP, we introduce an aircraft routing heuristic for use in real-world humanitarian relief operations. We show that our heuristic is able to produce similar results to the optimization, provides greater flexibility to account for realistic planning considerations, and solves within seconds.

Keywords: Haiti, Mixed-Integer Programming, Heuristics, Humanitarian Relief

1. Introduction

The Haitian Earthquake was the deadliest natural disaster ever recorded in the western hemisphere in a nation least prepared to handle it. On 12 January 2010, the 7.0 magnitude temblor caused nearly 112,000 fatalities. The United States' response was one of the largest disaster relief operations in its history. In the weeks following the disaster, the United States moved over 29,000 passengers and 18,000 tons of relief supplies. The scale of the relief effort required massive coordination between the 70 countries requesting the opportunity to land. The United States Air Force occupied approximately one third of these slot times. Despite this, it took several days for Air Mobility Command (AMC) to establish an efficient network to deliver relief aid to Haiti.

1.1 Problem Statement

At the start of the Haiti relief effort, there was no well-defined method to facilitate the efficient throughput of relief supplies. During the first 96 hours of the operation aircraft were sent to Haiti without strategy. Approximately 59% of AMC aircraft arrived late, and at less than full capacity. Furthermore, at the time of the disaster, there was only one operational airport in Haiti. Because the relief effort required international cooperation from nearly 70 nations, AMC was afforded only one third of the available landing slots. This meant that aircraft arriving late were often forced to divert to neighboring airports, delaying the arrival of life-saving humanitarian aid. This inefficiency revealed a need for a method by which decision makers could establish optimal networks for the expedient delivery of aid and the evacuation of people by reducing lateness and maximizing throughput. The purpose of this analysis is to produce a model that improves the delivery of humanitarian aid in the wake of a natural disaster using Operation Unified Response (OUR) as a baseline.

1.2 Related Work

Armacost et al. (2004) describes a modeling and algorithmic approach to an intractable network-design problem. In this case, integer programming formulation was used to optimize the United Parcel Service (UPS) delivery network to save an estimated $87 million over two years, and $189 million in the following decade. The most noteworthy modeling achievement accomplished by Armacost et al. is simultaneously determining the minimum-cost set of routes, fleet assignment, and package flows that satisfy constraints on various operating issues (aircraft types, aircraft range, aircraft capacity, airport capacity, etc.). This approach resonates with our model formulation in that we must also consider how to apply similar constraints, while meeting demands for aid, and proposing a schedule which minimizes lateness. Using some of the methods proposed by Armacost et al. (2004), we could potentially formulate our model in terms of aircraft routing decisions, and implicitly account for aid throughput. By following this example, the size of the problem can be reduced, and its tractability enhanced.

Baker et al. (2002) describes a linear programming model used for optimizing a global military airlift network for both cargo and passengers. The model takes a fleet of various aircraft, subject to many physical and policy constraints and models the deployment of these assets over time, to include aerial refueling, tactical aircraft shuttles, and crew constraints. Unlike flight scheduling models, this model's intended use is to analyze potential tactical and strategic issues in the military airlift network. The model has been used to assess questions such as "which routes are best suited for certain aircraft and war scenarios", "what was the performance of various combinations of aircraft type," and "how to improve aerial refueling operations". Since this model is aimed at analyzing policy questions, not as a tool for quick decision-making during a natural disaster, there are many aspects of the model that do not apply to our formulation. For example, in our network, supplies come from many bases and go to many bases, in contrast to the network in Baker et al. In addition, our model will not take aerial refueling into account. Aspects of the model that are useful include how to take into account different aircraft type, operating over time, and transporting both supplies and people.

Koepke et al. (2008) extends Bertsimas and Patterson's integer programming formulation (Bertsimas & Patterson, 1998) of the Multi-Airport Ground-Holding Problem to the US Air Force's air mobility network. Their formulation minimizes effects of system-wide disruptions while considering mission priority by providing a recommendation to delay certain aircraft on the ground. Airports can have a maximum number of aircraft that they can process at one time, called the maximum on ground (MOG). MOG constraints pose a difficulty in that they are often the primary factors which cause schedules to be infeasible. This integer program approach uses the MOG constraint to ensure that MOG limitations are met at each aerial port for every time period, just as our formulation does. A major difference between this formulations and our formulation is how time is accounted for. Koepke et al. (2008) uses two binary decision variables to indicate whether an aircraft has arrived and/or departed by a time period, rather than at that time period. Our formulation currently uses a decision variable that indicates whether the aircraft has arrived at a time period. Though our formulation does not currently use this binary method, improvements can be made through this approach as a way to better account for time—as this is one of the challenges that we faced.

1.3 Organization

The remainder of this article is organized as follows. In Section 2, we describe our methodology for data collection and model formulation. In Section 3, we discuss the numerical results and analysis for our optimization model. In Section 4, we describe a heuristic aircraft routing model that can be used to plan aircraft routing for humanitarian relief operations. We conclude in Section 5 with recommendations and suggestions for future work.

2. Methodology

2.1. Data Collection

Our data was provided by our client organization, Air Mobility Command. The original data set was in spreadsheet format and contained information about all of the missions during Operation Unified Response from 12 Jan 2010 to 9 Feb 2010. Below is the pertinent information that was provided for each mission:

- Mission number
- Aircraft type
- Scheduled departure time
- Actual departure time

- Scheduled arrival time
- Actual arrival time
- Departure airport codes
- Arrival airport codes
- Tons onloaded
- Tons offloaded

The data revealed that the delivery network became more organized over time, particularly after the first 96 hours of the operation. For our analysis, the first 96 hours of the operation started when the first plane with cargo landed in Haiti. Within the first 96 hours, there were 278 sorties flown and approximately 70% of the initial supplies were delivered to Haiti. The average capacity of flights landing in Haiti was approximately 37%. Our goal was to develop an optimization model that helps increase the throughput of aid, decrease the total number of sorties flown, and decrease the amount of time for the supplies to be delivered to Haiti.

2.2. Assumptions

The formulation of a mixed integer programming model that represents the airlift network relies on a number of assumptions:

1. We assume all information required for planning the operation (to include the location of available relief aid) is known with certainty at the start of the operation.
2. We assume that no unanticipated events will affect the airlift network, availability of aircraft, or delivery schedule (weather, unscheduled maintenance, accidents, etc.).
3. We assume that all aircraft operate consistently based on the averages of their listed capability (speed, range, etc.).
4. We assume that supply is immediately available and schedules will always be tenable, despite crew rest and maintenance restrictions.
5. We assume that cargo can always be loaded to aircraft capacity and that all payloads can be measured by pallets.

While these assumptions appear optimistic, it is important to point out that the model is intended only for planning in the early stages of a relief effort (first 96 hours).

2.3. Humanitarian Relief Model

We present our humanitarian relief model as a Mixed-Integer Program (MIP) based on the following:

Decision Variables

X_{ijta} (binary) = 1 if flight from airport *i* to airport *j* lands during time period *t* in aircraft type *a*
L_{ijta} (continuous [0,1]) = Aircraft Load (as a percentage of total aircraft capacity) on X_{ijta}

Parameters

C_{ijta} = Penalty of flying X_{ijta} (time interval of flight takeoff)
cap_a = Maximum aircraft capacity for aircraft type *a*
P = Percentage of supply required to get to Haiti during the time horizon
S_i = Total supply available at airport *i*
A_j = Initial number of aircraft at airport *j*
MOG_j = Maximum number of aircraft able to be on the ground at airport *j*
T = Last time period
d_{ij} = Distance between airport *i* and airport *j*
D_a = Maximum distance aircraft *a* can fly

Below is the mathematical representation of the model:

$$\min_{X,L} \sum_i \sum_j \sum_t \sum_a C_{ijta} X_{ijta} \tag{1}$$

$$\text{s.t.} \quad L_{ijta} \leq X_{ijta} \ \forall\, i,j,t,a \tag{2}$$

$$\sum_i \sum_t \sum_a cap_a L_{ijta} \geq \sum_i PS_i, j = \text{Haiti} \tag{3}$$

$$\sum_i \sum_a X_{ijta} \leq 1 \; \forall \; t, \; j = \text{Haiti} \tag{4}$$
$$0 \leq A_j + \sum_i \sum_{\tau=0}^{t} \sum_a X_{ij\tau a} - \sum_i \sum_{\tau=0}^{t+d_{jia}} \sum_a X_{ji\tau a} \leq MOG_j \; \forall \; j, t = [0, T] \tag{5}$$
$$S_j + \sum_i \sum_{\tau=0}^{t} \sum_a cap_a L_{ij\tau a} - \sum_i \sum_{\tau=0}^{t+d_{jia}} \sum_a cap_a \, L_{ij\tau a} \geq 0 \; \forall \; j, t = [0, T] \tag{6}$$
$$X_{ijta} = 0 \; \forall \; t, a, \; i = j \tag{7}$$
$$L_{ijta} = 0 \; \forall \; j, t, a, \; i = \text{Haiti} \tag{8}$$
$$d_{ij} X_{ijta} \leq D_a \; \forall \; i, j, t, a \tag{9}$$

(1) The objective function minimizes the overall cost of the operation by taking the sum product of all flights and the cost associated with each flight. Note that cost is not measured in dollars, but rather $C_{ijta} = t$.
(2) Maximum capacity constraint: ensures the load on an aircraft is less than the maximum load that type of aircraft can hold.
(3) Throughput constraint: ensures that all supplies (or percent of supplies) gets to the offload location
(4) Landing capacity constraint in Haiti. Due to single runway operations in Haiti, only one plane can land in each time interval.
(5) Maximum on Ground constraint: the number of planes able to park at an airport must be less than the MOG of that airport.
(6) Cargo constraint: ensures that the amount of cargo departing from each airport is less than or equal to the cargo available at that airport.
(7) Returning flight constraint: It ensures that none of the flights depart from and arrive at the same airport.
(8) Cargo at offload constraint: It ensures that once cargo arrives at the offload location it does not depart.
(9) Distance constraint: It limits the total distance that an aircraft can travel in one flight based on the aircraft type.

3. Optimization Results and Analysis

We programmed our model using the FICO Xpress Optimization Suite. We began with a small notional data set to validate the model. The "small model" approach had a narrow scope with only four airports (nodes) and one aircraft type. We ran the model over 100 time intervals, to represent 25 hours of the operation. Ultimately, the model scheduled 32 sorties between the 4 airports. The model took no longer than 12 time intervals to exhaust its initial supply, which is approximately equal to 6 hours assuming each time interval represents 30 minutes. All supply was at or en route to Haiti within the first 3 hours of the operation. The average aircraft capacity of flights arriving in Haiti was 100% (Table 1). The fact that the model exhausted its supply in such a short time period validates that it is accomplishing its objective function.

Following the "small model" approach, we proceeded with the "large model" approach. This model was "data-driven" utilizing the 16 airports that made the largest contribution to Operation Unified Response (by mission count). The model was still limited to 1 aircraft type, but we expanded to 192 -15 minute time intervals (representing 48 hours). The model scheduled 201 trips between the 16 airports. It took 103 time intervals to exhaust 80% of its supply (Table 1). In other words, 80% of the supply was at or en route to Haiti in just under 26 hours. The average aircraft capacity of flights landing in Haiti was 99.8%.

Table 1. Small and Large Model Parameters and Results

	# Airports	Time Intervals	# Trips	Time to Completion	Percent Supply Delivered
Small Model	4	100	32	12	100%
Large Model	16	192	201	103	80%

At this point we were comfortable applying our model on an even larger scale to more closely represent the Haiti Relief effort. We broadened the scope further, with 22 airports, and 384 -15 minute time intervals (representing 96 hours). The model scheduled 148 trips between the 22 airports, and exhausted 100% of its supply in 32 hours. The average aircraft capacity of flights landing at in Haiti was 99.2%. Figure 1 shows a comparison of the total supplies delivered to Haiti within the first 96 hours of OUR between the actual operation and our optimized model.

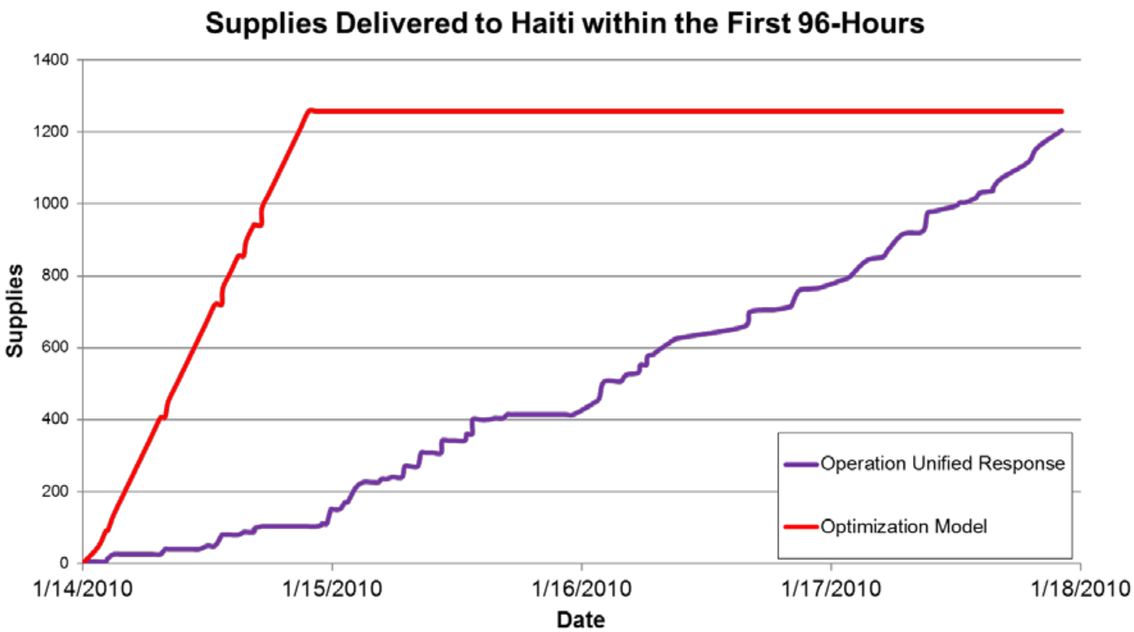

Figure 1. Comparison of Supplies Delivered to Haiti within the First 96 Hours of the Operation

As shown in Figure 1, our optimization approach would increase both the total amount of supplies delivered to Haiti as well as decrease the amount of time it would take to deliver that supplies. We observe that, due to Assumption 4, the optimization results are optimistic as they ignore aircraft ground time considerations. Ultimately, the proposed schedule and network from our model is more efficient than that used in OUR. Had it been used in the Haiti relief effort, it would have likely increased the throughput of relief aid to Haiti and saved additional lives. It is possible that using this model may have been impractical as already hours-long runtimes grow exponentially as more aircraft and airports are added to the model.

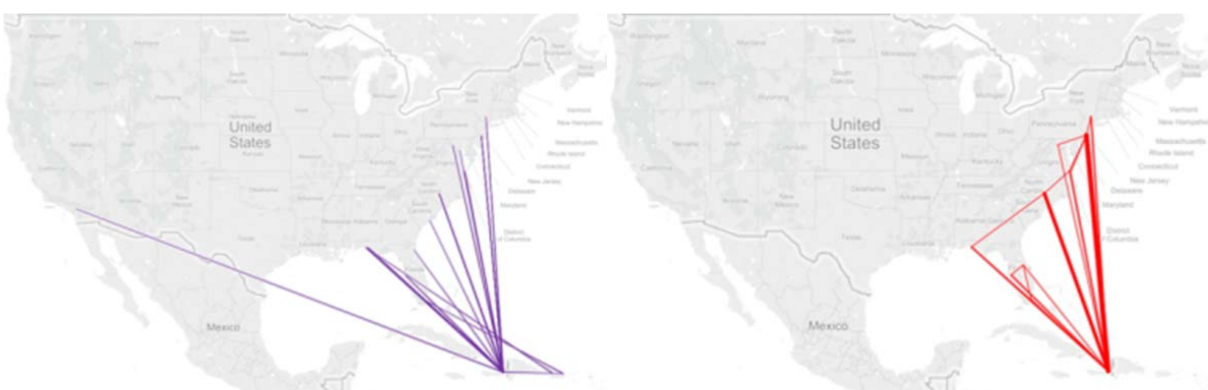

Figure 2. Comparison of Historical Network (Purple, Left) to Optimized Network (Red, Right)

Another relevant observation is the presence of transloading in the optimal solution (See Figure 2). In other words, the model scheduled aircraft to fly from more distant bases to closer, larger bases for the consolidation of supply. This observation supports the idea that transshipment is an efficient way to increase throughput to Haiti.

4. Aircraft Routing Heuristic

4.1. Motivating Observations

While the proposed optimization model yields a network that is much more efficient that the one executed during OUR, there are a few practical matters that need to be addressed in order to plan real-world humanitarian relief operations. First, aircraft missions are planned using standard ground times that account for, among other things, onload/offload, refueling, and crew changes (United States Air Force, 2011). Second, the model needs to be scalable with manageable runtimes in order to include, potentially all, airports within an area of interest and a diverse fleet of aircraft. Third, while we have observed that transloading is an element of efficient planning, real operations may dictate that a single transload base be used in order to make efficient use of support equipment. The optimization model does not identify a single transload base. Therefore, a heuristic approach to humanitarian airlift planning is imperative.

Our heuristic approach is based on four properties that we observed to be present in optimal solutions. First, aircraft landing at the relief zone offload location should be full whenever possible. Second, transloading makes for a more efficient operation. Third, the availability of landing slots at the relief zone offload location is the most limiting factor in delivering supplies. Lastly, MOG constraints (5) at supply and intermediate bases were never binding.

4.2. Heuristic Approach

An important intermediate step in developing our aircraft routing heuristic is to identify the transload base. Clearly, it is important that the base have sufficient MOG capacity, at least some planner defined *MinMOG* amount, in order to handle a large number of aircraft. Another requirement is that, in order to minimize multi-sortie missions, the base must be within the range of a single aircraft sortie of the offload location and also within single-sortie range of as many supply bases as possible. These rules give rise to the transload base function given in Algorithm 1.

```
transload base
      for each j in ListOfAirports
            if (MOG_j ≥ MinMOG AND d_{j,offload} ≤ aircraft range)
            then add j to ListOfTransloadAirports
      return j* from ListOfTransloadAirports that is within aircraft range
            of the largest number of airports
```

Algorithm 1. Transload Base Selection

Another required step is to identify intermediate, or en route, bases in the event that direct flights are not possible. A common approach to do this is to consider a “lens”, which is the overlap between the origin and destination (Sere, 2005). This approach is the basis for the intermediate base function given in Algorithm 2.

```
intermediate base (i)
      for each j in ListOfAirports
            if (j is within aircraft range of i)
            then add j to ListOfIntermediateAirports
      return j* from ListOfIntermediateAirports that is closest to offload
```

Algorithm 2. Intermediate Base Selection

The aircraft routing heuristic given in Algorithm 3 is based on the observation discussed in Section 4.1. It first attempts to schedule direct flights whenever possible. Next, it schedules flights to intermediate bases as necessary. When full flights are not possible, it uses transloading. It must also account for positioning empty aircraft in order to pick up cargo at supply bases. Lastly, it delivers all remaining supply that cannot be consolidated into a full aircraft load.

```
for t in [0,T]
      for each i in ListOfAirports
            if (aircraft is available at i AND
                  S_i exceeds aircraft capacity AND
                  i is within range of offload AND
                  landing slot t is available at offload)
            then schedule direct mission to offload
            if (aircraft is available at i AND
                  S_i exceeds aircraft capacity AND
                  i is outside range of offload AND
                  landing slot t is available at offload)
            then schedule mission to move supply to intermediate base (i)
            if (aircraft is available at i AND
                  S_i cannot fill aircraft)
            then schedule mission to move supply to transload base
            if (aircraft is not available at i)
            then pull empty aircraft from nearest base j where S_j = 0
            if (i = offload)
            then send empty aircraft to base j with largest supply to
                  aircraft capability deficit
            if (aircraft is available at i AND
                  S_i cannot fill aircraft AND
                  S_j = 0 for all other bases j)
            then schedule direct mission to offload
```

Algorithm 3. Aircraft Routing Heuristic

4.3. Heuristic Results and Analysis

We implement this heuristic in Microsoft Excel Visual Basic for Applications (VBA), a platform that is available to all planners. In order to validate our heuristic, we compare the heuristic results to the optimization model results using the same data and assumptions described in Sections 2 and 3. Most notably, we ignore aircraft ground times and assume that a homogenous fleet of aircraft is used. The heuristic produces, within seconds of computing time, an aircraft routing plan that is nearly identical (in terms of delivery time) to the optimization model's plan. There is a one-time-period difference in the delivery schedule that is caused by delivering remaining supply that cannot be consolidated (the final step in Algorithm 3). See the green delivery line in Figure 3.

Next, we remove several of the assumptions that were necessary for the MIP and use the heuristic with a mixed fleet of aircraft and require that ground times be observed. Including these more realistic considerations, we arrive at the blue delivery line in Figure 3. As expected, the heuristic delivery schedule that accounts for ground times takes longer to deliver supply to Haiti. The aircraft ground times are set to 1 hour at an offload location and 2 hours at all other bases. With these more realistic assumptions, the heuristic schedule delivers 100% of the supply in 50% of the time that it took in OUR.

It is also worthwhile to study the resulting airlift network. Figure 4 shows that the heuristic yields a substantial amount of transloading, with most occurring at the designated transload base of Charleston Air Force Base. Coincidentally, this is the same base that was used by airlift planners for transloading later on OUR.

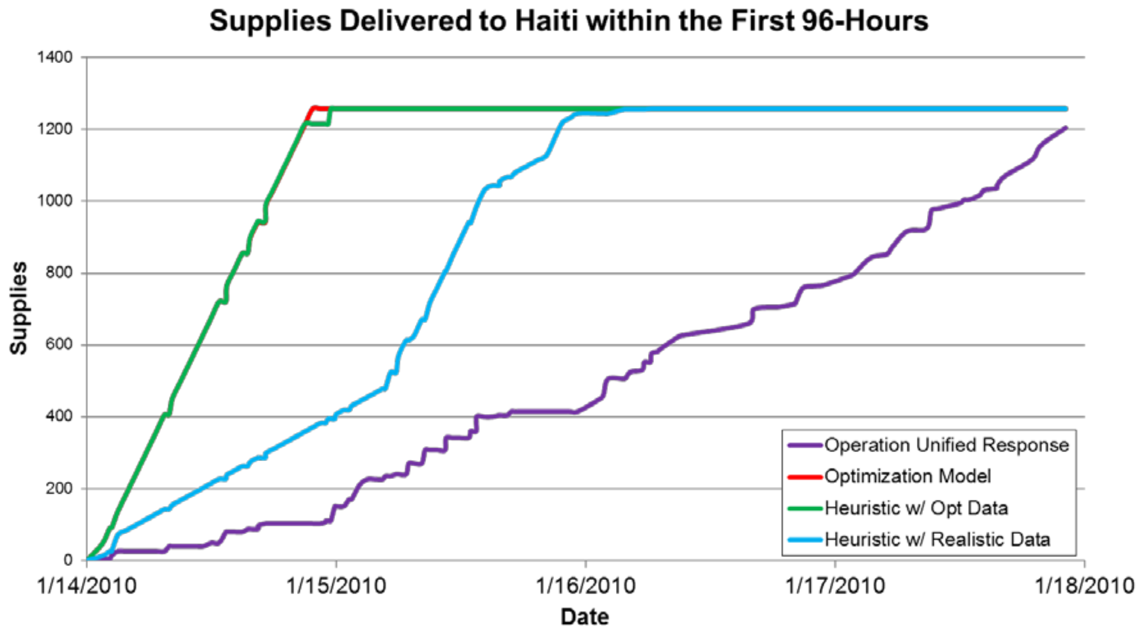

Figure 3. Comparison of Supplies Delivered to Haiti within the First 96 Hours of the Operation

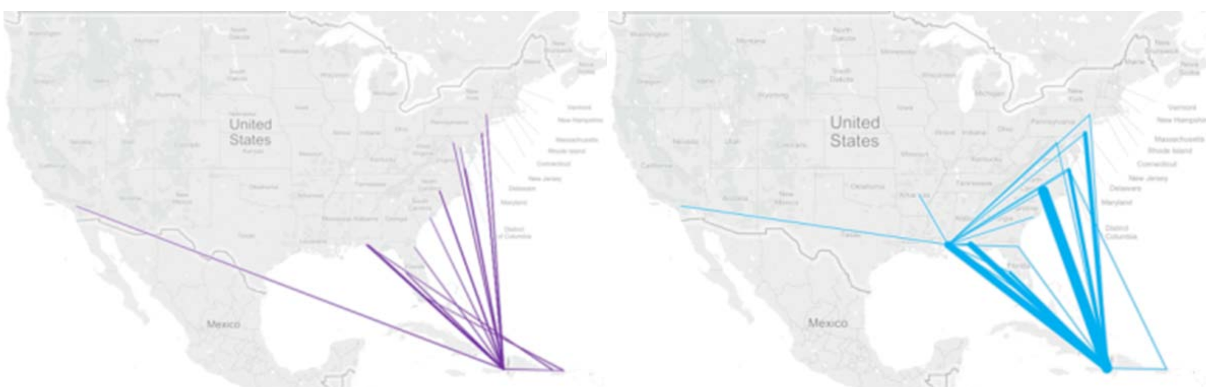

Figure 4. Comparison of Historical Network (Purple, Left) to Realistic Heuristic (Blue, Right)

4. Conclusions and Future Research

Our mixed integer programming model serves as an appropriate baseline to measure the success of AMC's method of scheduling flights in the early days of Operation Unified Response. The model proposes a schedule that theoretically delivers 100% of the humanitarian relief aid available at the start of the operation in under 32 hours. By comparison, the actual relief effort only delivered 70% of supply available at the start of the operation in the first 96 hours. This corresponds to a 30% increase in aid delivered in the first four days of Operation Unified Response. Furthermore, this additional aid is scheduled to be delivered in 1/3 of the time it took originally. The total number of sorties scheduled by our network is 148, which is 130 fewer sorties than those flown in the actual operation over the same period. In addition, the average capacity of aircraft landing in Haiti was 99.2% in our model. The integer programming model we present develops a schedule to expedite the delivery of humanitarian relief aid to Haiti, increasing throughput in the early days of disaster relief operations, and ultimately has the potential to save lives. However, it is important to note that our mixed integer programming model has significant limitations and does not capture all aspects of a realistic humanitarian relief effort due to assumptions that were made to improve tractability. Despite these limitations, the model results provide valuable insights that are used to develop a tractable heuristic that accounts for more realistic assumptions.

In order support realistic planning scenarios we describe a heuristic aircraft routing approach. We show that, under the same assumptions, this heuristic is able to produce results that are nearly equivalent to the optimal schedule generated using the mixed integer program. We then extend the heuristic by including considerations for standard ground times and a fleet of heterogeneous aircraft. By including these considerations, the heuristic accounts for more a wider range of realistic planning factors in comparison to the mixed integer program. Under these more realistic assumptions the heuristic is able to deliver 100% of the supply in 50% of the time it took during OUR. Moving forward, the heuristic based approach provides planners with a flexible tool for developing a humanitarian relief air network.

Mixed integer programs (MIPs) are powerful optimization tools, but ultimately cannot account for interdependent variance. In the future, we recommend the use of discrete event simulation to further analyze the airlift network. Simulation can better represent the highly interdependent components of the network, and account for unanticipated occurrences (unscheduled maintenance, weather, resource limitations, etc.) that the MIP fails to consider.

5. References

Armacost, A., Barnhart, C., Ware, K., & Wilson, A. (2004). UPS Optimizes Its Air Network. *Interfaces*. 34(1) 15-25.

Baker, S., Morton, D., Rosenthal, R., & Williams, L. (2002). Optimizing Military Airlift. INFORMS: Operations Research, 50(4), 582-602. Retrieved November 11, 2014, from http://pubsonline.informs.org/toc/opre/50/4

Bertsimas, D., & Patterson, S. "The air traffic flow management problem with enroute capacities." Operations research 46.3 (1998): 406-422.

Koepke, C., Armacost, A., Barnhart, C., & Kolitz, S. “An Integer Programming Approach to Support the US Air Force’s Air Mobility Network.” (2008). *Computer and Operations Research*, 35(6) 1771-1788.

Sere, M. (2005). Strategic airlift en route analysis and considerations to support the global war on terrorism (Master’s thesis, DTIC No. ADA435150). Fort Belvior, VA: Defense Technical Information Center.

United States Air Force. (2011). *Air Mobility Planning Factors* (Air Force Pamphlet 10-1403). Washington, DC: www.e-Publishing.af.mil.