A Framework for Collaborative Routing

prepared by the

Collaborative Routing Sub-Group

I. Introduction

Collaborative Decision Making (CDM) encompasses a new set of concepts, tools and procedures that have been applied to enhance the planning and control of Ground Delay Programs (GDPs) in the United States. The goal of Collaborative Routing (CR) is to apply the philosophy underlying CDM to address the management of the en-route airspace. While the "conventional wisdom" is that the major bottlenecks in the US National Airspace System (NAS) are associated with the limitations on the arrival and departure rates to and from airports, it has become increasingly evident that very significant delays and throughput degradations have arisen from en-route airspace problems and limitations. Most would agree that the principal source of problems is the convective weather activity centered around the summer months. Other problems are associated with demand surges that result from prior schedule disruptions possibly resulting from GDPs or other "upstream" control actions. A rather common set of en-route problems, generally associated with overloaded or closed airport arrival fixes, lie at the boundary between the domains of CR and GDPs.

En-route airspace congestion today is managed using a variety of actions generally under the control of one of three entities: the Air Traffic Control Systems Command Center (ATCSCC), the Air Route Traffic Control Centers (ARTCCs) and the Airline Operational Control Centers (AOCs). Localized problems and their solutions are typically handled by the appropriate ARTCC. Planning, managing and coordinating the response to major convective events are the domain of the ATCSCC, which has a severe weather unit specifically for this purpose. However, any action planned by the ATCSCC typically must be closely coordinated with appropriate ARTCCs. For example, when the ATCSCC plans major traffic reroutes, some ARTCCs will typically require that miles-in-trail (MIT) restrictions be placed on traffic entering certain areas to insure traffic levels do not overload sectors. The system-wide impact on throughput and delay is determined by the combination of the traffic reroutes and the MIT restrictions. The ATCSCC specialists and the ARTCC traffic flow managers have access to the Traffic Situation Display (TSD) which displays overall traffic status. However, no decision support tools are available to help design control actions or to evaluate their impact. In all cases, control actions or restrictions initiated by an ARTCC or the ATCSCC, leave leeway for the AOCs to react. For example, particularly if advance warning is given, the AOCs may be able to reroute flights away from congested areas. Thus, the overall response to an unfavorable situation is determined by the combined actions of the ATCSCC, the ARTCCs and the AOCs. In fact, one of the principal goals of CR is to better coordinate the individual responses of each of these entities.

In order to successfully design a CR strategy, the key CDM philosophical components must be adapted to the en-route airspace setting. A fundamental tenet of CDM is to allow, airspace Users, when possible, to independently relieve projected capacity-demand imbalances. By doing so, Users can design the most economic solutions. An absolute necessity to effectively accomplishing this objective is to provide airspace Users and managers with a common, accurate view of anticipated airspace demand and the associated traffic restrictions. To achieve this objective, airspace managers and Users must view NAS status based on a common information source. Furthermore, in order to have a timely and accurate estimate of demand, it is important that airspace Users provide early intent information, specifically, by filing planned trajectories in a timely fashion. While it is important that Users take responsibility for doing this, they cannot be expected to do so unless resource allocation processes have built-in incentives for doing so or, at the very least, there are not implicit penalties associated with early filings. Studies of recent airspace disruptions have shown that a significant portion of traffic in impacted areas had not originally filed through those areas. Thus, a second important component of improved demand prediction is to insure that reroute decisions made after a flight is airborne are done so with knowledge of the overall system impact. All of the above suggest the need for a common database and information display that provides current NAS status, including both traffic flows and control actions in place, that predicts future demand and capacity and that provides early notice of planned control actions. Users should be given the opportunity and the incentive to implement a proactive and positive response to any planned control action. Any required resource rationing should be based on fair allocation principles and should reward airlines that provide early intent information and/or take actions to relieve airspace congestion. To accomplish these goals decision support tools must be developed which allow for flexible viewing of NAS demand and capacity and which analyze the impact of any existing or planned NAS restrictions. Tools for collaborative resource allocation are also required. Of course, such tools must be available to all three parties that participate in problem resolution, namely, the ATCSCC, the ARTCCs and the AOCs.

II. Proposed Architecture

The type of problem we address here is a readily identified, demand-capacity imbalance in the en-route environment, particularly those associated with a weather front (e.g., convective weather, excessive demand). Presumably, this would be short-lived (less than 12 hours) situation, though the same principles could be applied to longer-term problems.

There are eight (8) components to the architecture, as described below. A complex CR scenario would be resolved by following the components in the order that they are presented but there will instances in which the components can be employed independently, in small groups or within a feedback loop to other components.

1. Early Problem Identification
This is process of identifying regions of the NAS that will most likely be subject to a capacity degradation based on forecasts made more than 4 hours in advance. The most typical scenario is the identification of areas where substantial convective weather activity is predicted. It will typically be is too early to have precise demand information, however, the predicted capacity degradation would be severe enough to cause disruptions to normal traffic flows.

2. Collection of Early Intent Information
This is the process of collecting user intent information early enough to allow for strategic problem identification and planning. The agreed upon definition of "early" is 4 hours in advance of flight departure. In this case, User intent is captured in the form of flight plan filings as well as more typical GDP information such as notification of flight timing changes and flight cancellations. In addition, under CR, a richer flight plan information structure is planned which will embody acceptable alternatives. Incentives will have been built into the resource allocation procedures to reward submission of early filings.

3. Situation Assessment/Advisory
TFM will routinely identify any anticipated traffic flow constraints and NAS performance degradations such as:

Miles-in-trail restrictions
Delays associated with certain air routes
Congestion associated with certain sectors
Blocked or degraded portions of the airspace

Additionally, advisories will be issued for any planned major initiatives.

Unlike the early problem identification, these AOC predictions and actions are based on both demand and capacity forecasts and thus will take into account both capacity degradations and demand surges or combinations thereof. Furthermore, at this point the information used should have a reasonably high level of accuracy.

4. User Response
Users adjust their flight plans, make cancellations, and form alternate route preferences in response to the announced conditions and advisories. This stage affords the Users the opportunity to move flights out of impacted areas or to alleviate problems totally so that major TFM initiatives become unnecessary.

5. Resource Rationing
Resource rationing will be required when the user response was not sufficient to eliminate a problem, thus requiring the imposition of a major TFM initiative similar to today's severe weather avoidance programs. In addition, it could be that certain "minor rationing" takes place to address small areas of congestion. An example, might be the use of alternate departure routes to avoid local thunderstorm activity.

The underlying principal for all rationing scheme are to allow the users to explicitly or implicitly make decisions requiring an economic tradeoff. The details of how resource rationing will proceed are the subject matter of on-going R&D. However, we now outline certain major components and principals upon which the overall scheme will be based.

A richer flight plan information structure will be developed to allow for the specification of alternatives together with criteria for trading off among them. The Users will take their own economic factors into account in specifying this flight plan information. The following is simple example of this flight plan structure:

> [flight-plan1; flight-plan2/del]
>
> The meaning of this information is that in the event that departure delay associated with flight-plan1 are greater than del longer than the departure delays associated with flight-plan2, the flight-plan2 should be employed.

It is anticipated that resource rationing will proceed in two phases. In the first phase aggregate traffic flow rates will be associated with certain air routes, or more generally, traffic classes. In the second phase, individual flights will be assigned departure times, routes (or reroutes), etc.

In the first phase, two criteria will be taken into account in setting aggregate flow rates:

> Maximization of airspace throughput
> Fair balance among major traffic flows (e.g. balance among east-to-west traffic, west-to-east traffic, traffic ascending from a terminal area, traffic descending into a terminal area etc.)

The assigning of routes and departure times to individual flights, while being more complicated that the assignment of controlled departure times for CDM-based GDPs, has many similarities and thus, GDP experience and principles can be applied. The GDP process, ration-by-schedule rations resource based on flight position in the OAG schedule. We believe a variant of this principle can be applied in the en-route airspace, but at the same time a second criteria must be brought to bear, namely, the flight plan filing time. This second criterion is important in order to encourage Users' to provide early intent information. The rationing scheme will involve forming two groups of flights:

> Group1: those flights whose flight plans were filed 4 or more hours in advance.
> Group2: those flights whose flights plans ere filed less than 4 hours in advance.

All Group1 flights will receive priority over all Group2 flights. Resources will be rationed among Group1 flight based on a variant of ration-by-schedule. Resources will be rationed among Group2 flights using a priority scheme that gives "weight" to both the file time and the schedule time. Thus, even among Group2 flights there is a penalty associated with filing later, so that filing 2 hours in advance is better than filing 1 hour in advance. The appendix contains a description of how schedule based rationing might proceed in the en-route setting.

6. User Adjustments

Users perform re-routes, cancellations and substitutions in response to the allocation of resources. This is comparable to the cancellation and substitution phase of the GDP process. This gives Users the chance to reallocate resources allotted to them and to mitigate potential damages associated with the demand-capacity imbalance. The fact that the individual flight departure time and route assignment has a similar structure to the current GDP process should allow for the development of similar procedures. We do recognize that this is a more complex setting so that significant R&D is required.

7. Maximization of Resource Usage

TFM re-evaluates the situation and reacts to User adjustments by executing compression-like algorithms, designed to maximize use of airspace. Again, the more complex en-route setting implies the development of such procedures represents an R&D challenge.

8. Revisions/Review

TFM has the opportunity to revise the ‘program’ parameters, such as en-route flow rates and new capacity/demand estimates.

III. Comments on the Architecture

Note that the first four components (steps) of the process can be considered preliminary to the en-route resource allocation because they reflect an on-going demand/capacity monitoring mode. In fact, these steps can follow an iterative loop as in Figure 1. This cycle continues until one of two things happens: either (1) the event is alleviated through some combination reduced demand and preserved capacity or (2) the event necessitates formal intervention on the part of TFM in the form of resource allocation. Steps 1 – 4 will invoke steps 5 – 8 only if it becomes necessary to pursue an advisory with resource allocation.

The last four steps (5 – 8) comprise the heart of the en-route resource allocation mechanism. This is analogous to the GDP resource allocation: first, scarce resources are allocated by TFM (RBS), then Users are given a chance to reallocate their resources among their flights (substitutions and cancellations), TFM re-evaluates the situation and makes appropriate revisions or takes up any slack in the system imposed by User adjustments (compression).

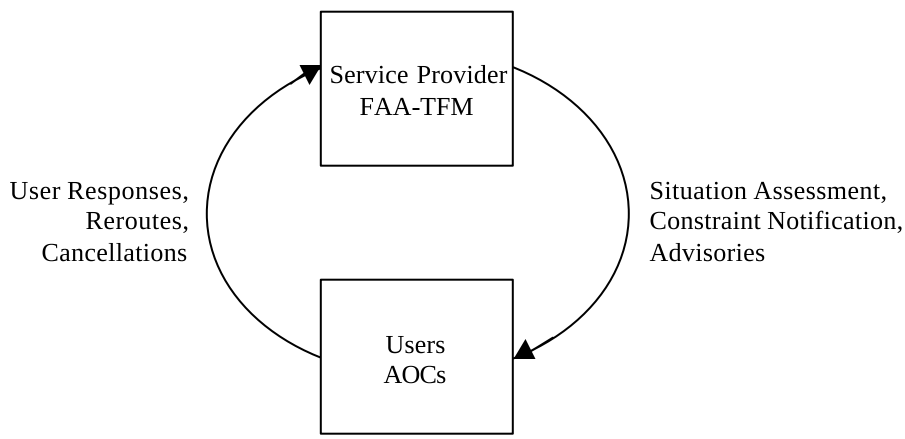

Figure 1: Cycle of Feedback during Monitoring Mode

The proposed architecture has the following desirable features:

- The categorization of traffic into classes (flow-paths) based on proposed routes is consistent with current TFM practices.
- The intuitive notion of a 'slot' and slot ownership carries over from a GDP.
- The collaborative processes employed are natural extensions to those used by SWAP and CDM GDP-E.
- The overall paradigm is similar to the Track Program used over the Pacific.
- Economic tradeoffs can be specified by the Users.
- Users are given an incentive to file early intent and preference information.

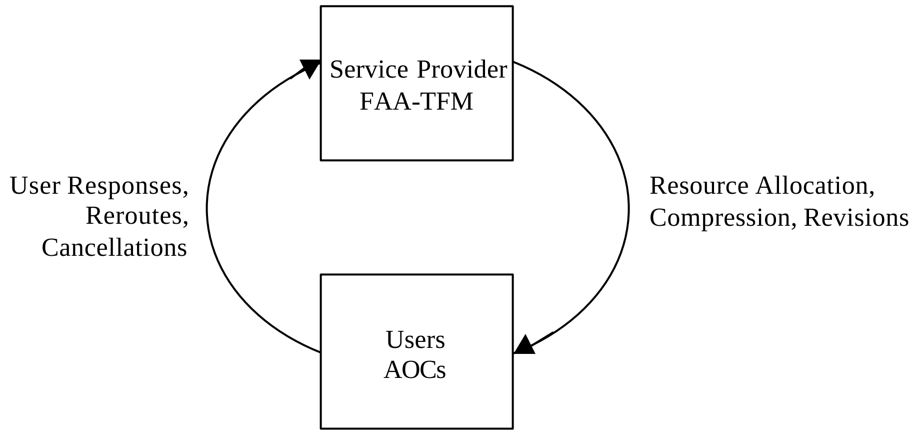

Figure 2: Constraint-Dissemination User-Preference Loop