# LIMITED DYNAMIC RE-ROUTING (LDRR)

*Carolyn Sorensen, ISA Software, Paris, France*

*Diana Liang, Office of System Architecture and Investment Analysis, ASD 100*

*Ian Crook, ISA Software, Paris, France*

## Abstract

The goal of this research is to assess the impact of the Limited Dynamic Re-Routing (LDRR) concept on the service provider (ATC), on overall NAS safety, and on the user by examining collaborative decision-making between ATC and the airline operations center (AOC) via the flight deck, assuming today's equipage for dynamically re-routing the aircraft. In planning a more optimal flight plan, the AOC considers the current route and delays along with the known SUA and weather situation; however, the AOC has no knowledge of the current traffic situation or unanticipated weather movement or SUA activity. Once a re-route request has been made, ATC will attempt to grant the request unless the modification results in a subsequent safety or separation violation.

In the end, based upon the individual circumstances for each flight, some flights benefited from LDRR while others showed little difference or a negative effect from re-routing with limited awareness. Indicators for the flights performing worst under optimization suggest that more foreknowledge of future impacts, i.e. better distribution of important ATM and weather data to all participants, would allow these flights to follow more optimal paths,

## The Concept

In the near future, limited dynamic re-routing will involve collaborative decision-making between ATC and the airline operations center (AOC) via the flight deck. The AOC will be able to re-plan flights as they progress through the airspace system to make use of a more optimized route in response to unforeseen situations, or to avoid congested areas.

The goal of this research was to assess the impact of the Limited Dynamic Re-routing concept on the service provider (ATC), on the NAS, and on the individual user. The focus was to explore the early stages of collaborative decision-making, where no additional equipage is anticipated in the cockpit.

At this stage of the concept the AOC, in collaboration with the service provider and the cockpit, can dynamically modify flight plans mid-route in response to unforeseen penalties that have induced undesirable delay. In calculating a more optimal flight plan, the AOC considers the current delay that has been induced on the flight, along with the known weather and SUA situation. However, the AOC has no knowledge of the current tactical traffic situation, unpredictable weather movement or any unanticipated SUA activity.

Although the basis for allowing the operator to review on-going flight plans is to allow them to optimize according to their own business objectives, in this analysis the optimization criteria were limited to improving fuel efficiency and reducing delay. The adapted flight plan was transmitted to the cockpit, which subsequently requested validation and acceptance from ATC. It was also assumed that no re-planning is permitted in the proximity of the origin or destination airport, where flights are subject to more specialized control rules and thus not available for dynamic re-planning.

Where possible, the ATC service provider responsible for the flight at the time a modification is requested attempted to grant the modified plan. Refusal of the request only occurred in cases where the requested modification resulted in an imminent separation violation. On the basis of the results, an initial assessment of the effectiveness and future potential of the concept and its impact on the different participants has been made.

A secondary, but equally important aspect of this research is to serve as a case study for the expansion of this prototype and its application in a large-scale, high fidelity simulation study of the Limited Dynamic Re-Routing (LDRR) proposed operational concept.

### Assumptions

The following key assumptions formed the basis for the behavior modeled in this scenario:

1. *Before departure, the AOC files a flight plan based on existing route structures and optimized, to the extent possible, with respect to expected adverse weather areas, known SUA activity and projected wind patterns.*

2. *The AOC/pilot will request a new route whenever the flight delay varies from their optimal projection by a certain threshold. This is most likely to occur when unanticipated external events take place, for example when a conflict avoidance maneuver is required, an unanticipated SUA closure occurs, or a weather front does not behave as predicted.*

3. *The AOC has limited situational awareness of unanticipated airspace restrictions and no knowledge of local traffic, but is normally aware of approaching weather patterns. Therefore, it was assumed that an optimized path presented by the AOC via the pilot would consider some, but not all, of the areas that should be avoided.*

4. *Requests for a new route will not be permitted within a given range of the origin and destination airports where a flight is considered to be under specific control rules.*

5. *A controller will attempt to grant the re-routing request, and will usually do so unless other traffic is impacted in a short-term conflict situation or the new route violates active procedural restrictions.*

## Scenario Development

ETMS flight plan data for a fair-weather, moderately heavy traffic day was used as the traffic sample. All flights traversing any of 4 central ARTCCs were extracted for the study. The 4 ARTCCs selected were Kansas City (ZKC), Denver (ZDV), Minneapolis (ZMP), and Cleveland (ZOB). This traffic sample captured the majority of the long-distance coast-to-coast flights as well as local and regional flights within the central core of the NAS. Flights to or from Canada or Mexico were kept, but all other international flights were dropped. RAMS Plus™ modeled the flight climb and descent profiles using a standard set of aircraft characteristics, consistent with those supported by the OPGEN model. The final Baseline sample included 15,423 flights.

All flights were modeled from departure airport to arrival airport, but since the study objective was to examine the enroute impacts of LDRR, RAMS Plus™ sectorization was configured with high altitude sector information for the four central ARTCCs only. Dummy high altitude sectors for the remaining areas of the CONUS were included for optimization purposes, but were not measured.

RAMS Plus™ was also provided with a set of controller taskload activities and event triggers which permit the comparison of controller taskload in both scenarios. New triggers were added to record pilot reroute requests (as returned from OPGEN), additional conflict searches necessitated by these requests, flight plan amendment tasks for accepted reroutes, and additional communication tasks.

## Modeling Approach

No existing standalone fast-time simulation model can represent all of the desired behavior described above nor capture the required measurements to support the analysis. The approach, therefore, was to take advantage of the distributed modeling approach that resulted from the RAMS-OPGEN Dynamic Link [1] project. One of the reasons this approach to simulation modeling has been developed is to cater for situations where no single model is available to study a given concept, however multiple tools and models exist, each providing unique functionality that collectively can represent the concept. The remainder of this section briefly describes the models used and their primary roles.

RAMS Plus™ served as the primary simulation tool. It was used to model the NAS and specifically the roles of the Air Traffic Controller, Special Use Area activity management, adverse weather management and dynamics and hence the source of the delay to traffic.

The Decision Tool (DT) served as the model of the AOC/pilot's behavior in choosing whether or not to request flight plan optimization. DT is a decision-making tool that accepts dynamic information and assesses a flight's current condition based upon pre-defined rules.

The FAA Optimized Trajectory Generator (OPGEN) was used to model the AOC trajectory planning tool. OPGEN accepts a flight's current 4D flight plan and calculates a new trajectory from the current time forward that is optimized according to defined criteria – in this case, optimizing for fuel burn first and then for desired arrival time. OPGEN creates flight trajectories that avoid known active restricted zones but it does not have any knowledge of other traffic nor of unanticipated restrictions due to unscheduled SUA activity or unexpected changes in weather patterns.

The ATMOS (Atmospheric) Weather Model is a distributed weather server that can be queried for wind information (speed and direction) at a given 3D location and time. RAMS applies the provided wind information to each flight profile to produce a wind-effected flight trajectory.

Enroute delays were introduced into the scenario using a series of airspace restricted zones. Some of these restricted zones were configured to be known to both controller and AOC models (RAMS Plus and OPGEN), and some areas were are known to the controller model only. The use of restricted zones and conflict avoidance maneuvers created sufficient delay in the model.

Since enroute delay can occur at varying points during a flight, it was decided that RAMS Plus would consider flight plan optimization at every sector crossing. At each crossing, RAMS Plus passed flight information to the Decision Tool, which was configured to recommend optimization if the current flight delay was 5 minutes or greater, if the flight was at least 200 nm from its origin or destination, and if the flight was at or above 24,000 feet. If the Decision Tool recommended optimization, flight data was then passed to OPGEN which returned a new partial flight profile from the current position to the destination. As the controller model, RAMS Plus then performed a conflict detection on the new profile and rejected it if a conflict within the current sector was found. If no conflict was found, it then checked for any additional restricted zones, representing unanticipated SUA or weather activity which the optimization model does not know about and which the newly optimized profile may cross. If SUA or weather activity was found, RAMS Plus inserted the necessary avoidance routes into the new profile and rejected the new profile if its final arrival time was more than 15 minutes later than the current profile's arrival time.

# Results

Results analysis focused on the following areas:

1. Overall effectiveness of the strategy, in terms of fuel usage and delays to traffic?
2. Service provider and safety impacts such as sector loading, controller taskload and the number and nature of conflicts?
3. What are the specific benefits to flights that have been allowed to be re-optimized, given the limited amount of information available to them?
4. What is the impact of allowing a subset of flights in the NAS to re-optimize on the remaining NAS users?
5. How effective is a partial solution where re-routing s based on an incomplete picture of the current (and future) situation?

## *Measures*

In order to consider the impacts of re-routing, several performance indicators and measures were examined, including:

**Fuelburn**, **Duration, Distance.** For each scenario, comparisons were made between initial (flight plan as filed, with calculated time and altitude at each point) and final (as flown) results. Fuelburn was calculated using an external fuelburn model. Total flight duration (in minutes) and distance traveled (in nautical miles) were also considered.

**Flight Delay.** Current flight delay in minutes was used during the re-routing scenario as a key determinant for optimization behavior. It is defined as the difference between the initial expected arrival time and the final arrival time.

**Conflicts**. The number of flights involved in conflicts, the number of conflicts, and the changes in the geometries of these conflicts are considered.

**Taskload**. Typical controller tasks and newly defined tasks for optimization were recorded. Cockpit/pilot taskload was not measured.

All flight-related metrics were calculated using the entire flight length, while sector- and controller-related metrics were based on the four ARTCCs only.

## *Overall LDRR Results*

As shown in Table 1, a slight deterioration in each flight performance measure was noted when the results for all flights in both scenarios were compared. A total of 39.56 hours of additional delay was recorded in the LDRR Scenario.

**Table 1. Total Flight Performance Results**

| **Measure (N=15,423)** | **Average % Change Baseline to LDRR** |
|---|---|
| Fuelburn | + 0.27 % |
| Flt Distance | + 0.09 % |
| Flight Delay | + 0.14 % |
| Avg Flt Delay | + 0.15 minutes |

Table 2 shows that while both total conflicts and the number of conflicts occurring within the 4 central ARTCCs were reduced in the LDRR scenario, the controller task count for these ARTCCs increased and there was a slight increase in conflicts occurring outside the 4 ARTCCs.

**Table 2. Service Proider Impacts: Conflicts and Taskload**

| **Measure (N=15,423)** | **Avg % Change** |
|---|---|
| Total No. Conflicts | - 0.6 % |
| No. Conflicts, Within 4 ARTCCs Only | - 1.9 % |
| No. Conflicts, Outside 4 ARTCCs | + 0.4 % |
| Total ATC Tasks, 4 ARTCCs Only | + 2.2 % |

A comparison of 15 minute instantaneous traffic counts showed a decrease (ranging from –1.4% to –4.7%) for each of the 4 measured ARTCCs.

These overall results present somewhat unexpected and apparently contradictory results and lead to further questions:

- *Why was there was no apparent improvement for traffic under LDRR?*
- *Why is there increased delay for users, but decreased flight counts and conflicts for the service provider?*

The remainder of this document will consider the possible influences, such as:

**Optimization Criteria**. Were the rules too stringent? Were too many optimization requests refused?

**Non-optimized traffic**. Were flights that were not optimized somehow penalized by LDRR?

**Optimized traffic**. Was the optimization process successful? Is it possible to optimize effectively with limited knowledge?

**Service provider**. Why did the 4 ARTCCs experience fewer simultaneous flights and conflicts, while the users experienced increased delays?

## *Optimization Criteria*

In order to be a eligible for dynamic re-routing in this scenario, a flight was required to meet the following criteria during the simulation:

- Total flight distance greater than 400 nm
- Cruise altitude of FL 240 or greater
- Five minutes or more of enroute delay experienced

Results showed that 11.1% (1710) of the flights in the baseline sample met these criteria during the simulation, and 92.6% of these eligible flights did request optimization at least once. (The remaining flights which met the criteria did not reach 5 minutes of delay until they were too close to their destination to permit rerouting.) Looking further, we also noted that of the 1584 flights that requested at least one optimization, 1448 had at least one optimization accepted and applied – representing 84.7% of all eligible flights.

Since each flight may have requested optimization on several occasions, we also considered the overall optimization acceptance rate. Overall, 72% of all requests for optimization were accepted and applied within the LDRR scenario. Of the remainder, only 2.3% were refused due to detection of a local conflict in the immediate area. In these cases, the appropriate controller tasks were recorded and the requesting flights remained on their current trajectory until the next sector boundary was reached. At that stage, the request process would be repeated provided the optimization criteria were still satisfied.

Finally, 25.7% of all optimization requests were rejected before the request was passed to the controller model. In some cases the optimization model could not optimize or chose to return the original trajectory. In other cases, it returned a fuel-optimized trajectory with a final arrival time more than 15 minutes later than the flight's current expected arrival time, or returned a trajectory containing excessive altitude changes (partial descents and climbs) prior to the final descent. These attempts were treated as AOC or cockpit decisions against requesting optimization at the current time, and were not passed to the controller model for further consideration nor counted in any of the workload or other measures.

Overall, these results show that most flights eligible for optimization did request optimization at least once, and that a high proportion (84.7%) of these flights did have at least one optimization accepted and applied within the simulation. And, while the optimization criteria and acceptance rules could have been made less – or more – stringent, there is no indicator within these results which might explain the increased traffic delays.

## *Impact on Non-Optimized Traffic*

The only difference between the Baseline and LDRR scenarios was the ability of some flights to request optimization. Therefore, we can consider any change in the results for flights which were not optimized to be a side-effect of LDRR.

To assess the magnitude of any such consequences, an analysis of the flight-based measures for flights that were not rerouted was performed. Table 3 presents the fuelburn, delay and distance impacts for the set of 13,975 flights that were not optimized in any way during the LDRR scenario. The results show a consistent but very slight improvement in each of these measures.

**Table 3. Flight Performance Results for Non-Optimized Flights**

| **Measure (N=13,975)** | **Avg % Change** |
|---|---|
| Fuelburn | - 0.02 % |
| Flt Distance | - 0.19 % |
| Flight Delay | - 0.04% |
| Avg Flt Delay | - 0.04 minutes |

An examination of conflicts and avoidance maneuvers showed a slight decrease in the number of conflicts (-0.52%), in the percentage of these flights involved in conflicts (-0.2%), and in the number of SUA avoidance maneuvers (-1%) for these flights. and provides the most likely explanation for the reductions shown in Table 3. These results clearly show that the optimization of a small subset of flights in the sample did not have a detrimental effect on other traffic. Furthermore, since there is a slight *decrease* in delay in this non-optimized set of flights, the *increase* in delay which is present in the total sample (Table 1) must be attributed to the set of flights which were optimized.

## *Impact on Optimized Flights*

Tables 4 and 5 present results for the set of 1,448 flights which were optimized at least once during the LDRR simulation. As shown in these tables, these flights were penalized by optimization and performed worse than average in fuelburn, delay and distance measures

**Table 4. Flight Performance Results for Optimized Flights**

| **Measure (N=1,448)** | **Avg % Change** |
|---|---|
| Fuelburn | + 1.47 % |
| Flt Distance | + 0.74 % |
| Flight Delay | + 1.04 % |
| Avg Flt Delay | + 2.0 minutes |

**Table 5. Conflicts and Avoidance for Optimized Flights**

| Measure (N=1,448) | Avg % Change |
|---|---|
| No. Conflicts | - 0.67 % |
| % of Flts in Conflict | - 0.4 % |
| No. SUA Avoidance Maneuvers | + 25.5 % |

These flights did show the highest delays in the baseline simulation, which was as expected given that the flights were required to have at least 5 minutes of delay before they qualified for optimization. It was somewhat surprising, however, that after optimization these flights had an average of 2 *additional* minutes of delay; that fuelburn and distance results were also worse, there were significantly more SUA avoidances; and yet, the number of conflicts was slightly lower. It was necessary to look at the optimization results in more detail to understand the results.

An examination of results for individual flights showed that some flights did benefit from optimization, showing considerably reduced delay in the LDRR scenario; while others showed little change or ended up penalized more than in the baseline case. No observable characteristics could be identified to explain why some flights improved and some worsened. It was necessary to examine separately the successfully and unsuccessfully optimized flights in order to judge the extent of the impacts and establish why the results varied from flight to flight.

Fuelburn, delay and distance measures were used to segregate the optimized flights into three subsets. A review of the percent change from baseline in these 3 measures revealed that many flights showed near zero change in one or more measures, and in many of these cases the change in one measure was slightly negative while another measure showed slightly positive. Because of this, only the results showing a change of at least +/- 0.25% (one-quarter of one percent) from Baseline to LDRR were considered "changed" in either direction (better or worse).

Flights showing an improvement of 0.25% or more in at least two of the three measures were considered "Successful" optimizations. Flights showing a deterioration in at least two measures were considered "Unsuccessful", and the remainder were placed in a "No Significant Change" group. Tables 6-8 present the flight measures results for these three subset groups, and Table 9 summarizes several additional flight and optimization characteristics.

**Table 6. Successfully Optimized Flights**

| Measure (N=261) | Avg % Change |
|---|---|
| Fuelburn | - 2.2 % |
| Distance | - 1.2 % |
| Delay | - 2.1 % |
| Avg Delay Per Flight | - 4.0 minutes |

**Table 7. No Significant Change Optimized Flights**

| Measure (N=442) | Avg % Change |
|---|---|
| Fuelburn | + 0.1 % |
| Distance | 0.0 % |
| Delay | 0.0 % |
| Avg Delay Per Flight | 0 minutes |

**Table 8. Unsuccessfully Optimized Flights**

| Measure (N=261) | Avg % Change |
|---|---|
| Fuelburn | + 2.9 % |
| Distance | + 1.7 % |
| Delay | + 2.5 % |
| Avg Delay Per Flight | + 5.1 minutes |

**Successfully optimized flights** performed worse in the Baseline, on average, than any other group. Baseline average delay was 7.4 minutes higher here than the average for unsuccessfully optimized flights, and a higher standard deviation of delay (11.1, in contrast to 6.8 for the unsuccessful group) points to a wider dispersion of final delay values for these flights.

Thus this group would appear to include many of the worst performing flights and in fact, the 75% quartile for baseline delay is significantly higher for this group (22.4 minutes) than for any other (13.2 for All, 11.5

**Table 9. Performance Indicators for Optimized Flight Subsets: LDRR Scenario**

| Measure | Successfully Optimized Flights | No Significant Change Flights | Unsuccessfully Optimized Flights |
|---|---|---|---|
| Median Initial Flight Distance (nm) | 1,331 | 1,177 | 1,457 |
| Average No. Flight Impacts (SUA and Conflict Avoidance) | 3.0 | 2.3 | 3.0 |
| % Change from Baseline: Total No. of SUA Avoidance Maneuvers | -1.7% | +3.1% | +51.0% |
| Average No. of Accepted Optimization Requests | 1.86 | 1.67 | 1.92 |
| Average No. of Refused Optimization Requests | .31 | .36 | .72 |
| % of Flights Encountering 1+ Impacts After Last Accepted Optimization | 22.6% | 13.8% | 29.8% |
| Median % Flight Time Remaining After Last Optimization Request | 26.3% | 23.2% | 20.7% |

for Unsuccessful, 11.9 for No Significant Change). The Successful flights, therefore, had the most to gain before optimization.

The average distance travelled for Successful flights was 126 nautical miles less than that for the Unsuccessful flights, but average impacts per flight was the same for both groups. Successful flights also experienced fewer than half as many refused optimization requests. With higher delays, equivalent impacts, shorter flights, and fewer refused optimizations, we can readily surmise that the Successful flights encountered more severe obstacles along their paths, obstacles that the optimization model was able to reduce or remove.

Finally, the Successful flights also benefited from a significant decrease in the number of conflicts, which may be largely attributed to lateral optimization moving some flights away from previously-encountered conflict points.

**Unsuccessfully optimized flights** showed only a small increase in the number of conflicts, but a 51% increase in the SUAs encountered (from Baseline to LDRR), indicating that their flight paths crossed more restricted zones which the optimizer did not know about. (As described in the Modeling Approach, if a path returned from the optimizer intersected additional SUAs, the controller model would insert avoidance routes for those SUAs into the path.) These flights also had the most refused optimization requests, which was most likely to occur when additional SUAs were inserted into the optimized path and the new total delay or distance exceeded the threshold.

The Unsuccessful group also had the highest percentage of flights with additional impacts after their last accepted optimization, and the smallest amount of flight time remaining after their last request, suggesting that more of these flights were affected by impacts near the end of the flight, and fewer were able to reduce delay sufficiently (to less than 5 minutes) before approaching their destination.

**No Significant Change flights** were the shortest in the optimized sample, on average. They made fewer optimization requests and were affected by fewer impacts.

In general, it would appear that these flights were affected by just sufficient impacts to reach the threshold for attempting optimization, but their ability to optimize was minimized by shorter flight lengths and/or known restricted zones which the optimizer could not significantly improve upon.

Thus, while some flights (18%) were able to improve upon their baseline results via optimization, 53% of flights attempting optimization finished with higher delays than they encountered in the baseline case.

Indicators for the flights performing worst under optimization suggest that more foreknowledge of future impacts, i.e. better distribution of important ATM and weather data to all participants, would allow these flights to follow more optimal paths, and would allow the users to be able to plan in accordance to their own business objectives more effectively.

## *Service Provider Impacts*

Table 2 showed a slight reduction (-0.6%) in the total number of conflicts between the two scenarios. Further conflict analysis showed variations in conflict geometry as well, i.e. differences in the angles of intersection between flights in conflict. Changes in either of these indicators may have implications on system predictability, safety, and controller workload.

Conflict geometry was analyzed separately, for both the Baseline and LDRR scenarios, for two groups of flights: the 13,839 flights which were never optimized; and the 1,448 flights which were optimized at least once. Figure 1 illustrates these results.

As discussed previously, the LDRR scenario showed only subtle impacts on the group of non-optimized flights, and as we might expect, the conflict geometry results for this group show only small variations in the distribution. (Variations in this group are due to conflicts between non-optimized and optimized flights in the LDRR scenario: as the optimized flights and conflicts shift, the angles of intersection will also vary.)

Differences in conflict geometry for optimized flights, however, were more interesting. In the Baseline scenario, the optimized flights show more overtaking and converging conflicts.

62.1% of conflicts to optimized flights were parallel, small or wide angle, same direction compared to 52.3% for non-optimized flights. This may be explained by recalling that the optimized flights had the highest delays, caused by SUA and conflict avoidance maneuvers; more flights were moved onto SUA avoidance routes, creating a higher traffic density and more overtaking conflicts in those areas.

In the LDRR scenario, these flights showed an increase in overtaking conflicts (parallel same) and an increase in parallel, small and wide angle opposite conflicts, which may be the result of the optimization model choosing more consistent paths around restricted zones.

An analysis of flight attitudes (cruise, climb, descent) during conflicts showed no significant change between the Baseline and LDRR scenarios for either group of flights. Optimized flights were involved in a higher proportion of conflicts where both flights were in cruise than non-optimized flights (76% versus 59%, in both scenarios).

All four measured ARTCCs showed an overall reduction in 15 minute instantaneous flight counts during the LDRR scenario. Average and total flight times in sector were also reduced. Flight time beyond the 4 ARTCCs, however, increased. Table 7-14 shows the change in total flight time within the 4 ARTCCs and for all other airspace (defined as total flight time minus ARTCC flight time).

Only a slight change in flight time in sector was found for flights which were never optimized. For optimized flights, a significant decrease was found within the 4 ARTCCs and a corresponding increase beyond. Since on average only 18% of optimized flights were able to *reduce* their overall flight duration while flight time remained constant or increased for other optimized flights (Tables 7-10 – 7-12), we can conclude that the reductions in sector occupancy are due not to increased flight efficiency, but rather to flights whose optimized trajectories spent less time within the 4 measured ARTCCs.

**Figure 1. Conflict Geometry**

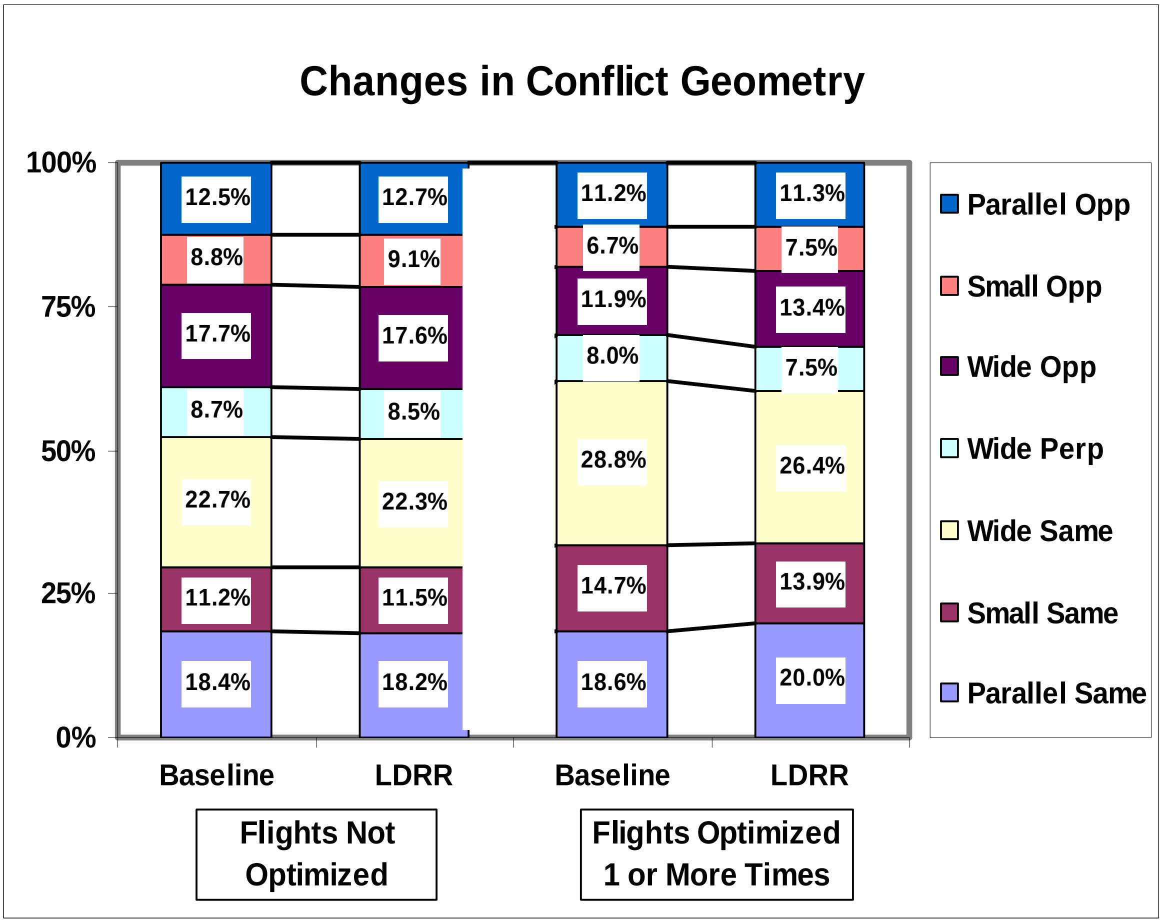


RAMS Plus recorded taskload events for the 4 measured ARTCCs using a set of ATC triggers representing typical events and controller tasks. Event triggers for optimization-specific events were added and recorded. For each event, the current time, sector, and flight were logged along with the relevant event information. A comparison of recorded events for optimized flights only, showed that task counts for typical sector crossing activities were 6.1% lower overall in the LDRR scenario. Point-outs were slightly higher. When the optimization-specific tasks which were present only in the LDRR scenario are added, the total task count had increased by 8.7%.

The noted reductions in sector occupancy, conflicts, and controller task counts (excluding the optimization-specific tasks) within these 4 central ARTCCs are consistent with each other, but we cannot apply them to the full NAS.

Because optimized flights were able to re-route outside of these 4 ARTCCs (spending more flight time, therefore, in other, un-measured ARTCCs), we cannot assume that controller tasks for the NAS were reduced overall.

## Conclusions

Based upon the individual circumstances for each flight, some flights benefited from LDRR while others showed little difference or a negative effect from re-routing with limited awareness. While some optimized flights (18%) were able to improve upon their baseline results via re-routing, with limited information, 53% of flights attempting optimization finished with higher delays than they encountered in the baseline.

Overall, the total number of conflicts in the scenario was reduced slightly as a result of optimization and this produced a small positive impact on the flights which did not attempt re-routing.

Indicators for the flights performing worst under optimization suggest that more foreknowledge of future impacts, i.e. better distribution of important ATM and weather data to all participants, would allow these flights to follow more optimal paths, and would allow the users to be able to plan in accordance to their own business objectives more effectively.

## References


[1] Bradford, S., Liang, D., Mondoloni, S (1998). Using Inter-Process to Expand Modeling Capabilities. The Second United States/European Air Traffic Management Research and Development Seminar, Orlando, December 1-4.


## Author Biographies

**Carolyn L. Sorensen** *works for ISA Software in Paris, France as an ATM analyst and software engineer. She is responsible for ATM analyses and future operational concepts simulation studies, and develops tools and models which assist with simulation preparation and results analysis. Previously, Ms. Sorensen worked for Crown Communications and served as the Technical Lead for the Sector Design Analysis Tool (SDAT) for the FAA's Air Traffic Airspace Planning and Management Branch (ATA-200) for several years. SDAT is an analytical model for airspace and traffic design and has been fielded to each ARTCC as well as FAA Headquarters, several regional offices, and the ATCSCC.*

**Diana Liang** *works for the Office of System Architecture and Investment Analysis for the Architecture and System Engineering Division of the FAA. She is responsible for the development of the NAS Architecture Tool and Interface called CATS-I, directing analyses in support of NAS Concept Validation, and the development of Modeling Tools and Fast-Time Simulations to support that validation. This work includes several models she is developing jointly with NASA and cooperative efforts with Europe via Eurocontrol. Prior to working for ASD, Ms. Liang worked in the Office of Energy and Environment for two years as the lead for the Emissions and Dispersion Modeling System (EDMS), updated the FAA's Air Quality Handbook and reviewed Environmental Impact Statements related to emissions. Ms. Liang holds a BS in Computer Science and is currently attending George Washington University.*

**Ian Crook** *has over 17 years experience in computing, specializing in the application of leading edge technologies to user oriented software systems. Specific skills include the use of Object Oriented methods, Artificial Intelligence, Discrete Event Simulation and Distributed Object Architectures.*

*Having spent his formative years working in the aviation manufacturing industry, specializing in the development of on-board aircraft control software, Ian spent five years developing telecommunication systems, before returning to the aviation industry in 1991. Since then, Ian has specialized in the design of ATM-oriented simulation systems.*

*Recent projects include the Eurocontrol Airspace Model (EAM), Capacity Analysis Facility (CAPAN), Reorganised ATC Mathematical Simulator (RAMS), Airspace-Airport Integrated Modeling System (AIMS), FAA RAMS-OPGEN dynamic link and the ISA Software/FAA ATMOS Weather server.*