Ninth USA/Europe Air Traffic Management Research and Development Seminar (ATM2011)

Integrating best-equipped best-served principles in ground delay programs

Andrew M. Churchill
Institute for Systems Research
University of Maryland
College Park, MD, U.S.A.
churchil@umd.edu

Michael O. Ball
R.H. Smith School of Business &
Institute for Systems Research
University of Maryland
College Park, MD, U.S.A.
mball@rhsmith.umd.edu

Alexander David Donaldson
R. John Hansman
Department of Aeronautics
and Astronautics
Massachusetts Institute of Technology
Cambridge, MA, U.S.A.
alexdon@mit.edu
rjhans@mit.edu

***Abstract*—Future air traffic management systems will consist not only of enhanced ground equipment, but also upgrades and new systems on board aircraft. To this end, an important tenet of future systems will be rewarding properly equipped aircraft. One method for doing so is through explicit prioritization of flights operated by equipped aircraft in traffic management initiatives. In this paper, the principle of Best Equipped, Best Served is examined as to its potential role in incentivizing equipage and enhancing the efficiency of Ground Delay Programs. To this end, several important policy questions pertaining to the direction of these benefits are examined. Then, three alternate allocation methods are described, incorporating aircraft equipage level as a criteria superseding scheduled arrival time. Then, a case study examining Newark Liberty International Airport, a critical and delay-prone node in the airspace system of the United States, is described. Several equipage scenarios are described, with particular attention paid to both the magnitude and distribution of the benefits realized from integrating Best Equipped, Best Served principles into Ground Delay Programs.**

Keywords-ground delay program, best-equipped best-served, BEBS, Next Generation Air Transportation System, NextGen

I. INTRODUCTION

To enable advanced air traffic management systems, both infrastructure improvements and aircraft system upgrades are required. Infrastructure improvements are largely financed by government entities are part of long-running plans to transition to newer systems. However, many costs to upgrade aircraft equipment are borne by aircraft operators. Thus, it is critical that appropriate incentive mechanisms be in place to encourage the installation of these advanced technology systems concurrent with the ground-based infrastructure upgrades.

Historically, the US air traffic management system has operated according to first-come, first-served principles. This means that flights are prioritized not according to any innate characteristics, but rather according to their arrival at some fix or other point in space. This is observed operationally in many arenas, including in providing departure and arrival clearances. At the same time, with the development of new operational concepts and decision support tools based on collaborative decision making (CDM), other flight prioritization rules have emerged, most notably ration-by-schedule (RBS).

A fundamental tenet of the FAA's Next Generation Air Transportation (NextGen) plan is *performance-based operations*. As a component of this concept, the FAA has proposed modifying existing flight prioritization principles to include use of the "best-equipped, best-served," in determining flight priority [1]. Under such a rule, those flights operated by aircraft equipped to achieve pre-established performance level are provided priority over those flights operated by unequipped or lesser-equipped aircraft. This concept has been endorsed by a federal advisory committee as an important means to help deliver the benefits of the Next Generation Air Transportation System [2].

II. RELEVANT POLICY QUESTIONS

Before presenting the details of this analysis, several important policy questions related to the use of BEBS are examined. While each of these is not answered definitively in this paper, each is examined to some degree.

- *Should a benefit (priority) be directed toward an equipped flight in cases where that flight does not produce an improvement in system performance?*

One could certainly argue that there is little justification for blindly encouraging equipage when no system benefit is produced. On the other hand, it may not be prudent to universally answering no to this question. For example, it could be that a certain percentage of aircraft must be equipped to achieve a benefit. Thus, in order to induce a sufficient number of aircraft to equip it might be necessary to provide a benefit to equipped flights even before the threshold is met, and any aggregate system benefits produced.

- *Should indirect benefits from equipped flights be proactively directed to specific groups of other flights?*

Some possible alternatives for directing such benefits include specifically directing such benefits to other equipped flights. Another alternative is to direct such benefits to the flight operator of the equipped aircraft. It should be noted that this particular approach may be attractive in cases where the

Research sponsored by the Federal Aviation Administration

flight operator cannot justify equipping the aircraft in question for the benefits that it directly produces but may find justification based on both the direct and indirect benefits. Both of these alternatives are explored in this paper.

- *Should non-equipped aircraft ever receive worse performance in order to provide a benefit to equipped aircraft?*

This is in fact, a generalization of the first question since if no system benefit were produced then it would be impossible to direct a benefit to an equipped flight without having a negative impact on a non-equipped flight. This question is posed separately because there may be many cases where the most practical way to direct performance benefits produced by equipped aircraft involves inducing a degradation in performance to at least some non-equipped aircraft.

- *Should system performance always be maximized?*

The overall goal of encouraging user equipage is to improve system performance. As we have identified in our experiments, it is quite possible that alternate approaches to BEBS yield different levels of system performance improvement. In particular, there can be tradeoffs between the total benefits achieved by equipped aircraft and the system-wide benefit/performance level. Thus approaches that do not necessarily maximize system performance but do a better job of encouraging equipage might be desirable.

III. User Equipage Options

In this section, several of the more important user equipage options to which BEBS prioritization methods might be applied are described.

- *RNAV/RNP – Area Navigation/Required Navigation Performance*

Together these capabilities allow aircraft to fly trajectories along an arbitrary set of (feasibly) defined 3-dimensional points in space. RNP trajectories can be flown within specific error tolerances (RNP level) and can include specifically defined curved segments. RNAV/RNP allows for the definition of more departure and arrival routes in and out of a metroplex of airports thereby increasing the capacity of such metroplexes.

- *GBAS – Ground Based Augmentation System*

Very high precision GPS- based navigation is enabled by GBAS, which requires that new ground infrastructure be put in place in the vicinity of an airport. This infrastructure provides for the increased accuracy of GPS-based navigation. Aircraft suitably equipped, can use the reference signals provided by the equipment on the ground to perform curved precision GPS Landing System (GLS) approaches (up to the Category I level) without the benefit of an Instrument Landing System (ILS).

- *ADS-B – Automatic Dependent Surveillance-Broadcast (-in/-out)*

Using GPS technology, an ADS-B-out equipped aircraft can continuously broadcast information on its position. This broadcast can both be monitored on the ground by air traffic controllers and in the air by other aircraft. Monitoring by other aircraft requires equipage with ADS-B-in. ADS-B-out is able to provide surveillance in areas not currently covered by radar, e.g. over large bodies of water. ADS-B-in provides very fast feedback to the cockpit on the location and movement of other aircraft thereby allowing a faster reaction time to traffic changes.

- *Datalink*

Datalink provides for two-way data communication between an aircraft and the ground. This has the potential to reduce voice frequency congestion, reduce pilot and controller workload and allow more complex messages and communication between air and ground automation systems.

Each of these capabilities has the potential to provide substantial improvements in NAS performance but also requires a substantial investment on the part of aircraft operators. The underlying costs involve both the cost of avionics and their installation and also the cost associated with idling an aircraft during the installation process. A general challenge for both NextGen and SESAR is finding a policy that makes these investments happen. While a mandate that requires flight operators to equip in certain ways to access the NAS or portions of the NAS may be appropriate, in general less intrusive mechanisms are used. The purpose of this paper is to analyze specific cases, where prioritization changes can be used to induce operators to make investments as appropriate.

IV. Ground Delay Programs

Ground delay programs are used when airport arrival demand is expected to exceed capacity for an extended period – typically several hours or more. When a GDP is imposed, flights destined for the subject airport are assigned controlled arrival times, typically necessitating that they hold on the ground at their departure airport. Ground delay is preferred over airborne delay, which would result from immediate departures, because it is requires less fuel and reduces air traffic controller workloads.

Ground delay programs are an operational implementation of the ground holding problem (GHP), as introduced by Odoni in [3]. Much of the literature in this area has focused on formulating this problem as an optimization problem, or on developing novel heuristic solution techniques. Most examinations of the GHP have focused primarily on considering various network characteristics ([4], [5], [6]) or on uncertainty in expected capacity values ([7], [8], [9]), but in each case, considering aircraft homogenously. In [10], Richetta and Odoni did consider multiple classes of aircraft in the context of aircraft weight and differential delay cost; however, each aircraft class used equal capacity. Of particular relevance to this work is the incorporation of CDM principles into GDP planning ([11], [12]). Specifically, specialized versions of the compression algorithm are implemented to focus certain indirect benefits to specific classes of flights.

In this research, the GHP is extended to consider two classes of aircraft, differing on their level of equipage, each of which has access to a different portion of airport arrival capacity. This models a future application of BEBS principles in a GDP, under which properly equipped aircraft receive priority access to airport capacity. The remainder of this paper

is devoted to examining several potential mechanisms for implementing this BEBS priority in GDPs. First, several allocation schemes are described in detail. Then, a case study examining Newark Liberty Airport (EWR) is described and the potential utility of these BEBS mechanisms examined.

V. MODELS FOR BEBS ALLOCATION

The methods we develop apply to the specific scenario where one portion of arrival capacity (the primary runway) is available to all aircraft and a second portion (the secondary runway) is only available to equipped aircraft. Thus, some (but certainly not necessarily all) equipped aircraft will be assigned to the second portion of capacity. In so doing not only will these equipped aircraft receive a benefit (reduced delay) but also space in the first capacity portion will be freed up thus providing a benefit to other aircraft. Several mechanisms are considered for addressing the scenario, by integrating best-equipped, best-served principles into the resource allocation systems used in ground delay programs. Each of these allocation schemes grants priority to equipped flights in a different manner, taking as input flight schedules and equipage status, and slot times and equipage requirements.

It is important to recognize that the modifications to the resource allocation schemes proposed here do not necessarily reflect wholesale departures from the current systems. Each mechanism represents a different approach to integrating BEBS principles with existing GDP flight prioritization, which is based on RBS. Thus, the basic approach to prioritizing many (usually most) flight remains the same.

Several input data are required for these allocation mechanisms, categorized broadly as pertaining to flights (Table I) or to slots (Table II). Flights are classed according to whether or not they are equipped with the appropriate equipment, and slots as to whether they are available to equipped flights. For both flights and slots, the data required for these allocation schemes are simple, and are consistent with the data used and available in operational systems.

The essential element of this research, and that which differentiates it from previous slot allocation methods, is the incorporation of the two classes of flights and slots – for flights, this is represented by its equipage status: equipped or unequipped. Likewise, slots are differentiated by their availability to equipped or unequipped flights, and are labeled then as universal (available to all flights) or enhanced (available only to equipped flights).

The objective of each algorithm is to develop a matching of flights to slots by setting the decision variable H_f equal to the slot s to which flight f is assigned. Although not strictly necessary because it may inferred from H_f, for algorithmic simplicity the inverse matching is also maintained, by setting I_s equal to the flight f assigned to slot s.

For each allocation method, assume initially that $H_f = -1 \ \forall f \in \Phi$ and $I_s = -1 \ \forall s \in \Sigma$. In addition, to guarantee that some feasible allocation of flights to slots always exists, assume that many slots exist with very late time markers, such that if no reasonable slot assignment exists, a flight can still be feasibly assigned, albeit with a large delay.

TABLE I. FLIGHT DATA

Variable	Definition	Description
Φ	-	Set of all flights
α_f	-	Scheduled arrival time of flight f
l_f	-	Equipage indicator of flight f: 1 = unequipped 2 = equipped
F^1	$\{f \in \Phi \mid l_f = 1\}$	Subset of unequipped flights
F^2	$\{f \in \Phi \mid l_f = 2\}$	Subset of equipped flights
Ψ		Set of airlines
a_f		Airline of flight f
A^a	$\{f \in \Phi \mid a_f = a\}$	Subset of flights controlled by airline a

TABLE II. SLOT DATA

Variable	Definition	Description
Σ	-	Set of all slots
τ_s	-	Arrival time for slot s
m_s	-	Equipage indicator of slot s: 1 = universal 2 = enhanced
S^1	$\{s \in \Sigma \mid m_s = 1\}$	Subset of slots available to all flights
S^2	$\{s \in \Sigma \mid m_s = 2\}$	Subset of slots available to equipped flights

The remainder of this section is devoted to formally outlining four allocation schemes for GDP's that include BEBS principles. The four methods differ as to the degree to which each implicitly grants priority to equipped flights. The relative magnitudes of these differences will be evaluated using a case study in the following section.

A. Selection and assignment procedure

The essence of each of the resource allocation procedures described in this section is the identification of the best flight to utilize a given slot. To that end, this procedure is described parametrically here, and is used many times in the remaining sections.

The first state of this procedure is to examine several sets of qualifications describing flights eligible for assignment to a given slot. For example, these qualifications may include a scheduled arrival time before some value and that a flight be properly equipped. In this case, the procedure identifies all flights meeting all of these criteria, as shown in the example in Figure 1, with Flights 2 and 3 meeting all criteria simultaneously. Then, assuming that this set is nonempty, the flight with the earliest scheduled arrival time is selected. Once the best flight has been identified, the records identifying flight-slot matchings are updated.

This procedure, named GP, is described below.

Procedure GP:

Inputs: $\{Z_i\}$, $\{\alpha_f\}$, $\{I_s\}$, $\{H_f\}$

Outputs: f^*, $\{I_s\}$, $\{H_f\}$

1. Identify feasible flights to assign to slot *s*, as set $U = \bigcap_i Z_i$
2. Identify best flight to assign to slot *s* as $f^* = \left\{ f \in U \mid \alpha_f = \min_{g \in U} \alpha_g \right\}$
3. If $f^* \neq \varnothing$, deassign flight f^* by setting $I_{H_{f^*}} = -1$, then assign it to slot *s* by setting $I_s = f^*$ and $H_{f^*} = s$

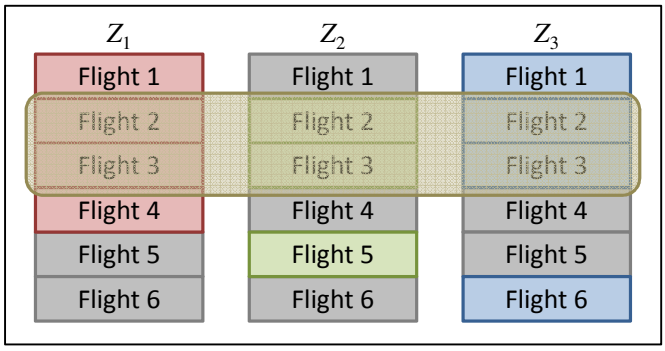

Figure 1. Identification of feasible flights

B. *Method 1: Conditional RBS*

The first method examined for incorporating BEBS principles builds on the practical implementation of the RBS principle. In this case, the list of all slots (for both equipped and unequipped flights) is examined sequentially. For each slot, the earliest flight that is not already assigned, is scheduled early enough to use the slot, and is properly equipped to use it is selected. In this method, all capacity is allocated in a single step – later allocation methods will extend this to a multistep process.

Procedure M1

0. Sort Σ according to earliest time τ_s
1. $\forall s \in \Sigma$, GP with
   $Z_0 = \left\{ f \in \Phi \mid H_f = -1 \right\}$,
   $Z_1 = \left\{ f \in \Phi \mid \alpha_f \leq \tau_s \right\}$,
   $Z_2 = \left\{ f \in \Phi \mid l_f \geq m_s \right\}$

C. *Method 2: Equipped exemptions*

The second allocation method examined is designed to incentivize equipage more strongly than in the first, by essentially providing exemptions from the GDP allocation process for equipped flights.

Again the procedure functions by examining slots sequentially. However, two stages of allocation are used. In the first, only equipped flights are considered, allowing their assignment to the earliest feasible slots without competition against unequipped flights. Once all equipped flights have been processed, then all unequipped slots are scanned again and unequipped flights assigned wherever possible.

Obviously this system provides a tremendous advantage to equipped flights by placing them as early as possible. In addition, it is likely, that some (perhaps many) equipped slots will go unused, as equipped flights are simply assigned to the earliest slot, irrespective of concerns about that slots equipage level.

Procedure M2

0. Sort S^1 and Σ according to earliest time τ_s
1. $\forall s \in \Sigma$, GP with
   $Z_0 = \left\{ f \in \Phi \mid H_f = -1 \right\}$,
   $Z_1 = \left\{ f \in F^2 \mid \alpha_f \leq \tau_s \right\}$
2. $\forall \left\{ s \in S^1 \mid I_s = -1 \right\}$, GP with
   $Z_0 = \left\{ f \in \Phi \mid H_f = -1 \right\}$,
   $Z_1 = \left\{ f \in F^1 \mid \alpha_f \leq \tau_s \right\}$

D. *Method 3: Sequential allocation with compression*

The third allocation method examined in this paper is the most complex. The objective in this allocation method is to most directly focus the benefits on those carriers that choose to equip some portion of their fleet. The motivation for this approach was discussed earlier, namely, the advantage obtained by the equipped flight might not be sufficient in and of itself to justify equipage but if additional benefits accrued to other flights in its fleet then the business case for the carrier could become positive.

To this end, the first stage of the procedure consists of allocating all flights, equipped or not, to the base level of capacity. This stage is consistent with the current RBS procedures, and establishes a baseline recognized as being fair.

After this initial allocation has been performed, the second stage begins. In this case, each enhanced slot is considered sequentially. If a flight can be feasibly reassigned while yielding some benefit to this slot, then the compression procedure (step 2b) takes effect. In this well-established procedure, each available universal slot is considered sequentially, beginning with the slot just having been deassigned from the flight moved into the enhanced slot. Flights of the same airline are given precedence (2.b.i); however when none is available, then flights of all other airlines are considered (2.b.ii). This process continues, moving up flights one at a time until reaching the end of the slot list. The process continues with moving an appropriate equipped flight into the next enhanced slot.

Procedure M3

0. Sort S^1, S^2, and Σ according to earliest time τ_s

1. $\forall s \in S^1$, GP with
   $Z_0 = \{ f \in \Phi \mid H_f = -1 \}$,
   $Z_1 = \{ f \in \Phi \mid \alpha_f \le \tau_s \}$

2. $\forall s \in S^2$

   a. GP with
      $Z_1 = \{ f \in F^2 \mid \alpha_f \le \tau_s \}$,
      $Z_2 = \{ f \in F^2 \mid \tau_{H_f} > \tau_s \}$

   b. If $g = f^* \neq \varnothing$, $\forall \{ t \in S^1 \mid I_t = -1,\ \tau_t \ge \tau_s \}$

      i. GP with
         $Z_1 = \{ f \in \Phi : \alpha_f \le \tau_t \}$,
         $Z_2 = \{ f \in \Phi : \tau_{H_f} > \tau_t \}$,
         $Z_3 = \{ f \in \Phi : a_f = a_g \}$

      ii. If $f^* = \varnothing$, GP with
         $Z_1 = \{ f \in \Phi : \alpha_f \le \tau_t \}$,
         $Z_2 = \{ f \in \Phi : \tau_{H_f} > \tau_t \}$

VI. CASE STUDY

To demonstrate the efficacy of the procedures defined in the previous section, a representative case study will be investigated using the traffic demand from an actual Ground Delay Program (GDP) at the Newark Liberty International Airport (EWR). In the example, the assumed equipment capability is a GBAS capable GBS navigation system which enables high precision GPS Landing System (GLS) approaches. The equipage scenarios vary based upon which carriers have equipped which aircraft, and to what degree.

A. Operational scenario

The Newark Liberty International Airport (EWR) is one of the 3 major airports servicing the New York metro area. EWR is one of the most delay prone airports in the US system due to high traffic demand and its constrained runway system shown in Fig. 2. The EWR runway system consists of 2 close parallel runways (4L/22R, 4R/22L) and a single crossing runway (11/29).

The Airport Arrival Rate (AAR) is dependent on the airport runway configuration and visibility conditions and varies significantly. In Low IFR or IFR conditions only a single runway is used for ILS approaches. Typically the ILS 4L or ILS 22R are used for arrivals with AARs between 28-34 arrivals per hour in Low IFR and 34-38 in IFR conditions [13]. In VFR conditions runway 11/29 can be used in coordination with 4R/22L to accept overflow arrival traffic increasing the AAR to 42-48+ [13]. Overflow aircraft on runway 11 are normally required to Land and Hold Short (LASO) of runway 4R/22L so the arrival rate to the primary runway is not reduced. When runway 29 is used for overflow the landings must be coordinated between both active runways.

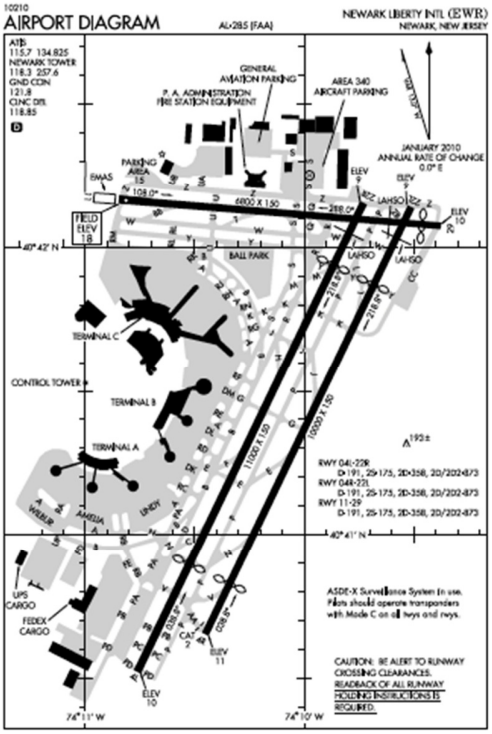

Figure 2. EWR Runway Layout [14]

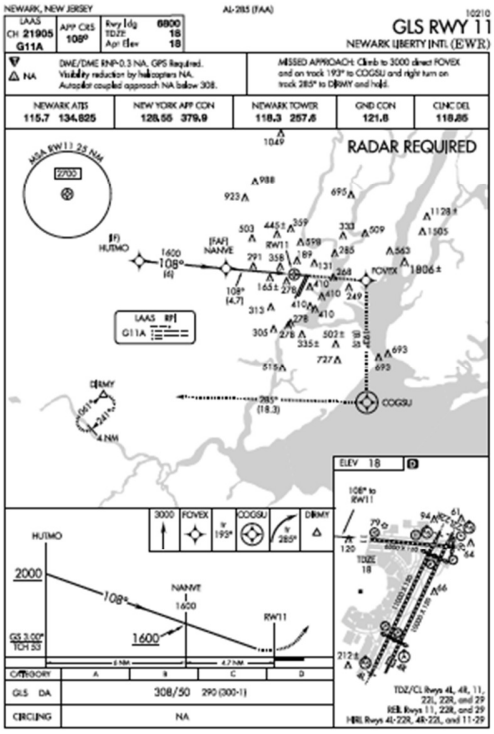

Figure 3. GLS Runway 11 Approach Procedure [14]

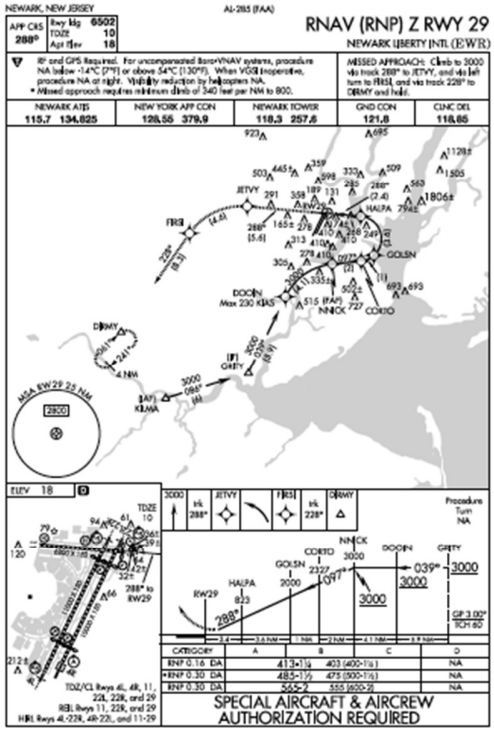

Figure 5. RNAV(RNP) Runway 29 Approach Procedure [14]

In the operational scenario considered for the case study, it is assumed that either GLS or low RNP (0.16, 0.30) capability would allow runway 11/29 to be used to accommodate overflow arrival traffic in IFR weather conditions. The current GLS RWY 11 procedure is shown in Fig. 3 and the current RNAV (RNP) Z RWY 29 is shown in Fig. 4. Note that the RNAV (RNP) Z RWY 29 procedure currently has a Special Aircraft and Aircrew Authorization Requirement (SAAAR). The GLS RWY 11 approach can be used down to ceilings of 300 ft and 1 mile visibility while the RNAV (RNP) Z RWY 29 can be used to ceilings of 400 ft and 1.5 mile visibility for RNP values of 0.16 and to ceilings of 500 ft RNP values of 0.30. The RNAV (RNP) to runway 29 requires a curved approach with a short final to avoid traffic at the New York LaGuardia (LGA) to the east and New Jersey Teterboro Airport (TEB) to the north. There is an additional RNAV (RNP) Y RWY 29 approach which is not shown for arrivals from the north which has a right turn to final.

In the operational scenario it is assumed that if runway 29 is used for overflow the Converging Runway Display Aid can be used to coordinate arrivals. If runway 11 is used for the overflow then it is assumed that arriving aircraft are capable of Land and Hold Short (LASO) procedures in IFR conditions.

The conditions for the operational scenario are assumed to be IFR with ceilings above 500 ft. The baseline AAR is assumed to be 34 arrivals an hour which is typical of IFR operations to the primary arrival runway 4R/22L. An additional 8 arrivals per hour are assumed from the overload runway 11/29 bringing up the potential AAR to 42 if appropriately equipped aircraft are available.

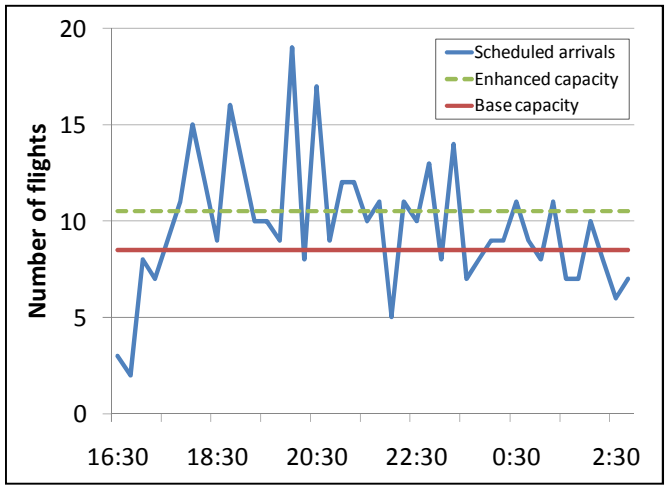

Figure 6. Flight schedule and airport capacities

The flight data used in this case study is drawn from flight schedules at EWR on June 8, 2007 from 16:30UTC until 03:00UTC on the following day as illustrated in Fig. 6. During this time period, 413 flights were scheduled, distributed among aircraft types as shown in Table III and among carriers in Table IV. As should be expected, the traffic mix at EWR is dominated by medium sized aircraft operating domestic routes for Continental Airlines, which maintains a very large hub operation there.

The controlling carrier, rather than the operating carrier, is shown as they are responsible for managing flights operating under a GDP, e.g. during a GDP, a large carrier such as Continental (the controlling carrier) may control the substitution process for certain affiliated carriers that serve as feeders for its hub. Further, it is likely that the controlling carrier would be the one that makes the decision to equip with the appropriate technology to access enhanced slots.

TABLE III. DISTRIBUTION OF AIRCRAFT TYPES

Class	Example types	Count
Heavy	A330, A340, B767, B777	40
Medium	A320, B737, MD80, DC9	219
Regional	E145, CRJ2, CRJ7	141
Other	LJ45, C550	13

TABLE IV. DISTRIBUTION OF AIRLINES

Code	Airline name	Count
COA	Continental Airlines	276
AAL	American Airlines	20
DAL	Delta Air Lines	11
UAL	United Airlines	11
USA	US Airways	11
NWA	Northwest Airlines	10
JBU	JetBlue	9
FDX	Federal Express	7
PVT	Privately operated	7
-	Other airlines	51

B. Equipage scenarios

Three future equipage scenarios are considered in this analysis, each reflecting a possible equipage decision by the carriers operating at EWR, as described below.

1) All COA RJ aircraft

In this scenario, only Continental Airlines is assumed to have equipped their entire fleet of RJ aircraft to access enhanced slots. This scenario is considered because the large COA presence at EWR could yield them significant benefits.

2) All COA, AAL, DAL RJ aircraft

The second equipage scenario consists of all RJ aircraft operated by Continental, American, and Delta having been equipped to access enhanced slots. This scenario reflects the situation in which, in addition to COA, these two other large carriers have equipped, potentially having done so primarily to yield benefits at their hub airports. However, they may also stand to gain at airports at which their presence is significantly smaller, such as EWR. Given that this scenario introduces the highest overall equipage level, it should be expected to yield the greatest aggregate benefit.

3) All AAL, DAL RJ aircraft

In the third equipage scenario, only American and Delta, each of which has a small operation at EWR, is assumed to have equipped their RJ aircraft. This reflects a situation under which the dominant hub carrier (COA) has not yet equipped, or has declined to equip, their RJ fleet. Obviously, under this equipage scenario, the expected aggregate benefit of introducing BEBS and enhanced slots would be the smallest.

4) Partial RJ equipage

The fourth equipage scenario examines the changes in performance with increasing levels of equipage. No particular carrier is posited as having equipped, but rather increasing fractions of RJ aircraft are assumed to have been equipped. At each fraction of equipped flights, several random instances are generated to ensure robust results. In this case, it is expected that the benefits should increase as the fraction of equipped aircraft increases, but should reach an asymptotic level at some fraction of equipped aircraft, as this represents the point at which all enhanced slots can be used.

C. Results

Several methods are used to describe the results of this case study, given the combination of the three allocation methods and the four differential equipage scenarios. The first three equipage scenarios will be compared simultaneously, then the variable equipage case will be examined separately.

The first analysis– the mean delay assigned under each equipage case – is shown in Fig. 7. Several trends are apparent in this figure. First, the equipage scenario with the greatest number of equipped flights, that under which three carriers equip yields the lowest delays. These delays however are not markedly lower than the first case with Continental only, as their RJ fleet is sufficiently large so as to use the entirety of the enhance capacity available in the case study. In addition, the third equipage scenario yields few benefits in the aggregate, given the very small number of flights able to access the enhanced slots.

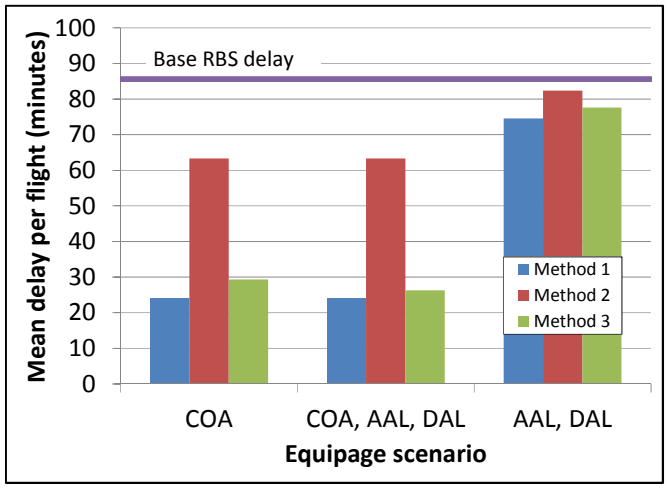

Figure 7. Mean assigned delays

Although the aggregate delays assigned under methods 1 and 3 are similar, their distribution between equipped and unequipped flights differs. In Fig. 8, the delay savings from the delays assigned under the base RBS procedure are shown. The dashed bars correspond to savings for equipped flights, while the solid bars correspond to savings for unequipped flights. It is clear that method 3 assigned greater savings for equipped flights than does method 1. This is consistent with the policy goal it sought to achieve.

However, it is also clear from 8 that the distribution of delay savings provided by method 2 – the exemption method – are significantly different from the other methods. The delay savings provided to equipped flights are the greatest for each equipage scenario, but this comes at a great cost. Namely, for the first two equipage scenarios, unequipped flights are actually disadvantaged, experiencing greater delays than under the base capacity using RBS allocation. For the third scenario with so few planes, the difference for unequipped flights is essentially zero.

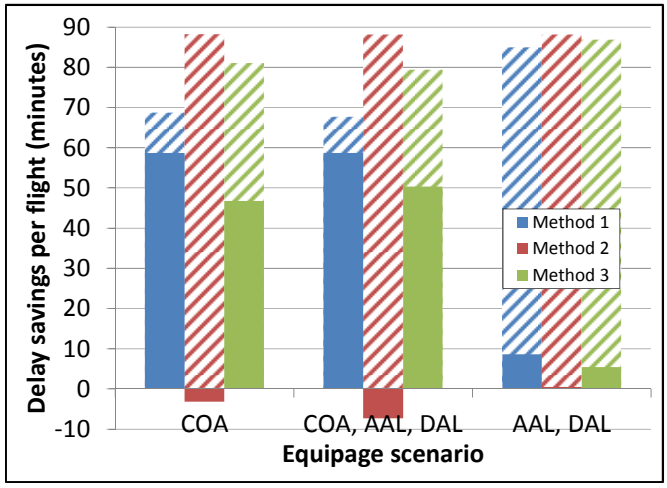

Figure 8. Mean delay savings for equipped and unequipped flights

However, as was raised in the policy section of this paper, the ability of BEBS to induce equipage depends more on the benefits received by the airline that equips. To that end, Fig. 9 depicts the delay savings for all flights operated by airlines that have equipped their RJ aircraft (dashed bars) and those that have not equipped any aircraft (solid bars). The trends observed here are similar to those in Fig. 9, in that airlines that choose to equip realize greater savings than those that do not choose to equip. Again, the application of allocation method 2 is detrimental to airlines that have no equipped any aircraft, as they experience greater delays than would have been assigned under the base RBS case.

The most important trend in Fig. 9 however lies in the great disparity in benefits realized by airlines that equipped versus those that did not. This benefit is most directly observable using method 3 – the sequential allocation method with compression – because it explicitly attempts to direct benefits to airlines that have equipped by using the compression procedure.

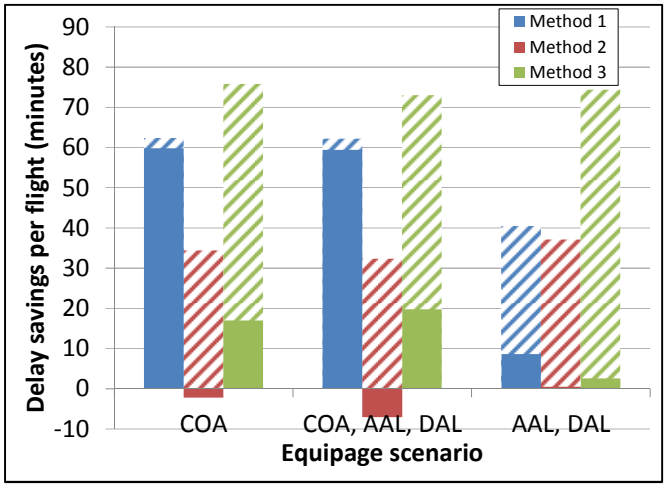

Figure 9. Mean delay savings for carriers operating equipped and nonequipped flights

The previous analyses have assumed that some discrete subset of airlines have chosen to equip their entire RJ fleet operating at EWR. However, going forward, it is likely that only some portion of eligible aircraft will be equipped to access enhanced slots during a GDP. Fig. 10 shows the variation in mean assigned delay under variable equipage conditions. Under this scenario, no particular airline is assumed to have unilaterally chosen to equip – any RJ is assumed to be equally likely to have been equipped.

The trend shown in this figure of decreasing mean assigned delay with increasing fraction of RJ equipage should be expected. The differences between methods are interesting. As with previous analyses, methods 1 and 3 are quite similar in aggregate assigned delay. Method 2 – using exemptions – is again significantly worse according to this metric. Also interesting in this figure is the marginally decreasing delay savings with increased equipage for methods 1 and 3. At around 70% equipage, little additional benefit is realized from equipping additional flights. This is due to all the enhanced slots having been utilized at that point. Given that there were 98 enhanced slots in this case study and 141 RJ aircraft eligible to be equipped, a saturation ratio of 69.5% should have been expected.

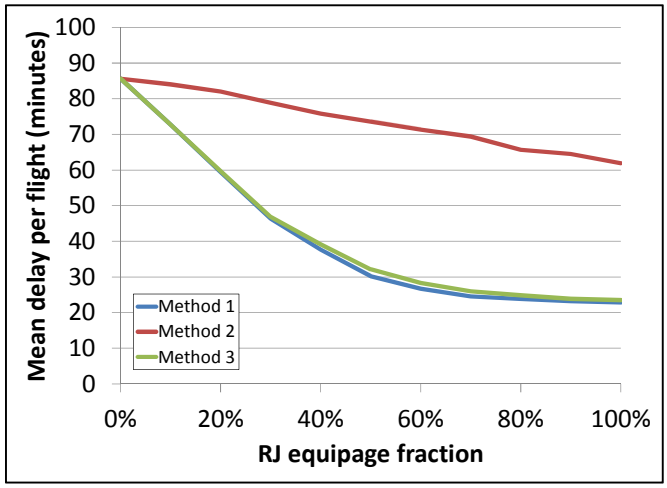

Figure 10. Mean assigned delays over range of equipage levels

Aggregate assigned delay is only one method of evaluating the variable performance of the allocation methods under increasing RJ equipage levels. In Fig. 11, the delay savings for both equipped (dashed lines) and unequipped (solid lines) flights are shown over a range of equipage levels. This figure shows that, at low equipage levels, the difference in delay savings between equipped and unequipped flights is very large. This should be expected, as those few equipped flights essentially get whatever enhanced slot they would like. Of course, as the fraction of equipped flights increases, this differential decreases and the two groups come nearer to one another in terms of mean difference from delay assigned under base capacity using RBS.

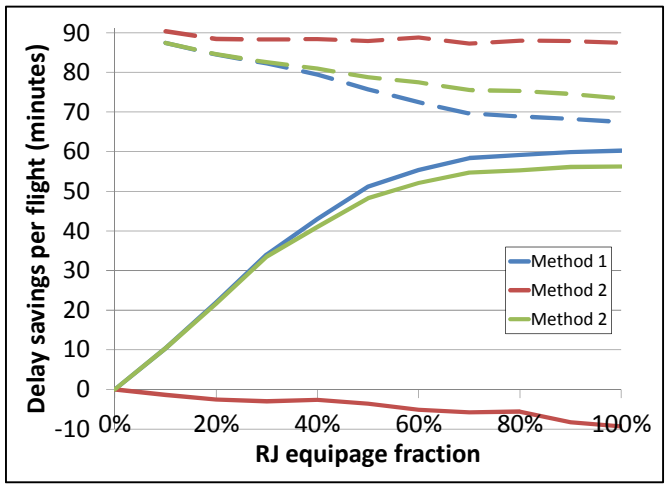

Figure 11. Mean delay savings for carriers operating equipped and nonequipped flights over range of equipage levels

The distributional differences of each of the three methods are also apparent across the entire range of possible equipage. Method 3 assigns greater delay savings to equipped flights than

method 1 does. In addition, the inequities of method 2 are once again apparent. The delay savings for equipped flights realized by this method are uniformly greater than the other methods, but this comes at the cost of assigning delay increases to unequipped flights.

Because exemptions may be a convenient and straightforward operational tactic for implementing BEBS principles in a GDP, it may be useful to examine more closely the equity impacts of this method. In Table V, the fraction of flights assigned increased delays, relative to what they would have received under RBS with the base capacity, is shown.

TABLE V. FRACTION OF FLIGHTS WITH INCREASED DELAYS

Equipage scenario	Method 1	Method 2	Method 3
A.	0	49.4%	0
B.	0	50.4%	0
C.	0	18.6%	0

As has been depicted in previous analyses, unequipped flights are only disadvantaged when the exemption method is employed. In the quite reasonable equipage scenarios 1 and 2, the impact to unequipped flights, in terms of increased delays is quite significant.

It is also instructive to consider the effect of varying the fraction of equipped flights, as in the final two figures. In 12, the fraction of unequipped flights experiencing delay increases is shown as a function of the fraction of the RJ fleet equipped to access enhanced slots. The fraction of unequipped flights impacted increases across the range of equipage levels. The aggregate impact of these delay increases is small, as depicted in Fig. 11. However, because the delays do increase, this may present a challenge to wide user acceptance for operational practice.

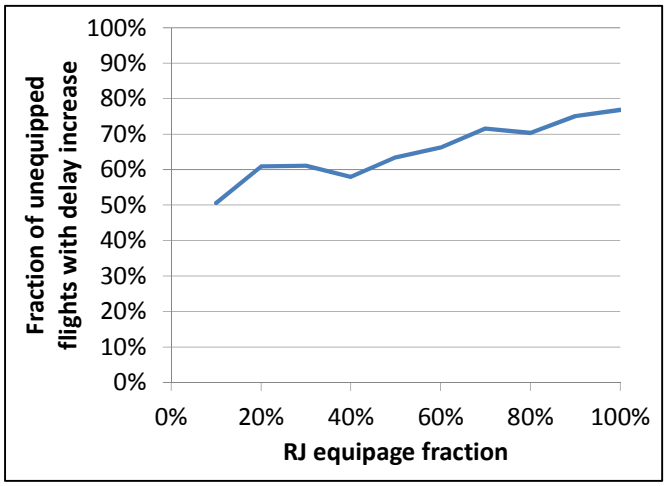

Figure 12. Fraction of unequipped flights experiencing delay increases over range of equipage levels

VII. CONCLUSIONS

This paper has examined several mechanisms for integrated best-equipped, best-served principles into ground delay programs. Several policy questions were introduced to motivate the research present here. These pertained primarily to the distribution of benefits realized by including BEBS principles and increased capacity.

Three allocation methods were described for including BEBS principles in a GDP. The first of these methods extended the existing RBS procedure used by simply selecting the earliest flight that is properly equipped for each slot. The second method reflects the desire to assign considerable benefit to equipped flights by granting them exemptions through the GDP. The final method began with the RBS allocation on the base capacity. Each enhanced slot was then added and compression performed after each addition. Thus, the benefits of introducing equipped flights to enhanced slots were most directly assigned to the airlines that chose to equip some portion of their fleet.

These three allocation methods were examined in concert with several realistic equipage scenarios using a case study for Newark Liberty International Airport. Method 3 was shown to be the most successful in assigning benefits to airlines that chose to equip the appropriate flights. However, this came at a slight cost in terms of aggregate assigned delay. According to this metric, the first method was most successful. The exemption procedure was quite successful in minimizing delays to equipped flights, but was overall quite inefficient because of the tremendous delays assigned to unequipped flights.

It is this last point that is critical in considering operational BEBS implementations. Simply exempting equipped flights to their desired arrival time, irrespective of the availability of an enhanced slot at that time, takes a considerable amount of base capacity away from unequipped flights and exacts a tremendous cost in efficiency terms. Thus, it is seems prudent that an operational implementation of BEBS principles in a GDP utilize a modified RBS procedure such as methods 1 or 3 to achieve a balance between efficiency – total assigned delay – and equity – distribution of delay savings to airlines that choose to equip appropriate aircraft.

ACKNOWLEDGMENT

The authors gratefully acknowledge the financial support of the Federal Aviation Administration, and assistance with GDP data provided by Metron Aviation.

REFERENCES

[1] Federal Aviation Administration, "NextGen Implementation Plan," 2010.

[2] U.S. Department of Transportation, "Future of Aviation Advisory Committee Reccommendations," 2010.

[3] A.R. Odoni, "The flow management problem in air traffic control," in *Flow Control of Congested Networks*, Amedeo R Odoni and G Szego, Eds. Berlin: Springer-Verlag, 1987.

[4] O. Richetta and A.R. Odoni, "Solving Optimally the Static Ground-Holding Policy Problem in Air Traffic Control," *Transportation Science*, vol. 27, pp. 228-238, 1993.

[5] P.B. Vranas, D.J. Bertsimas, and A.R Odoni, "The multi-airport ground-holding problem in air traffic control," *Operations Research*, vol. 42, pp. 249-261, 1994.

[6] D.J. Bertsimas and S. Stock Patterson, "The air traffic flow management problem with enroute capacities," *Operations Research*, vol. 46, no. 3, pp. 406-422, 1998.

[7] A.M. Churchill, "Coordinated and robust aviation network resource allocation," University of Maryland, Ph.D. Dissertation 2010.

[8] M.O. Ball, R.L. Hoffman, A.R. Odoni, and R. Rifkin, "A stochastic integer program with dual network structure and its application to the ground-holding problem," *Operations Research*, vol. 51, no. 1, pp. 167-171, 2003.

[9] A. Mukherjee and M. Hansen, "A dynamic stochastic model for the single airport ground holding problem," *Transportation Science*, vol. 41, no. 4, pp. 444-456, 2007.

[10] C.N. Glover and M.O. Ball, "Stochastic integer programming models for ground delay programs with weather uncertainty," in *Proceedings of 4th International Conference on Research in Air Transportation*, Budapest, 2010, pp. 217-223.

[11] T.W. Vossen and M.O. Ball, "Slot Trading Opportunities in Collaborative Ground Delay Programs," *Transportation Science*, vol. 40, no. 1, pp. 29-43, 2006.

[12] M. Wambsganss, "Collaborative decision making through dynamic information transfer," *Air Traffic Control Quarterly*, vol. 4, no. 2, pp. 109-125, 1997.

[13] http://www.fly.faa.gov/Information/east/zny/ewr/ewr_tm.htm

[14] FAA (NACO) Terminal Approach Procedures

AUTHOR BIOGRAPHIES

Andrew M. Churchill is a Research Associate in the Institute for Systems Research at the University of Maryland. Dr. Churchill received his B.S. degree in Aerospace Engineering, and M.S. and Ph.D. degrees in Civil and Environmental Engineering from the University of Maryland.

Michael O. Ball is the Associate Dean of Faculty and Research and the Orkand Corporation Professor of Management Science in the Robert H. Smith School of Business at the University of Maryland. He also holds a joint appointment with the Institute for Systems Research. Dr. Ball received his Ph.D. in Operations Research from Cornell University. He is co-Director of NEXTOR, the National Center of Excellence for Aviation Operations Research.

Alexander Donaldson is a Research Assistant in the MIT Department of Aeronautics & Astronautics. He received his B.S. and M.S degrees in Aeronautics and Astronautics from MIT.

R. John Hansman is the T. Wilson Professor of Aeronautics and Astronautics at the Massachusetts Institute of Technology where he is the Director of the MIT International Center for Air Transportation.