

Minimizing the Cost of Delay for Airspace Users

Stephen Kirby
ATM/RDS/NET
EUROCONTROL Experimental Centre
Bretigny-sur-Orge, France
Stephen.kirby@eurocontrol.int

Nadine Pilon
ATM/RDS/NET
EUROCONTROL Experimental Centre
Bretigny-sur-Orge, France
Nadine.pilon@eurocontrol.int

Abstract—In European air traffic flow management, regulation is the assignment of take-off times to prevent the over delivery of flights to sectors and airports. Delay is assigned to flights according to the principle of ‘first planned first served’ for entry into the regulated volume. The take-off times (and hence delays) take no account of the relative importance of flights from the airspace user perspective. The user driven prioritization process (UDPP) aims to provide airspace users with the opportunity to modify the sequence of flights in a regulation/hotspot to minimize their cost of delay. This paper reports on the first human-in-the-loop validation exercise to assess the UDPP concept. These preliminary results indicate that significant cost savings are possible for airspace users while UDPP non-participants are not affected to an unacceptable degree.

Keywords—UDPP, delay, airspace users, regulation, hotspot

I. INTRODUCTION

A. Regulation

When traffic is predicted to exceed capacity in a sector or on arrival at an airport flow managers in European control centres may ask the European Network Manager (EUROCONTROL) to regulate the volume concerned so as not to exceed the defined capacity. This is done by limiting the number of entries into the volume per unit time. The regulation is performed using an algorithm known as CASA (computer assisted slot allocation) [1]. CASA calculates entry times into the regulated volume and associated departure times (calculated take-off times, CTOT). Usually the CTOT will be later than the scheduled take-off time, hence regulations impose delay. Flights at the beginning of the regulation tend to experience little delay, but delays accumulate thereafter with flights in the middle often experiencing the highest delay; thereafter, allocated delays gradually subside to zero at the end of the regulation. CASA assigns delay according to the planned entry times of flights into the regulated volume. The algorithm has no awareness of the relative importance of flights for a given airspace user in that regulation. This is problematic for airspace users because each flight has a unique cost-delay curve: the cost of delay is not the same for every flight. Thus, for airspace users with several flights in a regulation CASA is unlikely to allocate delay such that airspace users’ costs are the minimum.

B. The Concept – User Driven Prioritization Process (UDPP)

The concept is being developed under the ongoing European SESAR work programme [2] (SESAR 1 from 2010 to 2016, SESAR2020 from 2016 to 2023). New terms are being used in SESAR: A flight that has left its blocks is ‘in execution’, and a volume where demand exceeds capacity is a ‘hotspot’. A regulation can be applied to a hotspot, but it doesn’t follow that a hotspot must have a regulation. The SESAR terms shall be used for the rest of this paper.

The UDPP concept is designed to address hotspots. When such a situation is identified by air traffic management, UDPP can be applied. The concept is designed to work for airport arrival or departure hotspots, or even hotspots at the airport, such as the de-icing area. There are plans to develop it further for the *en route* phase, but this is out of the scope of this paper. UDPP is designed to be applied in the short term planning phase (a few hours or days before execution), although it should be able to cope with a hotspot that has some flights in the short-term planning phase and some in execution.

There are two concept elements to the UDPP concept, which are complimentary:

- flight delay apportionment (FDA); and
- selective flight protection (SFP).

For FDA, an airspace user is allowed (but not compelled) to assign a numerical value $\{1, 2, 3, \dots, 9, Ba\}$ to some or all his flights. This can be done after the hotspot has been published (up to a suitable cut-off period) or before the hotspot is even known. FDA priority values can be changed as many times as necessary to match the changing priorities of flights. The simple idea is that a high priority flight should have less delay than a lower priority flight, which will have to accept more delay to balance out, and this can be achieved by assigning low numerical values to priority flights, and high numerical values to low priority flights. The clever among you may immediately think of assigning ‘1’ to every flight. Unfortunately, the airspace user would be saying “all my flights have equal priority”, which would result in none being prioritized. So FDA is a means for an airspace user to redistribute his total delay amongst his flights in the hotspot to match his priorities. One

final point: there is a ‘Ba’ (baseline) option too. When the hotspot is first published every flight in the hotspot will be assigned a baseline delay, which is the delay each will have if UDPP is unavailable. Assigning ‘Ba’ will force the FDA algorithm to ensure that the given flight will keep its baseline delay. The FDA algorithm is designed to ensure that an airspace user’s total delay in the hotspot after UDPP is the same as his total baseline delay. This is the ideal, and in practice some flexibility has been built into the FDA algorithm to ensure there are no gaps in the reformulated hotspot sequence as this could potentially decrease runway throughput.

SFP requires a little more explanation. The key idea is that an airspace user can heavily penalize one of his flights to allow several of his flights to be put back on time. This is achieved through a system of operating credits. Every flight is afforded 100 operating credits. In normal operating conditions (meaning without significant congestion) the operating index (OI) for any given volume will be 100 too. A flight that has at least as many operating credits as the operating index should be able to fly on schedule. However, when demand exceeds the available capacity and a hotspot is published a new, higher OI will be in force. Flights with just 100 operating credits will be subject to their baseline delays, calculated at the instant the hotspot is published. The value of the new OI will be directly proportional to the ratio of demand to capacity (in the so-called ‘stressed period’, i.e., the part of the hotspot where planned demand actually exceeds capacity). To illustrate the point, suppose that unfavourable meteorological conditions have reduced arrival capacity for airport X, and the planned demand is one third higher than the revised capacity. The OI increases from its usual 100 to 133. Any flight in the hotspot that wishes to arrive on time will need 133 operating credits (remember, all flights start with 100 credits). So, for this particular hotspot one flight could donate 99 credits and give 33 each to three others. The three with 133 would be put back on schedule, but the flight with now only one operating credit would be ‘suspended’, therefore put to the end of the hotspot. It is important to understand that the deprioritized flight, whilst in all probability heavily penalized with a long delay, would not be suspended in the conventional sense as it would still be in the hotspot sequence and still be able to fly. SFP creates temporary holes in the hotspot sequence, and injects suspended/protected flights elsewhere in the sequence. The algorithm by design attempts to eliminate all holes so as not to decrease runway throughput. The effect on the reformulated sequence is that many flights can be nudged forward or back a place or two, including those flights that belong to airspace users not taking part in UDPP.

To round off the overview of the UDPP concept it’s worth drawing attention to some principles to outline how we currently expect UDPP to work in practice. An airspace user may only apply FDA and SFP to *his* flights. Both FDA and SFP will probably be made available in a hotspot, not just one or the other. Participation in UDPP is voluntary: no airspace user will be forced to use it. Prioritizations can be submitted and revised as many times as necessary up to a suitable cut-off time defined by air traffic management and the airport operations centre (APOC) concerned. Airspace users may be submitting UDPP priorities simultaneously, or near

simultaneously – there are no turns. Finally, to be able to prioritize or protect a flight first the airspace user must deprioritize or suspend a flight further ahead in the sequence – this is referred to as ‘ration by effort’.

II. VALIDATION EXERCISE

A. The Setup – Platform and Prototype

Several validation exercises of increasing sophistication have been carried out to assess the UDPP concept. This paper concerns the latest exercise, which was carried out in 2016. The validation exercise used a human-in-the-loop platform, which provided five airspace user positions. Each position provided a fleet management view for an airspace user participant, permitting each to view the schedule of flights allocated to him and the impact on that schedule caused by a hotspot. The platform also provided airspace user participants the means to set priorities with FDA and SFP, and see the predicted affect via a what-if mode before submitting priorities for real. The platform also managed the publication of a hotspot, and the sequence of flights in the hotspot taking into account priorities submitted by airspace user participants.

A prototype was also used in some validation runs to provide airspace user participants an alternative means to set priorities with FDA and SFP. The main purpose of the prototype, however, was to provide a cost of delay estimate for each flight in the hotspot, and a total cost of delay per airspace user participant. The prototype was built by Sabre Airline Solutions, and part funded by the European SESAR work programme [2].

B. Validation Scenarios and Runs

Three different solutions (FDA only, SFP only, FDA and SFP) were assessed against two different departure capacity constraints (low visibility, de-icing – see Table I). These six solution scenarios were joined by two reference scenarios (one for low visibility, one for de-icing) to determine how each solution fared against the baseline of ‘doing nothing’. The two reference scenarios were simply the baseline delays calculated in the hotspot before any UDPP actions were made.

TABLE I. DESCRIPTIONS OF CAPACITYCONSTRAINTS

Capacity Constraint	Events		
	Start ^a	Phase 1 ^b	Phase 2
Low Visibility	No capacity constraint	Capacity constraint: 0730-1130	Capacity constraint: 0730-1130
	69 mov/hr ^c	39 mov/hr	30 mov/hr
	OI ^d = 100	OI = 150	OI = 194
De-Icing	No capacity constraint	Capacity constraint: 1545-1900	Not played
	69 mov/hr	21 mov/hr	
	OI = 100	OI = 191	

a. The starting conditions before any capacity constraints are imposed.
b. The phase denotes a change from the starting conditions, and represents a specific hotspot.
c. (Departure) runway movements per hour
d. Operating index. Normally this is 100, but can be increased to reflect demand exceeding the available capacity.

One traffic sample was used for the whole exercise. The sample covered 24 hours of data for Paris Charles de Gaulle (CDG) airport, France, for 3rd July 2015, although the sample was subject to some adaptations to suit the needs of the exercise. Modifications included the removal and modification of some flights, the de-identification of all remaining flights by reassigning new ICAO (International Civil Aviation Organization) and IATA (International Air Transport Association) codes, and the addition of fictional passenger and transferring passenger numbers.

Cost of delay data were generated by Sabre Airline Solutions for flights allocated to each airspace user participant. Several primary and reactionary delay factors were taken into account, such as: missed passenger connections, airport curfews, passenger goodwill, crew connections and maintenance activity.

Each validation scenario was run once only, and so the preferred terms ‘validation run’ or just ‘run’ are used hereafter to denote the run of a unique validation scenario.

The format for each run was:

- introductory briefing (2 minutes)
- performing the validation run (60-75 minutes)
- completing a questionnaire (5 minutes)
- short break (5 minutes)
- group debrief (30 minutes)

C. Airspace User Participant Positions

The platform was configured for four played positions and one reference position. Each played position was presented with only the flights allocated to that airspace user. The reference position was un-played, which included all the flights that weren’t under the control of the four played positions. 54% of the flights in the traffic sample belonged to the reference position and were never subject to any UDPP prioritization action. The reference position served as a means to measure the impact of UDPP prioritization actions by the four played positions on the flights of a non participant. (A non-participant refers to an airspace user who chooses not to use FDA and SFP, or is unable to do so.) The five positions were used in every validation run.

Strategic goals were suggested to participants before beginning the first run, and were requested to adhere to their goals for the entire exercise to provide comparability between runs. See Table II.

TABLE II. SUGGESTED STRATEGIES FOR PLAYED POSITIONS

Position Name	Operator	Type of Operation	Suggested Goal
Hub1	EEE	Short / medium haul operations	Protect single rotations (suspend multiple rotations / low load factors)
Hub2	LLL	Long haul	Protect A380 operations and 747 flights (commercially important and protects airport)

Hub3	XXX	Short haul operations primarily connecting to/from CDG hub with remote regional outstations	Protect single rotations where possible due limited rerouting options
Low Cost	HHH	Point-to-point	Ensure the following day starts on time without impact and aircraft are in place. Suspend flights (will be cancelled) to deliver this protection.
Other Airlines	OA	N/a - the reference position	N/a - the reference position

III. SUBJECTIVE DATA: RESULTS AND ANALYSIS

Airspace user participants were asked to complete a short questionnaire after each validation run, and then take part in a debriefing to discuss the run. The debriefing session allowed any interesting comments or themes from the questionnaires to be discussed in more detail. Some significant material was obtained, as expected, and generally added to and supported the results from the objective data that were recorded. This paper concentrates on the interpretation of the objective data.

IV. OBJECTIVE DATA: RESULTS AND ANALYSIS

A. FDA and SFP Algorithms

Fig. 1 shows the use of FDA priority values, in this case in the ‘FDA only, low visibility’ validation run. Each point on the graph represents a flight that has been assigned an FDA priority value by an airspace user participant. The figure shows that FDA priority values four, five, six and seven were hardly used by airspace user participants. The more extreme values were required to achieve the desired spread of delay across an airspace user’s fleet. Indeed, this was confirmed by verbal feedback from participants during and after the run.

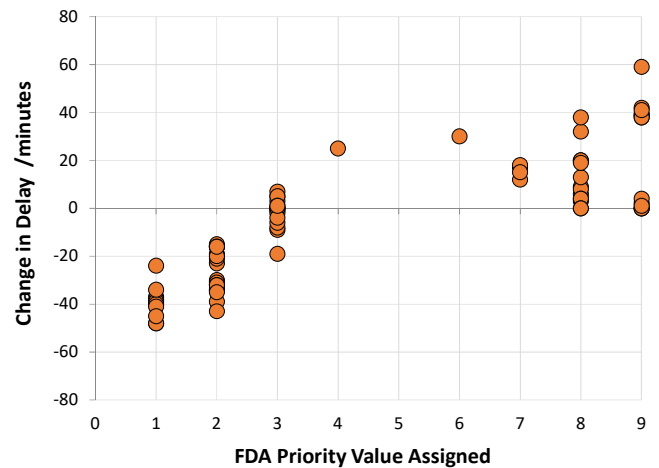


Figure 1. Use of different FDA priority values in the ‘FDA only, low visibility’ validation run, and the subsequent change in delay relative to the baseline delay. (Each point represents a flight that was assigned an FDA priority value at the end of the run.)

The effect of FDA priority value on the change of delay was broadly as expected: the lower the priority value the greater the delay reduction, and vice versa. Other runs that

permitted the use of FDA (either with or without SFP) produced similar results.

Fig. 2 shows the degree of volatility of total delay that was experienced by each airspace participant user during the ‘FDA only, low visibility’ run. The figure shows that the submission of new FDA priority values by one airspace user usually affected his own total delay and the total delay of other airspace user participants. The typical impact was no more than 2% of the total delay, either by increase or decrease. The unplayed (reference) position was hardly affected, which is good news: it is particularly desirable that UDPP non-participants should be protected from the UDPP actions of others.




Figure 2. Change of total delay per airspace user during the ‘FDA only, low visibility’ validation run. (A merge represents the submission of new priorities by one airspace user participant – indicated at the top of the figure.)

Table III shows the number of protections and suspensions carried out in the ‘SFP only, low visibility’ run. Of the 27 protected flights, 26 were assigned zero minutes of delay, and one got 2 minutes of delay. This is a basic check to confirm that the assignment of delay for protected flights was working correctly. Suspended flights received an average delay of 167 minutes, spanning a range from 26 to 317 minutes. The flight with a 26-minute delay was already near the end of the hotspot before being suspended, which explains why its delay was much lower than the average. Flights that were suspended near the start of the hotspot received the highest delays, as would be expected. These basic checks gave some confidence that the SFP algorithm was working correctly.

TABLE III. NUMBER OF SUSPENDED AND PROTECTED FLIGHTS IN THE ‘SFP ONLY, LOW VISIBILITY’ RUN.

Airspace User	Phase 1	
	Flights Protected	Flights Suspended
EEE	15	8
HHH	2	1
LLL	6	3
OA	-	-
XXX	4	2
Total	27	14

Fig. 3 shows the effect that the submission of SFP prioritizations had on airspace user participants. The column headings in the figure indicate which airspace user submitted the prioritizations. The figures shows that carrying out SFP actions increased that airspace user’s total delay (as expected). Whilst this might first seem counterproductive, first consider that UDPP is voluntary and then consider that an airspace user would only use SFP if it was in his operational/financial interest to do so, even if it were to increase his total delay. As we shall see shortly, the delay is less important than the cost of delay.

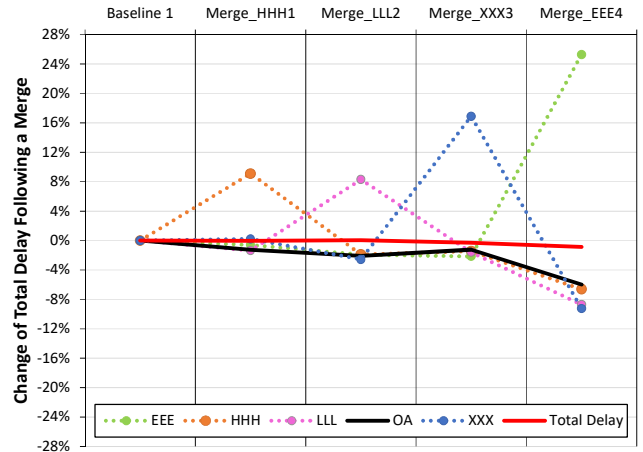


Figure 3. Change of total delay per airspace user during the ‘SFP only, low visibility’ validation run. (A merge represents the submission of new priorities by one airspace user position – indicated at the top of the figure.)

Two interesting observations can be drawn from Fig.3. First, the airspace user ‘OA’ experienced a 6% reduction of total delay, suggests that non-participant airspace users may reap some advantage too from UDPP. Second, and probably of less consequence, is that if SFP use is dominated by just one airspace user, his total delay will increase significantly. Just look at the reduction of total delay experienced by HHH, LLL and XXX airspace user participants when EEE submitted a large set of SFP prioritizations. The logical extension of this is that if SFP is used by many airspace users, and about in the same proportion, the total delays of these SFP users will be brought back down close to their baseline total delays. So SFP does not always have to result in a high increase in total delay for an SFP user because it will depend on the usage of other airspace users too.

B. Punctuality

For departure constraints, departure punctuality is a useful performance area. At least two definitions are widely used in Europe:

- $|AOBT^1 - SOBT^2| < 3$ minutes [3];
- $|AOBT - SOBT| < 5$ minutes [4].

¹ Actual off-block time.

² Scheduled off-block time, i.e., the departure time of a flight as printed on the passenger’s ticket.

Fig. 4 shows how FDA and SFP affect the punctuality of flights in a hotspot. FDA had little effect on improving punctuality, but SFP did bring significant improvement: 20% punctuality compared to about 0% in the baseline using the <3 minutes definition of punctuality. Whilst the data exclude any effects that pre-departure sequencing or final sequencing by the tower may have, nonetheless it does serve as a fair comparison between having SFP and doing nothing. The figures also shows quite nicely the effect on de-prioritized/suspended flights. Generally speaking, UDPP is less concerned with delay and more concerned with cost of delay, as we shall see shortly.

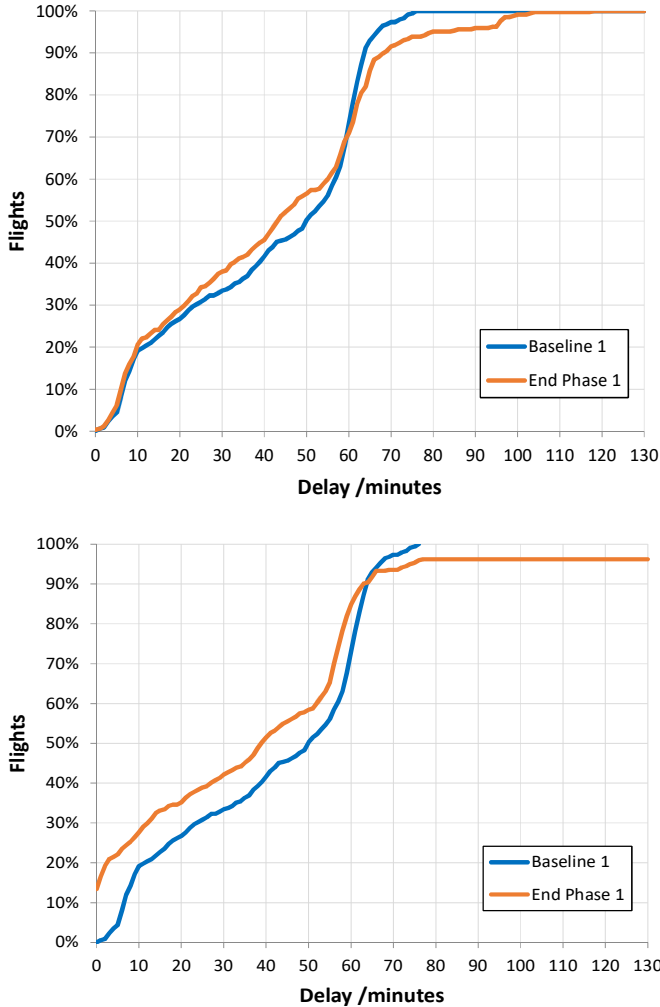


Figure 4. Cumulative percentage of delay for all flights in the low visibility capacity constraint with FDA only (top figure), and SFP only (bottom figure).

C. Cost of Delay

The most significant expected benefit from the UDPP concept is the reduction of cost of delay for airspace users. Depending on the airline, some dispatchers don't have visibility of the costs of delay, and will make operational decisions based on operational objectives. Even so, the operational decisions which guide the behaviour of dispatchers exist to reduce the cost of delay.

For selected validation runs airspace user participants were given the opportunity to prioritize flights using cost of delay information. Each flight in the hotspot had its own, unique cost-delay curve. Airspace user participants were able to prioritize their own flights using this information to minimize their costs of delay.

Fig. 5 shows the total delay and total cost of delay for the XXX (i.e., Hub3) airspace user in the 'FDA only, de-icing' validation run. Similar data were captured and plotted for the other airspace user participants in the same validation run, although are not presented here. The analysis of each leads to the same conclusions.

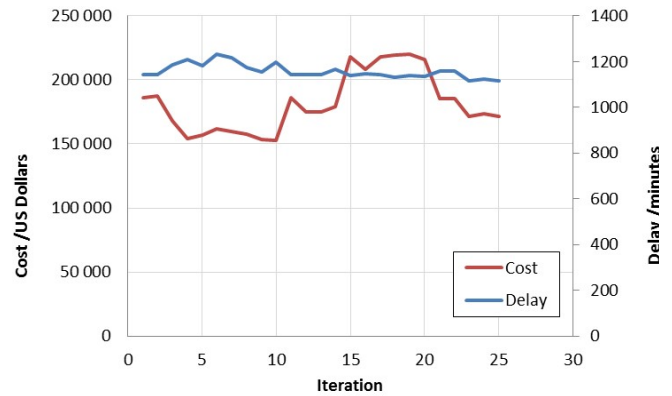


Figure 5. The total cost in US Dollars and the total delay in minutes for the XXX (i.e., Hub3) airspace user participant in the 'FDA-only, de-icing' validation run. The bottom axis shows the number of iterations (updates) of the sequence during the run.

The first observation from Fig. 5 is the decoupling of delay and the cost of delay. (This is the very feature that makes UDPP potentially beneficial.) The total delay for the XXX airspace user decreased by 2% compared to his baseline delay total. The lowest cost of delay achieved in the run was an 18% reduction from baseline (iteration 10 on the bottom axis), but the reduction at the end of the run was 7%. A second, important observation is that there is no smooth downward trend towards lower cost and an eventual minimum, as might be expected for an efficient process. Such a spasmodic graph suggests that a trial and error approach was used to minimize cost, and the airspace user didn't know when to stop prioritizing – it just wasn't obvious, and there wasn't any feedback to assist him.

Fig 6. shows the total delay and total cost of delay for the XXX (i.e., Hub3) airspace user in the 'SFP only, de-icing' validation run. Similar data were captured and plotted for the other played positions in the same validation run, although are not presented here. The analysis of each leads to the same conclusions.

The total cost of delay for the XXX (i.e., Hub 3) airspace user in Fig. 6 reduced by 11% compared to his baseline total delay. The cost of delay drops steadily, and reaches the smallest value at the end of the validation run. The problems of erratic total delay costs that were seen with FDA are not seen with SFP. Both these observations suggest that the process to reduce costs towards the minimum is effective, and was understood by this airspace user participant.

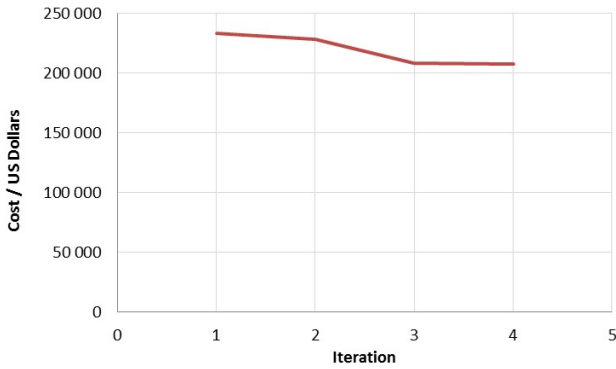


Figure 6. The total cost of delay in US Dollars for XXX (Hub3) airspace user in the ‘SFP only, de-icing’ validation run. The bottom axis shows the number of iterations (updates) of the sequence during the run.

D. Equity

Equity is an important performance area for airspace users taking part in the development of the UDPP concept. Unlike cost-effectiveness and punctuality, equity is a constraint rather than an intended area of performance improvement. A set of guiding principles has been developed and adopted to help guide the development of the UDPP concept [5]. The important question for this exercise arising from the principles was this: does UDPP discriminate against airspace users that choose not to participate?

Fig. 7 shows the effect of UDPP actions by airspace user participants on the flights belonging to the reference position, OA. Both FDA and SFP affected OA’s flights; some flights experienced a delay reduction, others an increase in delay. SFP caused more delay reduction for more OA flights than FDA. When FDA and SFP were available, the spread of delay change increased. At least 70% of OA’s flights in the hotspot either experienced no increase or delay or a reduction in delay. Of the 30% that experienced an increase, half of flights (for FDA, FDA+SFP solutions) had delay increases of no greater than 4 minutes.

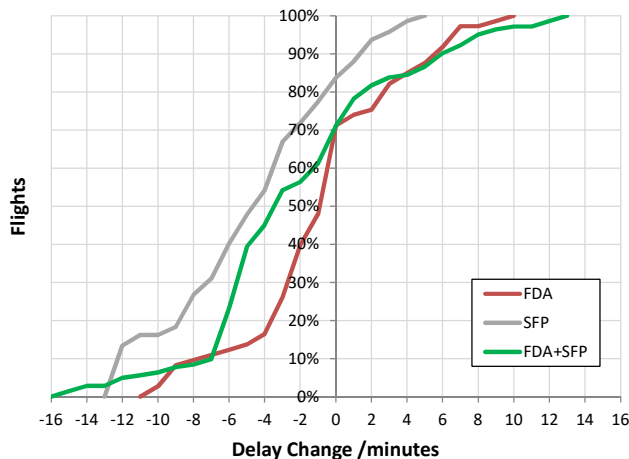


Figure 7. Change in delay for OA flights with the low visibility capacity constraint. (Calculated by comparing the end of phase I delay to the baseline delay for all OA flights.)

V. CONCLUSIONS AND FUTURE WORK

FDA and SFP brought measurable benefits to airspace users under the ideal conditions in this validation exercise, conditions such as the absence of network and airport operations’ constraints. The SFP concept element improved punctuality of flights in the hotspot. This is a subsidiary benefit. Both FDA and SFP brought cost savings to airspace user participants that used UDPP. Cost savings were typically in the range 10-20% of baseline delay cost.

The realism of the cost-delay profiles for flights used in the validation exercise is crucial to measure the cost benefits accurately. This is a difficult aspect because airspace users are reluctant to give up their commercially sensitive cost data. As far as the authors are aware, the cost profiles generated for the exercise are acceptable. The development of cost-delay profiles for future traffic samples will need close scrutiny from airspace users to ensure the required level of realism and accuracy is achieved.

Whilst equity is not a performance area where benefits are expected, it is a constraint that limits what UDPP can do, and is embedded in the FDA and SFP algorithms. The results in terms of equity are encouraging. The results show that, whilst there is an effect on airspace users when another airspace user performs UDPP actions, the effect is generally quite small. Indeed, more often than not the effect is a reduction of delay. SFP in particular had a beneficial effect on others’ flights.

The FDA and SFP algorithms generally behaved as expected. One surprise was that it was difficult to achieve lower costs savings with FDA ‘by hand’. Some automation may be required to support the dispatcher in selecting the most appropriate (cost-effective) set of FDA priority values for his flights in a hotspot. Research is ongoing to assess both algorithms more fully, and to make improvements where necessary.

This paper presents preliminary work. Whilst the results are valid and are useful indicators of potential performance in an operational environment, the UDPP concept must be subject to the full rigour of the validation process [6] to ensure it is fit for purpose. Future validation exercises will see the assessment of the concept in increasingly operationally realistic simulated environments: adding airport processes, introducing normal and abnormal network perturbations, allowing several airports to use UDPP simultaneously. The benefits and disbenefits for actors such as the airport operations centre and flow management controllers will begin to be assessed to compliment the further assessment of the benefits and disbenefits for the airspace user.

ACKNOWLEDGEMENT

This work was co-financed by EUROCONTROL and the SESAR Joint Undertaking (SJU) as part of the SESAR Exploratory Research programme of the European Union. Opinions expressed herein reflect the authors’ views only. Neither EUROCONTROL nor the SJU shall be liable for them or for any use that may be made of the information contained herein. Thanks to Sabre Airline Solutions for developing cost-delay curves for flights in the exercise – not an easy task.

REFERENCES

- [1] Computer assisted slot allocation (CASA), Software Requirements Document, Eurocontrol Network Manager, version 11.005, August 2006.
- [2] <http://www.sesarju.eu/discover-sesar/history>.
- [3] Guidance on KPIs and Data Collection Version 1 (2014), SESAR deliverable D85, Project B.05, edition 00.01.01, 09/12/2014.
- [4] <http://www.eurocontrol.int/sites/default/files/content/documents/official-documents/facts-and-figures/coda-reports/flad-december-2016.pdf>.
- [5] Improved flexibility and equity for airspace users during demand-capacity imbalance, 6th SESAR Innovation Days, hosted by TUDelft, 8-10 Nov 2016. (<http://www.sesarinnovationdays.eu/>)
- [6] European Operational concept Validation Methodology (E-OCVM), Volume I, Version 3.0, 2010. <https://www.eurocontrol.int/sites/default/files/publication/files/e-ocvm3-vol-1-022010.pdf>

AUTHOR BIOGRAPHIES

Stephen Kirby obtained his M.Sc. in Astronautics and Space Engineering from Cranfield University in 1994, and his B.Sc. (Hons) in Astrophysics from University College London (UCL) in 1993. He contributed to and led many safety related projects during his seven years at NATS, and participated in ICAO working groups for the North Atlantic region. He joined EUROCONTROL in 2008, and has spent the last six years working mostly within the European SESAR programme, first coordinating validation activities and then lately validating the user driven prioritization process concept.

Nadine Pilon obtained her PhD in operational research. She has been working for EUROCONTROL for more than 15 years, in safety, network and other areas of research. For the past six years she has been leading the UDPP project within the European SESAR work programme.