A Multi-Objective Generalized Random Adaptive Search Procedure for Resolving Airspace Congestion

Christine Taylor[1] and Craig Wanke[2]
The MITRE Corporation, McLean, VA, 22102

This paper extends previous work that examines the utility of employing a Generalized Random Adaptive Search Procedure (GRASP) to minimize the impact on the National Airspace while resolving congestion. As today's methods for managing congestion are mostly manual, and for which predictions of both capacity and traffic demand are uncertain, it is often difficult to find efficient, flight-specific resolution maneuvers for the hundreds of flights affected when congestion arises. This work continues to investigate improvements upon previous research that develops a non-optimal flight-specific congestion solution strategy based on a prioritized flight list by iteratively perturbing this prioritization to find better solutions. In this paper we examine the previous prioritization criteria considered and develop a new prioritization criterion that incorporates the two individual criteria. The research shows that combining the two criteria effectively bridges the difference between the results produced by the single criteria resulting in consistently good solutions across multiple metrics and objectives. Following this, we extend the optimization framework to include multi-metric objectives, examining the trade-off of congestion versus delay and how including inequity in the objective affects all three metrics. The research shows that significant benefits can be derived even when a metric is only lightly weighted in the objective, showing that multi-metric objectives are a significantly desirable formulation.

Nomenclature

ADM	=	Aggregate Demand Model
ATC	=	Air Traffic Control
CMA	=	Congestion Management Area
CRA	=	Congestion Resolution Area
ETMS	=	Enhanced Traffic Management System
GRASP	=	Generalized Random Adaptive Search Procedure
LAT	=	Look-ahead Time
MAP	=	Monitor and Alert Parameter
NAS	=	National Airspace System
TFM	=	Traffic Flow Management

I. Introduction

WITH the predicted increase in airspace demand over the next few years, new systems to aid in en-route congestion resolution are needed. Present day en-route congestion resolution occurs when the predicted demand exceeds the nominal capacity for a sector. Significant congestion problems typically involve hundreds of flights, and affect multiple en-route ATC facilities and U.S. National Airspace System (NAS) users (e.g. airlines). Currently, traffic managers in the NAS resolve congestion primarily through manual processes, relying on experience and limited traffic prediction data[1]. However, the scope of these congestion problems can cause difficulties for the human decision maker as they need to make effective and coordinated solutions in real time.

Identification of potential congestion relies on both the prediction of airspace demand and capacity over a several hour time horizon. In the NAS demand predictions are generated by the Enhanced Traffic Management System (ETMS) for most sectors over several hours, in 15 min time bins. These counts are based on filed flight

[1] Senior Simulation Modeling Engineer, M/S 450, AIAA Member.
[2] Senior Principle Simulation Modeling Engineer, M/S 450, Senior Member AIAA.

plans or schedule data, wind forecasts, and radar track reports for aircraft that are already airborne. However, as deviations from the predicted traffic demand arise frequently from such sources as, modifications in flight schedules or flight cancellations, the initiation of previously unscheduled flights, or the effect of weather on cruise altitude and air traffic control (ATC) intervention, it is necessary to view traffic demand as uncertain. The magnitude and characteristics of these uncertainties have been extensively described[2], measured[3], and modeled in the context of sector load forecasting[4-6].

Similarly, the nominal capacity of a sector is defined by a threshold called the Monitor/Alert Parameter (MAP), which provides an abstraction of the sector capacity, although does not capture the intricacies affecting controller workload[7-8]. Uncertainty is also inherent in the prediction of sector capacity. While ETMS generates alerts based on constant aircraft count thresholds, it is widely accepted that the real capacity of sectors depends on traffic complexity and weather, and should also be treated probabilistically.

Given the inherent uncertainty in the data provided to traffic managers, these predictions are of limited use, as they provide one data source for traffic managers who must develop ground delay or reroute initiatives to control congestion. There are no decision-support tools currently available to test traffic management initiatives, though extensive work has been done to develop such tools[9-11]. Furthermore, these initiatives typically affect flows of aircraft (e.g., rerouting all traffic between a pair of airports, or miles-in-trail spacing restrictions) rather than individual flights.

In order to alleviate some of these difficulties, pre-defined large-scale strategies, such as "National Playbook" routes[1] have become a standard method for addressing congestion. Given the prediction uncertainty arising from weather forecasts and flight plans and the effect of using large-scale resolution strategies, traffic managers tend to develop highly conservative decision making practices which can potentially result in unnecessary delays. As such, it is desirable to develop an automated support system that assists in the identification and resolution of congestion.

This paper presents research that is part of a broader effort to analyze traffic flow management (TFM) uncertainty through probabilistic modeling and to develop new decision support concepts that integrate this analysis in the congestion resolution process. Previous work in this area has investigated the use of a prioritized sort list for ordering flights that employs a single-pass deterministic greedy algorithm[12]. Further research has examined globally defined solutions through optimization with a genetic algorithm[13] and the Multi-Objective Hybrid Genetic Algorithm[14], which provide a 'global' solution to the problem but are computationally intensive. A potentially useful alternative that provides a compromise to both approaches was initially developed in Taylor et al[15] and this work continues to extend the concepts and address some issues presented.

Specifically, the focus of this paper is the development of a heuristic optimization algorithm that provides an improved congestion resolution strategy without significantly increasing the computational requirements needed to obtain these solutions. The heuristic algorithm, a Generalized Random Adaptive Search Procedure (GRASP), has been shown to aid in problems in which an ordering scheme must be defined but the exact nature of a desirable prioritization criterion is not fully understood[16-17]. The current work focuses on extending GRASP to a multi-metric objective framework for resolving airspace congestion. By analyzing the solution impact due to different relative weightings in the objective function and the application of an alternative prioritization criterion, further insight into the underlying behavior of the problem can be realized.

The benefits of the multi-objective GRASP approach are evaluated by way of an illustrative example problem. The next section discusses the general problem formulation that motivates the work presented in this paper. Section III presents the fundamentals of the GRASP algorithm, describes the different prioritization criteria used, details the metrics considered for inclusion in the multi-objective optimization and outlines the example problem considered. Section IV develops the new prioritization criterion considered in this paper and compares the results for single-metric optimization. Section V presents the results of multi-metric optimization with various weighting factors and examines how the inclusion of inequity in the objective function influences both delay and congestion. The conclusions and continuing work are discussed in Section V.

II. Problem Formulation

The problem under consideration is the resolution of predicted airspace congestion, in which congestion is simply defined as the condition that demand exceeds capacity. Given that both demand and capacity estimates are uncertain, we instead examine the probability of congestion, which identifies the probability that the expected demand exceeds the expected capacity. Therefore, to define congestion, estimates of both demand and capacity are required.

In this work, The Aggregate Demand Model (ADM) developed by Wanke et al[6] is utilized to predict the demand at different look-ahead times (LATs) for each sector. The ADM is a closed-form, statistical uncertainty model for en route sector demand predictions, suitable for both simulation and real-time applications. It is based on a comprehensive set of statistics for sector peak count prediction uncertainty collected over an 8 month period. The model forecasts peak traffic demand distributions based on four variables: the look-ahead time, the deterministic predicted peak count, the number of airborne flights in the peak count prediction, and the primary sector traffic type (departure, en-route, arrival, mixed).

The capacity of a sector is also a probabilistic quantity, given the uncertainty of weather and complexity on capacity. Although, an initial attempt has been made to define a probabilistic sector capacity metric[18], this work is still preliminary. As such, although only a rough abstraction of sector capacity, the MAP is the sector capacity value used in this work to identify areas of potential congestion.

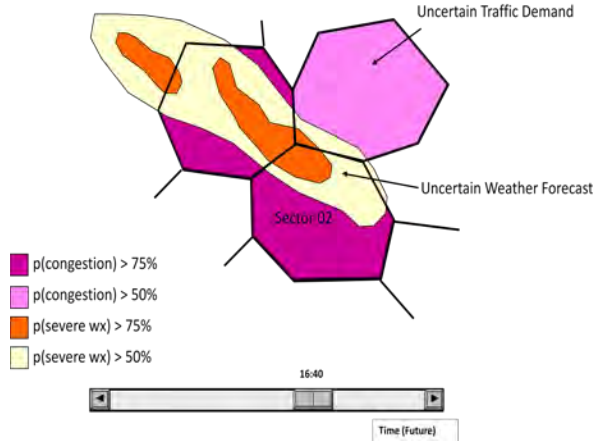

Figure 1. Probabilistic Sector Forecast

The resulting definition of the probability of congestion is obtained from the demand and capacity estimates as follows. The predicted traffic demand distribution and the predicted sector capacity distribution are convolved to determine predicted congestion distribution. The resulting congestion threshold is identified as the area of the convolved distribution where there is a probability that demand exceeds capacity.

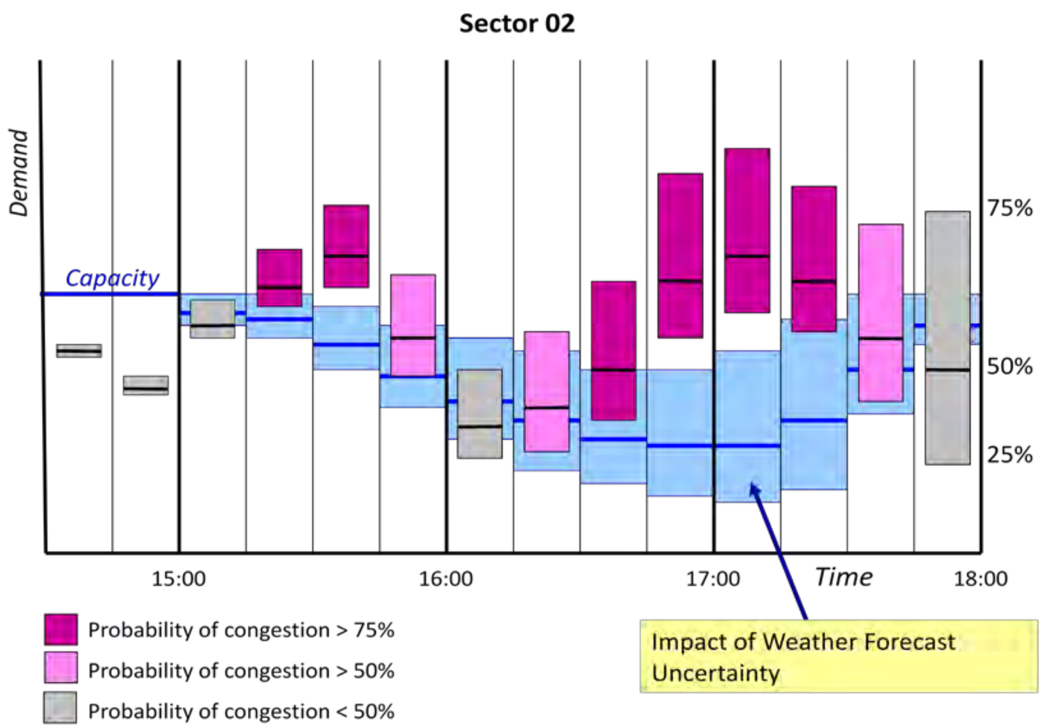

Figure 2. Detailed Congestion Predictions for Single Center

Figure 1 presents a notional probabilistic forecast of weather and congestion. The pink and purple codes indicate the probability of congestion for the three sectors shown. Sectors are classified into two categories: *congestion resolution area* (CRA) and *congestion management area* (CMA). The CRA consists of sectors in which congestion is predicted and resolution actions can be taken to reroute or ground delay flights as necessary. The CMA sectors correspond to neighboring sectors that must be monitored to ensure that resolution actions taken in the CRA do not adversely impact these neighboring sections.

To clearly understand the evolution of congestion with look-ahead time, the congestion probabilities can be decomposed into 15 minute intervals several hours into the future. Figure 2 presents this time-series information for

Sector 02 from Figure 1. The blue line represents the 50th percentile prediction of airspace capacity. As we can see, the uncertainty in the capacity estimates are small at short LAT, since the weather is not expected to impact that area until 30 minutes into the future (1500). For greater LAT, the weather is expected to reduce the capacity; however, the greater ranges reflect the uncertainty in the future position, size, and intensity of the weather. The gray, pink, and purple boxes reflect probabilistic demand predictions, for which the heights of the boxes, increasing with greater LAT, reflect the associated uncertainties in these predictions. The bottom, midline, and top of the boxes represent the 30th, 50th, and 80th percentile of the predicted demand distribution, in this example. The boxes are color-coded to reflect the probability that the actual demand exceeds the actual capacity.

Given the uncertainties in demand and capacity estimates, the resulting congestion management goal can be expressed as a maximum congestion probability, or "congestion risk." For example, a congestion management goal could be to take action such that there is a maximum congestion risk of 75\% for the next three hours. This is shown in Figure 2 for a single sector. The resolution problem, then, is to determine a set of flight-specific maneuvers that achieve this risk reduction while minimizing impact on airspace users. For the problem considered in this paper, the resolution actions can be the imposition of a ground delay, reroute, or both for a given flight. Other possible resolution actions that were not explored include cruise altitude changes, and imposition of time constraints on flight path waypoints (metering).

III. The GRASP Algorithm

The GRASP algorithm evaluates multiple potential solutions and returns the best solution found. Potential solutions are defined by randomly perturbing the prioritization scheme and resulting sort order of the candidate flight list, where the sorted candidate flight list determines which flights receive an undesirable resolution action in a given problem. GRASP was utilized to obtain benefits over a previously implemented deterministic greedy heuristic, but operates under a similar framework. Details of the greedy heuristic algorithm can be found in [Ref 12].

Figure 3 shows the overall implementation of the GRASP algorithm. As we can see from Figure 3, GRASP develops a sorted candidate flight list that modifies the absolute implementation of a priority order to account for unknown probabilistic impacts. Given that the sorted list is generated probabilistically, each iteration produces a

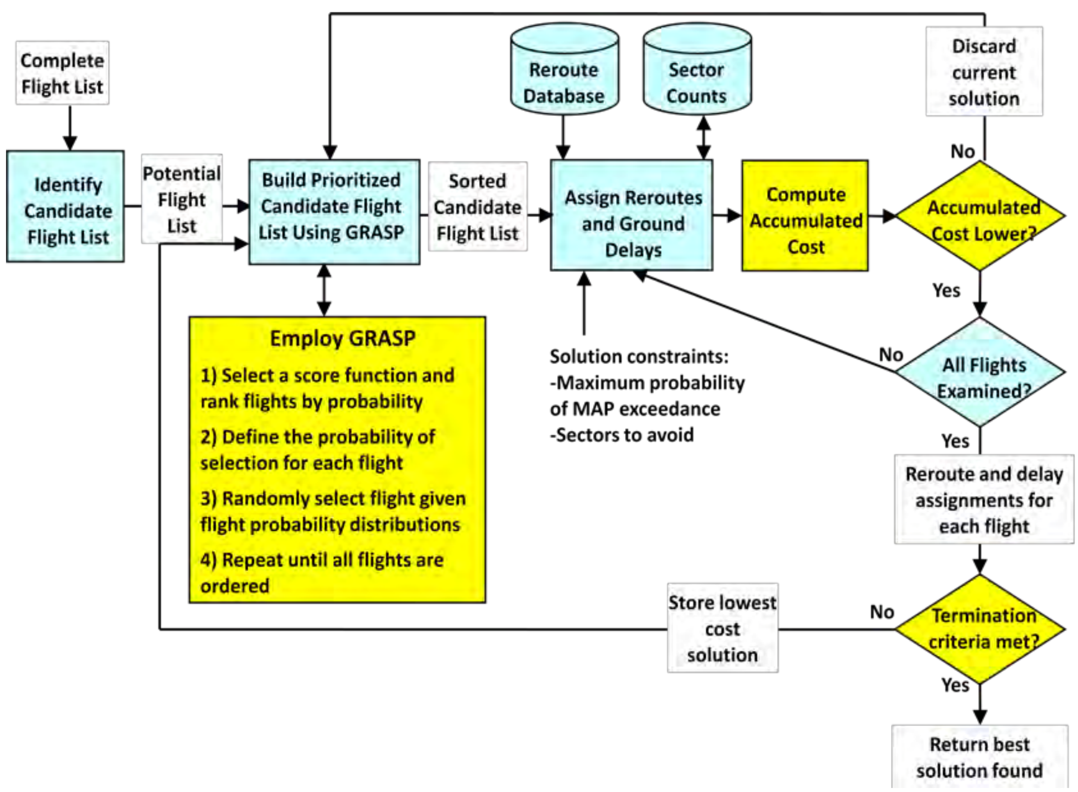

Figure 3. Overview of GRASP Congestion Management Algorithm

slightly different list which can change the overall impact on the system. Therefore, multiple iterations are performed until the termination criteria are met. A key aspect of the GRASP algorithm is how prioritization criteria are used to formulate the sorted candidate flight list.

A. Prioritization Criteria Selection

The GRASP algorithm uses prioritization criteria to influence the position of a flight in the sort order, in which priority in the sort order reflects a decreased likelihood of being assigned an undesirable resolution action. Therefore, it is desirable that the prioritization criterion selected align with the objective function, so that the ordering of the flights promotes good solutions.

This research considers two different prioritization criteria, namely time to congestion resolution area (CRA) and total time spent in the CRA. Using the prioritization criterion of time to CRA prioritizes flights entering the CRA sooner than other flights. This criterion represents how flights are traditionally prioritized in order to promote transparency and fairness in the decision process. Prioritizing by time spent in the CRA reflects that some flights have a greater impact on the system because they travel through the CRA longer. Therefore, if the flights that are in the CRA longer are prioritized lower and receive a resolution action, that resolution action may have a larger marginal benefit to the system than a resolution action for another flight that only minimally impacts the system. For each flight, the prioritization criteria is computed based on the original route of the flight and scaled to appropriate units as necessary.

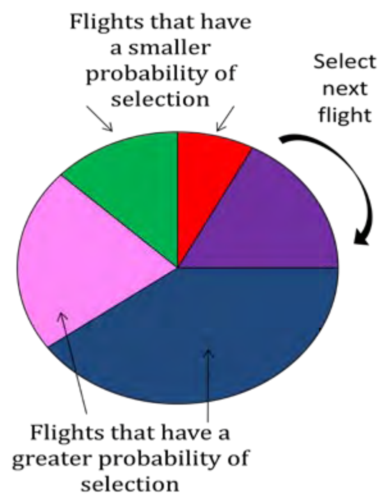

Figure 4. Illustrative Flight Probability Distribution

B. Development of a Sorted Candidate Flight List

A key aspect of the GRASP algorithm is in the construction of the sorted candidate flight list. Instead of strictly ordering the flights by the prioritization criteria, the GRASP algorithm utilizes a set of score functions to inform the sorted flight list construction. The score functions represent different weightings of the prioritization criteria and are constructed such that lower scores are more beneficial than higher scores for priority in the sort list.

For each prioritization criteria (pc) selected, multiple score functions are defined that vary the relative weighting of priority importance between flights. The score functions used to evaluate the prioritization criteria described above are:

$$\begin{array}{ll} S_1 = pc & S_2 = \sqrt{pc} \\ S_3 = pc^2 & S_4 = pc^{1.5} \\ \multicolumn{2}{c}{S_5 = \sqrt[3]{pc}} \end{array} \tag{1}$$

In each iteration of GRASP, one of the score functions is selected uniformly at random from the set of defined score functions. Using the selected score function, a score is calculated for each flight in the candidate list; however this score does not automatically determine the ordering of the flight in the sorted candidate list. Instead, the GRASP algorithm uses these scores to determine the probability of selection of a given flight for priority ordering. That probability of selection is defined as

$$\boldsymbol{P}(\boldsymbol{f}) = \frac{e^{-S_i^f}}{\sum_f e^{-S_i^f}} \tag{2}$$

in which f is a flight in the candidate list, S_i^f is the score of flight f evaluated for score function i, and $P(f)$ is the corresponding probability of selecting that flight. As the summed probability over all flights is one, the probability of each flight can be viewed as the relative proportion of selection.

Figure 4 illustrates this process through a simplified example. Each sector of the circle represents the probability of selection of a given flight. For each slot in the sorted candidate flight list, the wheel is "spun" and the next flight in selected. These probabilities are computed by defining the prioritization criterion and GRASP randomly selecting a score function.

Flights with high probabilities of selection are more likely to be selected early in the process and therefore have a high priority on the sort list. However, the relative probability values do not determine the final sort order in the flight list. Instead, a flight is selected when a randomly generated number falls within the probability sector of the flight. This process is repeated until all flights have been ordered in the candidate list.

C. Measuring Solution Quality

Multiple metrics are available to assess the quality of the resolution strategies developed, depending on organizational priorities. For this research, three metrics were defined: total delay, probability of congestion, and inequity. The goal is to minimize the value of the objective function, defined as one or more of those metrics.

The metric of total delay (J_d) selects the solution that provides the minimal increase in total system delay while attempting to meet the congestion target. Both ground delay (g) and airborne delay (a) contribute to the total delay. Only positive airborne delay is considered here to ensure that resolution actions that reroute flights to arrive at their destinations earlier than originally planned are neither penalized nor rewarded. The mathematical expression of delay is the sum of those delays:

$$J_d = \sum_f \left[g_f + \min(a_f, 0)\right] \tag{3}$$

The metric of congestion compares potential solutions on the basis of which option produces the minimum sector congestion. Under all metrics, the congestion target is provided and resolution maneuvers are selected to best meet this target; however this target cannot always be met given previous decisions and the desire to meet other objectives. The congestion metric (J_c) reflects the system level goal of reducing sector congestion, in which the impact on individual flights is a secondary consideration:

$$J_c = \sum_{i=1}^{N_c} \left[(\lambda_a(P_i - P_m) + \lambda_b(P_i - P_m)^2)\lambda_c\left(\frac{L_m - L_i}{L_m}\right)\right] \tag{4}$$

The quadratic cost term for positive deviations of congestion from the target level gives extra weight when the congestion of a sector (P_i) deviates from the sector specific target level (P_m). A decreasing linear cost on look-ahead time is included to represent that congestion occurring at a given LAT (L_i) is less concerning, all things being equal, when it is closest to the maximum look-ahead time (L_m). The weighting factors λ_a, λ_b, and λ_c represent different relative weightings of the components in the metric, in which the value of these factors were previously chosen through an off-line experimental study[13]. The total congestion is added for every sector and look-ahead time combination (N_c) in which congestion is present.

The final metric considered is inequity. Equity, as defined in this research, is when the delay incurred by the assignment of resolution maneuvers is as evenly distributed as possible over the different NAS customers. Therefore, the inequity metric (J_e) is defined as the minimum standard deviation of delay between customers. The total number of flights affected by delay is represented by N_F and the number of customers with these affected flights is represented by N_{ar}. For the purpose of this analysis, all general aviation (GA) flights are grouped together as a single customer.

$$J_e = \frac{\sum_{k=1}^{N_{ar}} \left(\overline{Delay}_k - Delay_{N_{ar}}\right)^2}{N_{ar} - 1} \tag{5}$$

in which

$$\overline{Delay}_k = \frac{\sum_{f=1}^{N_F} [g_f + \min(a_f, 0)]}{N_F} \tag{6}$$

$$Delay_{N_{ar}} = \frac{\overline{Delay}_k}{N_{ar}} \tag{7}$$

D. Assigning Resolution Actions

Given the prioritized candidate list, the algorithm proceeds by selecting each flight in turn and evaluating its nominal path against the congestion risk target to determine what (if any) action needs to be taken. The flight is then assigned a route and the total system cost incurred for that flight is computed, based on the selected metric to be minimized.

Because the GRASP algorithm minimizes the total system cost, the cost incurred by the resolution action defined for each flight is iteratively added to the system cost, in order to measure the performance of the current solution. If the current solution's accumulated cost is greater than the minimum cost solution found up to that point, the iteration is terminated and a new priority list is constructed. If instead, after all flights have been evaluated, the current solution has less total cost, it replaces the best solution found. The best solution found over all iterations is returned when the termination criteria, in this case defined to be the maximum number of iterations, is met.

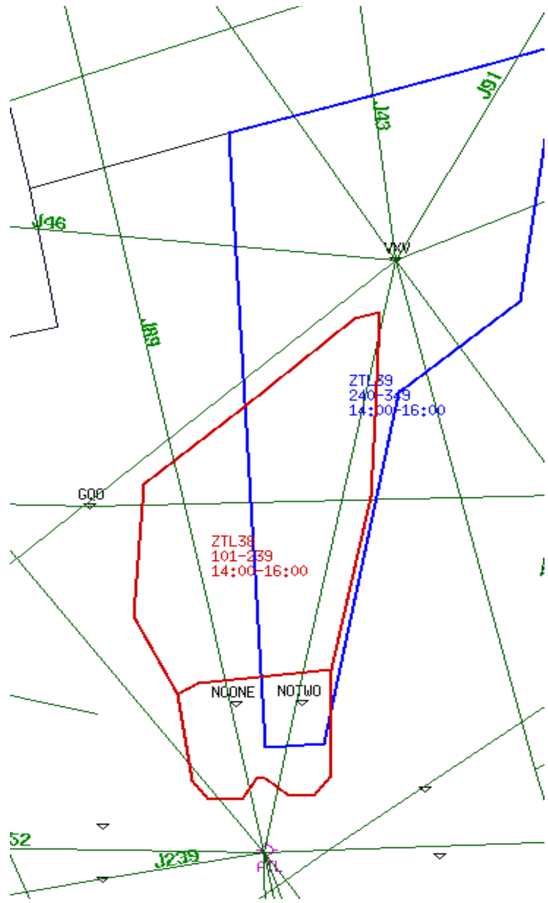

Figure 5. Geometry for Example Problem

E. Example Problem

In order to evaluate the performance of the GRASP algorithm, we constructed a sample and compared the results provided by each solution method to determine how effectively they resolve the congestion. The example congestion problem was developed from traffic patterns and predictions observed during January 2004. The traffic predictions used were from a TFM decision-support prototype[19] that uses a probabilistic traffic demand model, namely the Aggregate Demand Model (ADM) described in Reference 13, to determine the probability that the demand on a specific sector exceeds the normal Monitor Alert Parameter (MAP) value.

Figure 5 illustrates the geometry of the airspace structure around Atlanta Hartsfield International Airport (ATL), which is the example problem considered in this research. The CRA consists of two sectors: ZTL38 and ZTL39. ZTL38 is a low altitude sector, controlling altitudes from 10,000 up to 24,000 feet, and primarily handles the departure traffic over fixes NOONE and NOTWO. Departures from NOTWO transition into sector ZTL39, which is a high altitude sector handling traffic from 24,000 up to 35,000 feet. ZTL39 also handles a complex pattern of cruising traffic, as illustrated by the jet airways intersecting at the VXV navigational aid.

Figure 6 shows the congestion situation in the form of predicted sector counts. This is a probabilistic Center Monitor (CM) that uses the ADM to estimate peak traffic counts and generate alerts. In this version, the median predicted peak count for each sector is shown, thus compensating for prediction biases. Also, the purple and pink

	1400	1415	1430	1445	1500	1515	1530	1545
ZTL34 [12]	0*	2	4	4	8	9	6	2
ZTL36 [13]	0*	2	3	4	8	9	9	3
ZTL37 [13]	1	2	2	7	10	12	13	4
ZTL38 [13]	1	5	7	14	16	14	10	2
ZTL39 [13]	1	2	11	8	20	21	10	6
ZTL40 [18]	1	4	7	6	6	8	9	5

Figure 6. Probabilistic Congestion Display for Example Problem

alert colors are based directly on probabilities of congestion. A purple alert indicates a higher than 75% probability that the actual peak traffic demand will exceed the MAP. A pink alert indicates a 50% to 75% probability. The maximum congestion target for the example problem defined is 0.5, meaning that the goal is to obtain a solution that reduces congestion probabilities below 50%.

In Figure 6 both sectors 38 and 39 show congestion alerts between for the time periods between 45 and 90 minutes LAT (between 1445 and 1530), with a maximum probability of congestion of 0.96. Flows through these two CRA sectors are analyzed to determine all possible sectors to which flights in the CRA can exit, and from which sectors flights can enter. The 29 adjacent sectors identified constitute the congestion management area (CMA) for this problem and have a maximum probability of congestion of 0.61.

The goal of the GRASP algorithm is to determine a solution that meets the congestion target of 0.5 in the CRA, without adversely impacting the congestion in the CMA and with minimum total cost. For each prioritization criteria and objective function defined, the GRASP algorithm is run twenty-five times, with each run consisting of 100 iterations, and the results presented are averaged over all runs.

IV. Prioritization Criteria Development

In previous work[15] we showed that the two prioritization criteria, time to CRA and time spent in the CRA, provided improved solution quality as measured by the three metrics considered. However, there wasn't a single prioritization criterion that provided high quality solutions for each metric. Instead, the choice of prioritization criteria varied. As such, FY09 research delved deeper into the underlying dynamics of the prioritization criteria and score function selection to develop alternative criteria.

A. Impact of Score Functions

The value of implementing a variety of score functions and randomly choosing a score function for an iteration lies the change of the relative probability of selection in the sorted order for a given flight. Modifying the relative probability of selection in effect changes how likely flights are to swapped in the sorted list order, potentially leading to improved solutions.

The score functions described in Eqs. (1) provide five different probability distributions for each of the two prioritization criteria examined. The probability distributions are computed for each prioritization criteria using the original flight information of the 64 flights in the example problem described in Section III.E. The curves are generated by applying a fit to the histogram of probabilities generated using the given prioritization criteria and the chosen score function.

Using score function #5 a few flights have high probabilities of selection and therefore will most likely be selected first in the sort order. The majority of flights has relatively the same probability of selection and therefore can easily be interchanged, thereby producing vastly different sorted lists. Alternatively, using score function #3, there is a larger selection of flights with higher probabilities, yielding a greater interchange of orders at the top of the list. These flights however, are more likely to be chosen before the flights with lower probabilities. Effectively, this probability distribution bins the flights in the sort order, allowing only localized movement in the list.

Figure 7 shows the probability distributions associated with the score functions when selecting the prioritization criteria of time to CRA. Examining Figure 7 shows that each score function produces a different probability distribution; however they all do share the general characteristic that a high proportion of flights have a probability near 0.015. For example, taking the distribution of score function #5 (S_5), we see a high peak at around .015. Alternatively, examining score function #3 (S_3), we see a smaller peak at around the same probability, however instead of falling off as drastically at higher probabilities, a smaller but significant number of flights have a probability of 0.02. The probable impact of these different probability distributions in determining the sort order is as follows.

Examining the probability distributions for the score function when using the prioritization criteria of time spent in the CRA reveals a different picture. Figure 8 shows a more uniform probability distribution for each of the five score functions, where the general shape is the same for all, with only small shifts in height, width and location of the peaks. Selecting the prioritization criterion to be the time spent in the CRA produces different ordering than time spent in the CRA; however, both have been shown in previous research[15] to improve the solution quality as compared to solutions produced through deterministic orderings, but more improvements can be gained. Furthermore, we are searching for a single prioritization criterion that consistently produces high quality solutions under multiple objective functions.

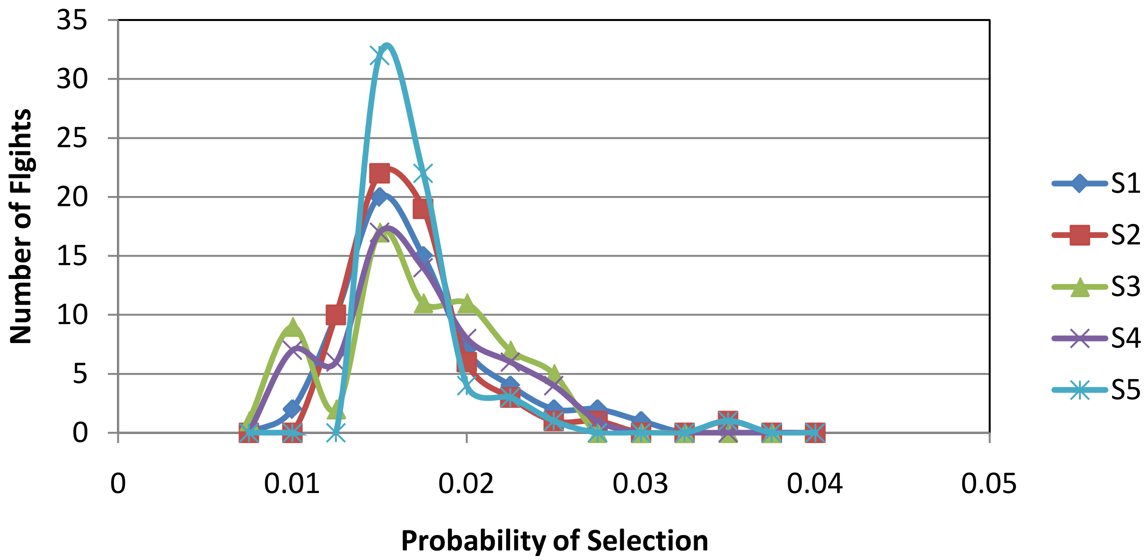

Figure 7. Probability Distributions of Score Functions for Time to CRA

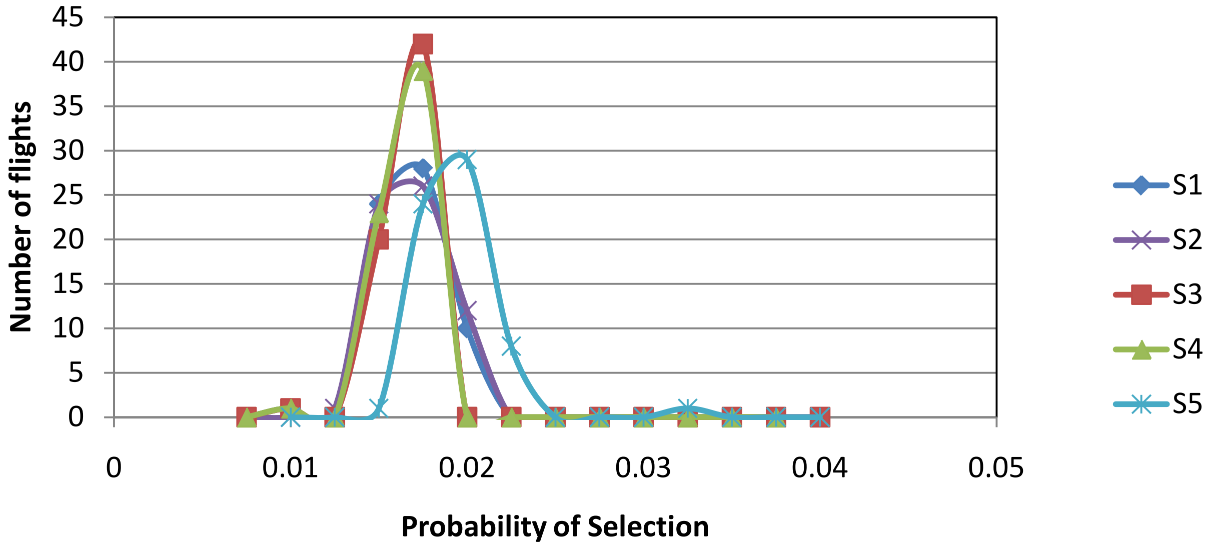

Figure 8. Probability Distributions of Score Functions for Time Spent in CRA

B. Combined Prioritization Criteria

Given the resulting probability distributions for the original score functions examined under each of the prioritization criteria, new score functions were defined that vary the weighting of the two prioritization criteria together when developing the sorted lists. The new score functions considered are:

$$\begin{array}{ll} S_{1*} = pc_1 * pc_2 & S_{2*} = \sqrt{pc_1} * pc_2 \\ S_{3*} = pc_1 * \sqrt{pc_2} & S_{4*} = \sqrt{pc_1} * \sqrt{pc_2} \\ S_5 = pc_1^2 * pc_2 & S_5 = pc_1 * pc_2^2 \end{array} \quad (8)$$

where $pc_1$is the time to CRA and pc_2 is the time spent in the CRA. These six score functions were chosen out of the many potential combinations of the two prioritization criteria to provide a variety of different probability distributions. Figure 9 provides the probability distributions for the new score functions. Examining Figure 9

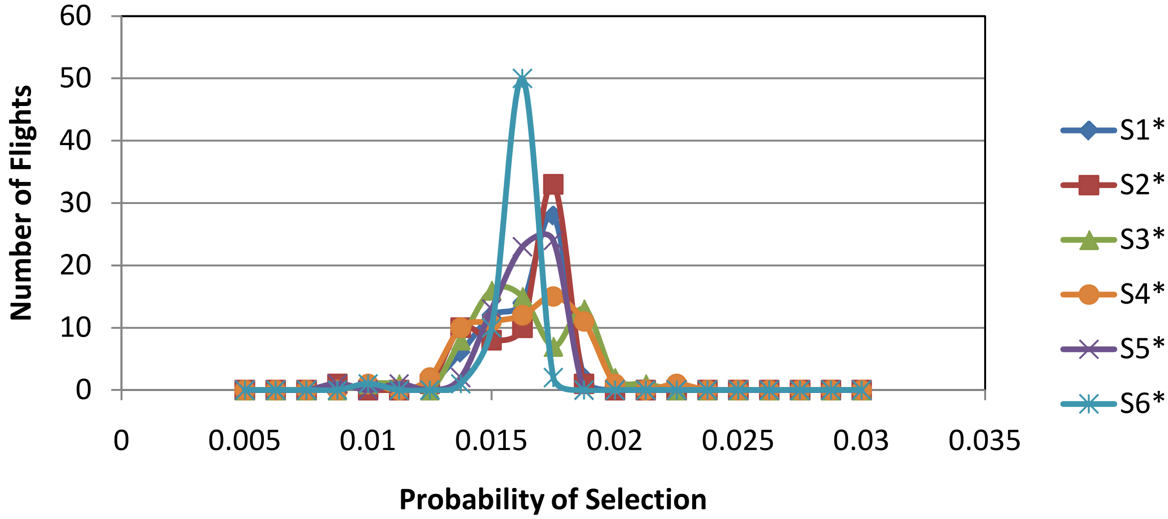

Figure 9. Probability Distribution for Combined Prioritization Criteria

shows how the score functions produce significantly different distributions, but are related to the distributions shown in Figure 7 and Figure 8.

The six new score functions were implemented in GRASP for the single-metric objectives of minimum delay, congestion, and inequity. The results are compared to the metric values produced by employing the original prioritization criteria using the original score functions. Again, the results presented are the averaged results for 25 runs of 100 iterations each. For clarity, future references to prioritization and score functions will be used as follows. "Time to CRA" refers to using the prioritization criterion of time to CRA with the original 5 score functions. "Time spent in CRA" refers to using the prioritization criterion of time spent in the CRA with the prioritization criteria using the original score functions. "Combined Prioritization" refers to using both the "time to CRA" and the "time spent in CRA" as the prioritization criteria in the six new score functions.

For the objective of minimum delay, Figure 10 shows the impact on each of the three metrics. Although only delay is minimized, all metrics are evaluated. Examining Figure 10 shows that the "combined prioritization" provides the lowest delay and the median valuation point for congestion and inequity of the three prioritization criteria. This improves upon the previous findings where "time spent in CRA" provided lower delay but much higher congestion and inequity than "time to CRA". Thus, for minimum delay, the combined prioritization may

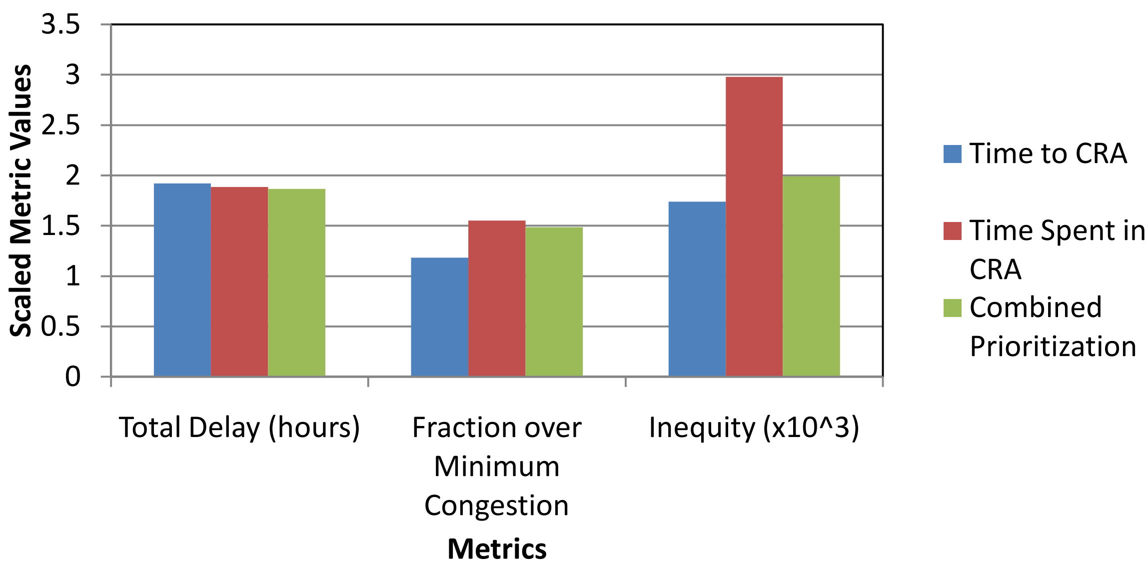

Figure 10. Prioritization Criteria Comparison for Minimum Delay

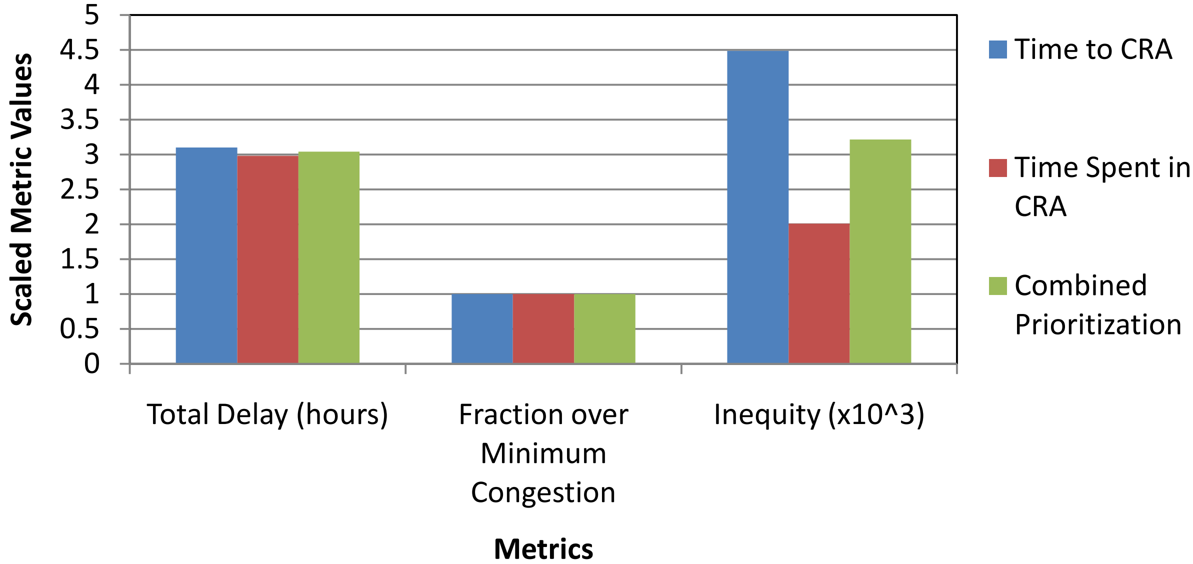

Figure 11. Prioritization Criteria Comparison for Minimum Congestion

provide an improved prioritization criterion.

For the objective of minimum congestion, Figure 11 shows the impact on each of the three metrics from implementing the three prioritization criteria. Examining Figure 11 shows that the "combined prioritization" provides the same minimum congestion as the other two prioritization criteria. Additionally, it provides the median valuation in the other two metrics. Unlike the minimum delay condition, here "time spent in CRA" produces the minimum congestion and the lowest delay and inequity, thereby decreasing the necessity of finding an alternative prioritization criterion for this objective function.

For the objective of minimum inequity, Figure 12 shows that all three prioritization criteria provide an extremely low inequity value, with "time to CRA" providing the lowest value. "Time to CRA" also provides the lowest congestion but the highest delay. The lowest delay comes from "time spent in CRA". The "combined prioritization" provides either the median or high value in each metric, and therefore may not be a desirable criterion for this objective.

The "combined prioritization" achieved a compromise between "time to CRA" and "time spent in CRA" for the single metric objective results. In two cases it produces the lowest metric value and in one case the highest metric

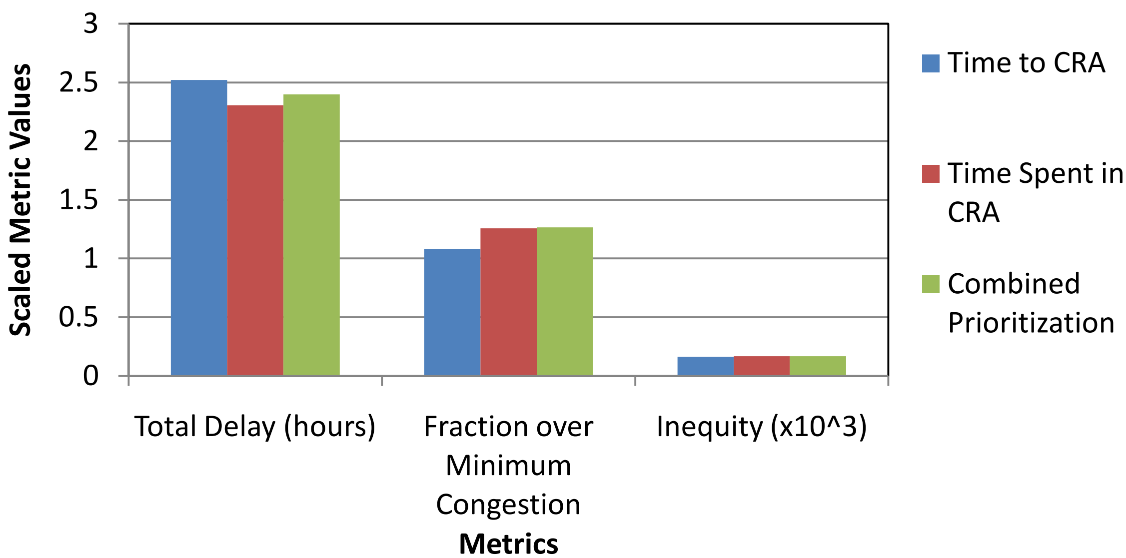

Figure 12. Prioritization Criteria Comparison for Minimum Inequity

value and in all other cases, the median metric value. Therefore it provides a more reliable prioritization criterion overall; however it does not necessarily improve the solution quality over the original prioritization criteria. Further investigations into alternative prioritization criteria would be necessary to identify a prioritization criterion and score function pair that consistently produces better solutions for all metrics.

V. Multi-Metric Optimization

The previous section shows the impact of the three prioritization criteria when each metric is considered in turn as a single metric objective. In reality, however, achieving balance between multiple objectives is of premium importance. Therefore, in this section we examine the impact of the different prioritization criteria on multi-metric objectives.

A. Congestion versus Delay

The metrics of delay and congestion are often competing objectives in congestion resolution. In this analysis, we examine how the prioritization criteria respond to changes in the relative weightings between these two metrics. The multi-term objective considered is

$$J = k_1 \frac{J_d}{N_d} + k_2 \frac{J_c}{N_c} \tag{9}$$

where k_1 and k_2 are the weighting factors for delay and congestion, respectively, and J_d and J_c are the delay and congestion metrics defined in Section III.C. To provide meaningful relationships between the weighting factors, the metric values are in units of hours of delay and fraction over minimum congestion.

Table 1 shows the various weighting parameters used in the above multi-term objective function to evaluate the impact of varying importance of delay and congestion using the different prioritization criteria.

Figure 13 shows the trend of increasing delay and decreasing congestion shown from left to right, corresponding to the different weightings from delay only to congestion only. For all three prioritization criteria, the “delay only” solutions provide a small variation in delay (as was shown in Figure 10) and a large variation in the resulting congestion. For the “congestion only” solutions, all prioritization criteria have the same congestion (as is shown in Figure 11) and there is only a small variation in delay.

Table 1. Relative Weighting Factors Cases for Multi-Metric Optimization

Delay Only	$k_1 = 1, k_2 = 0$
Heavily Emphasized Delay	$k_1 = 10k_2$
Moderately Emphasized Delay	$k_1 = 3k_2$
Equal Delay and Congestion	$k_1 = k_2$
Moderately Emphasized Congestion	$3k_1 = k_2$
Heavily Emphasized Congestion	$10k_1 = k_2$
Congestion Only	$k_1 = 0, k_2 = 1$

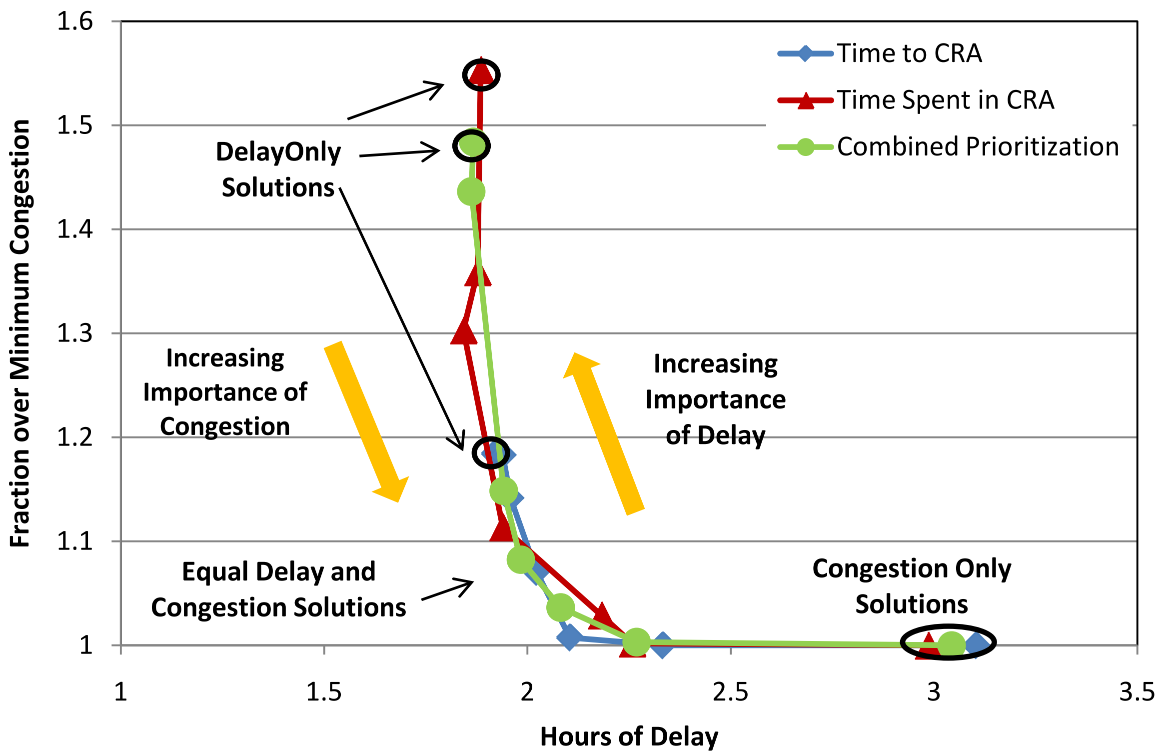

Figure 13. Congestion versus Delay

Figure 13 shows the trend of increasing delay and decreasing congestion shown from left to right, corresponding to the different weightings from delay only to congestion only. For all three prioritization criteria, the "delay only" solutions provide a small variation in delay (as was shown in Figure 10) and a large variation in the resulting congestion. For the "congestion only" solutions, all prioritization criteria have the same congestion (as is shown in Figure 11) and there is only a small variation in delay.

Moving from the "delay only" solutions to the "delay heavily emphasized" solutions shows small increases in delay for large decreases in congestion. Similarly, moving from the "congestion only" solutions to the "congestion heavily emphasized" solutions shows almost no increase in congestion for significant improvements in delay.

When delay and congestion are equally weighted, the three prioritization criteria provide similar congestion and

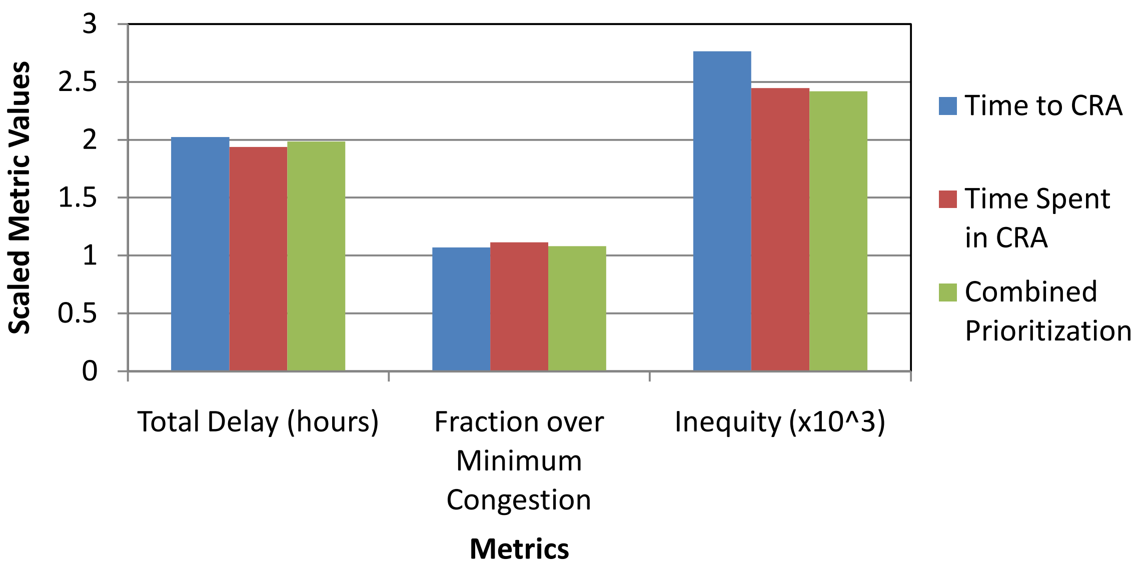

Figure 14. Prioritization Criteria Comparison for "Equal Delay and Congestion"

delay values, which is shown more clearly in Figure 14. Examining Figure 14 shows that for the “equal delay and congestion” case, “time spent in CRA” provides the lowest delay and highest congestion. The “time to CRA” provides the highest delay of the three prioritization criteria, but with the lowest congestion. The “combined prioritization” provides the balance between the two, as intended. Interestingly, the “combined prioritization” provides the lowest inequity of all three prioritization criteria, which depending on priorities might promote this prioritization over the alternatives.

B. Impact of Inequity

Understanding the trade-offs between delay and congestion is only part of the challenge in congestion management. Additionally, the delays assigned should not unduly impact any given NAS customer. The inequity metric defined in Section III.C provides one method for measuring the balance of delay distribution. This section investigates the impact of considering inequity directly within the multi-term objective by redefining the objective function as

$$J = k_1 \frac{J_d}{N_d} + k_2 \frac{J_c}{N_c} + k_3 \frac{J_e}{N_e}$$

where the new components of the objective function are k_3, the relative weighting of the inequity component of the objective; J_e, the inequity metric; and N_e, the scale factor of the inequity metric.

Figure 15 shows the minimization of the objective function using the same relationships (and naming conventions) for k_1 and k_2 as described in Table 1, and k_3 is set to be 0.1 or a tenth as important as the combination of delay and congestion.

Examining Figure 15 reveals that considering inequity, even with a small relative weighting to delay and congestion, changes the overall shape of the delay-congestion trade-off. As before, the same pattern emerges where heavily weighted delay provides a decrease in congestion with only a moderate increase in delay, as compared to the delay-only weighting (with the exception of the prioritization of time spent in the CRA, which actually decreases delay as well). Similarly, the heavily weighted congestion does not increase the congestion but does provide a decrease in delay over the congestion only solutions, albeit not as drastic as in the solutions where inequity is not

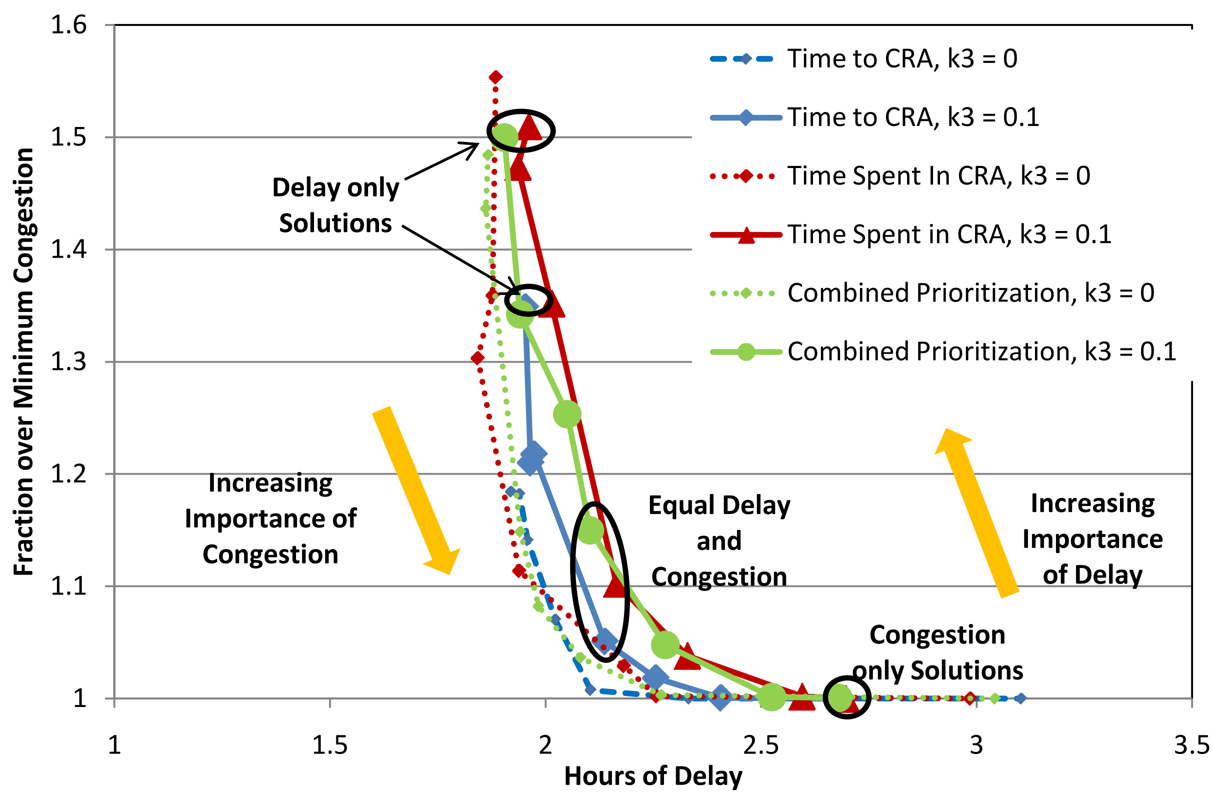

Figure 15. Congestion versus Delay with Inequity

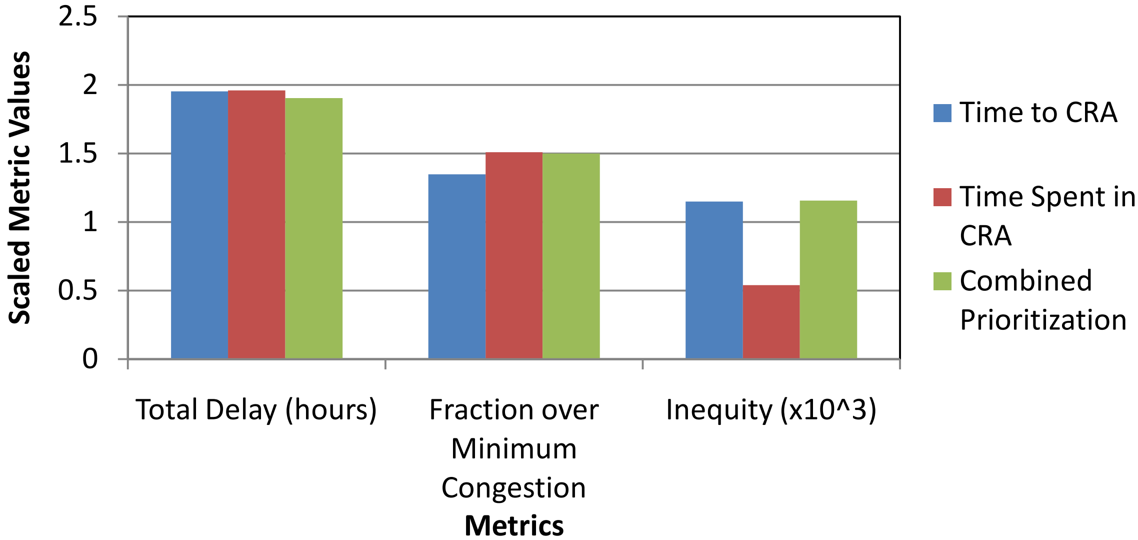

Figure 16. Prioritization Criteria Comparison for "Delay Only" with Inequity

considered. Finally, when considering the case where delay and congestion are equally weighted, we see a change in the relative trade-offs, where the combined prioritization provides the lowest delay, but the prioritization of time to CRA provides the lowest congestion.

To examine how the different objective function weighting factor combinations and prioritization criteria impact inequity in the solution, a more detailed view of the "delay only" with inequity, "congestion only" with inequity, and "equal delay and congestion" with inequity cases are presented. For the "delay only" with inequity case, Figure 16 reveals that the lowest delay is produced by the "combined prioritization" and the lowest congestion is produced by the "time to CRA". The "time spent in CRA" produces a slightly higher delay and congestion than the other two prioritization criteria, but significantly reduces the inequity. Decreases in inequity can be achieved through increases in delay as the additional delay can be evenly distributed to reduce inequity; however, the large decrease in inequity is disproportionate to the small increase in delay here, which suggests that the set of flights receiving delays was more equitable.

For the "congestion only" with inequity case, Figure 17 reveals that the "time to CRA" produces the minimum congestion, the lowest delay, and the lowest inequity. However, unlike the "delay only" with inequity case shown in Figure 16, the differences in metric values among the three prioritization criteria are very small for both delay and inequity.

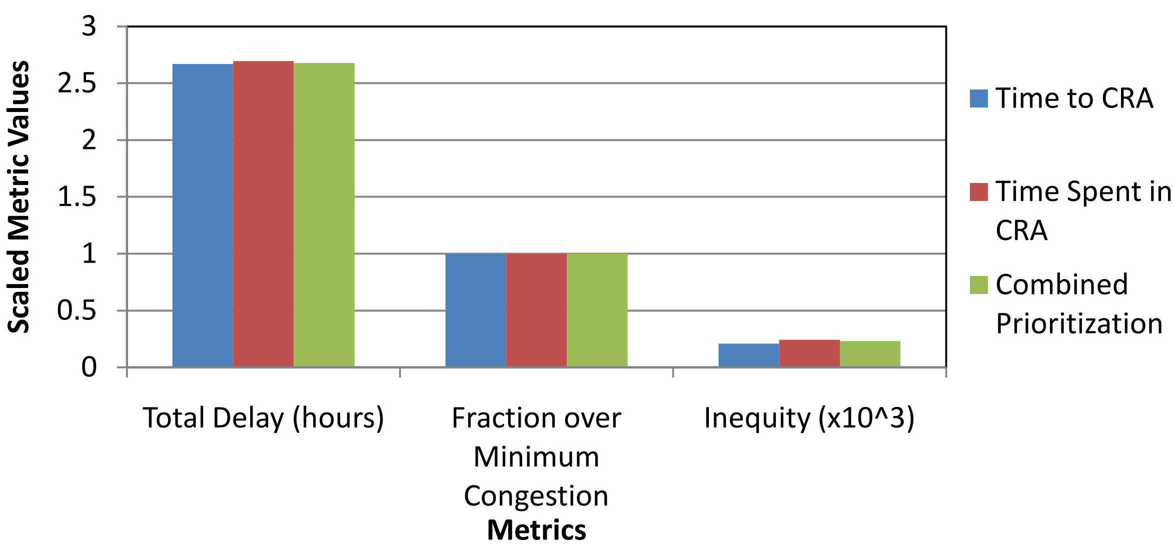

Figure 17. Prioritization Criteria Comparison for "Congestion Only" with Inequity

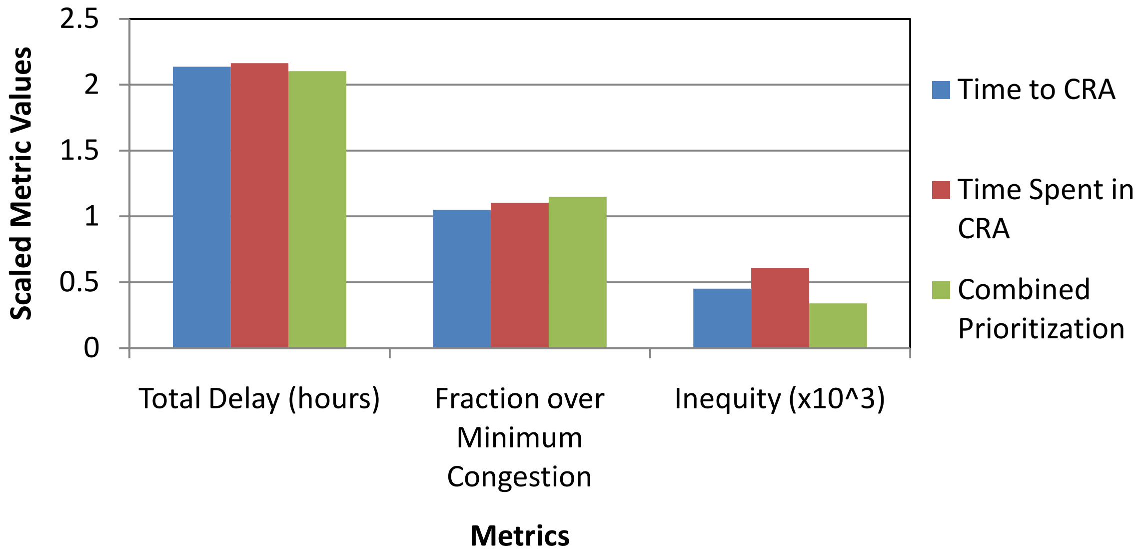

Figure 18. Prioritization Criteria Comparison for "Equal Delay and Congestion" with Inequity

When congestion and delay are equally weighted and inequity is included, Figure 18 shows that the "combined prioritization" provides the lowest delay and the lowest inequity, although the highest congestion. Here, instead of increasing the delay to decrease inequity, the congestion is increased, which describes a situation where a reduction in delay, and therefore a reduction in overall inequity, is achieved at the expense of more closely matching congestion targets.

In order to more clearly visualize the impact of inequity on metric values, the percent change in metric values from inequity not considered to inequity considered in the objective function are computed for each prioritization criteria and each weighting factor combination. Figure 19 shows the percent change in metric values for "time to CRA". A negative change shows improvement in the metric value when inequity is considered. Moving from left to right, Figure 19 also shows how these changes evolve as congestion is more heavily weighted. Examining Figure 19 reveals that considering inequity in the objective function always produces significant decreases in the inequity metric and that these reductions become larger as the emphasis of congestion is increased. The reduction of inequity

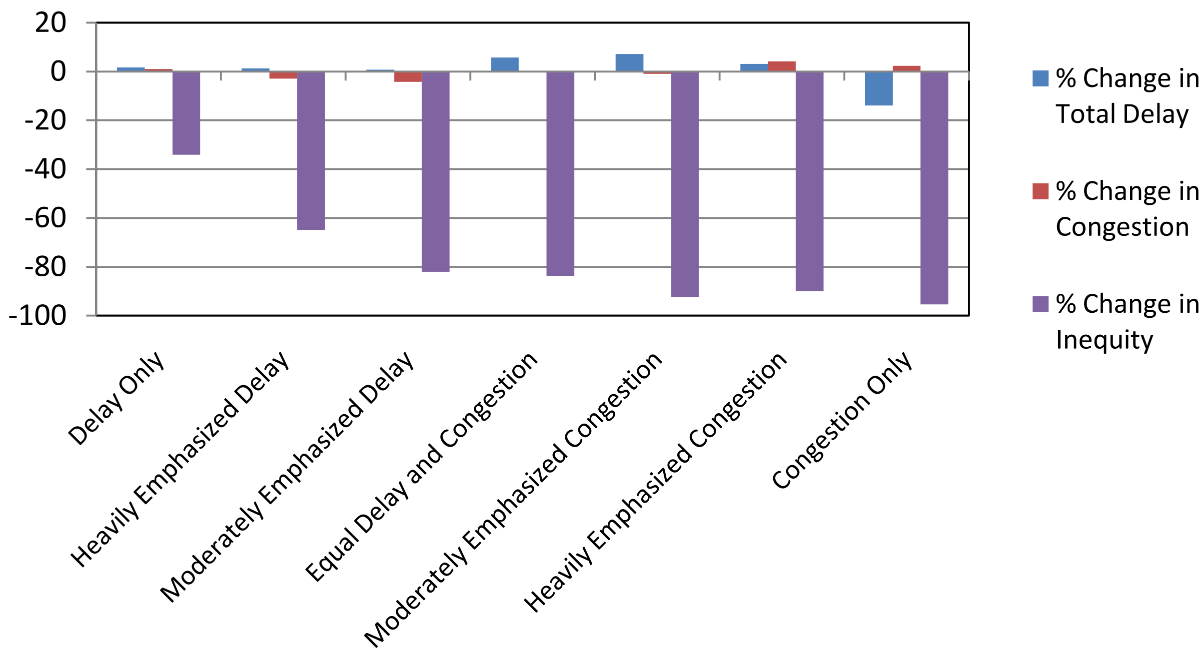

Figure 19. Percent Change from Including Inequity for "Time to CRA"

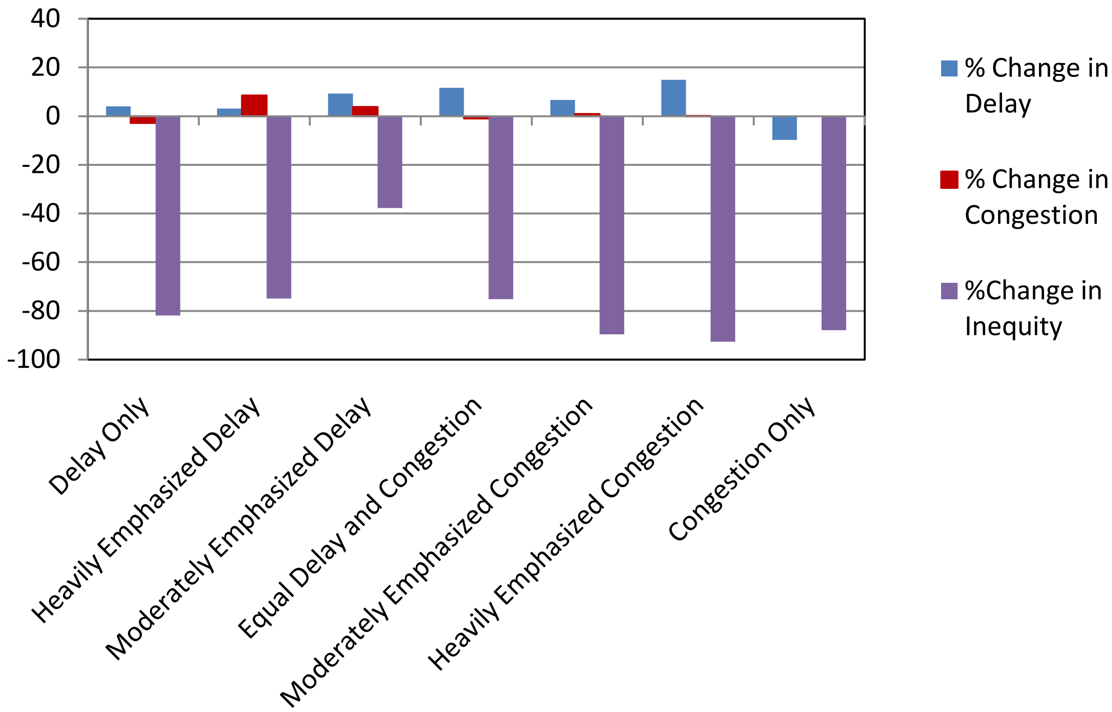

Figure 20. Percent Change from Including Inequity for "Time Spent in CRA"

is also paired with an increase in delay for every weighting factor except the “congestion only” case. The impact of inequity on congestion varies as congestion is more heavily weighted, in some cases causing an increase, in others a decrease. What is important to note is the magnitude of the changes in the metrics. Considering inequity in the objective decreases inequity between 34% and 95%, while providing maximum increases in delay and congestion of less than 8% and 5%, respectively.

Figure 20 shows the impact of inequity on metric values for the “time spent in CRA”. Again, we see a significant decrease in inequity; however the trend in decrease varies as congestion is emphasized. An increase in delay results from the consideration of inequity in every case, except for the “congestion only” case. And again, the impact on congestion varies. The decrease in inequity ranges between 37% and 93%, while the maximum increases in delay and congestion are less than 15% and 9%, respectively.

Figure 21 shows the impact of inequity on metric values for the "combined prioritization”. Again, significant decreases in inequity are shown, and the pattern of greater decreases in inequity as congestion is more heavily weighted emerges, like those shown in Figure 19 for “time to CRA”. Moving from left to right we see an increase in the change in delay, with the exception of the “congestion only” case. The impact on congestion from inequity varies. The decrease in inequity ranges between 41% and 93%, while the maximum increases in delay and congestion are less than 12% and 10%, respectively.

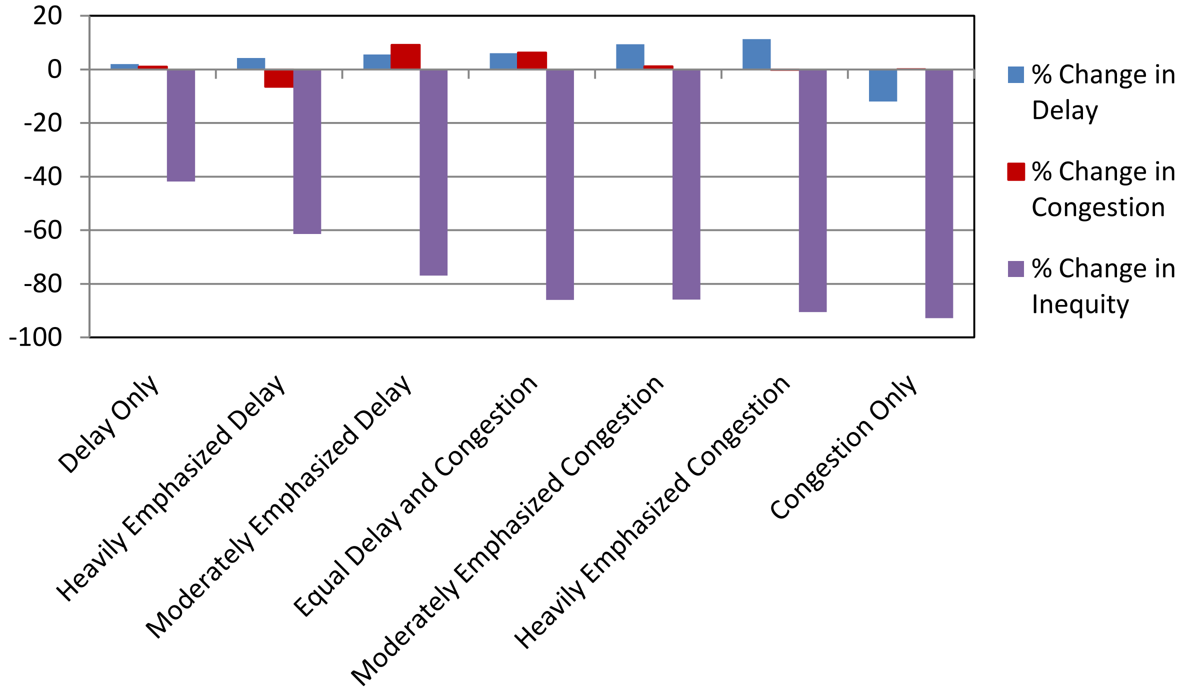

Figure 21. Percent Change from Including Inequity for "Combined Prioritization"

VI. Conclusions

In this research, the implementation of GRASP, a stochastic heuristic search procedure, was investigated for its merits in generating improved solutions for airspace congestion problems. Previous research[15] showed improved solution quality, as compared to results provided by a deterministic greedy heuristic, when optimizing single metric objectives within the congestion resolution framework. This research extended previous research by developing alternative prioritization criteria that provide heuristic for producing good solutions. Furthermore, the new and old criteria were used to generate solutions under a multi-metric objective, to analyze both the quality of the prioritization criteria and trade-offs in the relative priorities of different metrics as they impact airspace congestion resolution performance.

The choice of prioritization criteria, in combination with the selected score functions, defined the relative probability of a flight being selected. By defining the new set of score functions which used the combined prioritization criteria, new distributions were identified that provided a compromise between the individual criteria. The "combined prioritization" almost always provided the median metric value of the three prioritization criteria. Unlike "time to CRA" and "time spent in CRA" which would provide the best or worst metric values depending on the specific case, the "combined prioritization" could be consistently employed, regardless of objective function choice, to obtain reasonably good solutions. However, there were cases where the combined prioritization criteria did not perform as well. This suggests further research is needed to identify an alternate criterion or another factor to be included in the combined prioritization criteria that will improve solution quality in all cases identified.

This research considered the delay-congestion trade-off in depth, examining multiple relative weightings of delay and congestion using the prioritization criteria selected. The results show that considering heavily emphasized objectives provided almost the same performance of the prioritized metric and significantly improved results for the other metric. This point was further emphasized in the results where inequity was considered. Although, considering inequity in the objective sometimes yielded an increase in delay and congestion for the given weighting, a significant reduction in inequity was obtained. Strikingly, the consideration of inequity in the "congestion only" case produced a reduction even in delay, further illustrating both the inherent connection of these metrics and the need to evaluate solutions using a multi-metric objective. Further research analyzing alternative weightings of inequity may reveal that a lower weighting on inequity can provide almost identical congestion and delay results while still providing a reduction in inequity.

The analysis and conclusions presented in this research are based on a sample problem that is small in size by design to permit an in-depth analysis of the problem. However, further research into the impacts of GRASP on solution quality using larger, more complex problems is desirable. As the design space increases, the differences in solution quality from different prioritization criteria could become much larger, providing a more clear selection in the choice of heuristic and relative metric weightings.

Acknowledgments

The authors would like to thank Stephen Zobell, James DeArmon, Lixia Song, Sandeep Mulgund, Neera Sood, and Claude Jackson of the MITRE Corporation for their contributions to the research described here. Stephen Zobell developed the greedy heuristic referred to in this study.

Notice

The contents of this material reflect the views of the authors and The MITRE Corporation and do not necessarily reflect the views of the FAA or the DOT. Neither the Federal Aviation Administration nor the Department of Transportation makes any warranty or guarantee, or promise, expressed or implied, concerning the content or accuracy of these views

References

1. *Severe Weather and Route Management 2005,* Federal Aviation Administration - Air Traffic Control System Command Center, http://www.fly.faa.gov/Operations/Strategic_Planning/svrwx_handbook.html
2. Krozel, J., D. Rosman, and S. R. Grabbe. "Analysis of En Route Sector Demand Error Sources," AIAA 2002-5016, AIAA Guidance, Navigation, and Control Conference and Exhibit, Monterey, CA, USA, 5-8 August 2002.
3. Wanke, C., M. Callaham, D. Greenbaum, A. Masalonis, 2003. "Measuring Uncertainty in Airspace Demand Predictions for Traffic Flow Management Applications," 2003 AIAA Guidance, Navigation, and Control Conference and Exhibit, Austin, TX, 11-14 August 2003.
4. Meyn, L. A., 2002. "Probabilistic Methods for Air Traffic Demand Forecasting," AIAA 2000-4766, AIAA Guidance, Navigation, and Control Conference and Exhibit, Monterey, CA, USA, 5-8 August 2002.
5. Mueller, K. T., J. A. Sorenson, and G. J. Couluris. "Strategic Aircraft Trajectory Prediction Uncertainty and Statistical Sector Traffic Load Modeling," AIAA 2000-4765, AIAA Guidance, Navigation, and Control Conference and Exhibit, Monterey, CA, USA, 5-8 August 2002.
6. Wanke, C., S. Mulgund, D. Greenbaum, and L. Song, 2004: "Modeling Traffic Prediction Uncertainty for Traffic Management Decision Support," AIAA Guidance, Navigation, and Control Conference and Exhibit, Providence, RI, 16-19 August 2004.
7. Kopardekar, P., and Magyarits, S., "Dynamic Density: Measuring and Predicting Sector Complexity," *21st Digital Avionics Systems Conference*, Irvine, CA, 2002.
8. Masalonis, A., Callaham, M., and Wanke, C. "Dynamic Density and Complexity Metrics for Real-Time Traffic Flow Management," *5th USA/Europe Air Traffic Management Research and Development Seminar,* Budapest, Hungary, 23-28 June 2003.
9. Wanke, C., Rhodes, L., Connolly, K., DeArmon, J., and Topiwala, T., "Progressive Planning: A Decision-Support Framework for Collaborative TFM," *4th AIAA Aircraft Technology, Integration, and Operations Forum,* Chicago, IL, September 2004.
10. Sridhar, B., G. Chatterji, S. Grabbe, and K. Sheth, "Integration of Traffic Flow Management Decisions," *AIAA Guidance, Navigation, and Control Conference*, Monterey, CA, August 2002.
11. Wanke, C., N. Taber, S. Miller, C. Ball, and L. Fellman, "Human-in-the-Loop Evaluation of a Multi-Strategy Traffic Management Decision Support Capability," 5th *USA/Europe Air Traffic Management R&D Seminar*, Budapest, Hungary, June 2003.
12. Wanke, C, Zobell, S., and Song, L. "Probabilistic Airspace Congestion Management," *5th Aviation Technology, Integration, and Operations Conference*, Arlington, VA, 26-28 September 2005.
13. Sood, N., Mulgund, S., Wanke, C., and Greenbaum, D., "A Multi-Objective Genetic Algorithm for Solving Airspace Congestion Problems," *AIAA Guidance, Navigation and Control Conference and Exhibit*, AIAA Hilton Head, SC, August 2007.
14. S. Mulgund, Wanke, C., D. Greenbaum, and N. Sood, 2006: "A Genetic Algorithm Approach to Probabilistic Airspace Congestion Management," *AIAA Guidance, Navigation, and Control Conference and Exhibit*, Keystone, CO, August 2006.
15. Taylor, C. and Wanke, C., "A Generalized Random Adaptive Search Procedure for Solving Airspace Congestion Problems," *AIAA Guidance, Navigation, and Control Conference and Exhibit,* Honolulu, HI, August 2008.
16. Feo, T. and Resende, M., "Greedy Randomized Adaptive Search Procedure," *Journal of Global Optimization,* Vol. 6, No. 2, March 1995.
17. Taylor, C., Song, M., Klabjan, D., de Weck O., and Simchi-Levi, D., "A Mathematical Model for Interplanetary Logistics," *Logistics Spectrum,* Vol. 41, No. 2, January 2007
18. Song, L., D. Greenbaum, and C. Wanke, "Predicting Sector Capacity for TFM Decision Support," *6th AIAA Aviation, Integration, and Operations Conference,* Wichita, KS, 25-27 Sept. 2006.

19. Megenhardt, D. L., et al., "NCWF-2 Probabilistic Forecasts," *11th Conference on Aviation, Range, and Aerospace Meteorology*, Hyannis, MA, 4-8 October 2004.