U.S. Department of Transportation
Federal Aviation Administration

Flight Management System Weather Input Optimizer Final Report

Continuous Lower Energy, Emissions and Noise (CLEEN) Program

Submitted by General Electric

DOT/FAA/AEE/2014-05

The Continuous Lower Energy, Emissions and Noise (CLEEN) Program is a Federal Aviation Administration NextGen effort to accelerate development of environmentally promising aircraft technologies and sustainable alternative fuels. The CLEEN Program is managed by the FAA's Office of Environment and Energy.

The report presented herein is the final report deliverable submitted by General Electric for a project conducted under the CLEEN Program to develop the Flight Management System Weather Input Optimizer (FWIO), a tool that provides optimal weather data for a given flight plan. This project was conducted under FAA other transaction agreement (OTA) DTFAWA-10-C-00046. This is report is report number DOT/FAA/AEE/2014-05 by the FAA's Office of Environment and Energy.

CLEEN
Benefits Analysis
For the FMS Weather Input Optimizer (FWIO)

Prepared by:

GE Aviation Systems LLC
3290 Patterson Avenue, SE, Grand Rapids, MI 49512-1991, USA

Table of Contents

Table of Figures

Table of Tables

1 Introduction

Flight prediction algorithms are sensitive to the accuracy of their input weather information amongst other errors. Errors and inaccuracies in the weather model adversely impact predictions by introducing errors in predicted groundspeeds, which in turn raise the time and fuel costs for a flight. Successful implementation of NextGen and Single European Sky ATM Research (SESAR) Air Traffic Control initiatives require increased flight plan prediction accuracy to enable Trajectory Based Optimization (TBO) airspace solutions.

The Flight Management System (FMS) Weather Input Optimizer (FWIO) selects weather locations and values to use in flight predictions that minimize performance-based cost functions for each phase of flight. The purpose of this document is to present results of this optimization process versus simply minimizing weather residual or no optimization process for perfect weather forecasts, as well as simulated forecast errors.

2 Summary

The benefits analysis was performed using thirty-two (32) flight plans over 68 weather forecasts from 2011. The first analysis isolated FMS weather modeling errors from other error sources; in particular, weather forecasting errors. Utilization of the FMS Wind Input Optimization (FWIO) tool to determine wind and temperature inputs to the FMS produced average savings of 24 pounds of fuel per flight in descent for guidance maneuvers (73% reduction in guidance maneuver fuel usage), and increased predictions accuracy in other phases by over 70%.

In addition to the “perfect forecast“ analysis, the wind and temperature forecast error models have been statistically characterized by comparing available weather data recorded using the Data Acquisition Replay Tool (DART) to National Oceanic and Atmospheric Administration (NOAA) and AirDat forecasts. In determining a characteristic error, only a small set of flight recordings (22 flights corresponding to collected AirDat data and 40 flights corresponding to collected NOAA data) were obtained. In this limited set of data, the AirDat wind forecast has a 7.3 knot bias error with a standard deviation of 11.3 compared to 15 knot bias error with standard deviation of 13.3 for the NOAA forecast. In this sample set, AirDat’s wind forecast is approximately 35% better than NOAA (bias + one sigma noise). The AirDat temperature forecast has a 0.95 °C bias error with a standard deviation of 1.01 °C compared to a 0.47 °C bias error with a standard deviation of 0.27 °C for the NOAA forecast. In this sample set, NOAA’s temperature forecast is approximately 48% better than AirDat (bias + one sigma noise). Both wind forecast sources had considerable wind error (possibly skewed by small sample size), but had relatively low temperature error. For the purposes of this study, simulated errors were produced from the statistical characteristics of the forecast error and were applied to each of the 68 weather forecasts used.

When the simulated forecast error was included, benefits of using the FWIO tool were reduced from values for perfect forecast scenarios. However, the combined effect of reducing forecast error (using AirDat forecast) with reducing modeling error (using the FWIO tool) is significant for all phases of flight in each cost metric. See Table 1 for a summary of the benefits.

Table 1. Summary of FWIO Benefits on Average per Flight[1]

	FWIO vs. non-optimized (both using same forecast)			FWIO with AirDat vs. non-optimized w/ NOAA
	Perfect Forecast	NOAA Forecast	AirDat Forecast	
Descent Cost Savings (lbs - sec)	20 (71%)	13 (21%)	13 (28%)	21 (35%)
Descent Fuel Savings (lbs)	24 (73%)	16 (23%)	17 (28%)	27 (39%)
Cruise Temporal Prediction Accuracy Increase (s / NM)	0.014 (71%)	0.004 (1%)	0.005 (1%)	0.079 (28%)
Climb Distance Prediction Accuracy Increase (ft / s)	7.6 (80%)	0 (0%)	1.2 (6%)	8.4 (33%)

Due to the small weather collection sample size, the measured weather errors may not accurately represent the true error across many geographical regions and dates, and a more general error may be significantly lower. As an attempt to remove this potential inaccuracy, the study was repeated on a case with 1/3 of the error; FWIO benefits results with the modified forecast error are presented in Appendix A.

Overall, when combining the effect of using AirDat's forecast (instead of NOAA) with the FWIO optimization (instead of a non-optimized weather representation), descent guidance fuel usage is reduced from 70 lbs. down to 43 lbs.; a net savings of 27 lbs. per flight. It is important to note that these benefits are presented *on average per flight,* but on an individual flight basis using the FWIO tool with AirDat weather forecast, the benefits could yield higher costs. The largest single flight saving observed was 609 lbs. of fuel, and the smallest was 260 lbs. of additional cost (-260 saved).

3 Background Information

The weather optimization tool minimizes different cost functions for each phase of flight. In each phase, the cost function is designed to increase flight predictions accuracy and/or directly reduce cost to the airline. It is thought that increasing the accuracy of flight predictions will allow for improved performance of the NextGen or SESAR systems allowing denser air traffic patterns. Specifically, these systems plan to use 4D trajectory information (latitude, longitude, altitude, and time) to schedule arrivals and optimize traffic through the airspace.

Each of the cost functions is minimized for a single flight plan. The tool does not attempt to choose an optimal cruise altitude or cost index, but simply provides the optimal weather data given the flight plan (optimality is defined by the cost function).

4 Benefits Analysis

The following analysis was performed using 68 weather forecasts provided by AirDat. The 1800 Greenwich Mean Time (GMT) forecast was selected each day in an attempt to limit correlation between forecast databases. In cases where the 1800 GMT forecast was unavailable, the available forecast closest to 1800 (1200, 0600, 0000) was used. The choice to use one weather forecast per day was an attempt to remove bias in overall statistical results from analyzing similar forecasts. Even with the one-per-day selection, there is likely some forecast bias from selecting consecutive days within the same season. The weather was not categorized by severity; it is assumed that the sample size of 68 days is large enough to cover a wide range of weather types within the individual season (Spring 2011). To mitigate this potential source of bias, flight routes were selected spanning the western United States in all directions and lengths. Future studies to characterize weather forecast error would benefit from a wider timeframe of data collection spanning multiple seasons of the year.

1 Numbers reported are the delta between FWIO and non-optimized results, percentages shown are the percentage savings of using the FWIO versus no optimization for various forecast sources.

For each weather set, eight routes were predicted, each with four cruise altitudes, for a total of 32 routes per day. Table 2 shows the Alaska Airlines company routes and cruise altitudes used in this study. For all routes, the cost index is fixed at 25. The altitudes listed below were selected to span a range of standard cruise altitudes. Altitudes for the Seattle/Spokane routes are lower due to the short flight range.

Table 2. List of Flight Plans Used in Analysis

Route	Cruise Altitude 1	Cruise Altitude 2	Cruise Altitude 3	Cruise Altitude 4
SEAMSP1	FL280	FL310	FL340	FL370
MSPSEA1	FL280	FL310	FL340	FL370
GEGSEA2	FL240	FL260	FL280	FL300
SEAGEG1	FL240	FL260	FL280	FL300
SEASFO1	FL280	FL310	FL340	FL370
SFOSEA1	FL280	FL310	FL340	FL370
ORDSEA1	FL280	FL310	FL340	FL370
SEAORD1	FL280	FL310	FL340	FL370

4.1 Perfect Forecast

The hypothetical perfect forecast is defined as a forecast with zero error; the forecasted data exactly matches the true weather. For each flight plan, three different optimizations were tested assuming this perfect forecast. Further analysis to include the effects of forecast error was performed on two of the optimization options (paragraph 4.2).

Optimization 0 (op0) is the "no optimization" baseline; inputting wind from the weather forecast at each waypoint in cruise and at three fixed altitudes in descent (5000, 10000, 15000 ft.). The descent winds are retrieved from the weather database using the arrival airport latitude/longitude.

Optimization 1 (op1) is a version of the optimization algorithm that minimizes wind and temperature residual. This optimization is intended for comparison only, and is omitted from the forecast error analysis (paragraph 4.2).

Optimization 2 (op2) is the FWIO tool that minimizes the performance based cost functions outlined above. This algorithm uses detailed knowledge of the prediction functions to directly increase prediction performance. This algorithm selects weather locations and weather values at the selected locations that minimize the optimization cost function.

It should be noted again that analysis presented in this section (the hypothetical perfect forecast) only looks at savings due to weather model error, and does not address other factors that may impact the accuracy of flight predictions such as meteorological forecast errors, trajectory integration error, or unexpected air traffic control input. The benefits presented in this section in regards to percentage improvement-only factor in weather modeling error, which may only be a small piece of the overall error tree. The combined effect of model error and forecast error is shown in paragraph 4.2.

4.1.1 Descent

The descent phase optimization performance is measured by the following metrics:

- Cost function value -- 80% fuel cost, 20% time cost
- Error in predicted descent time ("continuous" reference weather model versus B737 U11-based weather model)

In the scenarios examined, the average cost associated with performing (non-idle thrust) guidance maneuvers to maintain the predicted trajectory was reduced when using the weather optimization tool (op2) versus no optimization (8 lbs. - 37 lbs. fuel savings depending on route; maximum savings for SFOSEA1 – FL370). The FWIO tool (op2) shows slight improvement over less sophisticated optimization methods (op1) (1 lb. – 11 lbs. fuel savings depending on route; maximum SEAORD1-FL370). Averaging all of the flight plans and weather forecasts, the total assessed cost using the FWIO tool is 8 lbs., compared to 28 lbs. without doing any optimization, and 13 lbs. with less sophisticated optimization. This represents a 71% assessed cost savings when using the weather optimization tool versus a 53% savings from less sophisticated optimization methods. The distribution for this averaged data is shown in Figure 1.

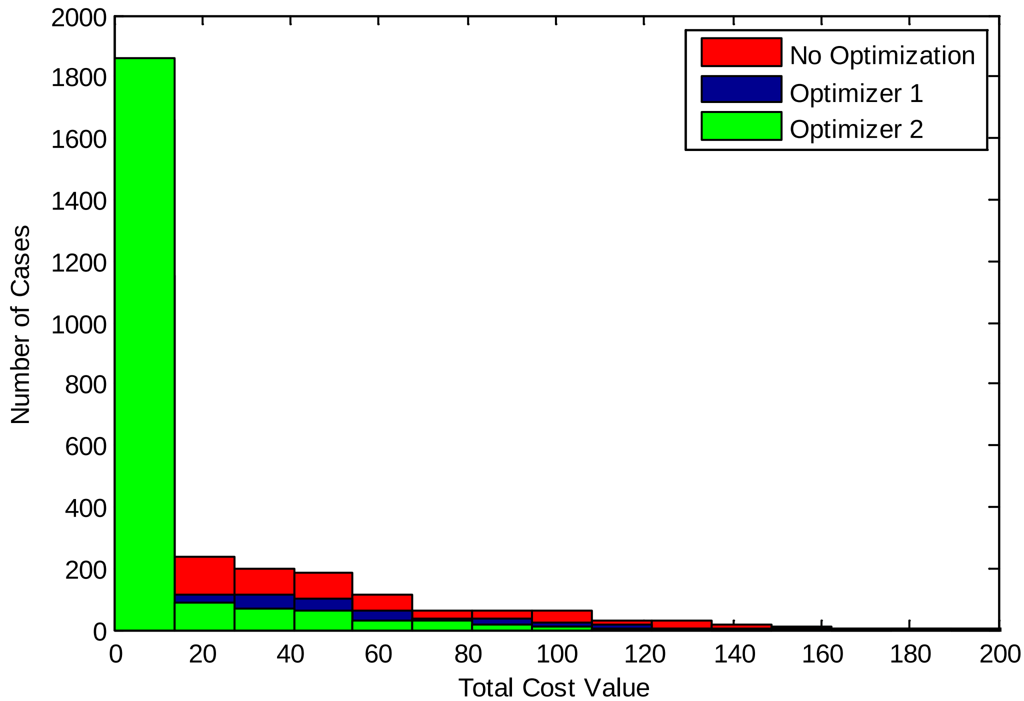

Figure 1. Descent Cost Value Histogram for All Optimization Options. Note the lower concentration of high cost cases for the FWIO tool (Optimizer 2)[2]

In Figure 1, all three optimizers have the same number of tested cases: 2176. It appears that optimizer 2 has more overall cases because it is plotted in front of the optimizer 1 and with no optimization bars. Figure 1 shows that there are more cases with high cost for no optimization and optimizer 1.

[2] All distributions shown have values in the 0-20 bin. Due to the higher concentration of "Optimizer 2" cases in this bin, the "No Optimization" and "Optimizer 1" values of the histogram are hidden behind the "Optimizer 2" bar.

While 28 lbs. of fuel does not seem like a significant amount, multiplied by the number of flights occurring each year amounts to a significant airline cost savings.

These fuel usage values are much lower than airlines' reported fuel cost. This could be attributed to other error sources not analyzed here (such as temperature effects and performance model errors). It is interesting to note that both optimization options provide approximately the same reduction in time cost compared to no optimization (5 seconds of time error). This is likely due to the heavy weighting of fuel cost (80%) in the overall descent cost function. When the weather optimization tool (optimization 2) is applied to the overall cost function, the weighting shifts the focus to the fuel component. Due to the termination tolerance used in the FWIO, it is possible (and likely) that no further iterations are performed to reduce time cost once fuel cost has been driven to zero. On the other hand, optimization 1 operates on a wind residual based cost function, which is directly related to time of flight error. As shown by the fuel cost results, this is sub-optimal in terms of fuel usage, but provides a relatively low time cost solution. Based on these results, it may be possible to further tune the FWIO descent cost function for iterations where fuel cost is zero and apply higher weighting to time cost. Additionally, FWIO users have the ability to change the fuel to time cost weighting depending on their specific needs.

In addition to the direct cost savings in fuel usage shown by using the weather optimization tool, there are indirect benefits to the overall air traffic management system through increased flight predictions accuracy. The predicted time of flight and predicted fuel usage accuracy was assessed using error, defined as the difference in the predicted quantity using the FMS weather model with weather from each optimizer, compared to a common reference of a flight prediction model that uses "pseudo-continuous" weather. Early arrival is counted as a negative time of flight error, and late arrivals are counted as a positive time of flight error. Time of flight accuracy is best represented by the standard deviation of this error quantity. These parameters are of particular importance to required time of arrival and arrival scheduling applications.

Figure 2 shows the overall distribution of all cases using each optimization algorithm. This plot shows significant benefit to using either of the weather optimization algorithms in terms of reducing the standard deviation of the time of flight error, with the FWIO showing a slight improvement over other optimization methods.

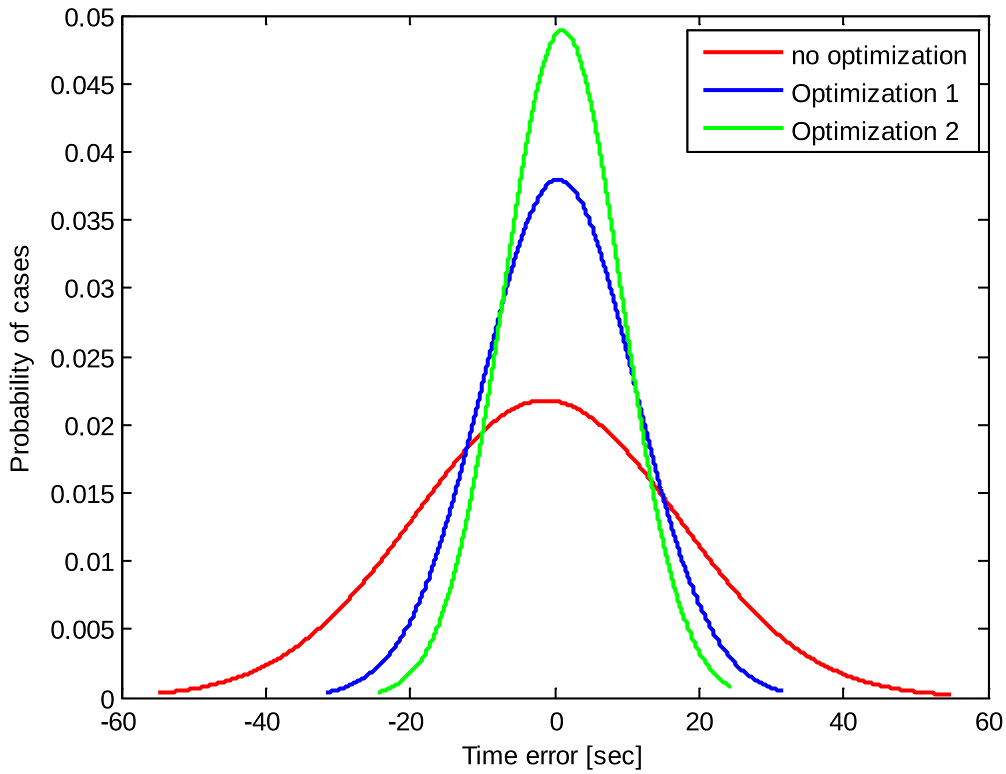

Figure 2. Distribution of Error in Predicted Descent Phase Time of Flight [sec] for All Cases Combined

This analysis shows that by using the weather optimization tool, the predicted time of flight error is reduced to less than 16 seconds in 95% of the cases, compared to 21 seconds using optimization 1, and 36 seconds for no optimization. This amounts to an improvement in temporal descent predictions accuracy of 55% when using the weather optimization tool. It is possible to improve all three methods by increasing the number of data points allowed in the descent weather model.

The results for predicted fuel usage error (Figure 3) show similar results to the predicted time error.

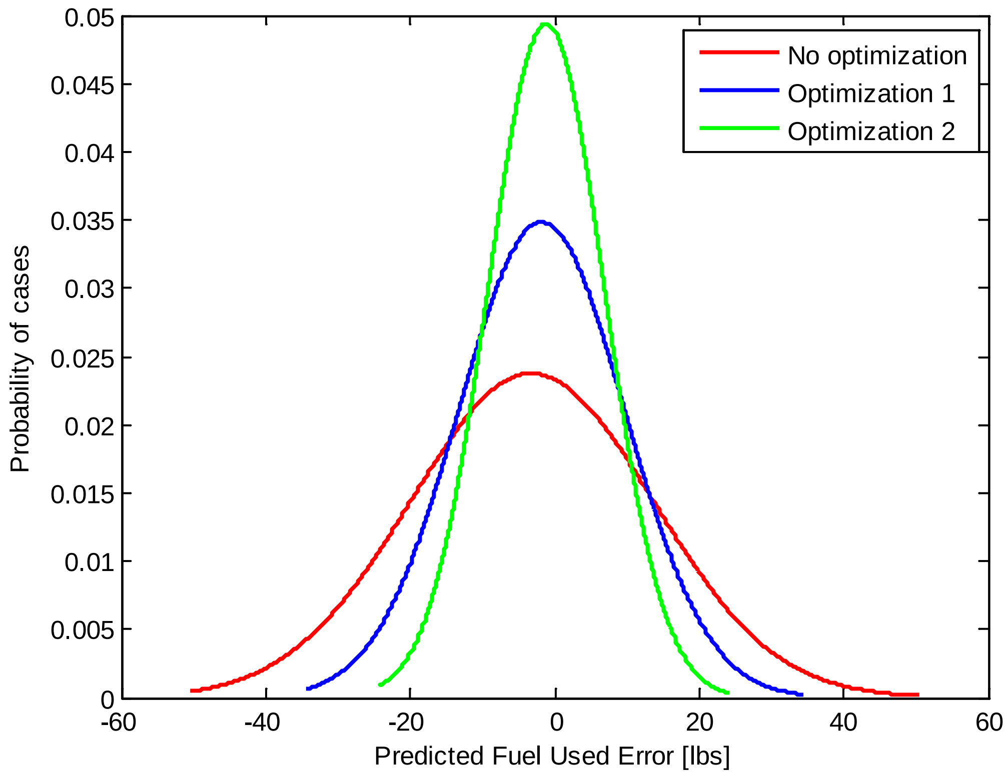

Figure 3. Distribution of Error in Predicted Descent Fuel Usage [lbs.] for All Cases Combined

4.1.2 Cruise

The cruise phase optimization performance is measured by a temporal cost function (prediction temporal accuracy per nautical mile in cruise). Unlike the descent cost function, this metric applies to the predicted path. Because the aircraft flies at constant altitude at constant thrust setting, the fuel usage is directly proportional to this metric. In this analysis, the optimization tools were not allowed to select weather locations, and were constrained to only top of climb, top of descent, and cruise waypoints. This choice was made because the routes used for analysis have many cruise waypoints, and in the majority of cases, additional weather locations were unnecessary. Further analysis should be performed allowing additional weather location selections for routes with long cruise legs (such as a DIRECT-TO) situation, as there is a potential for even greater accuracy improvement than shown in this study.

The FWIO weather values allowed higher temporal accuracy in flight prediction. On average, the FWIO values yielded 0.005 seconds of error/nautical mile, compared to 0.0196 seconds of error/nautical mile without any optimization. Not only did the results show an average of 4x reduction in mean and standard deviation of the cost function averaging all routes; each individual route showed some improvement.

4.1.3 Climb

The climb phase optimization performance is measured by a distance based cost function (feet of error/second in climb phase).

Similar to cruise phase, this cost function metric applies to the predicted path. In this analysis, the optimization tools were not allowed to select additional weather locations beyond what can be entered into the current B737 U11-based FMS, and were constrained to only ground level (current wind) and top of climb. Unlike cruise, there is no method to apply additional weather locations in the FMS, so no additional benefit is expected beyond the results presented here.

The FWIO tool yields approximately 5x reduction in mean and standard deviation of the climb cost function compared to no optimization for all routes. Averaging all cases, the FWIO tool predicted top of climb location within 2.022 feet per second of flight time, versus 9.594 feet per second of flight time without any optimization.

4.2 Forecast with Error

The analysis above shows that using weather from the FWIO can provide significant reduction in flight cost compared to non-optimized weather inputs for perfect forecast scenarios. True weather measurements in flight have been acquired from DART data recorders, which allow statistical characterization of the forecast error for use in the FWIO tool.

To examine the effect of forecast errors, the statistical forecast error models (shown in Appendix B) were applied to the predicted forecast to create hypothetical “truth” scenarios. The FWIO cost function value was then recomputed using the estimated weather (obtained through optimization of the original forecast) and the hypothetical “truth”. See Figure 4 for a graphical depiction of this process. Similar to previous analyses, eight routes with four altitudes each were analyzed for weather forecasts on 68 different days. For each of these route/altitude/day combinations, the error model was applied independently three times to generate three different “truth” weather cases for a small Monte Carlo run, in which the average of the three resulting costs was reported out. Due to processing time constraints, a larger Monte Carlo study was infeasible at this time.

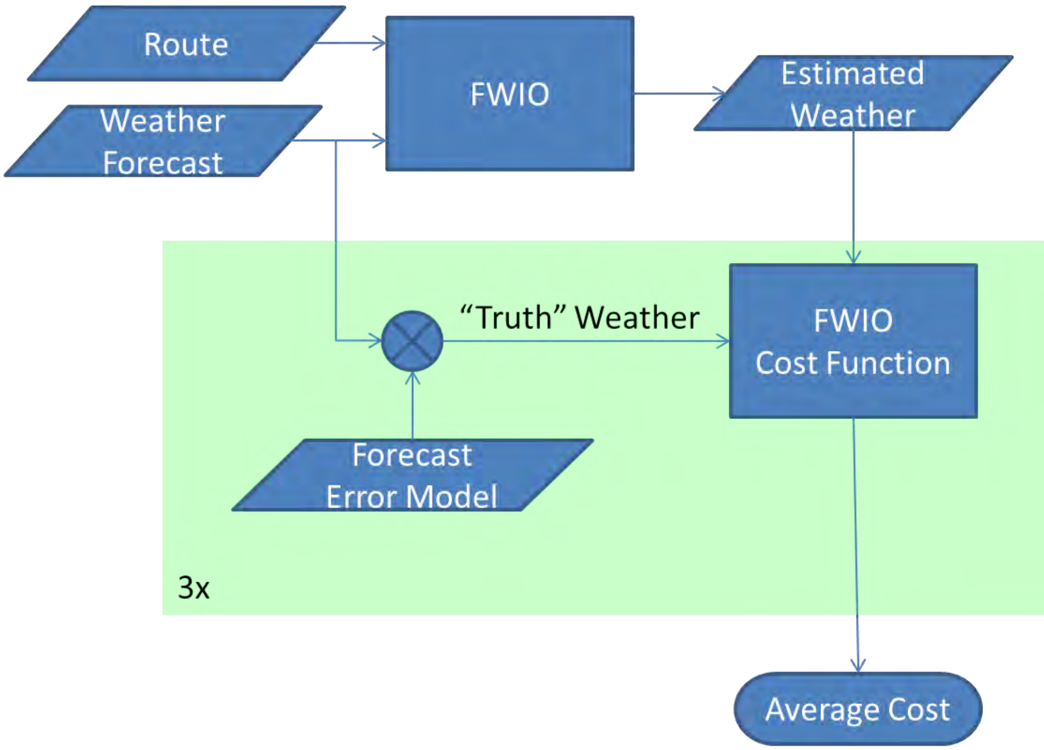

Figure 4. Flowchart for Simulation of "Truth" Weather Scenarios

4.2.1 Descent

When forecast error is taken into consideration, the fuel usage for guidance maneuvers is increased compared to scenarios with a hypothetical perfect forecast. The average descent cost with NOAA error levels increased from 28 to 61 lbs. for no optimization, and from 8 to 48 lbs. with optimization. Similarly, with AirDat errors, the cost is increased from 28 to 52 lbs. for no optimization and from 8 to 39 lbs. with optimization.

Overall, the cost savings of using the FWIO went from 20 lbs. with no error down to 13 lbs. with either error model. The original reported cost savings of (71% reduction in descent cost when using the FWIO versus a non-optimized forecast) has been reduced (now 21% for the NOAA error and 28% for the AirDat error – a much smaller fraction of the overall guidance maneuver cost). This reduction is to be expected because the optimization attempts to match the predicted forecast, not the true weather. When there are errors present, the true value may lie closer to the inaccurate forecast weather estimation, resulting in lower cost than even the optimized value. In addition, the forecast error from either NOAA or AirDat can be large in comparison to the model error. The FWIO tool can only be as good as the input forecast, and can only remove the model error component.

When comparing the combined effect of using the better forecast (AirDat) with the effect of using a better model (FWIO tool), the overall cost decreases from 61 lbs to 39 lbs; an overall savings of 21lbs; 13 can be attributed to the FWIO, and 8 attributed to the better forecast.

Comparing the average fuel usage for guidance maneuvers with forecast error to the perfect forecast analysis, the average cost with NOAA error levels increased from 32 to 70 for no optimization, and from 8 to 54 with optimization. Similarly, with AirDat errors, the cost is increased from 28 to 60 for no optimization and from 8 to 43 with optimization. In other words, using the FWIO tool will save 16 to 17 lbs. of fuel per flight in the presence of forecast error, down from 24 lbs. with a perfect forecast. Overall, when forecast error is included, fuel cost is higher, but can still be reduced by using the FWIO tool to generate weather input data.

When combining the effect of the better forecast (AirDat) with the better model (FWIO) descent guidance fuel usage goes from 70 lbs. down to 43 lbs., a net savings of 27 lbs. per flight. It is important to note that these benefits are presented on average per flight; but on an individual flight basis using the FWIO tool with AirDat weather forecast could yield higher costs. The largest single flight saving observed was 609 lbs. of fuel, and the smallest was 260 lbs. of additional cost (-260 saved).

4.2.2 Cruise

In the presence of forecast error, the optimization tool produced a weather estimate that resulted in a time accuracy improvement of 0.004 - 0.005 seconds/nautical mile on average. The perfect forecast analysis presented above showed an improvement of 0.014 seconds/nautical mile with no forecast error. It should be noted that both analyses used routes with many enroute cruise waypoints; resulting in very small model errors for both the optimized and non-optimized cases. In turn, the large forecast error component in this analysis tends to dominate the overall error, and thus the cruise time accuracy.

Reducing this forecast error component is therefore crucial to increasing temporal accuracy in cruise. Comparing the NOAA versus AirDat resulting errors, an average accuracy increase of 0.079 seconds/nautical mile travelled (28%) can be seen. Although this accuracy increase dwarfs the potential savings from using the FWIO tool for the routes in this study, routes with few cruise waypoints (such as a DIRECT-TO) have significantly higher model error which can be removed by the FWIO tool.

4.2.3 Climb

The average climb accuracy is not significantly improved (nor reduced) from use of the FWIO optimization tool in the presence of forecast error. Similar to cruise, the forecast error component tends to dominate over the modeling error component and wash away any potential savings. By using a better forecast model (AirDat) the average distance error can be reduced from 25.5 ft./second of flight in climb to 17.1 ft./second of flight; an accuracy improvement of 33%.

Appendix A FWIO Benefits with One Third Forecast Error

A factor of one third was selected to scale down all standard deviations and biases in the wind and temperature models to represent the case where the measured errors in the DART sample set represented a three sigma case. This is likely a lower bound (best case) estimate of benefits, with the true benefits lying between the results presented here (see Table 3), and the results presented in prior sections.

	FWIO vs. non-optimized (both using same forecast)			FWIO with AirDat vs. non-optimized w/ NOAA
	Perfect Forecast	NOAA Forecast	AirDat Forecast	
Descent Cost Savings (lbs - sec)	20 (71%)	17 (45%)	17 (45%)	17 (45%)
Descent Fuel Savings (lbs)	24 (73%)	20 (47%)	20 (46%)	21 (47%)
Cruise Temporal Prediction Accuracy Increase (s / NM)	0.014 (71%)	0.007 (7%)	0.006 (9%)	0.031 (31%)
Climb Distance Prediction Accuracy Increase (ft / s)	7.6 (80%)	3.2 (27%)	4.7 (44%)	5.8 (49%)

Table 3. Summary of FWIO Benefits using One Third Forecast Error on Average per Flight[3]

Note that with these smaller errors, the descent costs for either forecast are approximately equivalent, but cruise and climb prediction accuracy is improved with the AirDat forecast. This implies that for these reduced error levels; model error dominates in descent phase and forecast error dominates in climb and cruise phase.

[3] Numbers reported are the delta between FWIO and non-optimized results, percentages shown are the percentage savings of using the FWIO versus no optimization for various forecast sources.

Appendix B Forecast Error Characterization and Results

The weather forecast error is defined as the delta between measured weather and predicted weather. This error has been calculated at five second intervals for the entirety of each DART flight for temperature (Deviation from the International Standard Atmosphere (DISA)), north component of wind, and east component of wind to form a sequence of errors.

$$e = W_{DART} - W_{forecast}$$

The error for a given flight can be characterized in two components: bias and a "noise" component.

$$e = bias + noise$$

For each flight, the bias can be estimated by taking the mean of the sequence of errors. Using this estimated quantity, a zero mean noise component can be isolated.

The wind and temperature errors have very different "noise" characteristics. In fact, the "noise" is more truthfully described as simply the remaining error. In the case of temperature, this more closely resembles low standard deviation white noise. In the case of wind, the noise exhibits low frequency oscillatory behavior. This makes physical sense because atmospheric temperature is fairly well known, and can be described with a simple altitude model (loose dependence on latitude/longitude). Atmospheric wind is dependent on complex pressure systems moving around the surface of the earth. As the aircraft flies through various pressure systems, it may encounter winds earlier or later than expected. The wind error is generally high amplitude peaks and valleys. For this reason, the wind and temperature "noise" components are addressed with separate models.

In order to process the limited data set with the most possible data, each of the flight phases (climb/cruise/descent) are analyzed together as one set of forecast errors. In truth, each phase of flight possesses its own unique characteristics. For instance, measured temperature is typically higher than truth in climb and lower in descent due to thermal properties of the sensor (measurement lags truth). In descent, there is typically higher wind measurement error than other flight phases due to higher aircraft yaw. In addition, descents are typically flown into a headwind, which may produce a wind measurement bias. Some of these effects can be seen in the example plots shown above. With a larger data set, it may be possible to separate these effects out to obtain a more accurate representation of the true weather. For the purposes of characterizing the forecast errors in this study, these factors are assumed to be small, and DART measurements are used as truth data.

B.1 Temperature Error Model

As described in the previous section, a "noise" component can be extracted from the overall temperature error by subtracting out the bias term.

The resulting signal resembles a random walk or fractional Brownian motion, but upon further examination, there is no physical reason the error should grow uniformly with time[4] (assuming fresh forecast for all data points). Ruling out these types of errors, the best match is correlated Gaussian noise.

[4] As stated in paragraph 4.2.1, different flight phases possess different error characteristics, and higher error may be expected in climb or descent.

New simulated noise signals can be reproduced using the statistical measurements from DART flights. These resulting simulated signals are again smoothed using the one minute moving average filter. A few examples of simulated signals are shown in Figure 5.

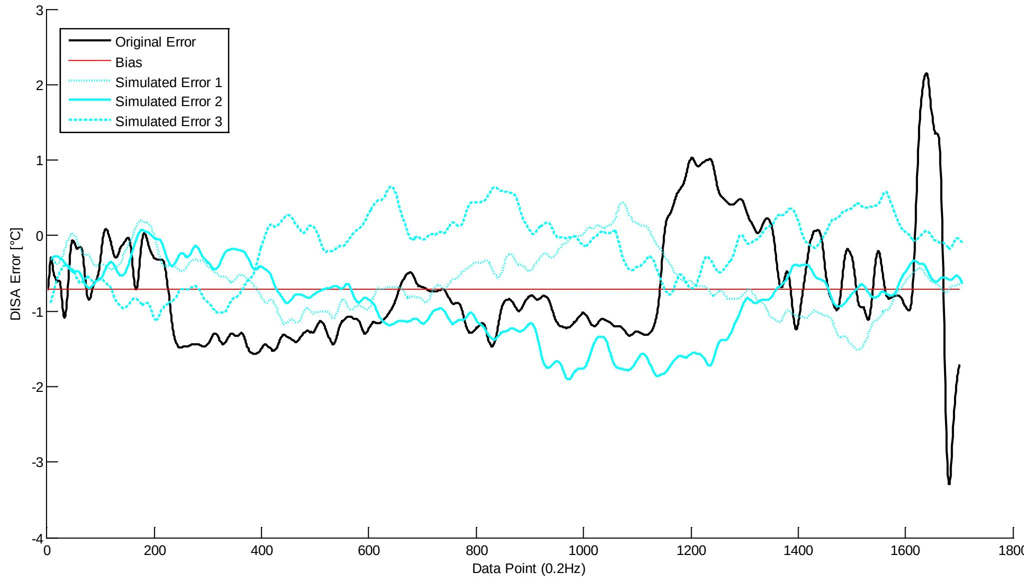

Figure 5. Example of Simulated DISA Error Signals

B.2 Wind Error Model

The remaining "noise" component of the wind error is initially filtered using a five minute moving average filter with zero phase lag to remove quantization and measurement errors. An example of filtering results is shown in Figure 6. It should be noted here that the five minute filter time window was selected to smooth the large spikes in the "noise" signal. It is thought that these spikes represent either measurement error (wind is not directly observed onboard the aircraft) or short-term wind gusts. For purposes of this analysis, short wind gusts and any unique characteristics by flight phase are neglected.

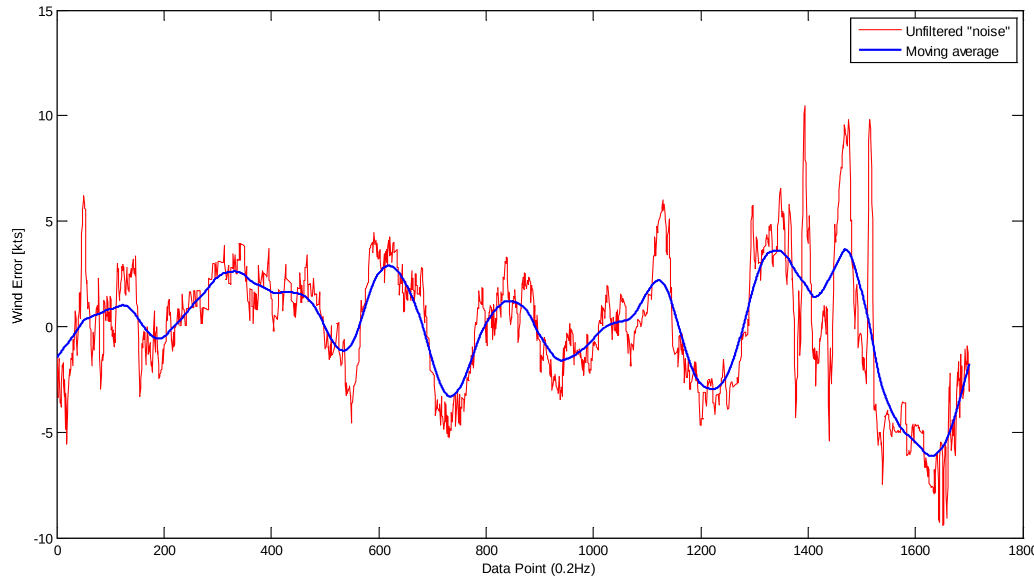

Figure 6. Example of Wind Moving Average Filter

The resulting wind "noise" signal is somewhat similar to the temperature "noise" signal, but exhibits much higher amplitude oscillations with lower frequency. The signal is short term highly correlated and long term uncorrelated. Because of the nature of this data, if the correlated Gaussian noise model is applied to these results, a very high correlation coefficient (0.9999) is obtained. This correlation tends to produce simulated signals that do not exhibit any oscillations. The model does not provide a means for long term versus short term correlation. Again, a random walk model can account for long and short term correlations through the Hurst parameter, but this model is ruled out due to its increasing error over time.

The wind "noise" signal can be modeled using an extrema model. In this model, each of the local extrema is identified in the "noise" signal. The rate of occurrence of these extrema can be modeled using a gamma distribution, which 4is a one-sided (0→∞) statistical distribution commonly used for waiting times. The Probability Density Function (PDF) of the gamma distribution is defined as:

$$y = f(x|a,b) = \frac{1}{b^a\Gamma(a)} x^{a-1} e^{\frac{-x}{b}}$$

Where Γ (x) is the gamma function

The defining parameters of the gamma distribution (a, b) can be estimated using numerical techniques.

In addition to the location of the extrema, the distribution of the values at these extrema is needed. These values tend to follow a correlated Gaussian noise distribution. The correlation coefficient and standard deviation of the extrema can be extracted from the signal data.

Simulated noise signals are produced for this study based on the DART weather statistics. A linear interpolation is used between extrema locations, and then a five minute moving average filter is applied to the data. This filter tends to apply a smooth transition between the extrema without introducing higher amplitude errors that would come from a spline or cubic interpolation fit. The moving average of an oscillatory signal tends to have reduced amplitude from the original signal. To remedy this effect, the resulting moving average is scaled upwards to have the same standard deviation as the original noise signal. Figure 7 shows a few examples of simulated wind errors.

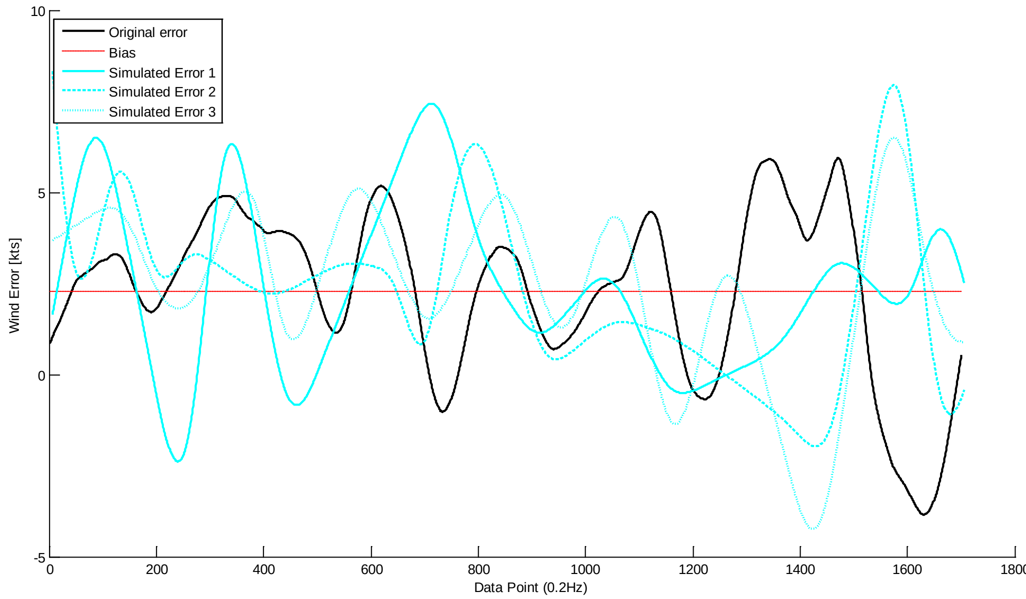

Figure 7. Example of Simulated Wind Component Error Signals

B.3 Forecast Error Model Statistical Results

The error models above were applied to the set of 99 DART 737-800 flights from mid-August through the end of September. The majority of these flights travelled in the US Pacific Northwest, and up and down the west coast. Of the 99 potential flights from Summer 2011, corresponding NOAA weather data was obtained for 40 flights, and AirDat data was obtained for 22 flights (with a very small overlap between the AirDat and NOAA data corresponding to the flight samples). The difference in sample size is due to different geographic collection regions between the collected NOAA and AirDat forecasts and data acquisition gaps. This extremely small number of samples in a focused region is undesirable for a large scale statistical analysis. The aggregate results presented below should be used with caution, as they may not accurately represent the entire population of weather error. Future forecast error characterization studies should use a much larger sample set across multiple seasons, or even years.

B.3.1 Wind Error Model

The statistical summary of the error model parameters for North wind component and East wind component are presented in Table 4 and Table 5 respectively. One interesting thing to note here is that the NOAA model is fairly accurate in the North/South direction, with only 2 knot bias and 4.6 knot standard deviation of the remaining error. Extrema occur at moderate intervals with low amplitude. On the other hand, for the East/West direction, the bias is much higher at 15 knots and standard deviation of the remaining error jumped to 11 knots. The extrema in the east west direction occur less frequently with high amplitude. The overall wind error in the NOAA model tends to stretch along the East/West direction. The AirDat error tends to be more omnidirectional in nature, with -4 knots of bias North/South and 6 knots East/West. The standard deviation of the remaining error is eight knots in either direction.

In normal weather patterns, the jet stream passes through the Pacific Northwest primarily from the West to East, so it is possible that the error seen is simply error in predicting jet-stream intensity. In other areas of the globe, the error may stretch in different directions, or could be circular in nature.

Regardless of the shape of the errors, they can be directly compared by taking the magnitude of the combined error vector. On average, the NOAA wind forecast has 15 knots of total bias error, compared to 7.3 from AirDat. The NOAA wind forecast has 13.3 knots of noise error, compared to 11.26 from AirDat. Most of the AirDat error model quantities have higher standard deviation, which could be attributed to the smaller sample size (22 flights versus 40).

Table 4. Wind North Component Error Model Parameters

Parameter	Mean Value (NOAA)	Standard Deviation (NOAA)	Mean Value (AirDat)	Standard Deviation (AirDat)
Bias –kts	1.92	3.49	-4.21	10.52
Extrema Correlation	0.02	0.34	0.30	0.33
Extrema Gamma (a)	2.03	1.37	2.52	1.62
Extrema Gamma (b)	56.78	23.83	46.31	26.26
Noise Standard Deviation – kts	4.61	1.39	7.38	5.26
Extrema Standard Deviation – kts	4.99	1.48	7.86	6.27

Table 5. Wind East Component Error Model Parameters

Parameter	Mean Value (NOAA)	Standard Deviation (NOAA)	Mean Value (AirDat)	Standard Deviation (AirDat)
Bias – kts	14.89	11.24	5.98	9.84
Extrema Correlation	0.30	0.37	0.08	0.46
Extrema Gamma (a)	2.11	1.72	2.30	1.52
Extrema Gamma (b)	72.42	42.95	51.67	31.34
Noise Standard Deviation –kts	11.03	4.87	7.28	7.46
Extrema Standard Deviation - kts	12.34	5.79	8.06	8.49

B.3.2 Temperature Error Model

The statistical summary of the DISA error model parameters is shown in Table 6. Both forecast sources are fairly accurate in their temperature predictions. NOAA has a slightly more accurate prediction, with an average bias of -0.5 °C compared to -1.0 °C for AirDat. The remaining noise has a standard deviation of 0.7 °C for NOAA and 1.3 °C for AirDat. Although this shows that NOAA DISA error is ~50% smaller, both forecasts have extremely small error on average.

Table 6. DISA Error Model Parameters

Parameter	Mean Value (NOAA)	Standard Deviation (NOAA)	Mean Value (AirDat)	Standard Deviation (AirDat)
Bias - °C	-0.47	0.27	-0.95	1.01
Noise Standard Deviation - °C	0.70	0.16	1.30	0.60
Noise Correlation Coefficient	0.998	0.001	0.999	0.001