2014 UKACC International Conference on Control
9th - 11th July 2014, Loughborough, U.K

Accounting for the Effect of Ground Delay on Commercial Formation Flight

Thomas Kent and Dr. Arthur Richards
Department of Aerospace Engineering
University of Bristol
Bristol, UK

***Abstract*—This paper explores the impact ground delay can have on designing and executing rendezvous operations such as formation flight. For a given formation pairing the route is fixed, speed policies are then generated to compensate for uncertainty in take-off times. Value Iteration is used to solve both the deterministic and stochastic Dynamic Programming problem for an entire state-space. The final optimal policies determine what course of action aircraft should take for any realization of delay. Finally a comparative case study shows that even with delay, formation flights can reach expected savings of 6.1% against flying solo.**

***Index Terms*—Delay, Routing, Formation Flight, Value Iteration, Dynamic Programming.**

I. Introduction

With a predicted future increase in fuel prices and demand for high speed air travel [1], ways to optimize commercial flight are of growing importance. Formation flight, whereby aircraft fly in close proximity to reduce induced drag and decrease fuel-burn, is one way commercial aircraft can save fuel. Recent studies have shown this to be a promising option [2]–[11]. Although many studies show a positive tradeoff between deviating routes, in order to meet up with other formation members, it is clear that any rendezvous would need to be carefully timed and coordinated. Therefore the impact delay could have on such an operation is of major interest.

Delays were estimated to cost the European airline industry 1.25 billion Euros in 2010 [12], with weather and airport operations contributing significantly. Such delays will always be a possibility and any commercial flight is at risk from being affected. However, when trying to design rendezvous operations such as formation flight, timing becomes a significant factor.

With many commercial aircraft flying between 300-450 knots during cruise, missing the rendezvous location by a minute can mean spatially missing it by 10-15kms. Therefore as the level of delay increases so too does the distance required to 'catch-up' in order to reach the formation. Any such catch-up manoeuvre will cause loss of performance compared to the ideal formation flight. The combination of this speed change along with the section of the formation fuel saving 'lost' means that any attempt to regain formation needs to be carefully costed and weighed against other possible solutions.

Delay can occur at any stage of the flight for a number of reasons. A significant proportion however, occurs at the airport (due to factors such as airport congestion [13]) known as ground delay. For the purpose of this paper, we assume that the take-off time is uncertain but that all subsequent operations are perfect. The approach taken could potentially extend to uncertainty in en-route flight such as turbulence, but this is left for future work. The main focus of this work is therefore addressing the impact of a delay in take-off has on other formation members.

Through the use of a state-space model this paper looks firstly to develop a Value Iteration problem. A Dynamic Programming (DP) algorithm is used to solve the deterministic region of the flight (i.e. when all aircraft have taken off). Then through Stochastic Dynamic Programming (SDP) [14] the uncertainty of ground delay is assessed via the assignment of optimal policies for any realization of delay. Finally the results of a case study of 210 transatlantic flights are compared and the results discussed.

II. A State Space Approach

The optimal formation routes are taken to be those found by the authors elsewhere [10], [11] (this previous work also outlines an estimated fuel-burn model to cost formation routes). However, the approach of this paper would still be applicable to any other choice of formation route. In order to reduce the complexity of this problem, formation routes are considered to be fixed, thus once an aircraft commits to a formation it must fly the geographical route regardless. Furthermore, this paper studies only formations of two aircraft while assuming a constant fuel burn discount of 10% during formation [6]–[9]. The concept presented could extend to formations of more than two aircraft, although the curse of dimensionality could cause significant increase in computation time.

A. Formation and Non-Formation States

Fixing the geographical route removes the dimensionality of varying the longitude and latitude locations. Instead only the current distance each aircraft is along its own path is varied. The location of each aircraft is therefore reduced to being implicitly defined by a one-dimension state variable. For the two aircraft formation case, where a Flight F_1 and a Flight F_2 take-off from two distinct airports. We are said to be in a state $(x_1, x_2) \in S$ if Flight F_1 and Flight F_2 are x_1 and x_2 km along their respective paths. The specific longitudes and latitudes can then be recovered from the already defined route. A subspace $SF \subseteq S$, the formation section, is defined

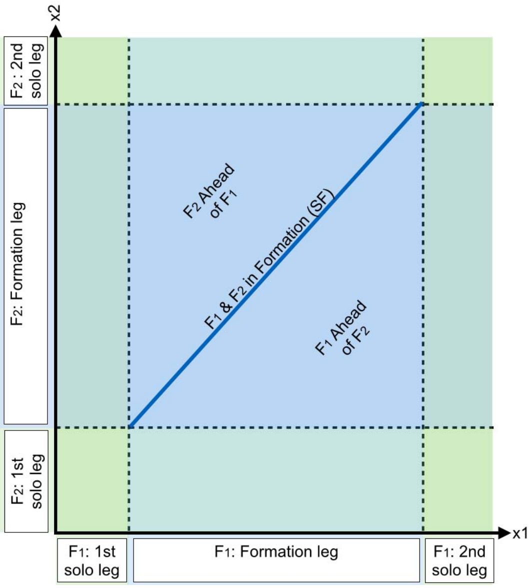

Fig. 1. State Space representation of a formation route and corresponding regions

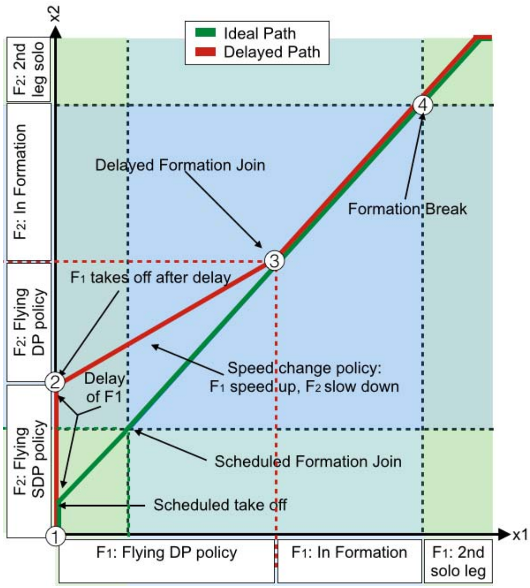

Fig. 2. Solution example with delay. Between 1 and 2 F_2 flies the SDP policy as F_1 is delayed. Between 2 and 3 both aircraft fly the DP policy, to meet at 3 and fly in formation until 4 where the break and fly solo.

to be the possible states that result in both aircraft being at the same geographical location and can therefore fly in formation (The diagonal line in Fig 1). A solution can transition from the non-formation state S to the formation state SF and begin to receive the fuel-reduction benefits of formation flight.

B. Moving Through the State Space

A graphical representation of this space is shown in Fig. 1. From any point in this space, a positive horizontal movement means Flight F_1 has moved along its trajectory and a positive vertical movement means Flight F_2 has moved along its trajectory. Any other movement, is a combination of the two, and so both are travelling along their paths at some given speed. A movement describes a gradient $m = \Delta x_2/\Delta x_1 = V_2/V_1$ defined by the two aircraft velocities V_1 and V_2. Assuming no time constraints, during solo flight each aircraft will fly at its best speed (i.e. the speed that minimizes its total fuel burn). The ratio of the aircrafts' solo-speeds describe a single nominal solo-gradient to move through the state-space m_{solo}. Therefore if it is decided that the aircraft should fly solo, then the path through the state-space will attempt to closely follow this gradient.

If an aircraft misses their rendezvous time (and so are not yet in formation) then this implies one of the aircraft is further along the route than the other. Thus if one aircraft is ahead, by adjusting the speeds (a decrease for the aircraft that is ahead and an increase for the aircraft behind) a 'catch-up' is performed and the formation joins at a later time (as outlined in Fig. 2). These new speeds define a new 'catch-up' gradient m_c, the intersection of the line described by this gradient and the formation state SF is the point at which the 'catch-up' is completed and a rendezvous is made.

For any given state one can define a reachable region, which are the future states which can be realistically reached through the ratio of possible velocities each aircraft can fly. These are defined by the gradients $m_A = \frac{V_1 min}{V_2 max}$ and $m_B = \frac{V_1 max}{V_2 min}$ and are the extremes of the aircraft speeds. The two lines defined by the two gradients m_A and m_B along with the current state (x_1, x_2) define this region,

$$R(x_1, x_2) = \{(x_1', x_2') : m_A \geq \frac{x_2' - x_2}{x_1' - x_1} \geq m_B\}. \quad (1)$$

The shape of this region will be defined by the relative efficiencies of each aircraft. In principle this region extends infinitely, however it is only necessary to look a small amount ahead.

The cost of reaching any point in this space will vary with the aircrafts' velocities. For a given m the velocities V_1 and V_2 of flight F_1 and flight F_2 respectively must satisfy the relationship $V_1 = m \times V_2$. Importantly, this relationship dictates the ratio but not the individual speeds. A one dimensional search over V_1 (or analogously V_2) can find the optimal speed

$$V_1^* \iff V_2^* = \frac{V_1^*}{m}, \quad (2)$$

which satisfies the ratio and minimizes the total cost. These two aspects combined allow for the calculation of the minimum cost to get between all possible reachable states (along with the corresponding speeds).

The following section explores the use of Value Iteration and Dynamic Programming (DP) to find solutions for optimally moving through this state-space.

III. Deterministic Dynamic Programming

The principle of Dynamic Programming (DP) is to reduce a complex problem into a sequence of smaller, simpler, subproblems, working sequentially backwards from an end goal to starting point. That is, if it is known how to get optimally from a state $n-1$ to the final state n then you need only work out how to optimally get from $n-2$ to $n-1$. The starting point can then be reached by recursively working backwards through all the states to arrive at a final solution.

This paper uses the method of Value Iteration, often used in reinforcement learning [15], to solve the DP problem. Starting from the final state and working backwards, at each new state the 'cost-to-go' to all reachable states is calculated. The best course of action is then decided, for example choosing the next state based on a minimal cost. By working backwards (backwards induction) each new state will be assigned a best cost and best action to take based on the previously calculated costs.

The assumption is made that once both aircraft are in the air the problem is deterministic. Therefore the boundaries of the state-space described in section II, (i.e. when $x_1 = 0$ or $x_2 = 0$), are the areas where one of the aircraft has yet to take off. Therefore uncertainty has an impact and so stochastic methods will be required. The interior i.e. when $x_1, x_2 > 0$, however, can effectively be solved via this kind of deterministic DP. Therefore the entire state-space can be split into two subsets:

$$S = \{(x_1, x_2) : x_1, x_2 > 0\}, \tag{3}$$

$$\hat{S} = \{(x_1, 0)\} \cup \{(0, x_2)\}, \forall x_1, x_2, \tag{4}$$

S which can be solved deterministically and $\hat{S}$ which will require the use of a Stochastic DP outlined in the following section. The state $\hat{S}$ is stochastic in nature and so the amount of delay can not be known before it occurs. In section IV it is further divided into cases for each aircraft being delayed, that is, a case for each of the axes x_1 and x_2 of S.

It is useful to note here that the use of a DP for the interior of this state-space is not entirely necessary and other, continuous methods can be used. However this work chooses to use a DP approach to develop a framework which in future could also include the effects of uncertainty en route (for example risks of adverse weather at certain locations).

A. Problem Formulation

For the deterministic formation flight problem the state s consists of how far along the path each aircraft is, the state-space $S \subseteq (0, x_{1max}] \times (0, x_{2max}]$ is then the finite set of possible discrete states s. A control $u \in U$ is the decision of which state $s' = u(s)$ to move to next. The finite set of all applicable controls U is assigned a reachability function $U_r : S \to \mathcal{P}(U)$, where $\mathcal{P}(U)$ denotes the power set of U, and $U_r(s)$ is the set of all controls which can be applied at the state s. The possible choices of $u \in U_r$ are defined by the reachability region R of equation (1). The cost function $C : S \times U \to \mathbb{R}^+$ returns the cost $C(s, u)$ of executing a given control u at a given state s. That is, the fuel the aircraft burn by applying $u(s)$ and moving from state s to s'.

At each step the system is at a distinct state $s \in S$ and can follow out any applicable action $u \in U_r(s) \subseteq U$ for a cost $C(s, u)$. A policy $\pi : S \to U$ is then, a mapping from the state-space S to the space of actions U describing which action to take at each state s to get to the next state s'. Finally, given a set of goal locations G and an initial state s_0 a solution is obtained by finding the policy $\pi \in \Pi$ which is optimal, denoted π^*.

An example of a possible solution from a state $s \in S$ is shown in Fig. 2 where s is point 2. The deterministic problem only applies between the points $2-4$ (i.e when both aircraft have taken off). At point 2 flight F_1 and F_2 apply the control u to fly the best speeds to meet at 3. Then a formation is made and flown between 3 and 4. The path between point 1 and 2 is stochastic and will be covered in section IV.

How well a policy π performs is based on the cost-to-go function J:

$$J_\pi(s) = C(s, \pi(s)) + \sum_{s' \in S} J_\pi(s'). \tag{5}$$

Defined as the cost to reach the next state s', by applying $\pi(s) = s'$ plus the (already calculated) cost to finish from s'. Finally, for any policy to be optimal it must satisfy the Bellman equations [14],

$$\begin{aligned} &J^*(s) = 0 \text{ if } s \in G, \text{ otherwise,} \\ &J^*(s) = \min_{u \in U(s)} \left[C(s, u) + \sum_{s' \in S} J^*(s') \right]. \end{aligned} \tag{6}$$

The corresponding optimal policy $\pi^*(s)$ at a state s can then be deduced from equation (5) by choosing the sequence of controls $u_1, u_2, \ldots, u_k$ such that $J_\pi(s)$ is minimized.

B. Value Iteration

With this in place, Value Iteration is used to solve the Dynamic Programming problem. The concept (first introduced by Bellman [16]) involves an initialized value function (usually zero) which is then iteratively updated with the best currently-found value as the algorithm progresses. That is, at each state s encountered, the minimal value

$$J(s) \leftarrow \min_{u \in U_r(s)} [C(s, a) + \sum_{s' \in S} J(s')], \tag{7}$$

is assigned to the value function for that state to create successively better solutions. Therefore the Value Iteration algorithm starts at a final state and iterates backwards until a sufficient solution is found.

IV. Stochastic Dynamic Programming

When uncertainty is introduced into the problem Stochastic Dynamic Programming (SDP) is needed. Without certainty about the next state, at each stage of the flight there needs to be a best action to take given any possible realization of the uncertainty. Therefore one does not know the absolute best solution, rather a set of policies to follow (based on the best expected outcome). Having costs which are sensitive to probabilistic events can be 'risky' and therefore the underlying solutions can also be sensitive. This methodology is key to trying to put a 'cost' on a possible solution with the intention of eventually making the solutions robust.

TABLE I
EXAMPLE r AND p VALUES FOR FOUR US AIRPORTS

Airport IATA	Late r	Late p	Early r	Early p
BOS	0.755	0.024	2.702	0.365
JFK	0.779	0.021	2.462	0.373
PHL	0.809	0.024	4.164	0.453
SFO	0.825	0.021	3.286	0.417

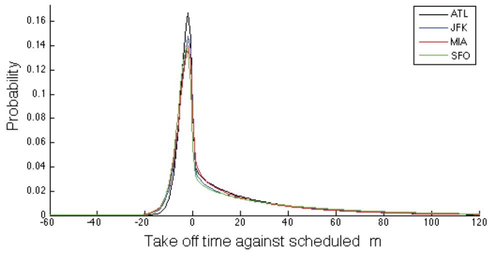

Fig. 3. Example Negative Binomial Distributions of the four US airports in table I

The SDP is very similar to the DP problem. The main difference is that the cost-to-go at each state is minimized over an expected cost and so cannot be guaranteed. It is necessary to first define the stochasticity of the problem using probability density functions.

A. Probability Density Functions of Airport Delay

In order to cost the probabilities of a particular delay occurring Probability Density Functions (PDFs) have been fitted to historical data. For a set of predefined airports the scheduled and actual take-off times were recorded for the month of October 2013. Discrete Negative Binomial Distributions (NBDs) were chosen as a reasonable fit to the data (given more data, better statistical models could also be used).

Given a succession of independent Bernoulli trials, each having a probability of success p and probability of failure $1-p$, then the number of trials needed in order to observe a given number r of successes defines a NBD. For some r and p the NBD is then defined as

$$f(k|r,p) \equiv \mathbb{P}(X=k) = \frac{\Gamma(r+k)}{\Gamma(r)\Gamma(k+1)}(1-p)^r p^k,$$
$$\text{for } k \in \mathbb{N}^0, \text{ where } \Gamma(x) = \int_0^\infty e^{-t}t^{x-1}dt. \tag{8}$$

The Gamma function $\Gamma(x)$ is an interpolation of the binomial coefficients, allowing for non-integer values of r (Table I gives some typical r values).

Some aircraft may also take-off early (early-delay), but not to the same extent, in fact it is common for aircraft to take-off between 0-15 minutes early. The distribution of this early-delay can also be modelled with a NBD, however it usually has a slightly different form to that of 'late-delay'. The resulting final PDF is therefore the normalized combination of two NBDs producing an antisymmetric curve peaking at 0 minutes (scheduled time) and rapidly decreasing at either side. The majority of the probability of take-off lies, as you might expect, within 30 minutes either side of scheduled take-off time, however it is not uncommon for aircraft to be delayed upwards of 60 minutes. Table I contains example r and p values used for the four US airports Boston Logan Int. (BOS), John F. Kennedy Int. (JFK), Philadelphia Int. (PHL) and San Francisco Int. (SFO). The corresponding NBDs are plotted in Fig. 3 for time difference t_D in minutes of observed take-off time against scheduled take-off time for $t_D \in [-60, 120]$.

B. Problem Formulation

As defined in equation (4) the stochastic region of the problem, $\hat{S}$, is constrained to the axes of S. The total flight duration, for each aircraft $i \in \{1,2\}$, is bounded by $T_{i_{min}} = \frac{x_{i_{max}}}{V_{i_{max}}}$ and $T_{i_{max}} = \frac{x_{i_{max}}}{V_{i_{min}}}$. For each aircraft i the stochastic state-space is temporally defined over t as:

$$\hat{S}_i \subseteq \cup_{i\in\{1,2\}}\{x_i \times [T_{i_{min}}, T_{i_{max}}]\}. \tag{9}$$

A state $s=(x,t) \in \hat{S}$ consists of both a position x along a particular axis and the time t it is at that location. A control $u(s)$ from the set of permissible controls $U_r : \hat{S} \to \mathcal{P}(U)$ is then the decision of how long to take to get to the next state s' (the deterministic problem was time invariant and the control was instead spatial). As each of the axes x_i are discretized into N_i spatial states $\{0, x_{i1}, x_{i2}, \dots, x_{iN_i}\}$, spatially we need only look ahead by one step.

Furthermore at each spatial state x_{ik} there is a feasible time range at which the aircraft could be in. This temporal region of reachability at time t is defined by the minimum and maximum velocities of the aircraft, that is $t' \in \left[t + \frac{\Delta x_{ik}}{V_{i_{max}}}, t + \frac{\Delta x_{ik}}{V_{i_{min}}}\right]$.

Finally the stochastic state-space $\hat{S}$ is augmented with a binary state $\delta \in \{0,1\}$ which defines whether or not the other aircraft has taken off yet. Therefore a solution begins in $\hat{S}$ and as soon as the other aircraft takes off the solution transitions to the deterministic solution and the DP solution on S is used. Therefore at each state $s=(x,t,\delta) \in \hat{S}$ a decision $u(t) \in U_r$ is made of how many time-steps t to take to reach the next discrete way point x', if the other aircraft takes-off ($\delta = 1$) the solution transitions to S and the DP is used. Finally given a set of goal locations G and initial state s_0 a solution is obtained by finding policies $\pi \in \Pi$ which are optimal.

The example of Fig. 2 outlines a possible solution path. The stochastic problem is to assign a policy between points $1-2$. As soon as both aircraft have taken off (point 2) the solution is deterministic and the DP can then be used between $2-4$.

Finally in the stochastic problem, the performance of any policy π is determined by its expected cost, governed by a given probability distribution f. Given the current state $s = (x,t,\delta)$ and the control applied gives the next state $u(s) = s' = (x',t',\delta')$. Let us define the probability functional $\mathcal{F}$, that the other aircraft takes off at s', as

$$\mathcal{F}(s,s') = \mathbb{P}(s',\delta) = \begin{cases} 1 & \text{if } \delta = 1, \\ f(t,t'|r,p) & \text{otherwise.} \end{cases} \tag{10}$$

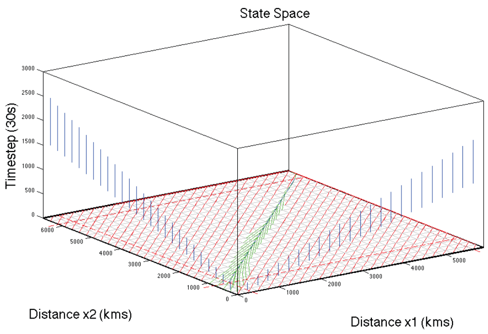

Fig. 4. Combination of deterministic and stochastic state-spaces S and $\hat{S}$.

With this in place the expected value function is defined as,

$$\mathbb{E}(s,s') = \sum_{s'\in S} [\mathcal{F}(s)J_\pi(x',t',1) + (1-\mathcal{F}(s))J_\pi(x',t',0)]. \tag{11}$$

where the corresponding cost to go function is,

$$J_\pi(s) = C(s,\pi(s)) + \mathbb{E}(s,s'). \tag{12}$$

This final Bellman equations are therefore

$$\begin{aligned} J^*(s) &= 0 \text{ if } s \in G, \text{ otherwise,} \\ J^*(s) &= \min_{u\in U(s)} \left[\mathbb{E}(s,s')\right]. \end{aligned} \tag{13}$$

Then for a given formation pair a final solution would consist of two parts. An optimal expected cost to go $J^*(s_0)$, from the initial state s_0, and the corresponding optimal policy π^* to follow until the total delay is realized.

V. Results

This section presents some results obtained by applying the methodology of sections II-IV for pairs of flights wishing to join formation. The DP and SDP have been implemented in Matlab which then decides the optimal policy each pair of flights should follow in order to minimize their expected fuel burn. The resolution of both the spatial and the temporal discretization has a large impact on both the quality of solution and the runtime. The results in this paper use a resolution in space of about 250km and about 30s in time, giving a good balance between run time (around 35s per pair) and solution quality (a choice of speed in increments of around 2 kph) .

A. Single Formation Example - BOSFRA & PHLMXP

The formation between Flight F_1 from Boston Logan International (BOS) to Frankfurt (FRA) airport and Flight F_2 from Philadelphia International (PHL) to Milan Malpensa (MXP) Airport is now explored. It has been chosen as it is a somewhat typical formation flight, achieving a very reasonable 7.35% fuel burn saving against solo flight when there is no delay. Neither of the two flights need to deviate much in order to rendezvous and need only coordinate timings in order to achieve this saving.

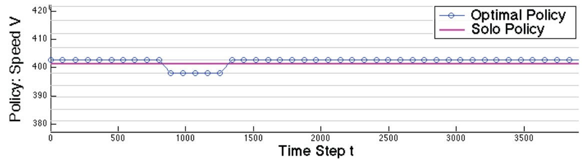

(a) Optimal Policy for Flight F_2 to follow until Flight F_1 takes off

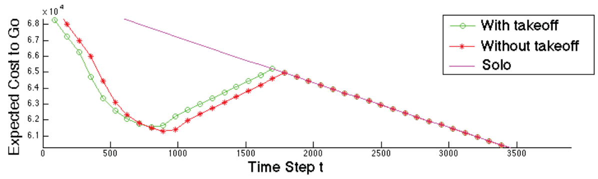

(b) Expected Cost to finish if Flight F_1 takes off at next point or not

Fig. 5. RDP solution for BOSFRA & PHLMXP for Flight F_2

Fig 5(a) shows the optimal 'no take-off' policy, that is, what to do at each step if the other aircraft doesn't take-off. The policy says Flight F_2 is to fly close to its nominal solo speed and then slow down in order to 'wait' for Flight F_1. This slow-down begins (as in Fig. 5(b)) as the expected cost, if F_1 takes off, becomes less than the expected cost if it does not. The policy is largely influenced by the deterministic part of the solution. This is because the window, in which rejoining formation is preferable, is reasonably small (similar to the green region in Fig. 4). The SDP solution policy attempts to stay within this window, by slowing down. The corresponding expected saving of this example is 6.44% (against solo flight) which is a significant saving in the presence of delay.

B. Transatlantic Case Study

In line with previous work by the Authors in [10], [11] a case study is now presented. A data set of 210 transatlantic flights between 26 US and 42 European airports is used. The aim is to compare the assignment of the formation pairings and whether using an expected cost (from the SDP solution) can reduce the impact delay may have on formations.

Firstly for each of the $21,945$ possible pairs (combinations of choosing 2 from 210 flights) a cost is calculated. Given that each aircraft may only belong to one formation, a Mixed Integer Linear Program (MILP), first outlined in [10], is used to optimally assign the aircraft into formation pairs to minimize the total cost.

The cost of each formation (and corresponding policy) is calculated by the SDP. The 'Expected value' cost (EC), is the probabilistically-weighted average of all possible values (i.e. the cost-to-go value J^* of Sec IV). While the 'Best-case' cost (BC) is the cost of a formation following the optimal SDP policy but no delay is realized (so aircraft take off on time). With these two costs, there are three scenarios to explore.

Case 1 - Delay Free: Without delay, there is no need to assign policies via an SDP. Aircraft fly at the best speeds and follow the schedule thus getting the maximum savings. The EC and the BC are therefore identical and around 8.5%

Case 2 - Fleet assignment based on best-case cost: Aircraft

TABLE II
TRANSATLANTIC CASE STUDY RESULTS

Assignment Scenario	Overall average saving (%) Best-case	Expected
Case 1	8.5%	8.5%
Case 2	8.2%	5.7%
Case 3	8.1%	6.1%

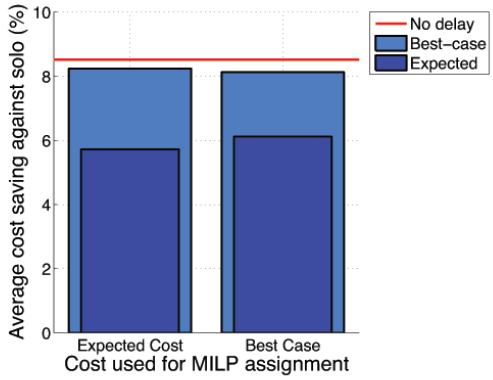

Fig. 6. Comparison of costs used for MILP assignment

fly the policy calculated by the SDP, the assignment of formation pairs is made by trying to attain the best possible saving, that is, by costing formations based on their BC. From this assignment, if all the routes go on to experience no delay, then an average saving of 8.2% (against solo flight) can be achieved (this is slightly lower than completely delay free as they must still follow their policy). However, if delays then occur (based on outlined PDFs) then the average ECs will be around 5.7%.

Case 3 - Fleet assignment based on expected cost: Aircraft fly the policy calculated by the SDP, the formation pairs are costed and assigned based on their EC. If all formations experience no delay then, a slightly lower average saving of 8.1% (against solo flight) can be achieved. However if delays occur (following the PDF), then the average ECs will be around 6.1%.

This highlights the importance of how a formation cost is assigned. As shown in Figure 6 choosing solutions based on the Best-Case scenario could yield better savings, but the overall expected savings are lower.

VI. CONCLUSIONS

This paper has presented a Stochastic Dynamic Program, in order to determine speed-policies for aircraft to follow, to compensate for any realization of delay. Secondly this method has provided a way to evaluate expected costs for proposed formation pairings, allowing fleet re-assigning to account for delays. While the case study results of section V show a reasonable impact to the potential fuel burn saving, formations can still expect to achieve around 6.1% on average.

The most important aspect to take from the results, however, is that how one costs the routes for the MILP assignment can have a large impact to possible savings. Naively assigning formations based on a best-case cost will likely result in suboptimal realizations due to the uncertainty. Future work must therefore include methods for costing formations to increase robustness. For example the work of Ref. [17], [18] explores costing solutions based on a function of their standard deviation. Therefore best-case potential may be sacrificed in order to minimize the variance in the achievable solutions.

ACKNOWLEDGEMENTS

The authors wish to thank the Engineering and Physical Sciences Research Council (EPSRC) for the ongoing support of this research.

REFERENCES

[1] SESAR Definition phase 1: Deliverable 1. In *Eurocontrol*, pages 0–71. Tech. Rep, 2007.
[2] Andrew Ning. *Aircraft Drag Reduction Through Extended Formation Flight*. PhD thesis, Stanford University, 2011.
[3] M Jake Vachon, Ronald J Ray, Kevin R Walsh, and Kimberly Ennix. F / A-18 Performance Benefits Measured During the Autonomous Formation Flight Project. *Instrumentation*, (September), 2003.
[4] Ronald J Ray, Brent R Cobleigh, M Jake Vachon, and Clinton St John. Flight Test Techniques Used To Evaluate Performance Benifits During Formation Flight. *NASATP2002210730*, (August), 2002.
[5] Jia Xu, Simeon Ning, Geoffery Bower, and Ilan Kroo. Aircraft Route Optimization for Heterogeneous Formation Flight. In *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference 20th AIAA/ASME/AHS Adaptive Structures Conference 14th AIAA*, number April, Reston, Virigina, April 2012. American Institute of Aeronautics and Astronautics.
[6] Rachel King and Ashok Gopalarathnam. Ideal Aerodynamics of Ground Effect and Formation Flight. *Journal of Aircraft*, 42(5):1188–1199, 2005.
[7] William Blake and David Gingras. Comparison of Predicted and Measured Formation Flight Interference Effects. *American Institute of Aeronautics and Astronautics*, 41(2):201–207, 2001.
[8] William Blake and Dieter Multhopp. Design, performance and modeling considerations for close formation flight. In *Proceedings of the AIAA Navigation Guidance and Control Conference*, number AIAA-98-4343, pages 476–486, 1998.
[9] P B S Lissaman and C A Shollenberger. Formation flight of birds. *Science*, 168(3934):1003–1005, 1970.
[10] Thomas E Kent and Arthur G Richards. A Geometric Approach to Optimal Routing for Commercial Formation Flight. (August):1–17, 2012.
[11] TE Kent and AG Richards. On Optimal Routing for Commercial Formation Flight. pages 1–11, 2013.
[12] Graham Cook, Andrew J and Tanner. European airline delay cost reference values. *EUROCONTROL Performance Review Unit*, (March), 2011.
[13] Husni R Idris, Bertrand Delcaire, Ioannis Anagnostakis, William D Hall, Nicolas Pujet, Eric Feron, R John Hansman, John-Paul Clarke, and Amedeo Odoni. Identification of flow constraint and control points in departure operations at airport systems. In *AIAA Guidance, Navigation and Control Conference*, 1998.
[14] Dimitri P Bertsekas. Dynamic Programming and Optimal Control 3rd Edition , Volume II by Chapter 6 Approximate Dynamic Programming Approximate Dynamic Programming. *Control*, II:1–200, 2010.
[15] A. Antos, C. Szepesvarf, and R. Munos. Value-Iteration Based Fitted Policy Iteration: Learning with a Single Trajectory. *2007 IEEE International Symposium on Approximate Dynamic Programming and Reinforcement Learning*, 2007.
[16] R E Bellman. *Dynamic Programming*, volume 11 of *RAND CORPORATION. Research studies*. Princeton University Press, 1957.
[17] L.F. Bertuccelli, M. Alighanbari, and J.P. How. Robust planning for coupled cooperative UAV missions. *2004 43rd IEEE Conference on Decision and Control (CDC) (IEEE Cat. No.04CH37601)*, pages 2917–2922 Vol.3, 2004.
[18] Mehdi Alighanbari, Luca F. Bertuccelli, and Jonathan P. How. A Robust Approach to the UAV Task Assignment Problem. *Proceedings of the 45th IEEE Conference on Decision and Control*, (1):5935–5940, 2006.