A Methodology and Results for Comparative Assessment of the Prediction Performance of the Collaborative Routing Coordination Tools (CRCT)

Kenneth S. Lindsay
The MITRE Corporation
Center for Advanced Aviation System Development
7515 Colshire Drive, McLean, VA 22102

Abstract

The Collaborative Routing Coordination Tools (CRCT) is the prototype of a set of tools to help the Federal Aviation Administration (FAA) to detect traffic flow problems in advance, to generate problem resolutions, and to evaluate the resolution strategies. CRCT does this by modeling four-dimensional aircraft trajectories and using them to predict demand for sector usage. A methodology was developed and used to compare the prediction performance of CRCT under various software and data configurations. The methodology can be and has been used for other tools (e.g., the Enhanced Traffic Management System (ETMS)) that predict sector demand. Several performance metrics were defined and tools were developed to calculate the metrics from prediction data and actual track data. The metrics assess how closely predictions match the sector loads that actually occur. This paper presents the methodology used to assess CRCT's prediction performance and the results of various comparative performance analyses. The analyses are based on CRCT runs against recordings of actual air traffic data on both good and bad weather days. Analysis results are compared for the two weather days and for three Air Route Traffic Control Centers (ARTCCs).

Introduction

The FAA is responsible for ensuring the safe and smooth flow of traffic in the National Airspace System (NAS). This entails monitoring traffic flows and NAS conditions and redirecting flows as necessary to make them safe and smooth and to balance capacity and demand while minimizing ground and airborne flight delays. To support this function, MITRE/CAASD has developed CRCT, an evolving prototype of tools to support early recognition of traffic flow problems, generation of problem resolutions, and evaluation of resolution strategies.

CRCT has been installed at the Kansas City (ZKC) and Indianapolis (ZID) ARTCCs and in the Air Traffic Control System Command Center (ATCSCC) where it is being evaluated. CRCT features that are deemed operationally mature are used to augment the capabilities of ETMS. Determination of the CRCT features that are to be implemented in ETMS is based on the results of field evaluations and on prediction performance assessments of CRCT. Flow constrained area (FCA) and rerouting are examples of CRCT capabilities that have been improved for and are being implemented in ETMS[1].

CRCT uses flight plans and adaptation data to model four-dimensional trajectories for aircraft[2]. The trajectories and other data are in turn used to predict and alert traffic management specialists when sector demand will exceed a preset threshold. When a traffic flow management (TFM) decision support tool (DST) such as CRCT predicts that the demand for a sector will exceed its threshold, the specialist will typically act to prevent the overload. As a result, the sector load that is actually realized will tend to be smaller than the predicted demand. The difference between predicted demand and actual sector load is therefore not a measure of true predictive accuracy if any intervention is allowed between the prediction time and the time for which the prediction is made[3]. In spite of this, such a measure is adequate when used to compare the relative performance of a TFM tool under various configurations. Measures of this type are used to identify CRCT capabilities that would improve the predictive accuracy of ETMS.

Previous prediction performance analyses for TFM and air traffic control (ATC) DSTs have compared predictions to actual events and allowed intervention between the prediction and event times. Actual recordings of field data (which contains controller interventions such as aircraft vectors) were used to assess the functional performance of the User Request Evaluation Tool (URET) prototype[4,5] and to verify that the production system met its accuracy requirements[6,7]. A similar methodology was used to assess trajectory prediction accuracy for the Center-TRACON Automation System (CTAS) Descent Advisor[8]. Comparative analyses of CRCT's prediction performance have also been done using CRCT predictions and recorded field data[3,9].

Prediction Performance Analysis

Prediction performance analysis of CRCT entails determining how closely CRCT predictions about aircraft behavior match the observed behavior of the aircraft, calculating operational metrics from these differences, and comparing the metrics for alternative CRCT data/software configurations. The analysis is based on recorded air traffic data, also known as scenario data. Since the scenario data includes controller interventions between prediction and event times, the analysis results tend to underestimate the true predictive accuracy of CRCT. CRCT predictions are made from trajectories that are modeled from flight plans and amendments contained in the scenario. Actual aircraft behavior is determined from track reports, which are also contained in the scenario.

Figure 1 illustrates the process that is used to conduct prediction performance analysis of CRCT. The process begins with a CRCT run against an air traffic scenario that consists of recorded ETMS data and optionally, recorded ARTCC Host data. Host data differs from ETMS data in that Host track updates occur on a 12-second cycle while ETMS track updates occur on a 1-minute cycle. After the CRCT run is made, a data extraction and analysis tool (DEAT) is run against the same scenario data. During the CRCT run, predicted sector transit records are collected and during the DEAT run, actual sector transit records are collected. Transit records and ETMS Monitor/Alert (M/A) predictions are input to the transit record analyzer (TRA), which calculates predicted transit errors for CRCT and for M/A. Finally, a data stratification and aggregation tool (DSAT) is used to stratify and aggregate TRA measures to produce operational performance metrics for multiple CRCT runs and/or ETMS M/A predictions. These results can then be compared to assess the relative performance of CRCT runs and/or ETMS M/A predictions.

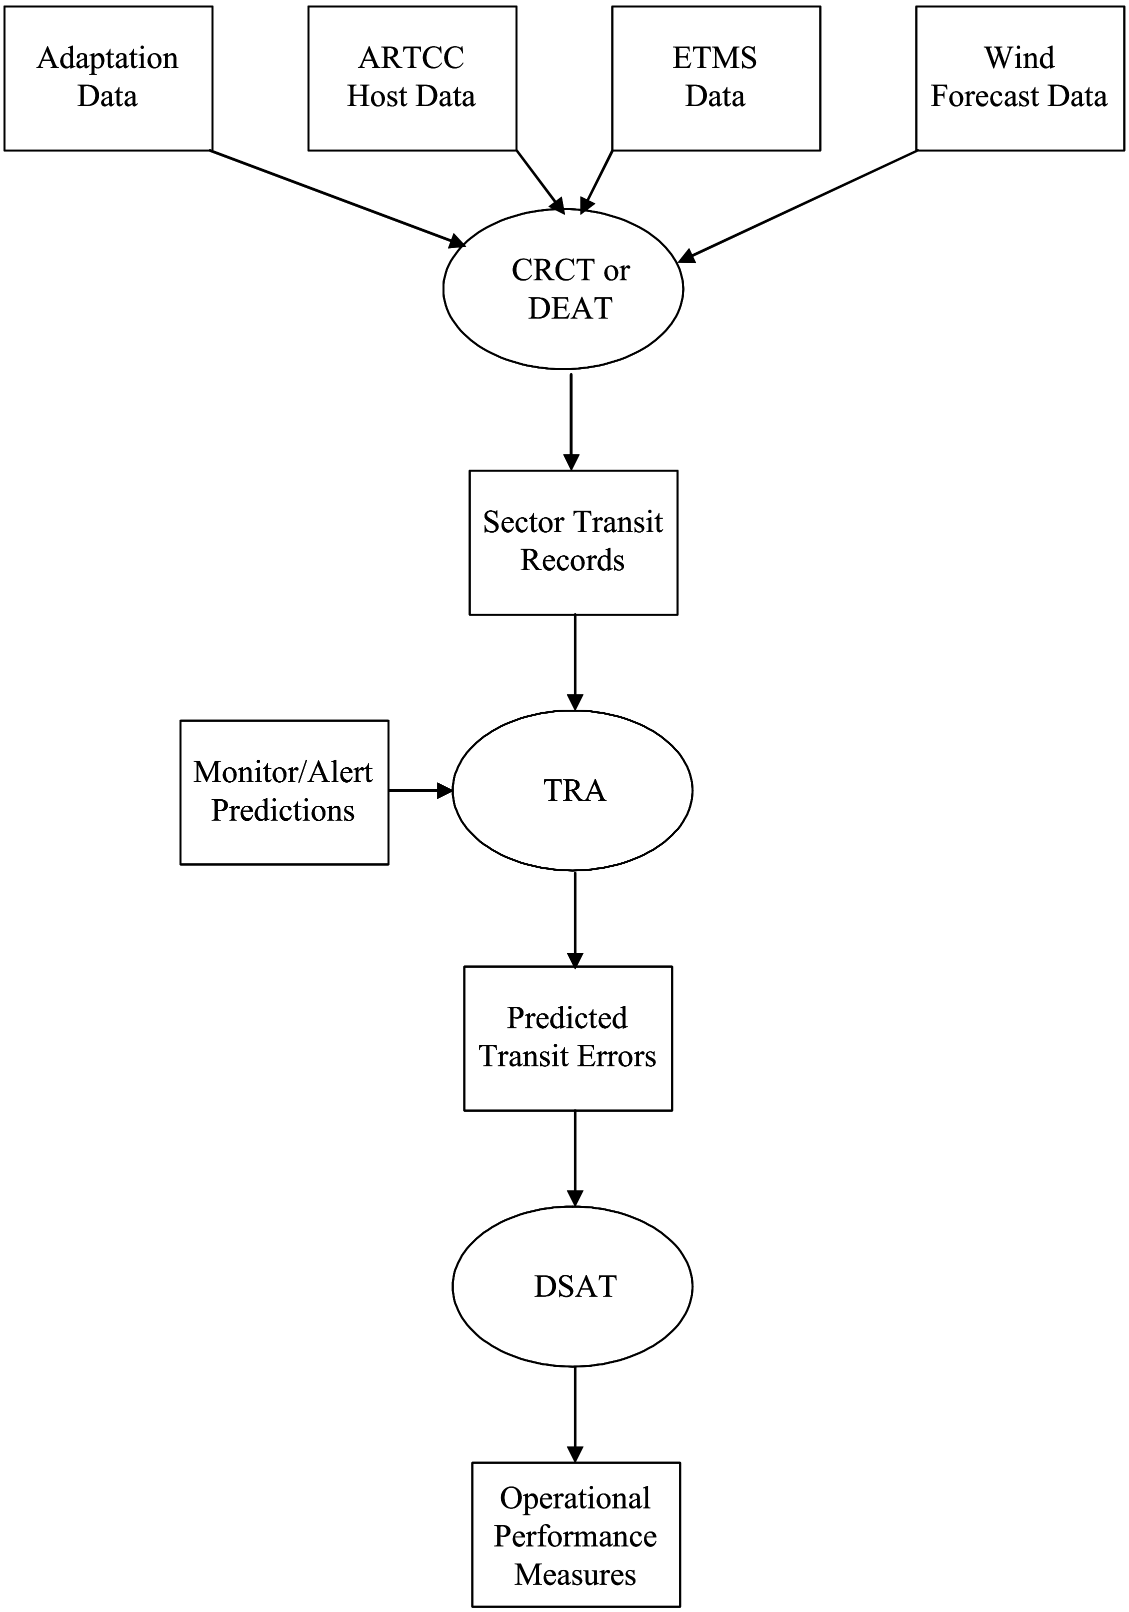

Figure 1. Prediction Performance Assessment Process, Tools, and Data

ETMS currently does not dynamically update sector configurations to reflect those specified in Host sectorization messages. Default sector configurations are used. As a result, the accuracy of M/A predictions can be degraded when sectors are reconfigured. CRCT, as configured in the field, does use Host messages to update the current sectorization. However, to ensure fair comparisons of CRCT and ETMS M/A, Host sectorization messages are not used during CRCT prediction performance analyses. In these analyses, default sector configurations are used by CRCT to generate predicted transit records and by DEAT to generate the actual transit records.

CRCT operational performance metrics characterize the accuracy of the CRCT predictions that are displayed to or noticeable by the user. These metrics consist of sector count error, sector entry error, sector dwell error, and hit rate. Predicted and actual sector transit records are used to calculate these measures. Sector transit records consist of the flight and sector IDs and the sector entry and exit times. Attributes such as flight type, sector type, and lookahead time are used to stratify the metrics. The lookahead time is the difference between the time at which the prediction is made and the time for which the prediction is made. The mean is determined as an aggregation of the operational performance measures. Each metric is described below.

Sector Count Error is the difference between the CRCT-predicted and the actual aircraft count for a sector for a 15-minute time bin. The mean absolute error is calculated for all CRCT-predicted sector counts. The error is stratified by lookahead time and can be stratified by sector type (low, high, and super-high altitude).

Sector Entry Error is the difference between the CRCT-predicted time and the actual time that an aircraft enters a sector. The mean or mean absolute sector entry error is calculated for all sector entry times predicted by CRCT. The error is stratified by lookahead time and can be stratified by flight type (active and proposed).

Sector Dwell Error is the difference between the CRCT-predicted and the actual time that an aircraft spends in a sector. The mean absolute error is calculated for all CRCT-predicted sector dwell times. The error is stratified by lookahead time and can be stratified by flight type.

There are two types of **Hit Rate**. Predictive hit rate is the fraction of flights predicted to enter a sector that actually do enter the sector. Actual hit rate is the fraction of flights that enter a sector that were predicted to enter the sector. The complement of the actual hit rate is the missed prediction rate and the complement of the predictive hit rate is the false prediction rate. The mean hit rate is calculated for all CRCT-predicted and for all actual sector transits. Hit rate is stratified by lookahead time and can be stratified by flight type.

Prediction Performance for CRCT with New Altitude Restriction Data

As illustrated in Figure 1, CRCT uses adaptation data to model trajectories. One component of this data is altitude restriction data. Altitude restrictions, derived from letters of agreement (LOAs) and from standard operating procedures (SOPs), specify constraints on aircraft altitude at center and sector boundaries. CRCT uses these constraints to model the vertical profile of the aircraft trajectory. Because of this, missing or inaccurate altitude restriction data will impact the predictive accuracy of CRCT.

Due to time constraints, the altitude restriction data that was initially used by CRCT was developed using various approximations and simplifications. The smallest number of restrictions was written that would get the most flights into the right sectors. For example, documented arrival restrictions to an airport that would result in flights entering a given sector were coded as a single restriction in the data, even if the restrictions were at different altitudes and for flights arriving from different directions. The coding for a restriction that applied to a major airport and its satellites did not list the satellites, meaning that the restriction will never be applied for the satellite airports. The altitude restriction data that is currently being used is a modification of the initial data and reflects several updates to LOAs and SOPs in ZID, ZKC, and the Memphis (ZME) ARTCCs. In addition, the initial CRCT altitude restriction data has been updated for the Atlanta (ZTL), Chicago (ZAU), Cleveland (ZOB), and Washington (ZDC) ARTCCs.

CRCT altitude restriction data has been updated again and the software that applies the restrictions has been modified. Restrictions have been crafted at a lower level of detail. Multiple restrictions are no longer coded as a single restriction if it means that accuracy would be compromised. The coding for a restriction for a major airport now explicitly lists the satellite airports to which the restriction also applies. ARTCC restriction data has been updated for all except the Boston (ZBW), Los Angeles (ZLA), Miami (ZMA), New York (ZNY), and Oakland (ZOA) ARTCCs. The altitude restriction application software has been updated to use additional criteria to determine whether or not to apply a restriction. Two of these criteria are eligibility altitude and restriction schedule. In addition, schedule time adjustments are made when required for daylight savings time.

A prediction performance analysis was performed on CRCT using no altitude restriction data, the current data, and the new data. CRCT runs were made for each of the three restriction data options. In all three cases, the trajectory modeling software was the same with the exception that the software used to process the new altitude restriction data contained new logic to use additional restriction data elements (e.g., eligibility altitude) in applying the restrictions. Operational performance metrics were generated and compared for ETMS M/A and for the predictions made during the CRCT runs.

ZID and ZKC scenario data, ETMS data, and RUC data were obtained for two days, one a good weather day and the other a bad weather day. Figure 2 illustrates weather radar images for one hour on 17 May 2002, the bad weather day, and for another hour on 22 May 2002, the good weather day. ZID and ZKC CRCT runs were made using both ARTCC and ETMS data. In addition, a Chicago ARTCC (ZAU) CRCT run was made using ETMS data only as Host data was not available. The ZAU run was made because ZAU has more transitioning flights than either ZID or ZKC and more of these flights are expected to be eligible for altitude restrictions. ZAU M/A predictions were not available for the two scenario days and as a result, ZAU operational performance metrics comparisons were made only for CRCT runs with the three restriction data options. A total of 18 CRCT runs were made for the various combinations of the three ARTCCs, the two weather days, and the three altitude restriction data options.

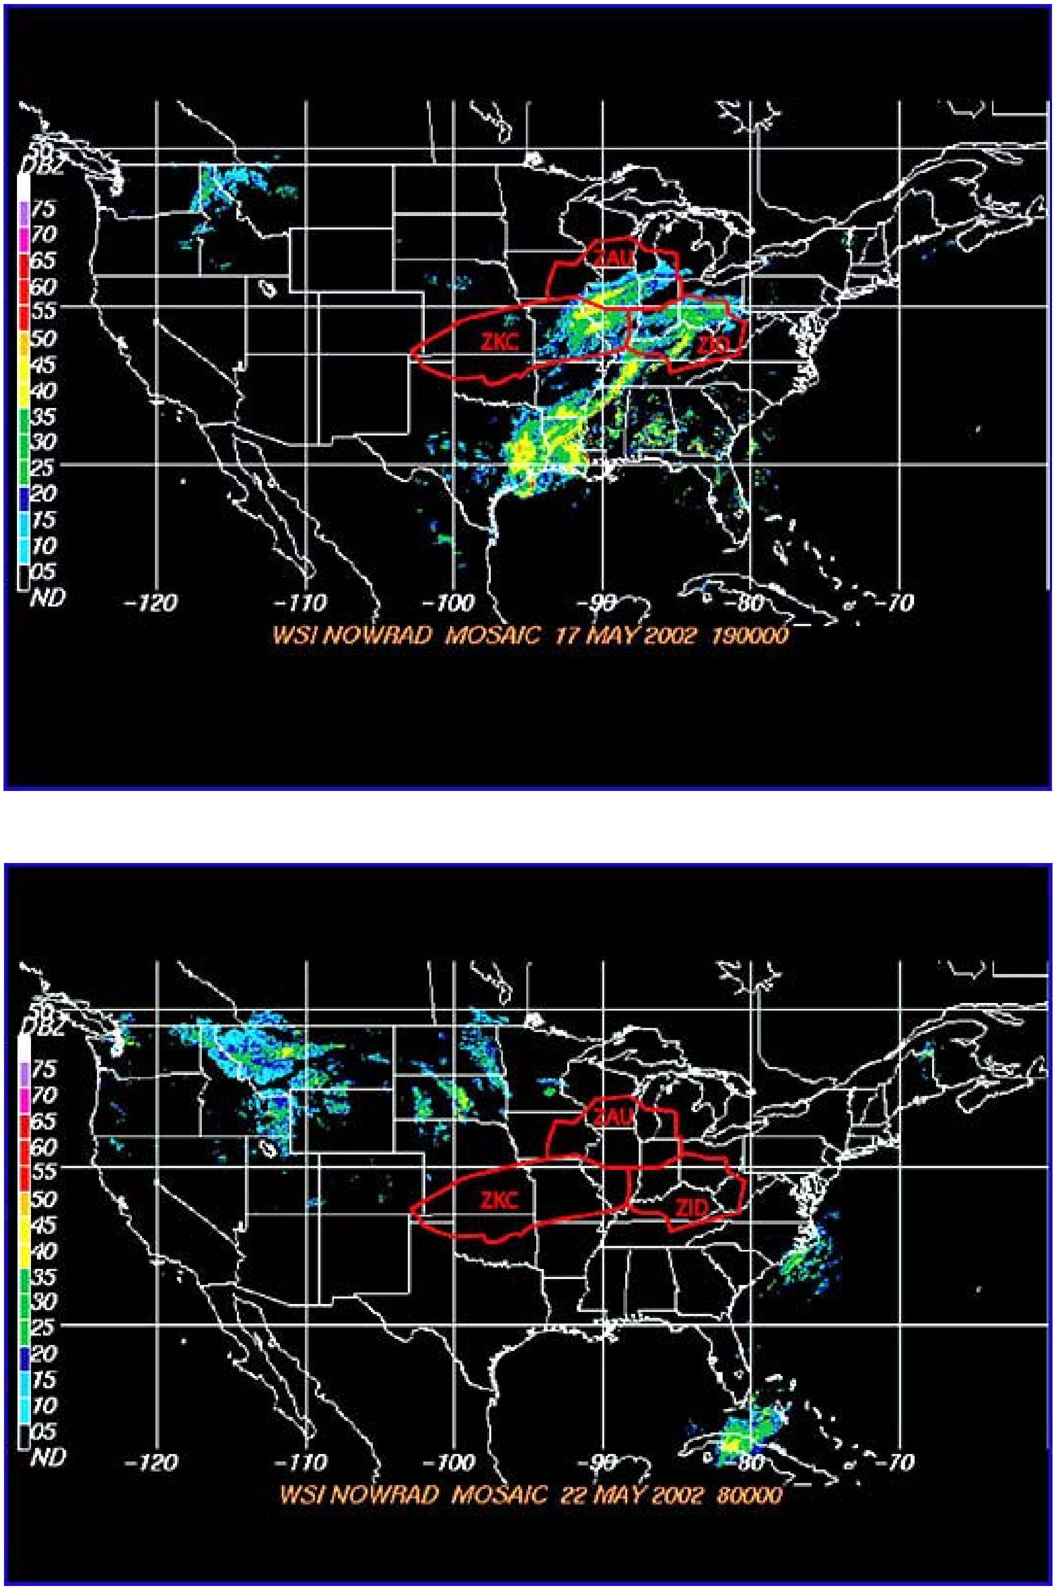

Figure 2. Weather Radar Images of Bad and Good Weather Days

On the bad weather day (17 May 2002), scattered convective weather existed in ZAU, ZID, and ZKC at the beginning of the scenario. The scattered areas expanded and consolidated into a single mass of bad weather that blanketed a large portion of the midwestern U.S. Some areas experienced thunderstorms. The mass moved slowly eastward and by the end of the scenario, it had left ZAU and ZKC, but still covered a large part of ZID. There were ground stop programs at seven major airports. One of the seven also had a ground delay program. All of the airports had flight cancellations. The CAN1 East and J6 Playbook reroutes were executed. On the good weather day (22 May 2002), the weather was clear in ZAU, ZID, and ZKC throughout the scenario. There were no ground stops, no ground delay programs, no cancellations, and no Playbook reroutes during the entire scenario.

Data was collected during the 18 CRCT runs and was used to calculate sector count error, sector entry error, hit rate, and sector dwell error. Figure 3 illustrates sector count error for the ZID CRCT runs, Figures 4 illustrates predictive hit rate for the ZAU runs, and Figure 5 illustrates actual hit rate for the ZKC runs. The prediction performance of CRCT with various altitude restriction data is about the same whether the current or the new altitude restriction data is used (Figures 3 – 5). This result holds for good and bad weather days in ZID, ZKC, and ZAU. Prediction performance was generally better when altitude restriction data (current or new) is used than when it is not (Figure 3). The notable exception to this is the ZAU CRCT runs for the bad weather day, which had a better sector count error and better predictive (Figure 4) and actual hit rates when no altitude restriction data was used. This is probably due to the fact that the ZAU altitude restriction data has not been refined as much as its ZID and ZKC counterparts. The bad weather probably contributes to the unexpected result.

CRCT prediction performance was generally better than that for M/A, with some qualifications. CRCT sector count and dwell errors were generally smaller than those for M/A (Figure 3). CRCT sector entry error was less than the corresponding M/A error for both weather days. Comparison of M/A and CRCT hit rate yields mixed results for ZID and ZKC. M/A predictive hit rate is higher for ZKC. For ZID, M/A hit rate is higher than the hit rate for CRCT runs with no restriction data. On the bad weather day when restriction data is used, M/A and CRCT predictive hit rates are comparable; on the good weather day, the CRCT hit rate is higher. Actual sector hit rate is higher for CRCT than for M/A in ZID runs. In ZKC runs, M/A hit rate is slightly higher on the bad weather day with the opposite being the case on the good weather day (Figure 5).

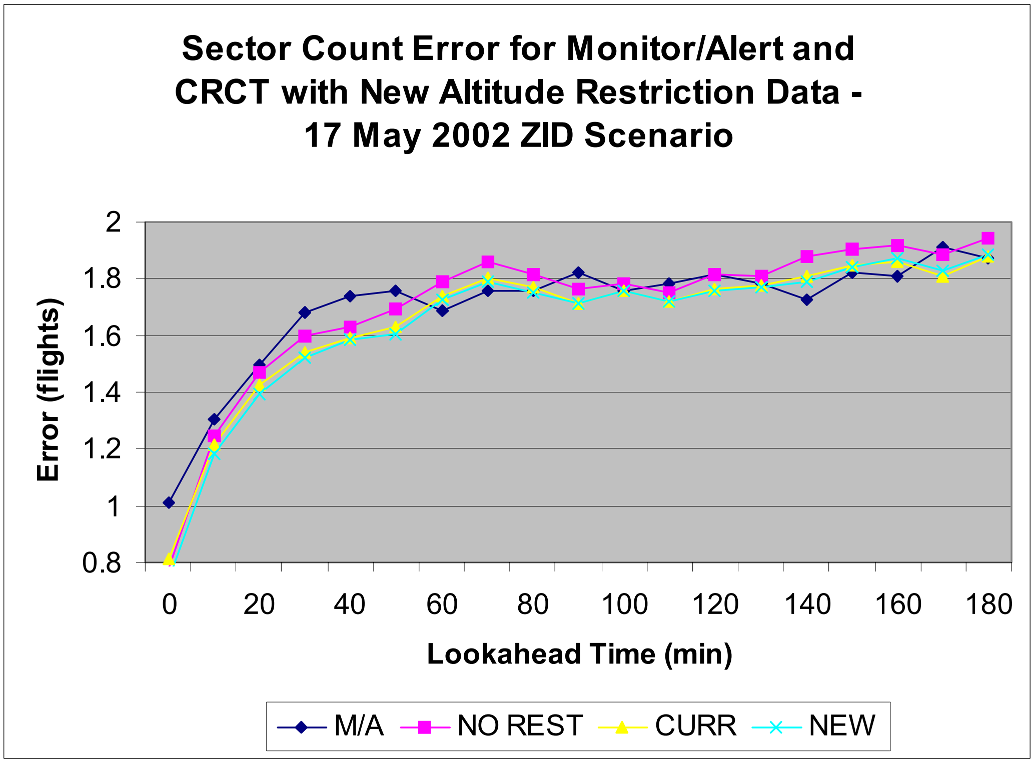

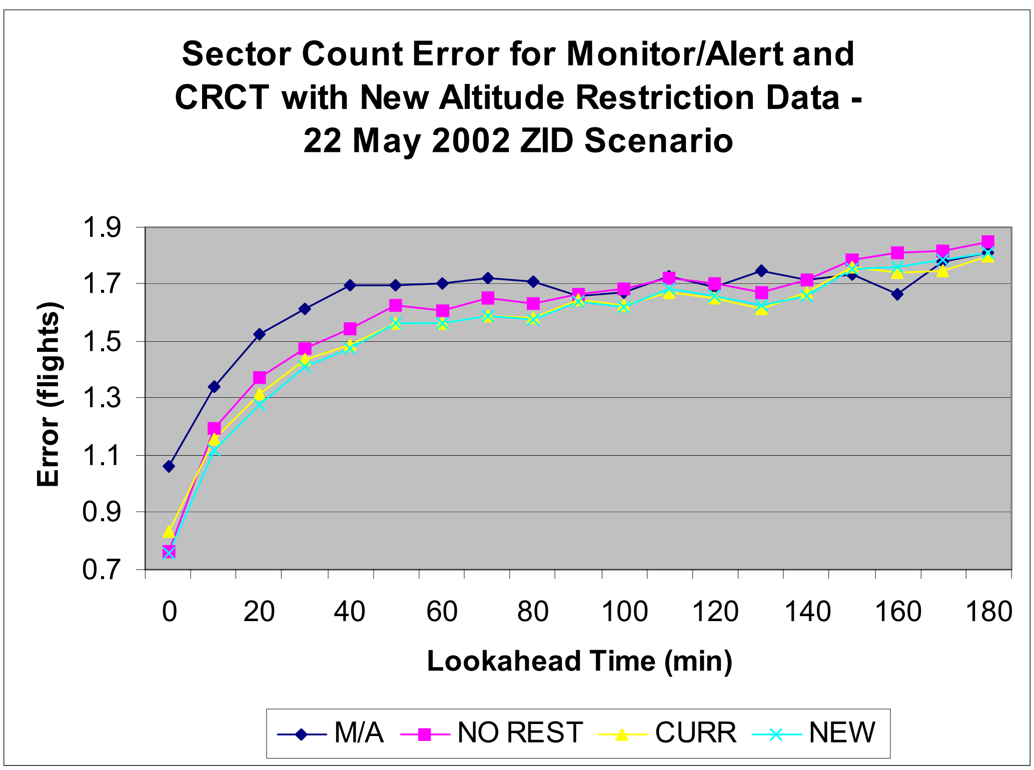

Figure 3. Sector Count Error for ZID Scenarios

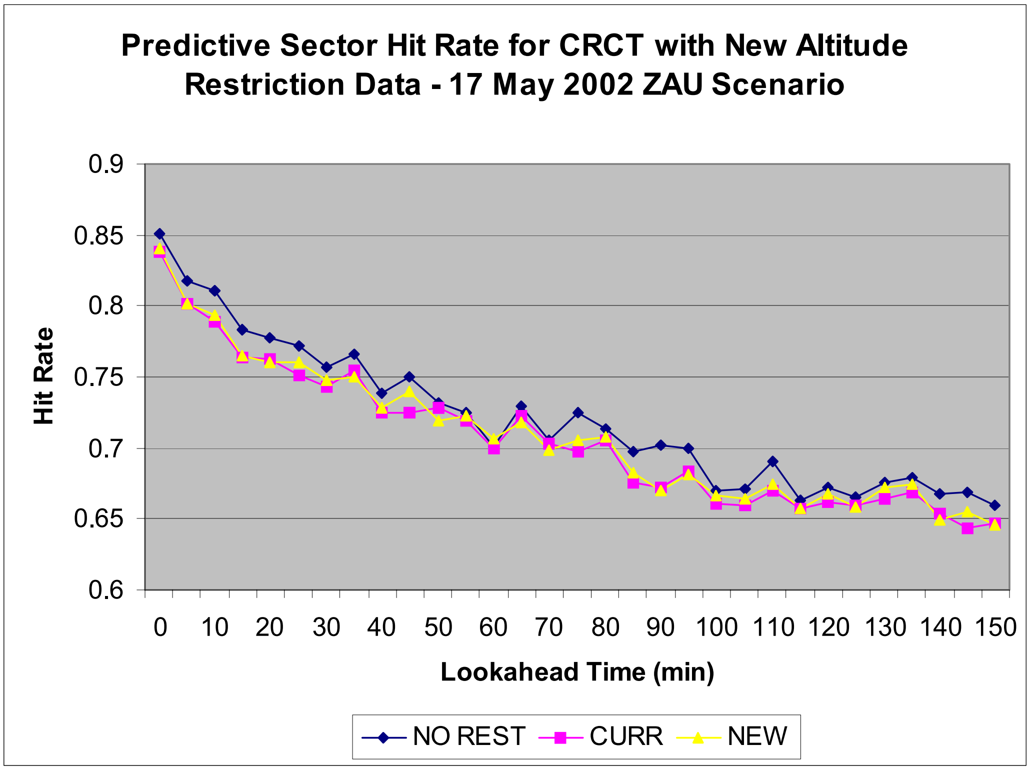

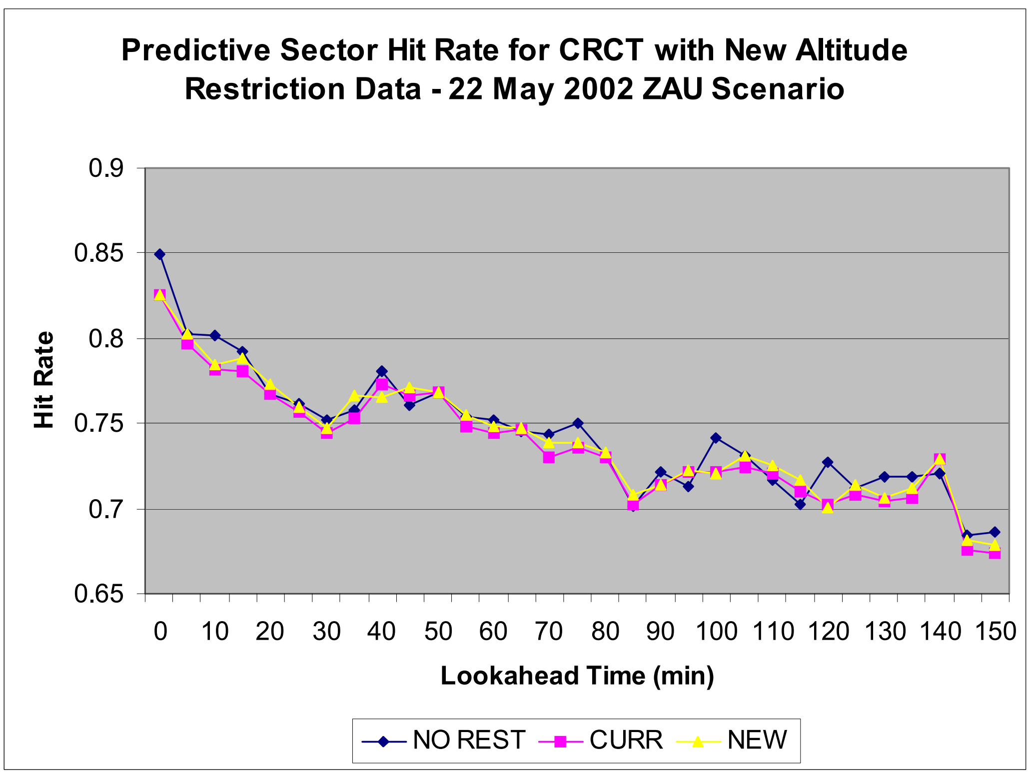

Figure 4. Predictive Hit Rate for ZAU with ETMS Data

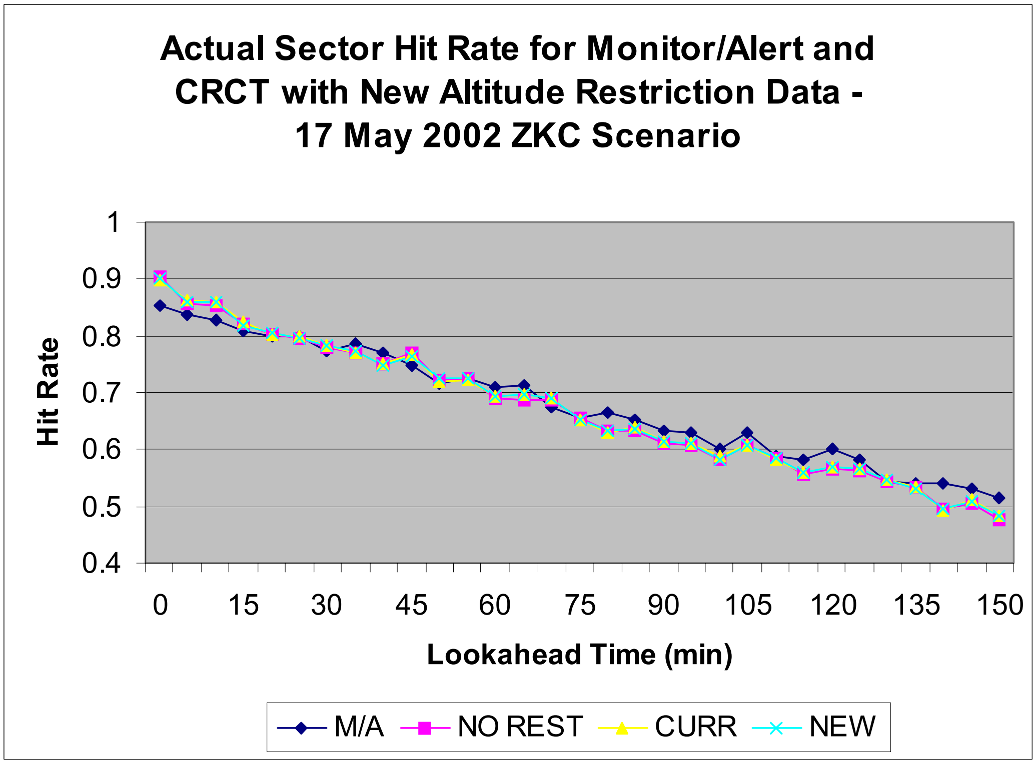

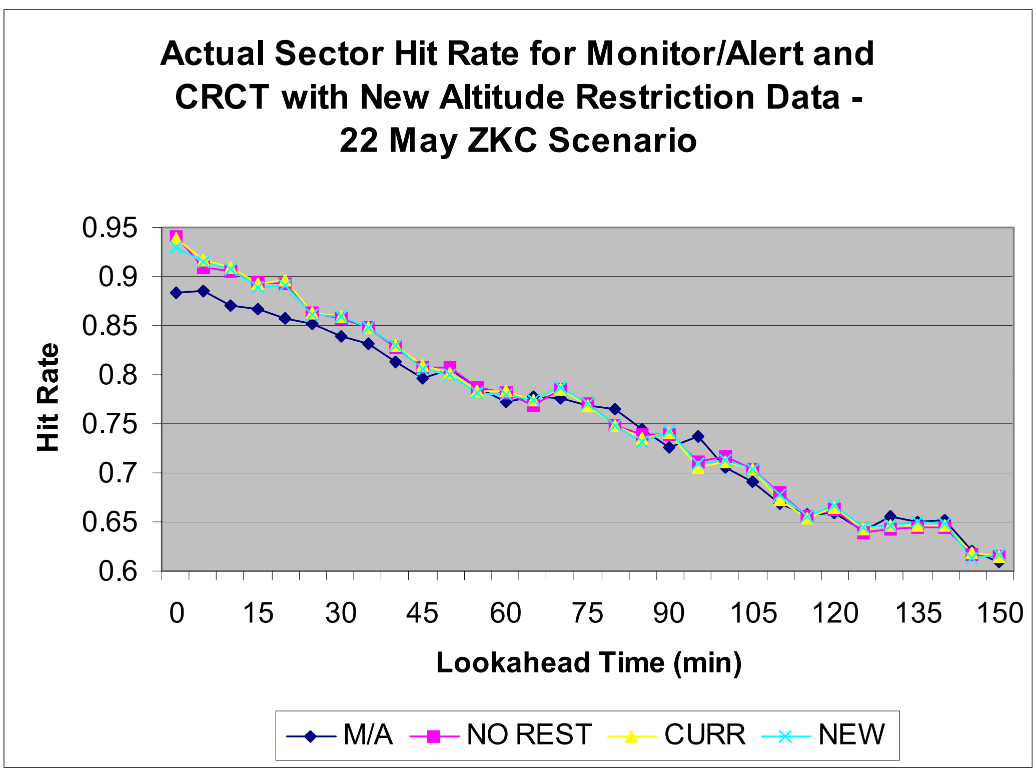

Figure 5. Actual Hit Rate for ZKC Scenarios

<u>Prediction Performance for CRCT with New Aircraft Performance Data</u>

CRCT predictions are based on the aircraft trajectories that it models. Various data sets are used to model a trajectory. One of these, aircraft performance data, characterizes the behavior of various aircraft types during vertical transition. This data specifies information such as climb and descent speeds and gradients. The accuracy of the aircraft performance data that CRCT uses affects the accuracy of the vertical profiles of the trajectories and this in turn affects the accuracy of CRCT predictions.

The current aircraft performance data being used by the ZID, ZKC, and ATCSCC CRCT systems was obtained from corresponding URET prototype data. The URET aircraft performance data was largely crafted from manufacturers' handbooks. As this data was not available from many manufacturers, default performance data was established for jet, turboprop, and piston aircraft for which specific data was not known. Subsequently, track data was analyzed to determine empirical aircraft performance data. This derived data was used to update URET aircraft performance for engine types for which such data was already explicitly available and to establish new data for engine types for which such data was previously unavailable.

The derived aircraft performance data resulted in modest improvements in URET trajectory modeling accuracy. However, unlike the initial URET aircraft performance data, the revised data was never adopted by CRCT. An analysis was performed to determine how the prediction performance of CRCT would be affected by the updated aircraft performance data and to determine if the CRCT data should be updated. Two versions of updated aircraft performance data were obtained from URET. The earlier version is implemented in the URET prototype and in the fielded URET Core Capability Limited Deployment (CCLD) system. The later version updates the earlier one for eventual use by the Problem Analysis and Resolution Ranking (PARR) capability of the URET prototype. The ZID and ZKC scenario data, ETMS data, and RUC data were obtained for two days, one a good weather day and the other a bad weather day. Twelve CRCT runs were made for the various combinations of the two ARTCC scenarios, the two weather days, and the three versions (current and two updates) of aircraft performance data.

On 5 March 2002, the good weather day, the weather was clear in both ZID and ZKC throughout the entire scenario. However, ORD did experience bad weather for a short time and as a result had short ground stop and ground delay programs. Some flights were cancelled following the ground delay program. No other facilities had a ground stop, a ground delay, or a Playbook reroute during the scenario. On 26 March 2002, the bad weather day, a north-south band of severe weather moved across ZID and ZKC during the scenario. There were ground stops at 13 major airports in the eastern and midwestern U.S. There were ground delay programs and cancelled flights at seven major eastern U.S. airports. There was no Playbook reroute during the scenario.

Data was collected during the 12 CRCT runs and was used to calculate sector count error, sector entry error, hit rate, and sector dwell error. Prediction performance results were almost identical for CRCT runs using the URET CCLD and URET PARR aircraft performance data versions. For this reason, discussion of CRCT prediction performance results will not include those for the URET PARR version of the aircraft performance data. Figure 6 illustrates sector count error for the ZKC CRCT runs, Figure 7 illustrates sector entry error for the ZID runs, and Figure 8 illustrates sector dwell error for the ZID runs. Updating the aircraft performance data almost always improved CRCT's prediction performance (Figures 6 – 8). CRCT sector count error was smaller when the URET CCLD data was used than when the current CRCT aircraft performance data was used (Figure 6). CRCT sector count error when the URET CCLD data was used was almost always smaller than the corresponding M/A error (Figure 6). The one exception was at later lookahead times for the CRCT run with the ZKC bad weather scenario.

CRCT sector entry error was smaller than that for M/A (Figure 7). The error was generally smaller when the current aircraft performance data was used than when the URET CCLD data was used. However, the difference between the two sets of errors was relatively small. CRCT predictive and actual sector hit rates were comparable when the two versions of aircraft performance data were used. CRCT predictive hit rate was as good as or better than the corresponding M/A rate for the ZID scenarios. The reverse is true for the ZKC scenarios. CRCT actual hit rate was as good as or better than that for M/A. Finally, sector dwell error was smaller for CRCT runs with the URET CCLD aircraft performance data (Figure 8). Sector dwell error when the current CRCT aircraft performance data was used was in turn smaller than the corresponding M/A error.

Comparison of the prediction performance metrics shows that using updated aircraft performance data usually improves the performance of CRCT. However, it is not clear that the improvement is consistent enough or large enough to make a compelling case for updating the aircraft performance data from the current version being used in CRCT to the URET CCLD (or URET PARR) version. Since metrics for CRCT were usually better than their M/A counterparts, it may be worthwhile to inspect ETMS aircraft performance data and the logic that applies it. It may be possible to improve M/A accuracy by making some logic and/or data changes.

Comparison to Previous Prediction Performance Analyses

FY01 analyses assessed the prediction performance of CRCT with and without the following features: modeling of altitude restrictions, use of ARTCC Host data, and adaptive modeling of departure delays[3]. To increase the level of confidence that can be placed in the previous results, in FY02 the analyses were repeated using additional scenarios. Both the FY01 and FY02 analyses found that modeling altitude restrictions decreased CRCT sector count error for ZID scenarios. Figure 3, the result of an FY02 analysis, illustrates this for the 17 and 22 May 2002 ZID scenarios. Although the FY01 analyses found that modeling restrictions had no marked effect on the error for ZKC scenarios, the FY02 analyses found that CRCT sector count error also decreased for ZKC scenarios when altitude restrictions were modeled. The probable cause for the difference is that, as of FY01, the ZID restriction data had been refined more than its ZKC counterpart and was therefore more accurate. In addition, the ZKC altitude restriction data had been updated between the FY01 and FY02 analyses.

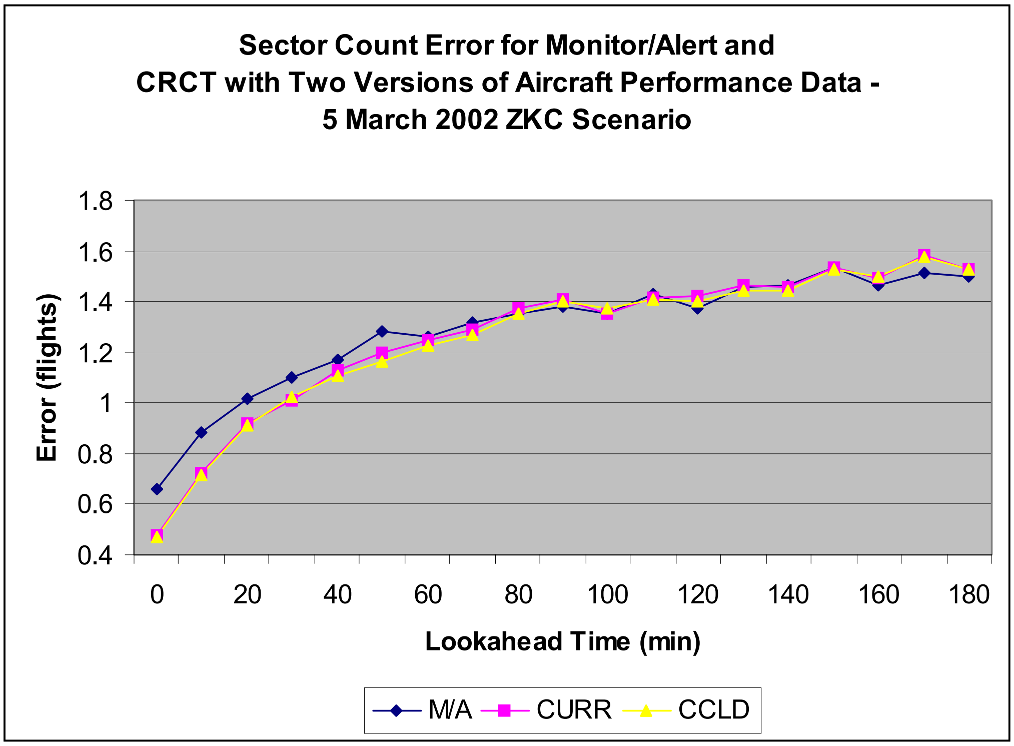

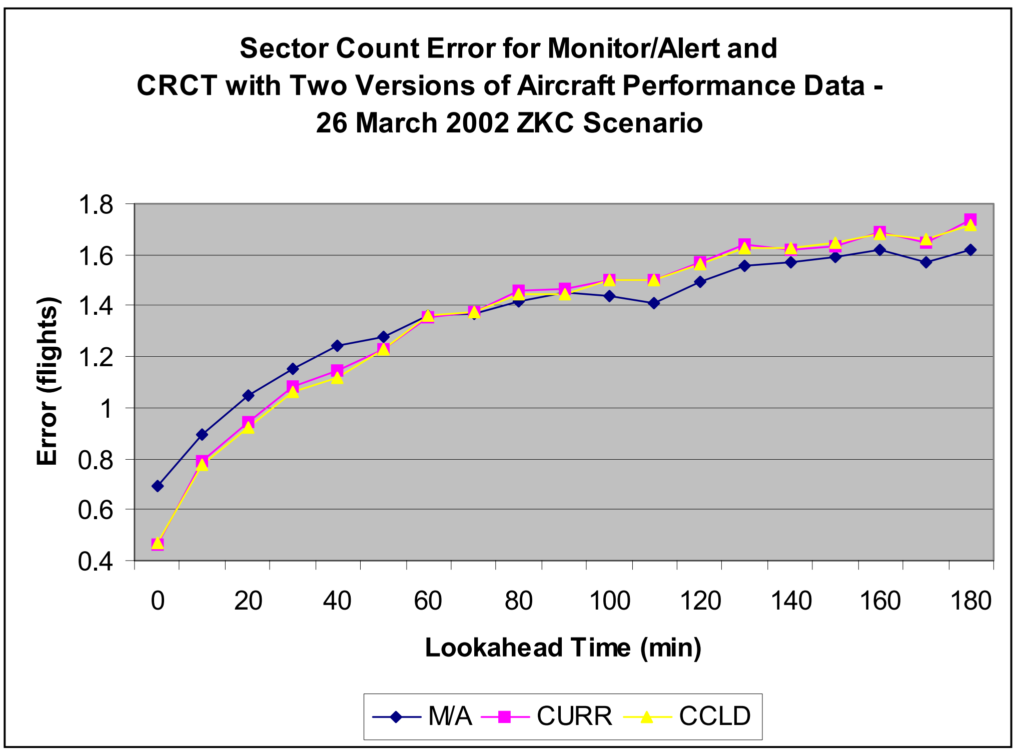

Figure 6. Sector Count Error for ZKC Scenarios

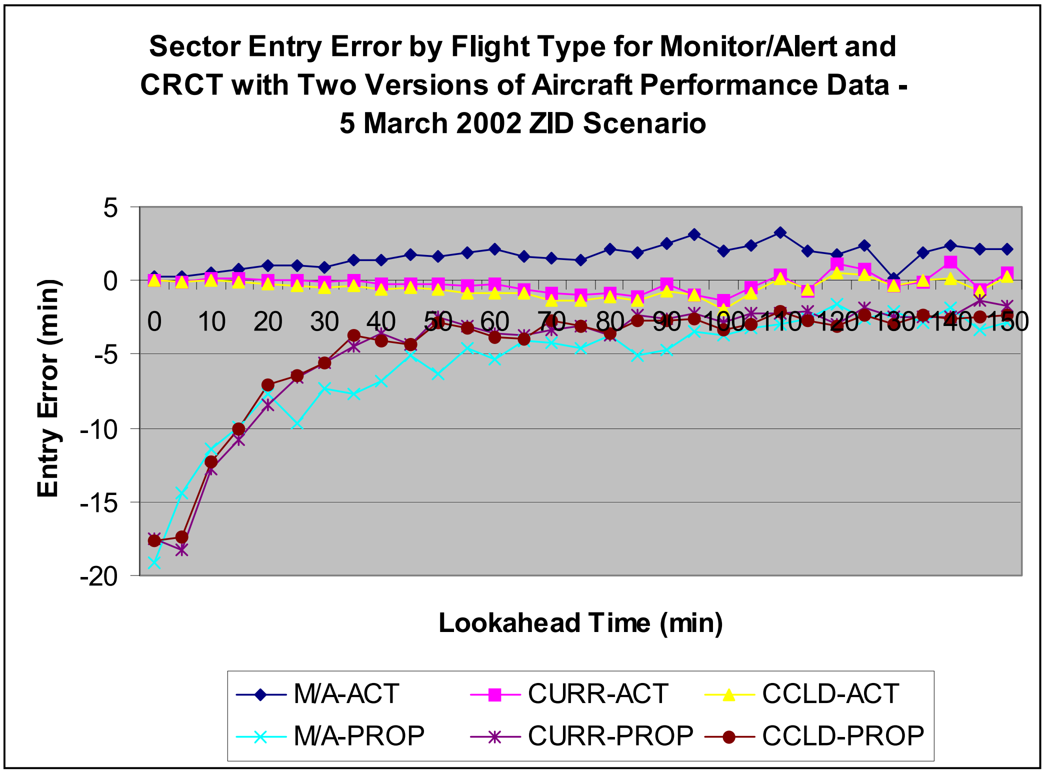

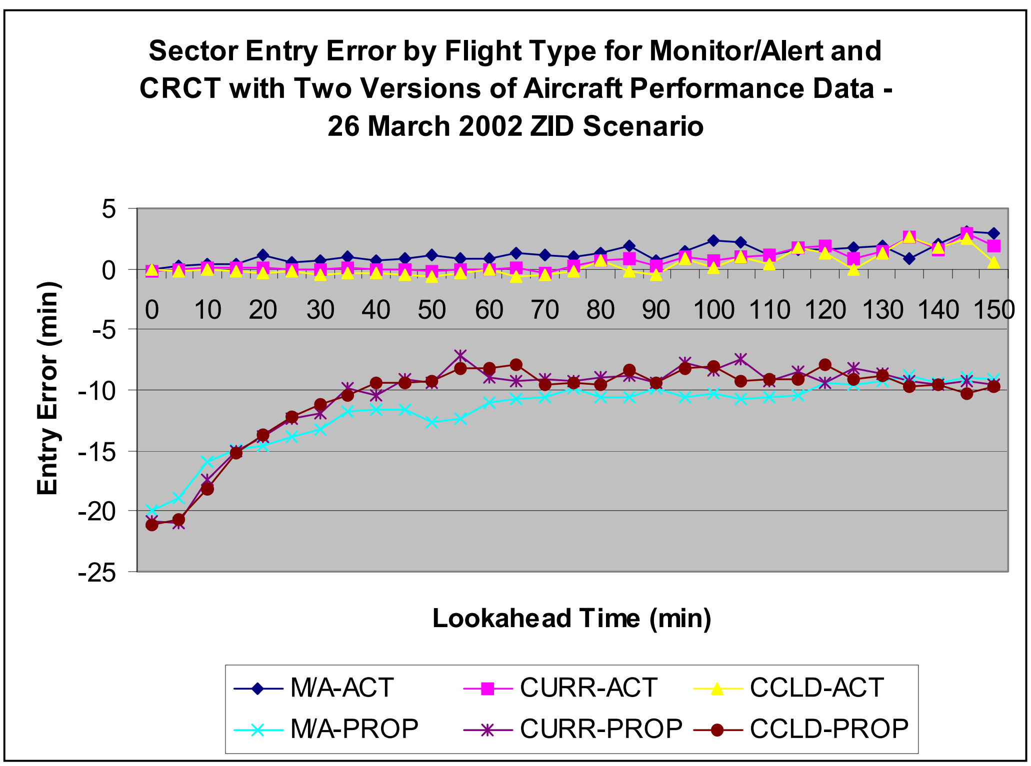

Figure 7. Sector Entry Error for ZID Scenarios

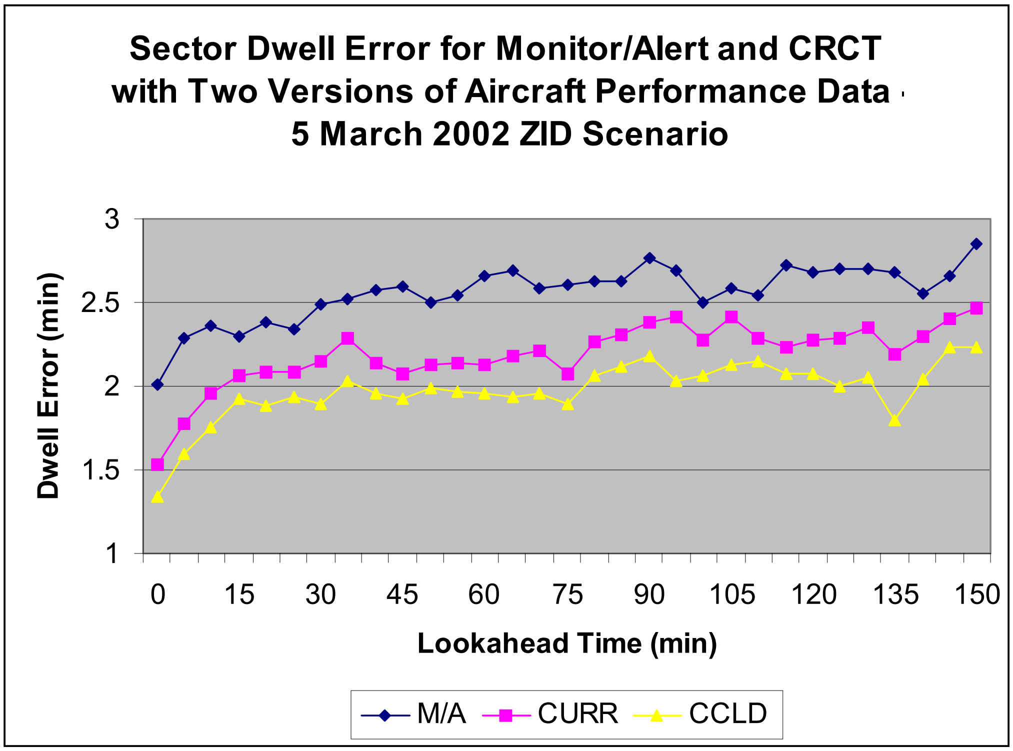

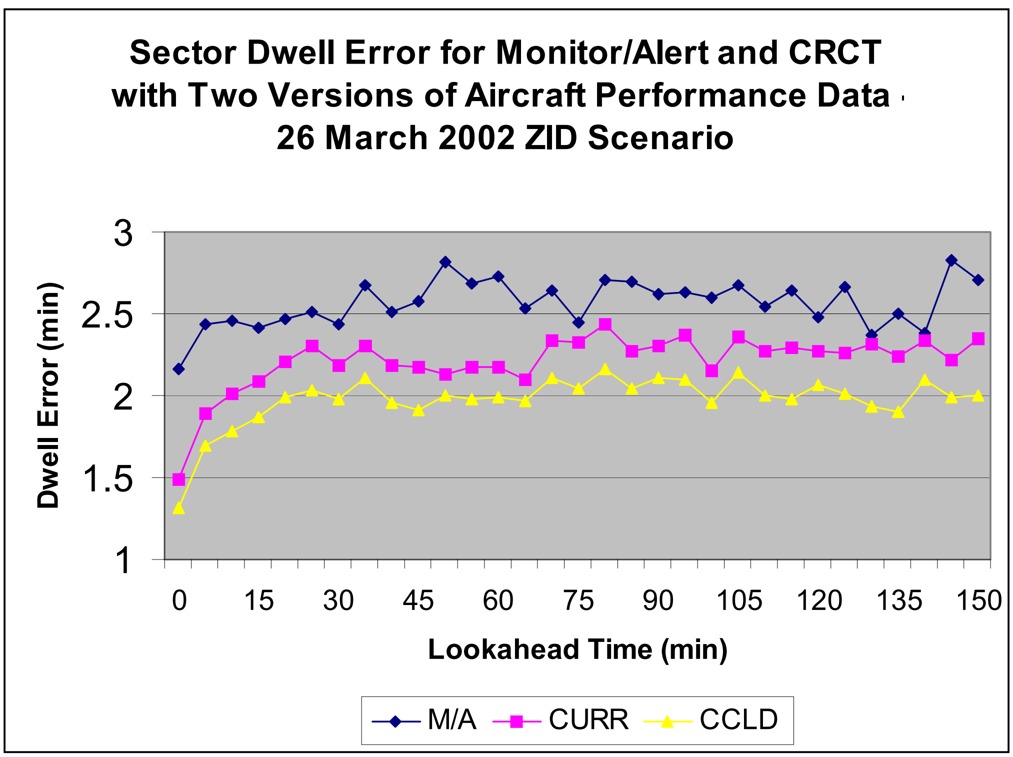

Figure 8. Sector Dwell Error for ZID Scenarios

CRCT optionally uses local ARTCC Host track data, which is updated every 12 seconds, to supplement ETMS track data, which is updated every minute. The more frequent Host position updates can mean earlier detection and remodeling of flights that deviate from their routes. An FY01 analysis found that using ARTCC Host track data provided a small benefit at early lookahead times but virtually no benefit at later times. In fact, using Host data slightly degraded accuracy at some lookahead times. This counterintuitive result was again found by an FY02 analysis, as shown in Figure 9, which illustrates sector count error for CRCT runs against the 17 and 22 May 2002 ZKC scenarios. In the figure, the curve labeled "BASE" represents sector count error when Host data is used and the curve labeled "NO HOST" represents error when no Host data is used. Although the cause of the unexpected result is not fully understood, the result itself suggests that using ARTCC Host track data does not significantly impact CRCT sector count error.

When aircraft are delayed beyond their scheduled departure times, optional CRCT logic can be invoked to determine delayed departure times. These delayed departure times are themselves iteratively delayed, as needed, if they are not met. An FY01 analysis found that this adaptive departure delay modeling improved CRCT predictions at lookahead times less than an hour. A similar FY02 analysis obtained the same result. This is illustrated again in Figure 9. In the figure, the curve labeled "BASE" represents sector count error when delay modeling is used and the curve labeled "NO ADDM" represents error when no delay modeling is used.

FY01 analyses found that the prediction performance of CRCT is better than that of M/A for ZID. CRCT and M/A have similar prediction performance for ZKC. This was seen in the metrics sector count error, sector entry error, and predictive hit rate. FY02 analyses found that the prediction performance of CRCT is generally better than that of M/A in both ZID and ZKC. Figure 9 illustrates FY02 analyses sector count errors for M/A and CRCT (curve labeled "BASE") runs against two ZKC scenarios.

Conclusions and Recommendations

The following conclusions and recommendations are based on the findings of various CRCT prediction performance analyses. Two sets of analyses were performed, one in FY01 and another in FY02. Results for the two sets of analyses are largely consistent and support the conclusions and recommendations.

- Although updating altitude restriction data did not significantly improve CRCT accuracy, using restriction data (current or new) resulted in better prediction performance than not using any restriction data. This suggests that modeling restrictions may improve M/A accuracy if it were implemented in ETMS.
- The new (URET CCLD or URET PARR) aircraft performance data improves CRCT accuracy. Further research is recommended to determine if and how the new performance data could be used by ETMS and if ETMS use of this data will improve M/A predictive accuracy.

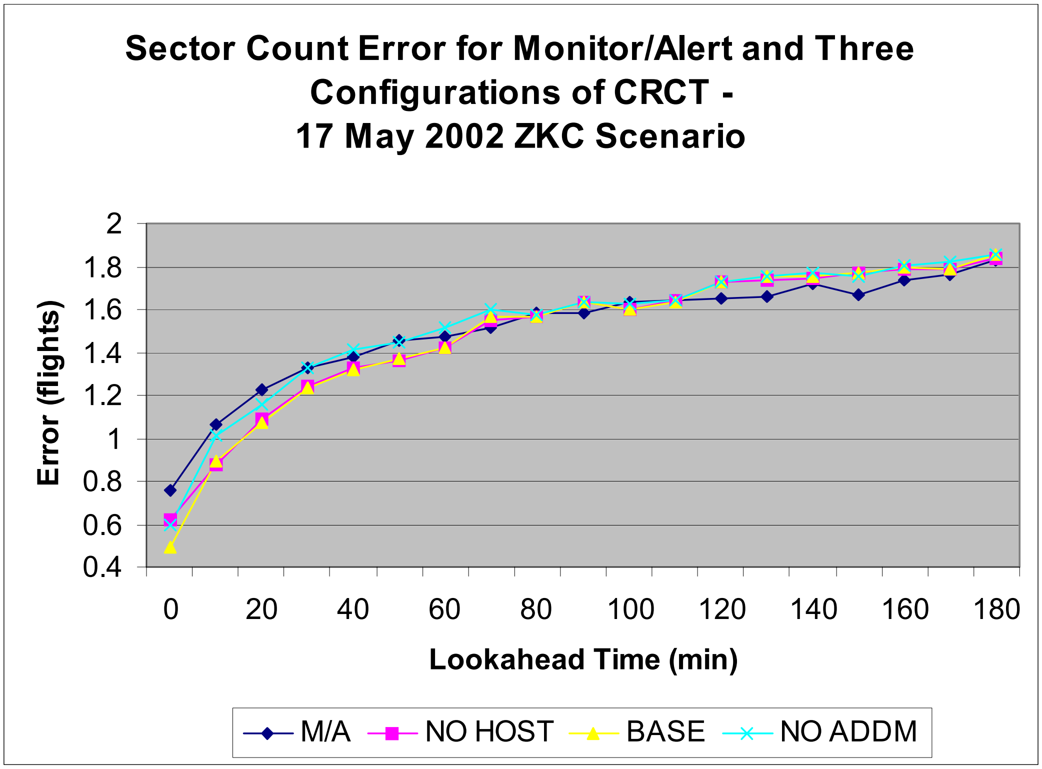

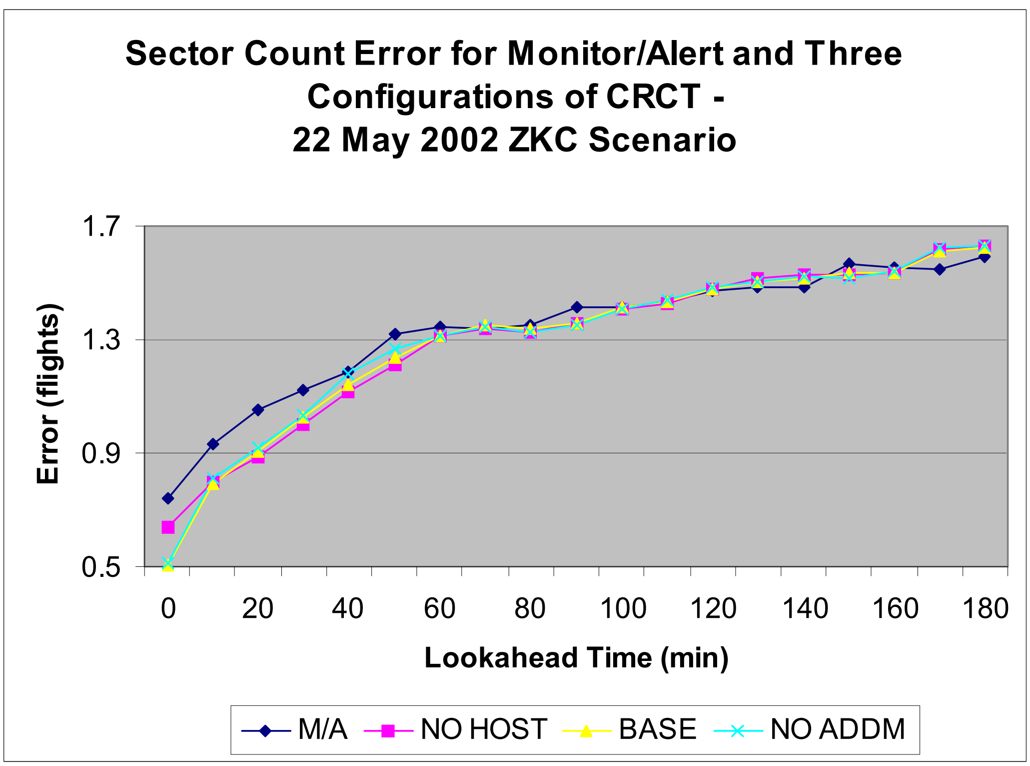

Figure 9. Sector Count Error for ZKC Scenarios

- Use of ARTCC Host track data (12 second updates) to supplement ETMS track data (1-minute updates) does not appreciably improve CRCT's prediction performance. In some cases, Host data actually degrades accuracy. The cause of this is not understood. Nevertheless, Host data is beneficial for other reasons (e.g., sectorization messages) and its continued use is recommended.

- To be consistent with ETMS M/A, Host sectorization messages are not used in CRCT prediction performance assessments. To better understand their impact on CRCT accuracy, it is recommended that additional assessments be made in which the sectorization messages are used.

- Adaptive departure delay modeling improves CRCT's prediction performance at lookahead times up to one hour. Further research is needed to determine how CRCT departure delay modeling compares to that for ETMS and if adoption of the CRCT approach in ETMS would improve M/A accuracy.

- On the whole, CRCT prediction performance appears to be better than that for M/A in both ZID and ZKC. There are cases; however, where M/A accuracy is better than CRCT accuracy, depending on metric, scenario, and ARTCC.

- Future prediction performance analyses are recommended to further validate FY01 and FY02 results and to investigate results that are not clearly understood. For example prediction performance analyses varying aircraft performance data found that CRCT predictive hit rate is higher than M/A in ZID with the opposite being true in ZKC.

References

1. Beaton, E. K., L. S. Rhodes, L. R. Rhodes, and R. Robbins, *CRCT Capabilities Detailed Functional Description,* MTR 00W302R1, The MITRE Corporation, McLean, VA, November 2002.
2. Holly, K., Y. Cabeza, M. Callaham, D. Greenbaum, A. Masalonis, and C. Wanke, *Feasibility of Using Air Traffic Complexity Metrics for TFM Decision Support*, MTR 02W55, The MITRE Corporation, McLean, VA, August 2002.
3. Masalonis, A. J., C. R. Wanke, D. P. Greenbaum, C. K. Jackson, E. K. Beaton, S. M. Mahmassani, D. Chaloux, E. A. Cherniavsky, and J. Forren, *Preliminary Analysis of CRCT Traffic Prediction Performance*, MP 01W62, The MITRE Corporation, McLean, VA, March 2001.
4. Brudnicki, D. J., K. S. Lindsay, and A. L. McFarland, "Assessment of Field Trials, Algorithmic Performance, and Benefits of the User Request Evaluation Tool (URET) Conflict Probe," *Proceedings of the 16th AIAA/IEEE Digital Avionics Systems Conference*, October 26-30, 1997.
5. Lindsay, K. S. and N. E. Rozen, Functional Performance of the User Request Evaluation Tool (URET) Delivery 3.4 Release 2, MTR 00W115, The MITRE Corporation, McLean, VA, December 2000.
6. Paglione, M. M., H. F. Ryan, R. D. Oaks, J. S. Summerill, and M. L. Cale, *Trajectory Prediction Accuracy Report: User Request Evaluation Tool (URET)/Center-TRACON*

Automation System (CTAS), DOT/FAA/CT-TN99/10, Federal Aviation Administration, Atlantic City, NJ, May 1999.

7. Cale, M. L., S. Liu, R. D. Oaks, M. Paglione, H. F. Ryan, and S. Summerill, "A Generic Sampling Technique for Measuring Aircraft Trajectory Prediction Accuracy," *4th USA/Europe Air Traffic Management R&D Seminar*, December 3-7, 2001.
8. Green, S. M. and R. Vivona, "Field Evaluation of Descent Advisor Trajectory Prediction Accuracy" AIAA 96-3764, *AIAA Guidance, Navigation, and Control Conference*, July 1996.
9. Lindsay, K. S., *Predictive Accuracy Analyses for the Collaborative Rerouting Coordination Tools (CRCT)*, MTR 02W62, The MITRE Corporation, McLean, VA, September 2002.