AF2008097

Combining Airspace Sectors for the Efficient Use of Air Traffic Control Resources

Michael Bloem* and Parimal Kopardekar†

NASA Ames Research Center, Moffett Field, CA 94035-1000

A greedy heuristic algorithm for combining under-utilized airspace sectors to conserve air traffic control resources is suggested and analyzed. Simulations using historical air traffic data suggest that systematically combining under-utilized airspace sectors within centers can lead to significantly fewer sectors and therefore a more efficient utilization of air traffic control resources. The efficiency gains are more modest when sectors are only allowed to combine with sectors in the same controller area of specialization. Parameters can be adjusted to ensure that this approach to combining sectors works well with operational procedures as well as to trade off efficiency and the likelihood of sectors exceeding their capacity. Quantitative analyses show that the likelihood of sectors exceeding their capacities can be decreased effectively without significant sacrifices of efficiency gains by changing the minimum capacity gap parameter.

I. Introduction

The problem of congestion in the National Airspace System (NAS) is serious and well known. As discussed in the Next Generation Air Transportation Integrated Plan, expected growth in air traffic will lead to excessive levels of congestion in many areas of the NAS in the near future. Congested airspace and airports could cost consumers up to $20 billion per year by 2025.[1] Meanwhile, in many parts of the NAS and at many times during a typical day, air traffic control (ATC) resources are over-allocated. For example, in Cleveland center, 75% of sectors are operating near or below half capacity even at the busiest times of the day. This excess capacity is particularly troubling considering the current and projected shortage of air traffic controllers. Aside from its troubling potential safety implications, a controller shortage could also lead to increased delays.

Dynamic Airspace Configuration (DAC) involves changing how the national airspace is divided into centers, sectors, or other airspace components dynamically depending on the state of the NAS. Recent research has produced numerous DAC concepts and tools that achieve goals such as balancing the workload between sectors and building sectors that conform to desired air traffic flows. This research uses theoretical tools such as integer programming,[2] computational geometry,[3] genetic algorithms,[4] and many others.

While this DAC research has led to some promising results, this body of research has two shortfalls. First, much of this research does not explicitly minimize the number of sectors (and therefore the required ATC resources). Secondly, many of these approaches would require that new sectorizations be implemented at least seasonally and up to multiple times each day. The air traffic control system is not able to implement new sectorizations with this frequency, nor will it be able to in the near term (within three years). Currently, implementing changes in sectors takes 6-18 months, even though most changes in sectors are only incremental. Training a controller on a new set of sectors requires six months to two years.[5] Increasing the frequency with which new sectorizations can be implemented will require significant improvements in automation. Therefore, many approaches suggested in DAC research cannot be implemented in the near-term.

In this paper, an algorithm for systematically combining current airspace sectors is presented and analyzed. This algorithm directly fills the two gaps in DAC research mentioned above. Firstly, this algorithm can be implemented in the near-term. Airspace sectors are already being combined during off-peak hours.

*Aerospace Engineer, Automation Concepts Research Branch, MS 210-10, NASA Ames Research Center, Moffett Field, CA 94035-1000.

†Principal Investigator, NextGen Airspace Project, MS 210-10, NASA Ames Research Center, Moffett Field, CA 94035-1000 and Senior Member, AIAA.

Moreover, the algorithm presented here uses measures of predicted sector capacity utilization that are readily available in air traffic control centers. Secondly, this research is explicitly and exclusively concerned with minimizing the number of sectors and therefore the quantity of ATC resources required to handle a given quantity of air traffic in a center.

The relevant sector utilization and capacity metrics and algorithm parameters are described, the details of the algorithm are explained, and the algorithm is classified in Section II. In Section III, a quantitative analysis of the performance of the algorithm is described, followed by the results of this analysis and a discussion of the results. Finally, Section IV contains some concluding remarks.

II. An Algorithm for Combining Airspace Sectors

II.A. Sector Utilization and Capacity Metrics

The algorithm presented in this paper for combining sectors is based upon a measure of predicted excess capacity in sectors. Such a measure has two components. The first is a predicted measure of the utilization of a sector. The second is a measure of the maximum possible safe utilization of a sector, which is also referred to as the capacity of a sector. These components share units and so the predicted excess capacity is the difference between the predicted utilization and the capacity.

The most simple and widely used sector utilization metric is the maximum instantaneous aircraft count over a 15 minute time interval. Also, each sector has a “Monitor Alert Parameter” (MAP) value that designates the sector’s capacity in units of instantaneous aircraft count. For this study the maximum instantaneous aircraft count will be used to measure the utilization of a sector and the MAP value will be used to indicate its capacity.

Using maximum instantaneous aircraft count as the measure of utilization has several advantages. It is simple, widely used, readily predicted in real-time for planning purposes, and the uncertainty in its prediction is relatively well understood.[6,7] However, in many cases this measure does not correlate well with the actual complexity of a sector as indicated by air traffic controller feedback. Therefore, more sophisticated and accurate measures of complexity have been developed.[8] These measures are not readily available nor predicted in real-time, so they were not considered. Furthermore, as more automation is added to the air traffic control system, measuring the utilization of a sector with regards to a human controller may become irrelevant. As the approach presented here is envisaged for implementation in the near term, however, only complexity measures based on human controllers and currently used by air traffic control centers are considered.

Finally, a mechanism for determining the capacity of a new sector that is the combination of two smaller sectors must be established. One possibility is to apply the method used by the FAA to determine MAP values.[9] This function sets MAP values as a function of average dwell time of aircraft in the sector. Another is to set airspace capacities based on the volume of the newly created airspace sectors and a model of controller workload.[10] A more conservative and straightforward approach used here is to simply compare the two MAP values of the sectors being combined and to set the capacity of the new combined sector to the larger MAP value. Future work may consider a more complex or less conservative determination of the capacity of a new sector.

II.B. Algorithm Parameters

The algorithm uses four parameters, which are summarized in Table 1. Altering these parameters will change the sectorizations resulting from the algorithm. Also, altering these parameters allows a user to configure the algorithm to fit with the operational procedures of the air traffic control service provider.

Table 1. Algorithm Parameters

Parameter	Name
t_c	combination time
t_d	combination duration
t_n	advance notice duration
g	minimum capacity gap

The first parameter is a vector that contains the times that a sector may combine during a day. This *combination time* parameter will be denoted as t_c. An air traffic control service provider may wish to only combine sectors during times of the day that are not busy or when a shift change of air traffic controllers is scheduled to occur. Alternately, allowing sectors to combine repeatedly during the day may lead to better capacity management and utilization.

The next parameter is the *combination duration* parameter (t_d). A user may determine that combining sectors for long periods of time or combining them until the next combination time (t_c) is desired in order to reduce confusion regarding sectorizations. Alternately, sector combinations of shorter duration may be desired if the resectorization was caused by a temporary change in capacity and utilization, as may be the case during poor weather. Again, allowing shorter and more frequent sector combinations would likely lead to a more efficient use of air traffic control resources.

With this algorithm, sector combinations may be scheduled any length of time before they are implemented. This allows the air traffic control service provider time to schedule employees or otherwise prepare for the combination. The *advance notice duration* parameter t_n is the time between sector combination planning and execution. Larger values for this parameter require sector combinations to be based on longer-term and thus less accurate predictions of sector utilization.

Finally, the parameter g is the *minimum capacity gap* parameter. When evaluating whether neighboring sectors can be safely combined, the predicted excess capacity of the hypothetical combined sector must be greater than the minimum capacity gap. When changing this parameter, the user will trade off between increased safety and increased efficiency.

The effect of changing these parameters on the efficiency and safety of the resulting sectorizations will be quantitatively analyzed in Section III. Further research may investigate the possibility of dynamically changing these parameters based on the air traffic context.

II.C. Algorithm for Combining Airspace Sectors

There are three main steps to the proposed algorithm:

1. Compute the predicted capacity gaps for all possible two-sector combinations in the center.
2. Combine the two sectors whose combination has the largest predicted capacity gap.
3. Repeat until the largest predicted capacity gap is smaller than the minimum capacity gap.

The algorithm is more precisely specified as shown in Algorithm 1.

Input: set of sectors ($s \in S$), set of neighbors for each sector ($N(s)$), capacity of each sector ($C(s)$), predicted utilization of each sector at each time ($U(s,t)$), t_c, t_d, t_n, g
Output: new set of sectors ($\hat{s} \in \hat{S}$), neighbors of each new sector ($\hat{N}(\hat{s})$), capacity of each new sector ($\hat{C}(\hat{s})$), predicted utilization of each new sector at each time ($\hat{U}(\hat{s},t)$)

$\hat{S} \leftarrow S$
foreach $\hat{s}_i \in \hat{S}$ **do**
  $\hat{N}(\hat{s}_i) \leftarrow N(\hat{s}_i)$
  $\hat{C}(\hat{s}_i) \leftarrow C(\hat{s}_i)$
end
while $t = (t_c - t_n)$ **do**
  foreach $\hat{s}_i \in \hat{S}$ **do**
    foreach $\hat{n}_j \in \hat{N}(\hat{s}_i)$ **do**
      for $u = t_c$ **to** $(t_c + t_d)$ **do**
        if $\max\{\hat{C}(\hat{s}_i), \hat{C}(\hat{n}_j)\} - (\hat{U}(\hat{s}_i,u) + \hat{U}(\hat{n}_j,u)) < nbrMinGap(\hat{s}_i,\hat{n}_j)$ **then**
          $nbrMinGap(\hat{s}_i,\hat{n}_j) \leftarrow \max\{\hat{C}(\hat{s}_i), \hat{C}(\hat{n}_j)\} - (\hat{U}(\hat{s}_i,u) + \hat{U}(\hat{n}_j,u))$
        end
      end
    end
    $sectMaxGap(\hat{s}_i) = \max_{\hat{n} \in \hat{N}(\hat{s}_i)}\{nbrMinGap(\hat{s}_i,\hat{n})\}$
    $sectMaxNbr(\hat{s}_i) = \text{argmax}_{\hat{n} \in \hat{N}(\hat{s}_i)}\{nbrMinGap(\hat{s}_i,\hat{n})\}$
  end
  if $\max_{\hat{s} \in \hat{S}}\{sectMaxGap(\hat{s})\} > g$ **then**
    $\hat{s}_{max} \leftarrow \text{argmax}_{\hat{s} \in \hat{S}}\{sectMaxGap(\hat{s})\}$
    $\hat{S} \leftarrow \hat{S} \setminus \{\hat{s}_{max}, sectMaxNbr(\hat{s}_{max})\}$
    $\hat{S} \leftarrow \{\hat{S}, \bar{s}\}$
    $\hat{C}(\bar{s}) \leftarrow \max\{\hat{C}(\hat{s}_{max}), \hat{C}(sectMaxNbr(\hat{s}_{max}))\}$
    $\hat{N}(\bar{s}) \leftarrow \{\hat{N}(\hat{s}_{max}) \cup \hat{N}(sectMaxNbr(\hat{s}_{max}))\} \setminus \{\hat{s}_{max}, sectMaxNbr(\hat{s}_{max})\}$
    for $u = t_c$ **to** $(t_c + t_d)$ **do**
      $\hat{U}(\bar{s}) \leftarrow \hat{U}(\hat{s}_{max},u) + \hat{U}(sectMaxNbr(\hat{s}_{max}),u)$
    end
  else
    exit
  end
end

Algorithm 1: Algorithm for combining under-utilized airspace sectors

Two versions of this algorithm were developed and analyzed. The difference between these versions of the algorithms relates to "areas of specialization" within centers. Currently, each sector in a center is part of an area of specialization; each center typically contains 4-8 areas of specialization. Controllers are trained on sectors within a particular area of specialization. All controllers are qualified to work each sector in an area of specialization, but controllers rarely are trained on sectors spanning multiple areas of specialization.[5] Therefore, a near-term implementation of this algorithm should consider that neighboring sectors could only be combined if they are in the same area of specialization. When sector combinations are restricted to sectors within the same area of specialization, the algorithm is referred to as the *restricted algorithm*. This restriction may be relaxed as automation enables controllers to control more sectors in a center. When sectors are allowed to combine with any neighbor, regardless of area of specialization, the algorithm is referred to at the *unrestricted algorithm*.

This is a recursive algorithm. The algorithm can also be classified as greedy because at each step it chooses the best combination: the combination that results in the most excess capacity in the new sector combination. This excess capacity can then be used in later steps of the recursion for further sector combinations.

The computational complexity of this algorithm is $O(n^2)$, where n is the number of sectors in the initial sectorization. Therefore the computation time will grow quickly as more sectors are considered. For the number of sectors in a typical center, however, this computation time is not prohibitive. The computational complexity is linear in the combination duration (t_d). The complexity also grows non-linearly with the number of sector combinations that are implemented, although this number of combinations is not known a priori.

There are several strengths of this algorithm. It produces a new sectorization that utilizes air traffic control resources at least as efficiently as the original sectorization. It only uses information currently available in ATC centers. There are several algorithm parameters that can be changed to increase the safety of the resulting sectorization. These parameters also enable users to tailor the algorithm to work with existing operational procedures. In fact, because it combines this configurability with the use of existing sectors and data, this algorithm allows for a systematically more efficient utilization of airspace that can be implemented in the near term. In practice the algorithm executes in less than a second when calculating hour-long combinations for a center with 20 sectors.

This algorithm has several weaknesses as well. While using existing sectors as the building blocks for a new sectorization allows for short-term implementability, it restricts the possible airspace configurations and therefore also restricts the efficiency of the resulting sectorizations. Moreover, while this approach will yield more efficient air traffic control resource utilization, it does so by eliminating unused capacity, not by increasing capacity where capacity is lacking. This algorithm also does not guarantee sector convexity in the resulting sectorization.

From a more technical perspective, a weakness of this algorithm is that it is a heuristic with no guarantee of optimality. The problem faced here can be mapped to a graph theory problem by considering sectors as vertices and each sector's set of neighbors as defining edges between vertices.[11] The optimal solutions for related problems in graph theory can be found in polynomial time. For example, this problem is similar to the maximum cardinality graph matching problem,[12] but with two significant differences. In this case the edges in the desired "matching" may include adjacent edges. In terms of combining sectors, this means that more than just two sectors can be combined. This difference actually makes this problem a hypergraph matching problem, which cannot be solved in polynomial time. The second difference is that in this problem many vertices cannot be matched due to capacity constraints, and furthermore the set of vertices that are forbidden from being matched due to capacity constraints changes depending on what vertices have already been matched. These complications induce a more challenging problem than the maximum cardinality graph matching problem.

III. Quantitative Analysis of the Algorithm

III.A. Sector Aircraft Counts and Predictions

Sector aircraft count predictions are a crucial input to this algorithm. Therefore, in order to analyze the behavior of this algorithm, an accurate and realistic description of sector aircraft count predictions is necessary.

The details of how sector aircraft counts are predicted is not of importance in this research, however. What does matter is that any quantitative analysis of this algorithm uses a realistic sample of predicted aircraft counts for each sector over time. Also, the accuracy of these predicted aircraft counts with respect to the observed aircraft counts must reflect the actual distribution of such predictions.

A realistic sample of sector aircraft count predictions over time is obtained by simulating air traffic in the Future ATM Concepts Evaluation Tool (FACET)[13] and simply counting the aircraft in each sector at each time. The flights scheduled for a day are simulated instead of playing back flight data to eliminate the effect of traffic flow management (TFM) actions on sector aircraft counts. TFM actions, such as miles-in-trail restrictions and ground delay programs, are used to prevent sectors from exceeding their capacities, among other things. Simulating actual traffic schedules allows for an analysis of aircraft sector counts that is unbiased by TFM actions.

The second of requirement for a realistic quantitative analysis of this algorithm is that the sector aircraft count prediction uncertainty be modeled accurately. Research by Wanke et al.[6,7] has analyzed and characterized this uncertainty. In this work, Wanke et al. determine what probability distributions best represent the inaccuracy of sector aircraft count predictions similar to those used in air traffic control centers. They use regression analysis to find equations that give the parameters of these distributions (such as the mean or variance) as functions of the magnitude of the sector aircraft count prediction, the number of departed and not departed flights in the aircraft count prediction, the type of airspace sector being analyzed, and the look-ahead time of the prediction. The distribution of the error in a sector count prediction is assumed in this research to be independent of the errors in neighboring sectors.

In order to understand how variations in the actual sector aircraft counts impact the performance of the algorithm, a Monte-Carlo analysis was performed. In this analysis, each uncombined sector count distribution is sampled 500 times for each 15-minute time step. Samples of the actual aircraft count for combined sectors are found by summing individual samples from the distributions for the sectors that make up the combined sector. These samples of the actual aircraft count for combined sectors are used to approximate the distribution of the realized sector aircraft count for combined sectors. In particular, they are used to approximate the likelihood of combined sectors exceeding their capacity.

Changing the airspace sectorization would impact sector aircraft counts and potentially impact sector aircraft count prediction uncertainty as well. The effect of changing the airspace configuration on air traffic flow is not considered in this analysis. As the new sectorization is made up of existing sectors, this effect is assumed to be negligible.

III.B. Simulation Details

The behavior of this algorithm was studied under a variety of circumstances. The algorithm was applied to two centers with different air traffic characteristics: Cleveland and Salt Lake City. Cleveland is an exceptionally congested center where TFM initiatives are often in place to control flows into the New York City airports or the Chicago airports, while Salt Lake City tends to be less congested.

The algorithm was also simulated with data from seven good or moderate weather weekdays and two bad weather weekdays. Days were classified as good, moderate, or bad weather according to the number of hours of weather-related delays in the NAS. The good or moderate weather days were 2 March 2007, 8 May 2007, 24 May 2007, 29 May 2007, 13 June 2007, 4 July 2007, and 24 July 2007 and the bad weather days were 27 July 2006 and 10 July 2007. The data are Aircraft Situational Display to Industry (ASDI) data. The location of the weather relative to Cleveland and Salt Lake City centers was not analyzed. Rather than simulating air traffic on bad weather days, FACET was used to play back actual flight tracks on those days. Playing back rather than simulating the data allowed for the impact of weather on sector counts to be preserved. No effort was made to estimate the reduction in sector capacity values resulting from weather.

Combination times occurred each hour, and each combination lasted for an hour. The advance notice duration was also set to an hour. Finally, the minimum capacity gap was set to three. Parametric studies, in which the impact of adjusting these parameters is investigated, were conducted as well. The results of these studies can be found in sub-section III.C.1.

The sectorization used for these simulations is from July of 2007. For simplicity, the algorithm is only applied to a two-dimensional airspace and sub-sectors are not considered. The algorithm is directly applicable to three-dimensional sectorizations, however. The sectors chosen for these simulations are those high sectors containing airspace at or above 29,000 feet. This analysis therefore focuses on higher altitude sectors containing mostly en-route crossing traffic (as opposed to ascending or descending traffic).

III.C. Discussion of Results

The results of simulations of good and moderate weather weekdays will be presented first. When the restricted algorithm is applied to these two centers, simulations show that modest gains in efficiency are possible. Table 2 summarizes the simulation results for this scenario.

The sector-hours metric is calculated by multiplying the number of sectors operating at each 15-minute time step by the length of the time step and then summing over all time steps (in each of the seven weekdays simulated). If one controller is working each sector, this metric is equivalent to controller-hours.

This metric is also proportional to the area under the respective curves in Figure 1. There are eight areas of specialization in both the Cleveland and Salt Lake City sectors, so eight is the minimum number of sectors possible in this scenario. This figure shows that most of the possible combinations are at night, but also that sector combinations are possible even during busy times of the day. More sector combinations are possible in Salt Lake City, so the reduction in sector-hours is more profound there than in Cleveland center.

The next metric is the average over all time steps of the median sector utilization at each time step from the seven weekdays. It measures what portion of the deployed air traffic control capacity is utilized. When sectors are combined as prescribed by the algorithm, this measure of capacity utilization increases moderately. Figure 2 shows the quartiles of the sector utilization data over all sectors in all of the seven weekdays. Most of the increase in utilization can be seen in the increase in the 100^{th} percentile of utilization

Table 2. **Efficiency and Safety Performance of Sectorizations Produced by the Restricted Algorithm**

Metric	Cleveland Uncombined	Cleveland Combined	Cleveland Change	Salt Lake City Uncombined	Salt Lake City Combined	Salt Lake City Change
Sector-Hours	1,771	1,510	−14.73%	2,898	2,023	−30.19%
Average Median Utilization (% MAP)	25.03%	28.31%	+13.09%	26.81%	34.65%	+29.22%
Expected Number of Over Capacity Sectors (Number of sectors)	0.03866	0.04941	+27.81%	0.2309	0.2651	+14.81%

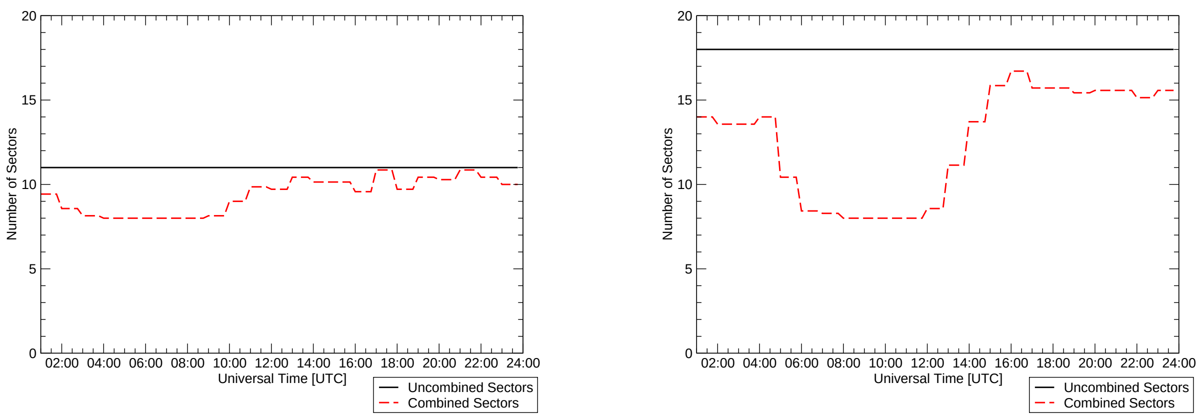

(a) (b)

Figure 1. Average number of uncombined sectors and combined sectors resulting from the restricted algorithm over the course of a day in (a) Cleveland center and (b) Salt Lake City center. These sector counts are averaged over seven good and moderate weather weekdays.

(the top whisker of each box-and-whisker pair in the figure) during the night time. The simulations also indicate that the efficiency of capacity utilization would increase even during busy daytime periods.

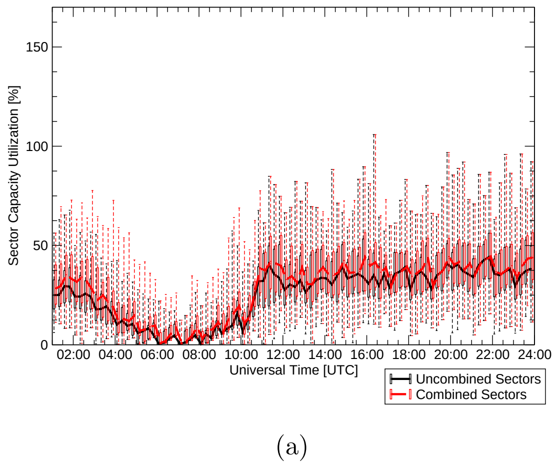

Figure 2. Distribution of sector utilization for uncombined sectors and combined sectors resulting from the restricted algorithm in (a) Cleveland center and (b) Salt Lake City center.

Finally, in order to capture the impact of combining sectors on the safety of the resulting airspace configuration, the expected number of airspace sectors exceeding the MAP value was computed. For both Cleveland and Salt Lake City, this number was usually well below 1, as can be seen in Figure 3. While combining airspace sectors does increase the likelihood of a sector utilization exceeding the capacity of the sector, both the increase and the resulting expected value are small.

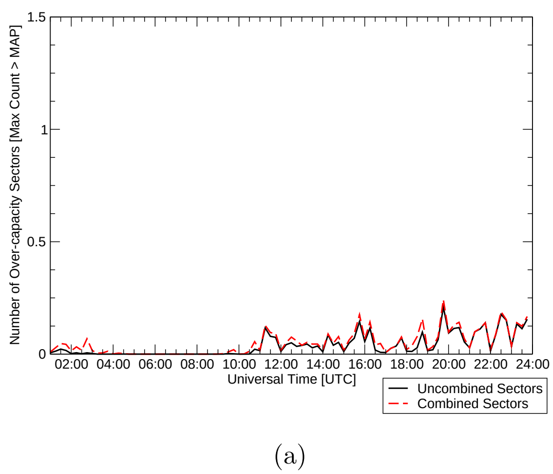

Figure 3. Expected number of over-capacity sectors in (a) Cleveland center and (b) Salt Lake City center when sectors are not combined and when sectors are combined according to the restricted algorithm. These expected values are estimated from a Monte-Carlo analysis of simulated air traffic data from seven good and moderate weather weekdays.

The unrestricted algorithm generates sectorizations that are significantly more efficient and yet still safe. Table 3 shows the impressive gains in efficiency and small losses in safety when the algorithm is simulated with data from seven good or moderate weather weekdays.

Allowing sector combinations between neighboring sectors not in the same area of specialization leads to much larger gains in efficiency than allowing sector combinations only within the same area of specialization. The sector-hours are reduced by more than 40% for both centers and the average median sector utilization increased by around 74% and 94% for Cleveland and Salt Lake City, respectively. Meanwhile, the expected number of sectors exceeding their MAP value at each time step does increase, although it should be noted

Table 3. **Efficiency and Safety Performance of Sectorizations Produced by the Unrestricted Algorithm**

Metric	Cleveland Uncombined	Cleveland Combined	Cleveland Change	Salt Lake City Uncombined	Salt Lake City Combined	Salt Lake City Change
Sector-Hours	1,771	1,031	−41.78%	2,898	1,517	−47.65%
Average Median Utilization (% MAP)	25.07%	43.51%	+73.52%	26.79%	52.06%	+94.35%
Expected Number of Over Capacity Sectors (Number of sectors)	0.03808	0.07999	+110.1%	0.2322	0.3162	+36.18%

that the MAP values of the combined sectors are conservatively set to a low value. Further analysis would be needed to determine if this increase is an acceptable sacrifice for these gains in efficiency. Figures 4, 5, and 6 show the number of sectors, the utilization of sectors, and the expected number of sectors exceeding capacity at each time step, respectively.

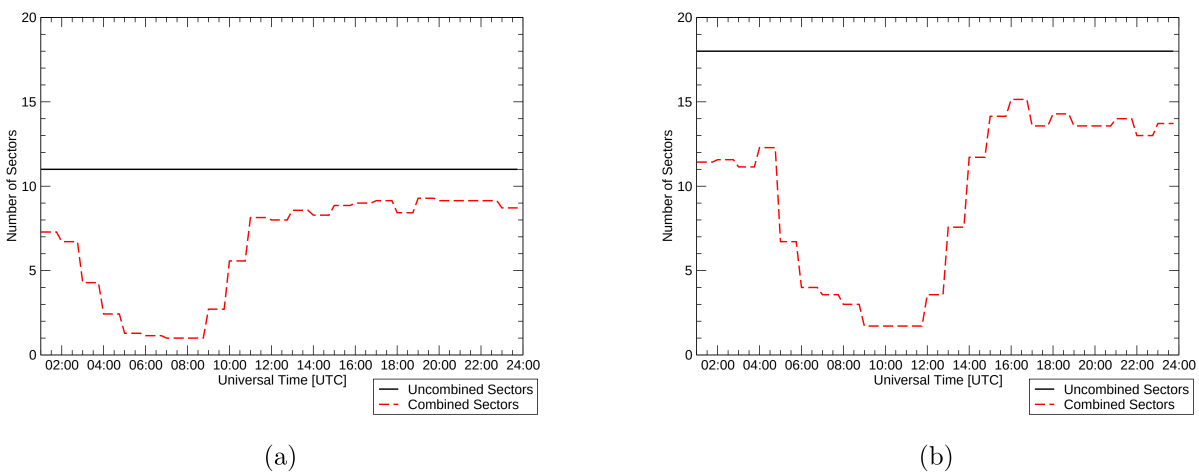

Figure 4. Average number of uncombined sectors and combined sectors resulting from the unrestricted algorithm over the course of a day in (a) Cleveland center and (b) Salt Lake City center. These sector counts are averaged over seven good and moderate weather weekdays.

The combinations produced by this algorithm also were compared to sector combinations that were actually executed in Cleveland center, as seen in Figure 7. Fewer combinations are implemented in practice than suggested by the restricted algorithm. The unrestricted algorithm would implement many more combinations that are implemented in practice. Overall, using the restricted algorithm would lead to a 9% decrease in sector-hours from the number of sector-hours used in practice. The unrestricted algoirthm would reduce the number of sector-hours required in practice by 38%.

The next results are from simulations where the algorithm is applied to actual air traffic data from two bad weather weekdays. For these simulations the air traffic is not simulated in FACET, but rather played back in order to preserve the impacts of weather on air traffic. A summary of the simulation results for the unrestricted algorithm is presented in Table 4.

The results for bad weather weekdays are not significantly different from those presented in Table 3 for good weather weekdays. The efficiency gains are slightly lower for the bad weather case and the increase in the number of expected over-capacity sectors increases slightly, but in general weather does not seem to impact the performance of the sectorizations proposed by the algorithm. More detailed analysis of bad weather days is required to better understand the behavior of this algorithm when applied during bad weather days. This analysis should take into account the location of weather relative to the centers being

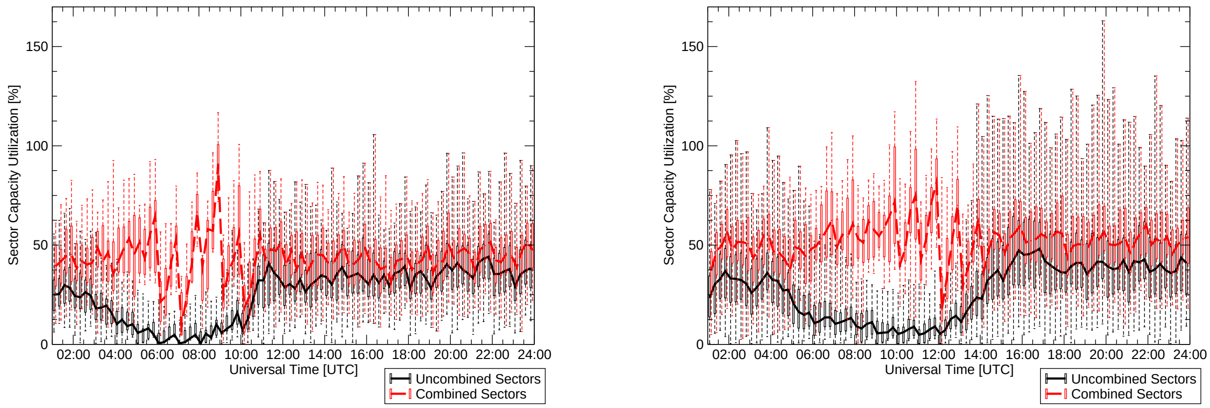

(a) (b)

Figure 5. Distribution of sector utilization for uncombined sectors and combined sectors resulting from the restricted algorithm in (a) Cleveland center and (b) Salt Lake City center.

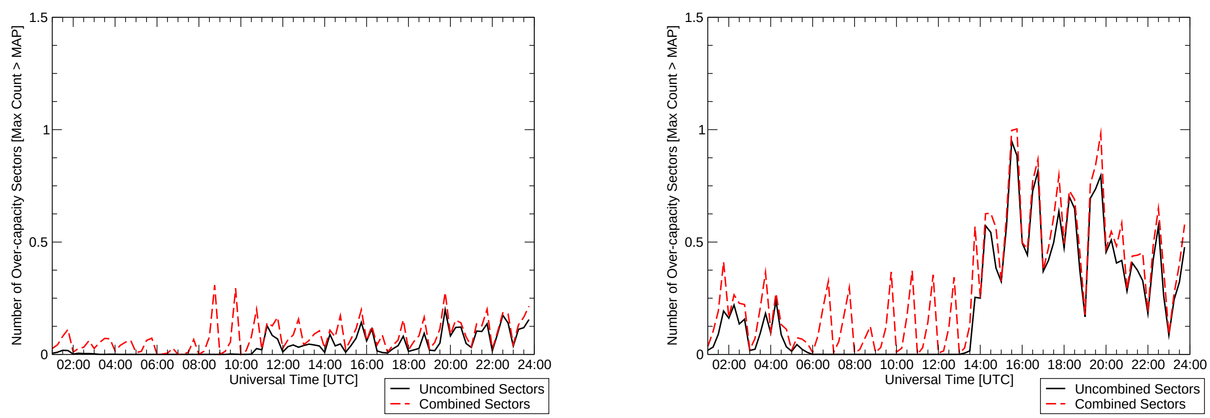

(a) (b)

Figure 6. Expected number of over-capacity sectors in (a) Cleveland center and (b) Salt Lake City center when sectors are not combined and when sectors are combined according to the restricted algorithm. These expected values are estimated from a Monte-Carlo analysis of simulated air traffic data from seven good and moderate weather weekdays.

Table 4. Efficiency and Safety Results of Simulations Combining any Neighboring Sectors in Bad Weather

	Cleveland			Salt Lake City		
Metric	**Uncombined**	**Combined**	**Change**	**Uncombined**	**Combined**	**Change**
Sector-Hours	506	310	−38.74%	828	451	−45.53%
Average Median Utilization (% MAP)	25.98%	44.83%	+72.58%	27.07%	50.57%	+86.80%
Expected Number of Over Capacity Sectors (Number of sectors)	0.03097	0.07640	+146.7%	0.1877	0.2516	+34.04%

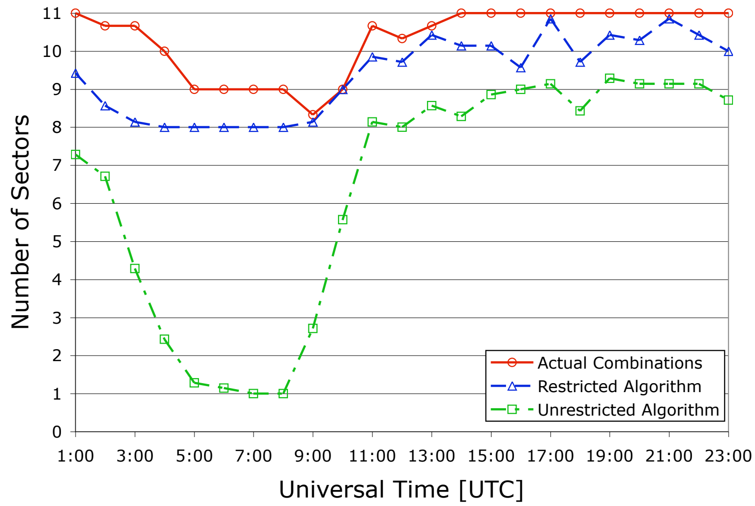

Figure 7. Average number of sectors over the course of a day in practice, when using the restricted algorithm, and when using the unrestricted algorithm.

investigated. Other metrics should also be considered because the reasons for combining sectors during bad weather may differ from the reasons for combining during clear weather.

III.C.1. Parametric Analysis

A parametric analysis was performed in order to determine the sensitivity of the performance of this algorithm in both centers to changes in the algorithm parameters. Each of the parameters in Table 1 were varied individually around the default parameter values described in Section III.B. The parameters t_c and t_d are so closely related that they were varied simultaneously between 4 and 40 (corresponding to sector combinations that last between 1 hour and 10 hours). The t_n parameter was assigned values between zero minutes and 4 hours. Finally, g was given values between 0 and 5.

The parameters were only varied around their default values, so there are many combinations of parameter values that were not investigated. Therefore this analysis does not give a complete picture of the impact of changing algorithm parameters; it only reveals one small part of this complete picture. For example, while the combination duration was allowed to vary widely, the effect of changing when the combinations are implemented (t_c) for each combination duration was not investigated. The time chosen to execute a combination could potentially make a large difference in the performance of the resulting sectorizations. Tuning the algorithm parameters will be an important aspect of any implementation of this algorithm.

First the impact of changing parameter values on the number of sector-hours will be investigated. Figure 8 shows the number of sector-hours as a percentage of the number of sector-hours when sectors are not combined as a function of variations of g, t_c, and t_d.

One result that Figure 8 makes obvious is that when sector combinations are constrained to sectors in the same area of specialization, the reduction in sector-hour usage as a percentage of the sector-hours used by uncombined sectors increases by between 20% and 25% almost regardless of the choice of algorithm parameters. Therefore, improving automation or changing air traffic controller training procedures to allow sector combinations across areas of specialization would have a dramatic positive impact on the utility of this algorithm.

When the capacity gap is set to zero aircraft, the sectorizations produced by the algorithm use only 45% of the sector-hours that would be consumed by entirely uncombined sectors over the course of this simulation. As the capacity gap increases to five, the utilization of sector-hours by the sectorizations produced by the algorithm increases linearly to around 60%. Allowing more aggressive combinations leads to larger efficiency gains. In reality, such "aggressive" combinations are probably relatively conservative because of how the capacity of combined sectors is set (see sub-section II.A).

Figure 8 (b) shows that the advance notice duration parameter value has no impact on the efficiency gains of this algorithm, as measured in sector hours. The predictions of sector hour counts utilized in this simulation are not dependent on the look-ahead time of the predictions. For example, the predicted number

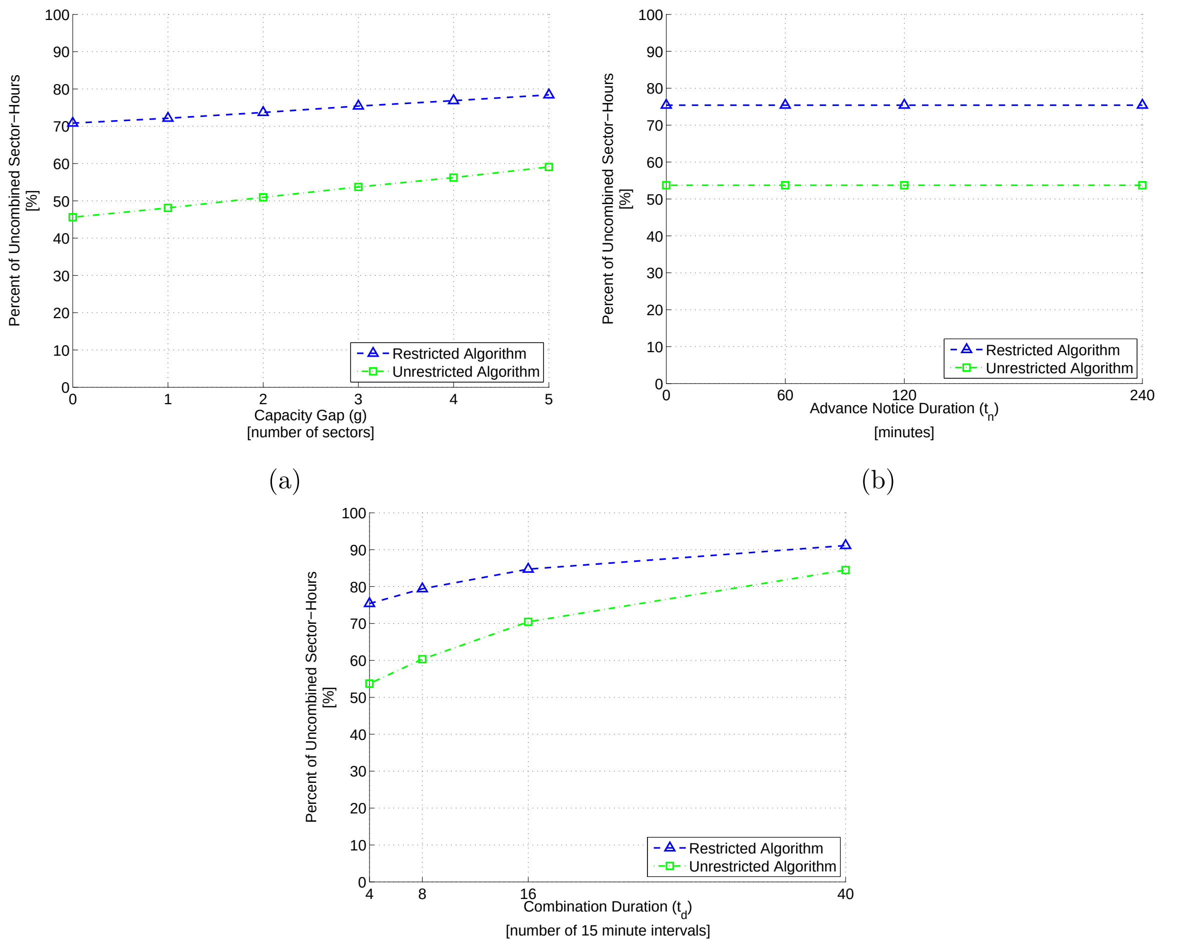

Figure 8. Effect of changes in (a) g, (b) t_n, and (c) t_d on the quantity of sector hours utilized by sectorizations produced by the algorithm expressed as a percentage of the sector-hours used when no sectors are combined.

of flights that will be in a given sector at a given time does not change depending on whether the prediction was made one hour in advance or ten hours in advance. In reality sector aircraft count predictions may be systematically higher or lower depending on look-ahead time, but in this research it is assumed that this is not the case, leading to the scheduling of identical sector combinations regardless of look-ahead time.

Finally, as the combination duration grows from 4 to 40 intervals of 15 minutes (from 1 hour to 10 hours), the efficiency of the sectorizations resulting from the algorithm decreases. When combinations occur at a frequency just higher than two times per day the efficiency gains are less than 20% even when any two sectors are allowed to combine. Note that this analysis does not attempt to find the optimal times for the combinations. Investigating changes in the combination times or allowing non-uniform combination durations could lead to further gains in efficiency in the sectorizations resulting from this algorithm.

A similar analysis was performed to determine the impact of changing parameter values on the expected number of over-capacity sectors. Figure 9 shows the worst-time expected number of over-capacity sectors. The expected number of over-capacity sectors at each time step in the two centers is determined by conducting a Monte-Carlo analysis as described in sub-section III.A. Here the highest expected number of over-capacity sectors from all of the 15 minute time intervals in the simulation is plotted because this worst-time value is what should be minimized to ensure a safe sectorization.

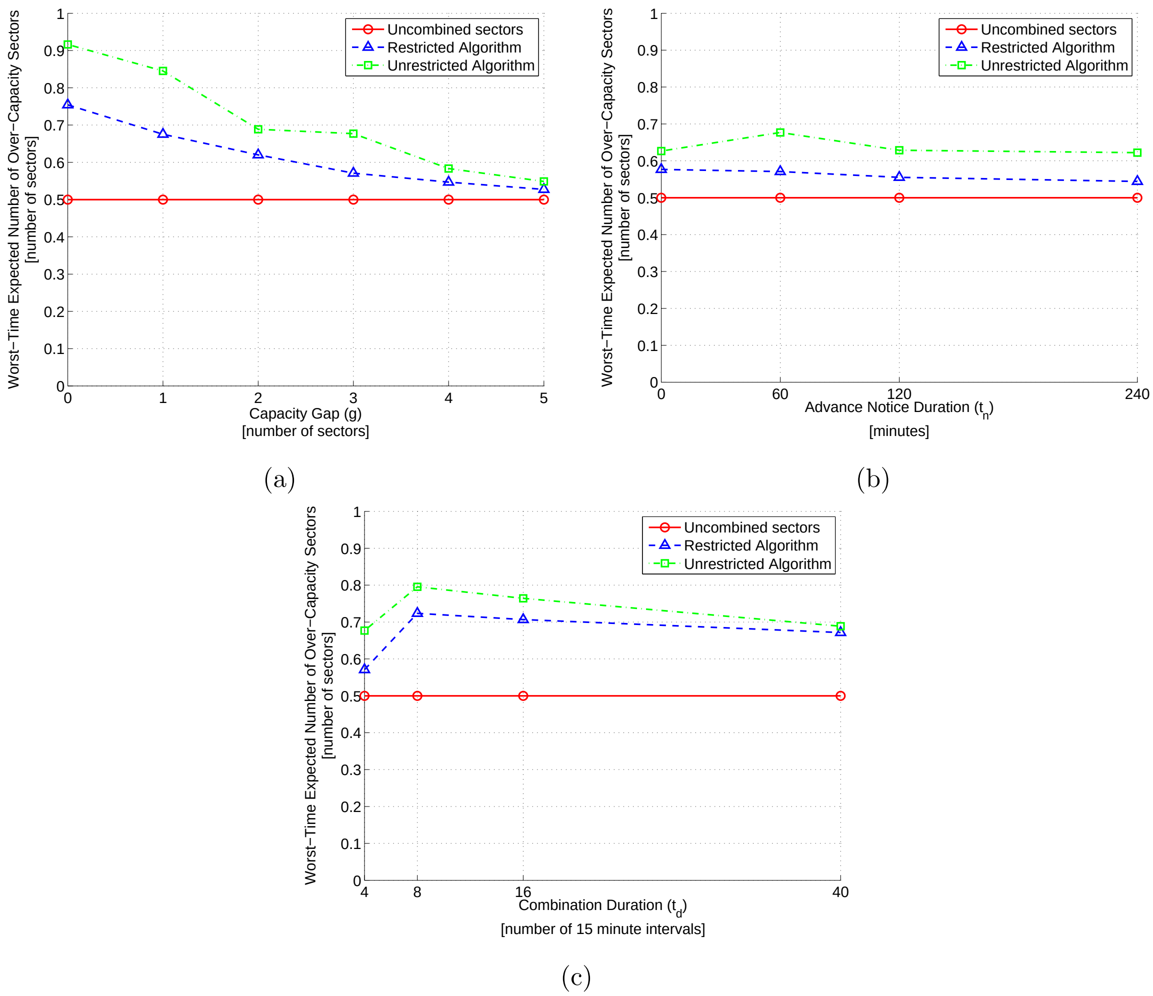

Figure 9. Effect of changes in (a) g, (b) t_n, and (c) t_d on the worst-time expected number of over-capacity sectors.

When no sector combinations occur, the worst-time expected number of over-capacity sectors in these two centers on these days is almost exactly 0.5. When sectors are aggressively combined with a zero capacity gap, this worst-time expected value increases to more than 0.9. Even when sectors are only allowed to combine with neighbors in the same area of specialization, the expected value increases to around 0.75. As the capacity gap increases to 5, the expected number of over-capacity sectors decreases to a value just slightly higher than the value that is achieved with uncombined sectors, as shown in Figure 9 (a).

As was the case for sector-hours usage, the advance notice duration has little impact on the worst-time expected number of over-capacity sectors. The reasons for this are the same as those given above for other metrics. Figure 9 (b) shows how this metric stays relatively constant in spite of changes in the advance notice duration.

Two competing factors influence the effect of combination duration on the worst-time expected number of over-capacity sectors. As the combination duration increases, fewer potential sector combinations are implemented because fewer meet the capacity gap requirement. This results in fewer combined sectors and therefore fewer over-capacity sectors. However, sector aircraft count predictions tend to be low, particularly as look-ahead time increases.[6] This leads to more sector combinations that appear to have sufficient excess capacity when they are planned and implemented but end up exceeding their capacity due to low sector-count predictions. The interaction of these two competing factors leads to worst-time expected over-capacity sector counts that first increase and then later decrease as combination durations increase.

This parametric analysis also allows for tradeoffs to be studied. The tradeoff between efficiency of ATC resource usage and likelihood of sector utilization exceeding capacity is particularly important when using this algorithm. Monitoring the same quantity of airspace with fewer and fewer sectors means that there is an increasing likelihood of sectors exceeding their capacity. This tradeoff is depicted graphically in Figure 10. Increasing the minimum capacity gap (g) decreases the worst-time expected number of over capacity sectors without sacrificing much of the savings in sector-hours. Therefore, if the likelihood of sectors exceeding their capacity becomes too large, the minimum capacity gap can be used to decrease this likelihood.

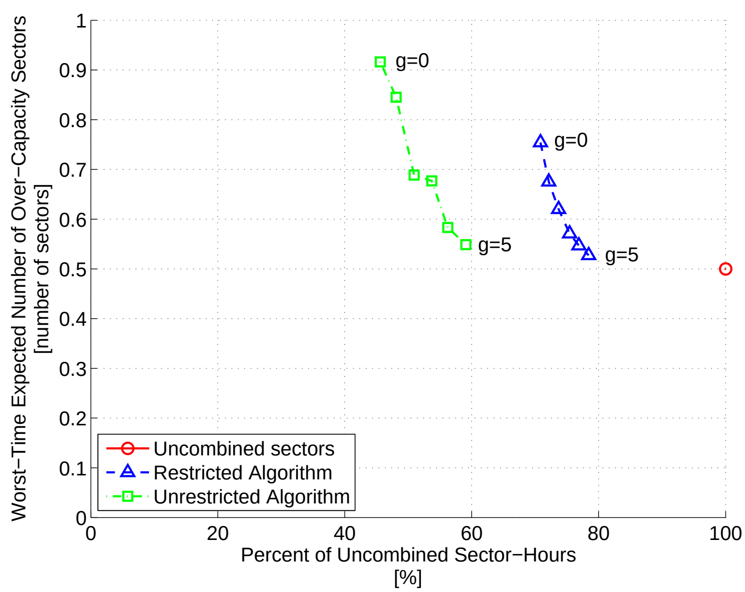

Figure 10. Curves showing the tradeoff between sector-hour usage and worst-time expected number of over-capacity sectors for the unrestricted algorithm and the restricted algorithm with several minimum capacity gaps. The values of these metrics for the uncombined sectors are also plotted as a point.

IV. Conclusions

This analysis has shown that a relatively simple greedy heuristic algorithm can systematically combine airspace sectors to significantly improve the efficiency of air traffic control resource utilization. When sectors are permitted to combine with any neighboring sector, the number of sector-hours is reduced by between 41% and 48% in simulations of different centers. This reduction of sector-hours is accompanied by an increase in the utilization of sectors and also an increase in the expected number of over-capacity sectors at any given time. When sectors are only permitted to combine with neighboring sectors that are in the same area of specialization, however, the efficiency gains are much more modest. This limits the short-term utility of an application of this algorithm. However, these results demonstrate that improvements in automation that would enable controllers to work any sector in a center would allow for large efficiency gains. Parametric studies show that altering the parameters in this algorithm, notably the capacity gap, allows a user to increase efficiency gains resulting from sector combinations by tolerating increased likelihood of sector utilizations exceeding sector capacities.

The analysis presented here should be followed by further research. The definition of the capacity of combined sectors used here is conservative and could be made more sophisticated and accurate, which would likely lead to more possible sector combinations. Further research from a graph theoretical perspective could help find an efficient optimal algorithm rather than a heuristic algorithm. Simulations should also be performed to characterize the performance of this algorithm when applied to three-dimensional sectorizations. This algorithm could serve as a simple alternative or baseline to which more advanced DAC algorithms can be compared. If this algorithm were used with a carefully designed initial sectorization and more advanced and appropriate complexity metrics, it may be a simple yet useful DAC tool in the near to mid-term.

Acknowledgments

The authors would like to thank Craig Wanke and Lixia Song of the MITRE Corporation for providing the details of their characterization of the uncertainty in sector aircraft count predictions. Also, thanks to Joe Mitchell of the State University of New York at Stony Brook for his assistance in understanding and classifying this problem and the algorithm. The authors are also grateful to Mark Evans of the Federal Aviation Administration for providing sector combination data and also for his detailed comments and suggestions regarding this work.

References

[1] "Next Generation Air Transportation System Integrated Plan," Tech. rep., Joint Planning and Development Office, 2004.

[2] Yousefi, A. and Donahue, G. L., "Temporal and Spatial Distribution of Airspace Complexity for Air Traffic Controller Workload-Based Sectorization," *Proc. of AIAA Aviation Technology, Integration and Operations Forum*, Chicago, IL, September 2004.

[3] Bau, A., Mitchell, J. S. B., and Sabhnani, G., "Geometric Algorithms for Optimal Airspace Design and Air Traffic Controller Workload Balancing," *Proc. of Workshop on Algorithm Engineering and Experiments*, San Francisco, CA, January 2008.

[4] Delahaye, D., Alliot, J.-M., Schoenauer, M., and Farges, J.-L., "Genetic algorithms for partitioning air space," *Proc. of 10th IEEE Conference on Artificial Intelligence Applications*, San Antonio, TX, March 1994, pp. 291–297.

[5] Evans, M., private e-mail communication, January 30, 2008.

[6] Wanke, C. R., Callaham, M. B., Greenbaum, D., and Massalonis, A., "Measuring Uncertainty in Airspace Demand Predictions for Traffic Flow Management Applications," *Proc. of AIAA Guidance, Navigation, and Control Conference*, Austin, TX, August 2003.

[7] Wanke, C., Song, L., Zobell, S., Greenbaum, D., and Mulgund, S., "Probabilistic Congestion Management," *Proc. of 6th USA/Europe Air Traffic Management Research and Development Seminar*, Baltimore, MD, June 2005.

[8] Kopardekar, P., Schwartz, A., Magyarits, S., and Rhodes, J., "Airspace Complexity Measurement: An Air Traffic Control Simulation Analysis," *Proc. of 7th USA/Europe ATM Seminar*, Barcelona, Spain, July 2007.

[9] Federal Aviation Administration, *Order 7210.3U Facility Operation and Administration*, February 2006.

[10] Welch, J. D., Martin, B. D., and Sridhar, B., "Macroscopic Workload Model for Estimating En-route Sector Capacity," *Proc. of 7th USA/Europe ATM Research and Development Seminar*, Barcelona, Spain, July 2007.

[11] Diestel, R., *Graph Theory*, Springer-Verlag, New York, 3rd ed., 2005.

[12] Galil, Z., "Efficient Algorithms for Finding Maximum Matching in Graphs," *Computing Surveys*, Vol. 18, No. 1, March 1986, pp. 23–38.

[13] Bilimoria, K., Sridhar, B., Chatterji, G., Sheth, K., and Grabbe, S., "FACET: Future ATM Concepts Evaluation Tool," *Proc. of 3rd USA/Europe Air Traffic Management R & D Seminar*, Napoli, Italy, June 2000.