

MIT Open Access Articles

Cost-Aware Congestion Management Protocols for Advanced Air Mobility

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Balakrishnan, Hamsa. 2022. "Cost-Aware Congestion Management Protocols for Advanced Air Mobility." International Conference on Research in Air Transportation (ICRAT).

As Published: <https://www.icrat.org/previous-conferences/10th-international-conference/papers/>

Persistent URL: <https://hdl.handle.net/1721.1/145406>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Cost-Aware Congestion Management Protocols for Advanced Air Mobility

Victor Qin and Hamsa Balakrishnan
 Massachusetts Institute of Technology
 Cambridge, MA, USA
 {victorqi, hamsa}@mit.edu

Abstract—Air traffic congestion management has traditionally relied on centralized optimization, which may not be practical for large-scale and on-demand applications. The emergence of advanced air mobility motivates the use of prioritization protocols, similar to rules of the road. We propose a cost-aware backpressure prioritization method for air mobility traffic management protocols, based on the second-price auction. We demonstrate using simulations of several advanced air mobility scenarios that our prioritization method increases economic efficiency and fairness across flights and aircraft operators.

Keywords— *advanced air mobility; congestion management; mechanism design; economics of AAM*

I. INTRODUCTION

Market analyses predict that the number of Advanced Air Mobility (AAM) operations will far exceed that of conventional aviation operations [1, 2, 3, 4]. It is generally expected that the scale and density of AAM operations will be such that traditional air traffic management (ATM) paradigms will no longer be sufficient. Conventional air traffic management algorithms aim to achieve efficiency through centralized optimization; however, these approaches do not scale well computationally. Furthermore, most operational concepts assume that AAM traffic will be managed by private service providers within a federated architecture, rather than by an air navigation service provider in a centralized architecture [5].

Several characteristics of AAM operations motivate the use of protocol-based, or rules-of-the-road, approaches to congestion management [6]. The first is, as mentioned before, the scale of demand, which poses a barrier to centralized traffic flow optimization. Secondly, many AAM applications (e.g., urban air mobility, drone deliveries) tend to be on-demand in nature, making long-term planning ineffective. Thirdly, competition between aircraft operators results in an unwillingness to share complete information on flights (e.g., their complete flight plans). However, most prior work on congestion management protocols, both for AAM and road traffic, assume that all flights have equal delay costs. In reality,

an urgent drone delivery flight may have a higher delay cost than a sightseeing tour, and should be appropriately prioritized when a region of airspace (sector) becomes congested. These distinctions between flights, expressed in the willingness to pay, should be considered in order to improve the economic efficiency of airspace allocation.

Auctions offer an effective method of eliciting information useful for flight prioritization, while maintaining the privacy of aircraft operators and efficiently allocating resources. An aircraft operator can signal information on how much they value a flight through their bid, while keeping considerations such as the destination or expected arrival time private. In this paper, we propose the use of auctions as a prioritization mechanism for congested air traffic situations. Specifically, we extend the congestion management protocol proposed in [6] to incorporate auction-based prioritization schemes that account for aircraft operator valuations of their flights.

A. Related work

Market-based approaches have been studied for strategic demand management and tactical deconfliction in the aviation context, including airport slot auctions [7], slot trading during Ground Delay Programs [8], and mobility permits for airspace sector access [9]. More recently, there have proposals to consider auctions and other market-based mechanisms for AAM airspace use [10, 11]. Auctions for congestion management have been studied primarily for road networks, including for congestion pricing in a downtown area [12] and for managing autonomous traffic in an intersection [13]. The latter idea was extended to account for bids from chains of cars with a proportional payment mechanism, along with a “wallet” that controls how cars bid as they traverse their trajectory [14].

Congestion control protocols have been extensively studied in the context of communication networks [15] and road networks [16, 17]. We refer the reader to [6] for a discussion of other examples of congestion control protocols. In particular, [17] used the concept of backpressure to formulate a con-

trol law across multiple intersections. Protocol-based methods have also been used for aircraft trajectory deconfliction using heading and velocity changes [18]. Recently, [6] proposed a congestion management protocol for AAM operations; they however ignored any delay cost variations across flights. In this paper, we augment this congestion management protocol with an auction-based mechanism in order to account for flight delay costs while prioritizing airspace access.

We aim to determine which flight has priority when a sector is contested, i.e., when multiple flights request access to a sector (which can only accommodate one flight) at the same time. There are several properties we desire of the resulting methods:

- 1) *Economic Efficiency*: We want to minimize the sum of cost of delays and maximize weighted throughput of flights throughout the system.
- 2) *Ex post rationality*: Flights should rationally want to participate in the system, and the mechanism should never make a flight worse off (i.e., the operator should not pay more than their valuation for the flight).
- 3) *Fairness*: Costs of delay incurred should be evenly spread across flights in the system. We distinguish between unweighted and economic fairness - they will be more formally defined in IV.

The second-price auction (VCG) satisfies the first two properties, and provides a basis with which to explore the third.

B. Contributions

In this work, we make two key contributions:

- 1) **Chained flight auctions**: We propose a method for building flight bids and running an auction for conflicts across multiple intersections.
- 2) **Cost-aware congestion management protocols**: We account for variable operating costs in our congestion management protocol, which allows us to achieve better economic efficiency.

Using multiple AAM traffic scenarios, we demonstrate that the proposed cost-aware prioritization mechanisms perform similarly in delay and fairness to other prioritization methods, while exhibiting superior performance on metrics of weighted delay and fairness.

Section II presents the problem setup. Section III discusses prioritization mechanisms, building from a basic auction into the full second-price backpressure (SPB) prioritization. Section IV presents results on the performance of cost-aware congestion management protocols in four simulated AAM environments, and discusses the implications. We conclude in Section V with some promising directions for further investigation.

II. FORMULATION AND BACKGROUND

A. Problem

We consider a discrete-time setting on a hexagonal grid of the set of sectors $G = \{s_1, s_2, \dots, s_N\}$ of capacity 1, with a flight being able to move to any sector adjacent to the one it is currently in. We assume that flights do not replan trajectories while in flight. At each time, the only congestion management action considered by a protocol is whether or not to allow a flight into a sector.

Each flight x will be at sector $S_t(x) \in G$ at time $t \in [0, T]$. To continue along their trajectory, flights will bid for the next sector in their path $B_t(x) \in G$, with a bid price $b(x) \in \mathbb{R}$. For ease of notation the subscript may be dropped for $S(x), B(x)$. We also assume that each flight has a fixed cost of operation $p(x) \in \mathbb{R}$, representing the cost per unit time for a flight to operate in the air. This can also be viewed as the “variable cost” in the economic sense, and an aircraft operator should be willing to pay up to $p(x)$ for that flight, as each unit of delay will cost it an extra $p(x)$ over its expected cost.

We assume that flight operators are truthful, that there is no collusion, and that there is no strategic bidding or deconfliction by the operators. This assumption means that we assume $b(x) = p(x)$, $\forall t$; we will use $p(x)$ throughout this paper. Relaxing this assumption is a direction for future research.

B. Protocol-based Congestion Management

The prioritization schemes discussed in this paper are implemented within the congestion management protocol given in [6]. The protocol first prioritized cycles of flights that gridlock the system, then resolves contested sectors by highest backpressure (see III.A for an explanation of backpressure). For each contested sector, the protocol prioritizes all flight requests, then gives access to all flights in priority order until the the sector capacity is reached.

The prioritization methods we discuss can maintain the reduced-information, decentralized principles embedded in the protocol from [6], by only requiring that only the bid amount for flight is passed between sectors.

III. COST-AWARE PRIORITIZATION METHODS

A. Setup

In this section, we illustrate methods for deconflicting flight bids for entering a single contested sector s , with assumed capacity of 1. These results can be shown to be generalized to multiple capacities as well. We define a chain of flights $X = \{x_0, x_1, \dots, x_k\}$ as a group of flights where $B(x_1) = S(x_0), B(x_2) = S(x_1) \dots B(x_k) = S(x_{k-1})$. We can define another chain $Y = y_0, y_1, \dots, y_l$ as contesting sector s with X if $B(x_0) = B(y_0) = s$. We will additionally indicate $p(X) = \sum_{x \in X} p(x)$, to represent the sum of bid prices of all flights in chain X .

There may be multiple contested sectors along a chain of flights. We resolve this by working from the perspective of the highest backpressure sector s^* . Let there be chains of flights $X^1 = \{x_0^1, x_1^1, \dots, x_j^1\}$, $X^2 = \{x_0^2, x_1^2, \dots, x_i^2\}$, \dots , $X^k = \{x_0^k, \dots, x_m^k\}$ bidding for s^* , where the sector of the last flight (eg. $S(x_j^1)$) is the termination of the chain or a subcontested sector. A subcontested sector is defined by the chain terminated at it (eg. $S(x_j^1)$ will be denoted as s^1). Let the chains contesting s^1 be defined as $X^{1,1} = \{x_0^{1,1}, x_1^{1,1}, \dots, x_j^{1,1}\}$, \dots , $X^{1,k} = \{x_0^{1,k}, x_1^{1,k}, \dots, x_m^{1,k}\}$. Further subcontests extend this notation.

For example, in Fig. 1, we can define the chains $X^1 = [0]$ and $X^2 = [2, 6]$ centered on the central green contested sector as s^* . We further define chains $X^3 = [1]$, $X^{3,1} = [3]$, and $X^{3,2} = [4, 5]$, with the subcontest at $S(1)$.

We now discuss the concept of (raw) backpressure. Backpressure is defined as the longest chain of flights queued behind a flight, plus the flight. In the example in Figure 1, Flight 0 has a backpressure of 1 because no flights are requesting its sector. Flight 2 has a backpressure of 2, because Flight 6 is requesting $S(2)$. Flight 1 has a backpressure of 3, because the longest chain directed back to flight 1 is the chain formed by Flights 1, 4, and 5.

This concept of backpressure can be distinguished from the definition of weighted backpressure, where we incorporate flight bids to the backpressure definition. We define weighted backpressure as the sum of bids from flights following a flight plus that flight's bid. For example, in Figure 1 Flight 2 would have a weighted backpressure of 7.

Our protocol implements a modified second-price Vickrey-Clarke-Groves mechanism (referred to as the second-price mechanism [19]) $\chi(X^1, X^2, \dots)$ to determine which flights get to proceed to their desired sector. Winning flights then pay an amount determined by the payment mechanism $\rho_{i,j}(X^1, X^2, \dots)$ for flight x_j^i , using a proportional method and splitting the cost of the winning price by their variable cost (similar to [14]). We will demonstrate this auction mechanism in successively more complex situations over the next few sections, starting with a simple second-price prioritization, introducing proportional payment with chains of flights and weighted backpressure, and culminating in describing a generalized prioritization method that can resolve conflicts between several chains of flights with multiple contested sectors.

B. Second-Price Prioritization (SP), without Backpressure

We first consider the simple case of the second-price mechanism that ignores backpressure, where we only consider the flights adjacent to the contested sector. Let there be chains of flights $X^1 = \{x_0^1, \dots, x_j^1\}$, $X^2 = \{x_0^2, \dots, x_i^2\}$, \dots , $X^k = \{x_0^k, \dots, x_m^k\}$ all attempting to enter sector s . Let $X = \{X^1, X^2, \dots, X^k\}$, $x = \{x_0^1, x_0^2, \dots, x_0^k\}$. We define the choice




Fig. 1: An example of a contested sector. The red sectors contain flights, with (id, bid) denoted in the sector. The green sector in the middle is being contested by 4 continuous chains.

and price mechanism as follows:

$$\begin{aligned} \chi(X) &= \operatorname{argmax}_{x_0^i} p(x_0^i) \\ \rho_{i,0}(x) &= \sum_{j \neq i} p_j(\chi(x/x_0^j)) - \sum_{j \neq i} p_j(\chi(x)) \end{aligned} \quad (1)$$

The choice function selects the winning chain by examining the bids first flight in each chain (the flight adjacent to the contested sector) and choosing the highest bid as the winner. The payment function then selects the second highest price among all first flights in each chain as the winner's payment, while all other flights pay nothing. This can also be seen as a method of prioritizing between chains of length 1, ignoring any backpressure or bids beyond the first flight.

We can use Fig. 1 as an example of the mechanism in (1). The mechanism examines the bids from flights 0, 1, 2. Flight 2 has the highest bid at $B(2) = 3$, while Flight 1 has the second highest bid $B(1) = 2$. Thus Flight 2 pays the second price $\rho(2) = 2$, while $\rho(1) = \rho(0) = 0$. This algorithm is tested as the SECONDPRICE algorithm in simulation.

The above mechanism is straightforward, and maintains many of the positive traits of the VCG mechanism (including efficiency, truthfulness, etc.). However, this ignores delays incurred by flights not adjacent to the contested sector that may be much more serious. In our previous example for instance, not selecting Flight 1 also delays Flight 3, which incurs a very large delay cost. This motivates the following prioritization mechanism, which accounts for weighted backpressure and proportionally distributes costs along the winning chain.

C. Second-Price Prioritization, with Weighted Backpressure and no Subconflicts

In this section we introduce the concept of proportional payment, in order to allow chains to bid together and for flights following in a chain to express their preferences. Let us consider the case when chains of greater than length 1 are bidding for a sector s . Let there be chains of flights $X^1 = \{x_0^1, x_1^1, \dots, x_j^1\}$, $X^2 = \{x_0^2, x_1^2, \dots, x_i^2\}$, \dots , $X^k =$

$\{x_0^k, \dots, x_m^k\}$; assume these chains do not contain subconflicts. To pick the winning chain and price, we choose the chain with the highest sum total of bids and divide the second highest sum total across the winning flights proportionally, where each flight pays a weighted fraction of the winning price. Let $X = \{X^1, \dots, X^k\}$. We define the mechanism as:

$$\chi(X) = \operatorname{argmax}_{X^i} p(X^i) \quad (2)$$

$$\rho_{i,m}(X) = \frac{p(x_m^i)}{p(X^i)} \left(\sum_{j \neq i} p_j(\chi(X/X^i)) - \sum_{j \neq i} p_j(\chi(X)) \right)$$

The choice function in (2) is similar to (1), but instead of only considering the bid of the first flight of the chain, it considers the sum of bids from all flights in a chain (the weighted backpressure). The chain is thus considered one total bidding unit. The payment function is modified so that if chain X^i wins, each flight x_j^i pays a fraction of the second price, where the fraction is defined by their bid $p(x_j^i)$ over the total chain's bid $p(X^i)$. This is shown in Algorithm 1.

Algorithm 1 Second-Price Simplified Backpressure Algorithm

Given: Chains $X = (X^1, \dots, X^k)$

Output: Winning chain X^i , prices $(\rho(x_1^i), \dots, \rho(x_m^i))$

- 1: $C = \operatorname{sort}(X, p(c))$ for $c \in X$
 - 2: $winner = C[0]$
 - 3: $price = p(C[1])$
 - 4: **for** flight $x_j^i \in winner$ **do**
 - 5: $\rho_{i,l} = \frac{p(x_j^i)}{p(winner)} price$
 - 6: **end for**
 - 7: **return** $winner, (\rho(x_1^i), \dots, \rho(x_m^i))$
-

This mechanism is now able to allow flights not directly adjacent to the main conflict to express preferences, and contribute to their chain getting priority and advancing. However, this mechanism is unable to deal with subconflicts that divide and split chains. In the next section, we can continue building on Algorithm 1 by treating every possible chain as its own unit, and resolve conflicts as such.

D. Second-Price Mechanism, with Weighted Backpressure and Subconflicts (SPB)

When we have multiple contests, the resolution of the highest backpressure sector has important implications for which subcontests must be resolved next. For example, if x_0^1 is allowed into s^* and the chain X^1 is allowed to proceed, then a subconflict s^2 at the tail of chain X^2 will not be resolved and we only can consider how to resolve s^1 . Alternatively, if sector s^2 is being contested by flights with high costs of delay (relative to those flights contesting sector s^1), by selecting X^1 we may pick a less efficient group of flights to

move. To resolve these issues, we propose selecting the highest total bid among all possible combinations of chains that could proceed, and then pricing by the second highest combination that does not involve elements of the selected winner. Possible combinations will be a continuous set of chains, such as X^3 and $X^{3,1}$ in Fig. 1; we select the continuous chain with the highest total bid.

Let X represent the set of all chains $X^1, \dots, X^k, X^{1,1}, \dots$, and let α be the collection of superscripts defining an flight's subchain. We can define the mechanism as follows:

$$\chi(X) = \operatorname{argmax}_{X^i, X^j, \dots} p(X^i) + p(X^{i,m}) + \dots \quad (3)$$

$$\rho_{\alpha,j}(X) = \frac{p(x_j^\alpha)}{p(X^j)} \left(\sum_{j \neq \alpha} p_j(\chi(X/X^\alpha)) - \sum_{j \neq \alpha} p_j(\chi(X)) \right)$$

The choice mechanism picks the continuous chain with the highest total bid. If chain X^i is chosen, only subchains $X^{i,m}$ that are contesting the sector at the tail of X^i can also be considered for inclusion into the winning flights. The combination of chains $X^i, X^{i,m}, \dots$ is collectively selected as the winner. The price mechanism is similar to the one presented in (2), except the second price that is being determined is the second continuous chain with the highest total bid $X^k, X^{k,n}, \dots$ that does not contain any element from the winning chain $\chi(X)$. This mechanism is tested as the SECONDBACK method in simulation. See Algorithm 2 for an implementation of (3).

Algorithm 2 Second-Price Backpressure (SPB) Algorithm

Given: Chains $X = (X^1, \dots, X^k, X^{1,1}, \dots)$

Output: Winning chain $(X^i, X^{i,m}, \dots, X^{i,m,\dots})$, prices $(\rho(x_1^i), \dots, \rho(x_j^{i,m,\dots}))$

- 1: $C = \{\}$
 - 2: **for** header chain $X^i \in X$ **do**
 - 3: $C = C \cup \text{AllChains}(X^i, X)$
 - 4: **end for**
 - 5: $C = \operatorname{sort}(C, p(c))$ for $c \in C$
 - 6: $winner = C[0]$
 - 7: $index = 1, price = 0$
 - 8: **while** $price = 0, index < \operatorname{len}(C)$ **do**
 - 9: **if** $C[index] \cup winner = \{\}$ **then**
 - 10: $price = p(C[index])$
 - 11: **break**
 - 12: **end if**
 - 13: $index++$
 - 14: **end while**
 - 15: **for** flight $x_l^{(i,\dots)} \in winner$ **do**
 - 16: $\rho_{(i,\dots),l} = \frac{p(x_l^{(i,\dots)})}{p(winner)} price$
 - 17: **end for**
 - 18: **return** $winner, (\rho(x_1^i), \dots, \rho(x_j^{i,m,\dots}))$
-

Algorithm 3 All Chains Helper Function (*AllChains*)**Given:** Chain X^i , Chains $X = (X^1, \dots, X^k, X^{1,1}, \dots)$ **Output:** Set of chains C^i

```

1:  $C^i = \{\}$ 
2: for subchains  $X^{i,k} \in (X^{i,1}, \dots, X^{i,m})$  do
3:    $C^i = C^i \cup \text{AllChains}(X^{i,k})$ 
4: end for
5: for subchain  $C_{sub}^i = \{X^{i,1}, \dots\} \in C^i$  do
6:    $C_{sub}^i = X^i \cup C_{sub}^i$ 
7: end for
8: return  $C^i$ 

```

Fig. 1 offers an example of the mechanism in (3). The green sector in the middle is being contested by 4 continuous chains $X^1 = [0]$, $X^2 = [2, 6]$, $X^3 = [1]$, $X^{3,1} = [3]$, and $X^{3,2} = [4, 5]$. We pick the chain with the highest total bid, which is $p(X^3, X^{3,1}) = 8$. The price paid by the entire chain is determined by the second highest total bid, which comes from $p(X^2) = 7$. This cost is then proportionally divided across all flights in $\{X^3, X^{3,1}\}$, such that flight 1 pays $\rho_1 = \frac{2}{8}p(X^2) = \frac{7}{4}$, and flight 3 pays $\rho_3 = \frac{6}{8}p(X^2) = \frac{21}{4}$. No other flights pay.

E. Properties

Under assumptions of truthfulness, the SPB mechanism is Pareto efficient in the one-step optimization. A brief sketch can be provided: we can build a tree with the highest backpressure contest sector at the root, each sector as a node, and each sub-contest sector as a branching node. An flight's current sector $S(x)$ is linked to its bid sector $B(x)$, and the price each flight bids is the weight (distance) of that link. SPB is then guaranteed to select the path from leaf to root that is the longest.

While SPB is not strategyproof in total, flight bids are not affected by bids from flights outside of their chain. Other than collusion (which is not within the scope of this paper), flights can only change their payment by underbidding with respect to other flights in their chain, so that they pay a smaller proportion of the price. Future work can improve strategyproofness in the method, although balancing this with efficiency and budget-balance requirements is impossible by the Myerson–Satterthwaite theorem [19, 20].

A possible alternative in the payment mechanism in (3) is to choose the second highest bid among all continuous chains, regardless if subchains are contained within the winning chain. This is possible, but further erodes the strategyproof properties of the second-price method because flights have an incentive to also underbid against other flights not associated with their possible winning chain.

IV. RESULTS

We evaluated our protocol against several other prioritization methods on several simulated traffic environments, inspired by

[6]. We measured delay and delay weighted by flight costs (economic efficiency), and show that the SPB protocol in (3) is economically more efficient and fair across multiple operators. We will first illustrate the different test cases used and define the metrics measured, then explain the results when considering independent flights and when grouping flights with arbitrary operators.

A. Comparison of protocols

We compared the SECONDPRICE and SECONDBACK prioritization methods against several other prioritization methods: RANDOM: A random flight is selected to proceed.

ROUND ROBIN: Flights take turns to enter sectors, with flights that have waited the longest getting priority to proceed, similar to how stop signs operate in road traffic conditions today. Deadlocks are broken randomly.

BACKPRESSURE: The flight with the highest (raw) backpressure (the longest following chain) is allowed to proceed. See [6] for a comprehensive explanation.

The backpressure method has been shown to be Pareto-efficient in the one-step optimization by [6], and is our main point of comparison for these results.

B. Scenario and grid design

We used a 7-radius (169 sector) hex grid for simulation. At time t , the protocol accepts requests bids from flights for sectors, then determines and gives approval to winners to enter their requested sector at time $t + 1$. Flights begin on the “ground”, and request access to the sector directly above their origin location. Once they receive approval, they move into the “air” and proceed to their destination sector. Flights “finish” their trajectory at the end of the timestep t_f of when they enter at their destination sector, freeing up the sector for an flight to enter and occupy at $t_f + 1$.

Trajectories are assumed to be the shortest path between origin and destination, and are given by a straight line from the origin to destination sector. The expected travel time for each trajectory is assumed to be the length of the shortest path, with each flight crossing a sector in one unit of time. Flights are initialized with a random cost of travel $p(x)$ between $[1, 10]$ in every scenario. Operators have the same expected value of average cost of travel across all flights.

Random flight scenario: The random scenario simulates 126 flights split across 3 operators traverse the grid. Origin and destination points were randomly and uniformly drawn across all hex points, and departure times were uniformly drawn from between 0 to 50.

Bimodal flight scenario: The bimodal scenario simulates a scenario where there may be peak demand times and origin/destination locations. Similar to the random scenario, 126 flights split across 3 operators traverse the grid. Origin and destination sectors were determined by assigning every

sector a probability in $[0, 1)$, with all probabilities summing to 1. Flight takeoff times were drawn between $[0, 50]$ with a probability generated by the equation $\mathcal{N}(40, 5) + \mathcal{N}(20, 8)$, where \mathcal{N} is the normal distribution.

Cross-flow flight scenario: The cross-flow scenario studies how protocols resolve a heavy amount of traffic through central sectors. Operators originate from 4 points along the top of the grid, and have 4 possible destinations on the opposite side. This creates a large amount of traffic in the central sectors, where many flights intersect. Departure times were generated using the above equation from the bimodal scenario. We test 3 operators with 30, 30, and 40 flights each.

Hub-and-spoke flight scenario: This scenario represents a package delivery system, where flights originate on the outskirts of the grid and move to destinations across the whole grid. Six operators with 25 flights each start from six origin “warehouses”, with start times determined by a Poisson process and destinations distributed uniformly across the grid.

C. Metrics and Numerical Results

Each scenario was tested with 100 random trials to obtain the results below. For flights, we measured raw and weighted total delay and raw and weighted standard deviation of total delay across all flights. Delay for each flight f is defined as the number of time units above the expected travel time $t_d = t_{arrival} - t_{expected}$. Total raw delay is then the sum of delay across all flights (4a). Standard deviation of raw delay is the standard deviation of delay across all flights (4b), with the total number of flights denoted as $N_{flights}$. These metrics measure the total system efficiency and system fairness, and we aim to minimize both (shown in the second row of Fig. 2).

$$T_{r,d} = \sum_{f \in flights} t_d \quad (4a)$$

$$\sigma_{r,flights} = \sqrt{\frac{\sum_{f \in flights} \left(t_d - \frac{T_{r,d}}{N_{flights}} \right)^2}{N_{flights}}} \quad (4b)$$

The weighted delay weights each unit of delay by the cost of each flight $p(x)t_d$; the total and standard deviation of weighted delay (5) measure economic efficiency and fairness. They are shown in the third row of Fig. 2.

$$T_{w,d} = \sum_{f \in flights} p(x)t_d \quad (5a)$$

$$\sigma_{w,flights} = \sqrt{\frac{\sum_{f \in flights} \left(p(x)t_d - \frac{T_{w,d}}{N_{flights}} \right)^2}{N_{flights}}} \quad (5b)$$

For operators, we start by defining raw and weighted mean delay as the sum of the delay incurred by each flight under that operator, either unweighted or weighted, normalized by

the number of flights under that operator (expressed in (6a)-(6b) respectively). Standard deviation of total delay takes the standard deviation of the mean delays of operators ((6c) and (6d)) and can be understood as a measure for unweighted and economic fairness across operators. We plot (6c) and (6d) in the fourth and fifth rows of Fig. 2 respectively.

$$\mu_{r,op} = \frac{1}{N_{op}} \sum_{f \in op} t_d \quad (6a)$$

$$\mu_{w,op} = \frac{1}{N_{op}} \sum_{f \in op} p(x)t_d \quad (6b)$$

$$\sigma_{r,ops} = \sqrt{\frac{\sum_{op \in ops} \left(\mu_{r,op} - \frac{1}{N} \sum_{op \in ops} \mu_{r,op} \right)^2}{N}} \quad (6c)$$

$$\sigma_{w,ops} = \sqrt{\frac{\sum_{op \in ops} \left(\mu_{w,op} - \frac{1}{N} \sum_{op \in ops} \mu_{w,op} \right)^2}{N}} \quad (6d)$$

D. Discussion

We begin by noting that random, round-robin, and backpressure metrics are similar to those found in [6]. For flights, we can see that while SECONDBACK slightly underperforms BACKPRESSURE in both raw delay and standard deviation of delay, it outperforms BACKPRESSURE in both metrics after weighting by the variable cost of each flight. This makes sense because BACKPRESSURE has been shown to be optimal in the unweighted case in [6], but its cost-agnostic approach leads it to suffer after weighting by variable costs. SECONDPRISE is clustered with the other protocols methods as it ignores backpressure, but it outperforms the ROUND ROBIN and RANDOM protocols after weighting. This shows that adding second-price considerations to the prioritization protocol does indeed improve economic efficiency. Notably, SECONDBACK and SECONDPRISE compared to BACKPRESSURE have high raw standard deviation of delay for the hub-and-spoke scenario (e.g., warehousing and delivery services), but much lower weighted standard deviation of delay.

When measuring operator fairness, we present the summed standard deviation of mean operator delay across protocols. We find that the sum of raw standard deviation of mean operator delays is minimized by the BACKPRESSURE protocol. SECONDBACK performs similarly to BACKPRESSURE. When considering weighted standard deviation, we find that SECONDBACK shows significant improvements compared to BACKPRESSURE in all scenarios considered. This shows the potential of SECONDBACK in ensuring economic fairness.

V. CONCLUSIONS

We present a cost-aware congestion management prioritization method that ensures economic efficiency and reduces flight costs incurred across the system. We demonstrated its compatibility with a proposed AAM congestion management

protocol, and used simulation results to show that it outperforms existing prioritization methods in economic efficiency and fairness, both across all flights and between operators. These prioritization methods account for variable delay costs among flights while maintaining the advantages of prioritization protocols.

The tradeoffs between truthfulness, efficiency, and revenue in the AAM context are interesting questions for future research. The design of incentives (for example, for aircraft operators to replan and adjust their trajectories), and the incorporation of time-varying operator behavior will further increase the ability of cost-aware mechanisms to efficiently prioritize AAM traffic.

REFERENCES

- [1] K. Balakrishnan, J. Polastre, J. Mooberry, R. Golding, and P. Sachs. *Blueprint for the Sky: The roadmap for the safe integration of autonomous aircraft*. Tech. rep. Airbus UTM, 2018.
- [2] Booz Allen Hamilton. *Urban Air Mobility (UAM) Market Survey*. Tech. rep. <https://ntrs.nasa.gov/citations/20190001472>. NASA, Nov. 2018.
- [3] Crown Consulting. *Urban Air Mobility (UAM) Market Survey*. Tech. rep. <https://ntrs.nasa.gov/citations/20190002046>. NASA, Nov. 2018.
- [4] Parker D Vascik, Hamsa Balakrishnan, and R John Hansman. “Assessment of Air Traffic Control for Urban Air Mobility and Unmanned Systems”. In: 10th International Conference on Research in Air Transportation. 2018, p. 9.
- [5] Deloitte Consulting LLP. *UAM Vision Concept of Operations (ConOps) UAM Maturity Level (UML) 4 – v1.0*. Tech. rep. NASA, Dec. 2020.
- [6] Christopher Chin, Karthik Gopalakrishnan, Hamsa Balakrishnan, Maxim Egorov, and Antony Evans. “Protocol-Based Congestion Management for Advanced Air Mobility”. In: Fourteenth USA/Europe Air Traffic Management Research and Development Seminar. 2021, p. 10.
- [7] Michael O Ball, George L Donohue, and Karla Hoffman. “Auctions for the Safe, Efficient, and Equitable Allocation of Airspace System Resources”. In: *Combinatorial Auctions*. Ed. by Peter Cramton, Yoav Shoham, and Richard Steinberg. The MIT Press, Dec. 9, 2005, pp. 507–538.
- [8] Thomas W. M. Vossen and Michael O. Ball. “Slot Trading Opportunities in Collaborative Ground Delay Programs”. In: *Transportation Science* 40.1 (Feb. 2006), pp. 29–43.
- [9] Luca Corolli, Tatjana Bolić, Lorenzo Castelli, and Desirée Rigonat. “Tradable Mobility Permits for the Strategic Allocation of Air Traffic”. In: 6th International Conference on Research in Air Transportation. 2014, p. 8.
- [10] Brent Skorup. “Auctioning Airspace”. In: *North Carolina Journal of Law & Technology* 21 (2019), pp. 79–109.
- [11] Sven Seuken, Paul Friedrich, and Ludwig Dierks. “Market Design for Drone Traffic Management”. In: *arXiv:2110.13784 [cs]* (Oct. 26, 2021).
- [12] Dušan Teodorović, Konstantinos Triantis, Praveen Edara, Yueqin Zhao, and Snežana Mladenović. “Auction-Based Congestion Pricing”. In: *Transportation Planning and Technology* 31.4 (Aug. 2008), pp. 399–416.
- [13] M. Vasirani and S. Ossowski. “A Market-Inspired Approach for Intersection Management in Urban Road Traffic Networks”. In: *Journal of Artificial Intelligence Research* 43 (Apr. 24, 2012), pp. 621–659.
- [14] Dustin Carlino, Stephen D. Boyles, and Peter Stone. “Auction-based autonomous intersection management”. In: *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*. 2013 16th International IEEE Conference on Intelligent Transportation Systems - (ITSC 2013). The Hague, Netherlands: IEEE, Oct. 2013, pp. 529–534.
- [15] Steven H Low, Fernando Paganini, and John C Doyle. “Internet congestion control”. In: *IEEE control systems magazine* 22.1 (2002), pp. 28–43.
- [16] Ayesha Atta, Sagheer Abbas, M Adnan Khan, Gulzar Ahmed, and Umer Farooq. “An adaptive approach: Smart traffic congestion control system”. In: *Journal of King Saud University-Computer and Information Sciences* (2018).
- [17] Jean Gregoire, Xiangjun Qian, Emilio Frazzoli, Arnaud de La Fortelle, and Tichakorn Wongpiromsarn. “Capacity-aware back-pressure traffic signal control”. In: *arXiv:1309.6484 [cs]* (Apr. 22, 2014).
- [18] Inseok Hwang, Jegyom Kim, and Claire Tomlin. “Protocol-Based Conflict Resolution for Air Traffic Control”. In: *Air Traffic Control Quarterly* 15.1 (Jan. 2007), pp. 1–34.
- [19] Yoav Shoham and Kevin Leyton-Brown. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, 2008.
- [20] Roger B Myerson and Mark A Satterthwaite. “Efficient mechanisms for bilateral trading”. In: *Journal of Economic Theory* 29.2 (Apr. 1983), pp. 265–281.

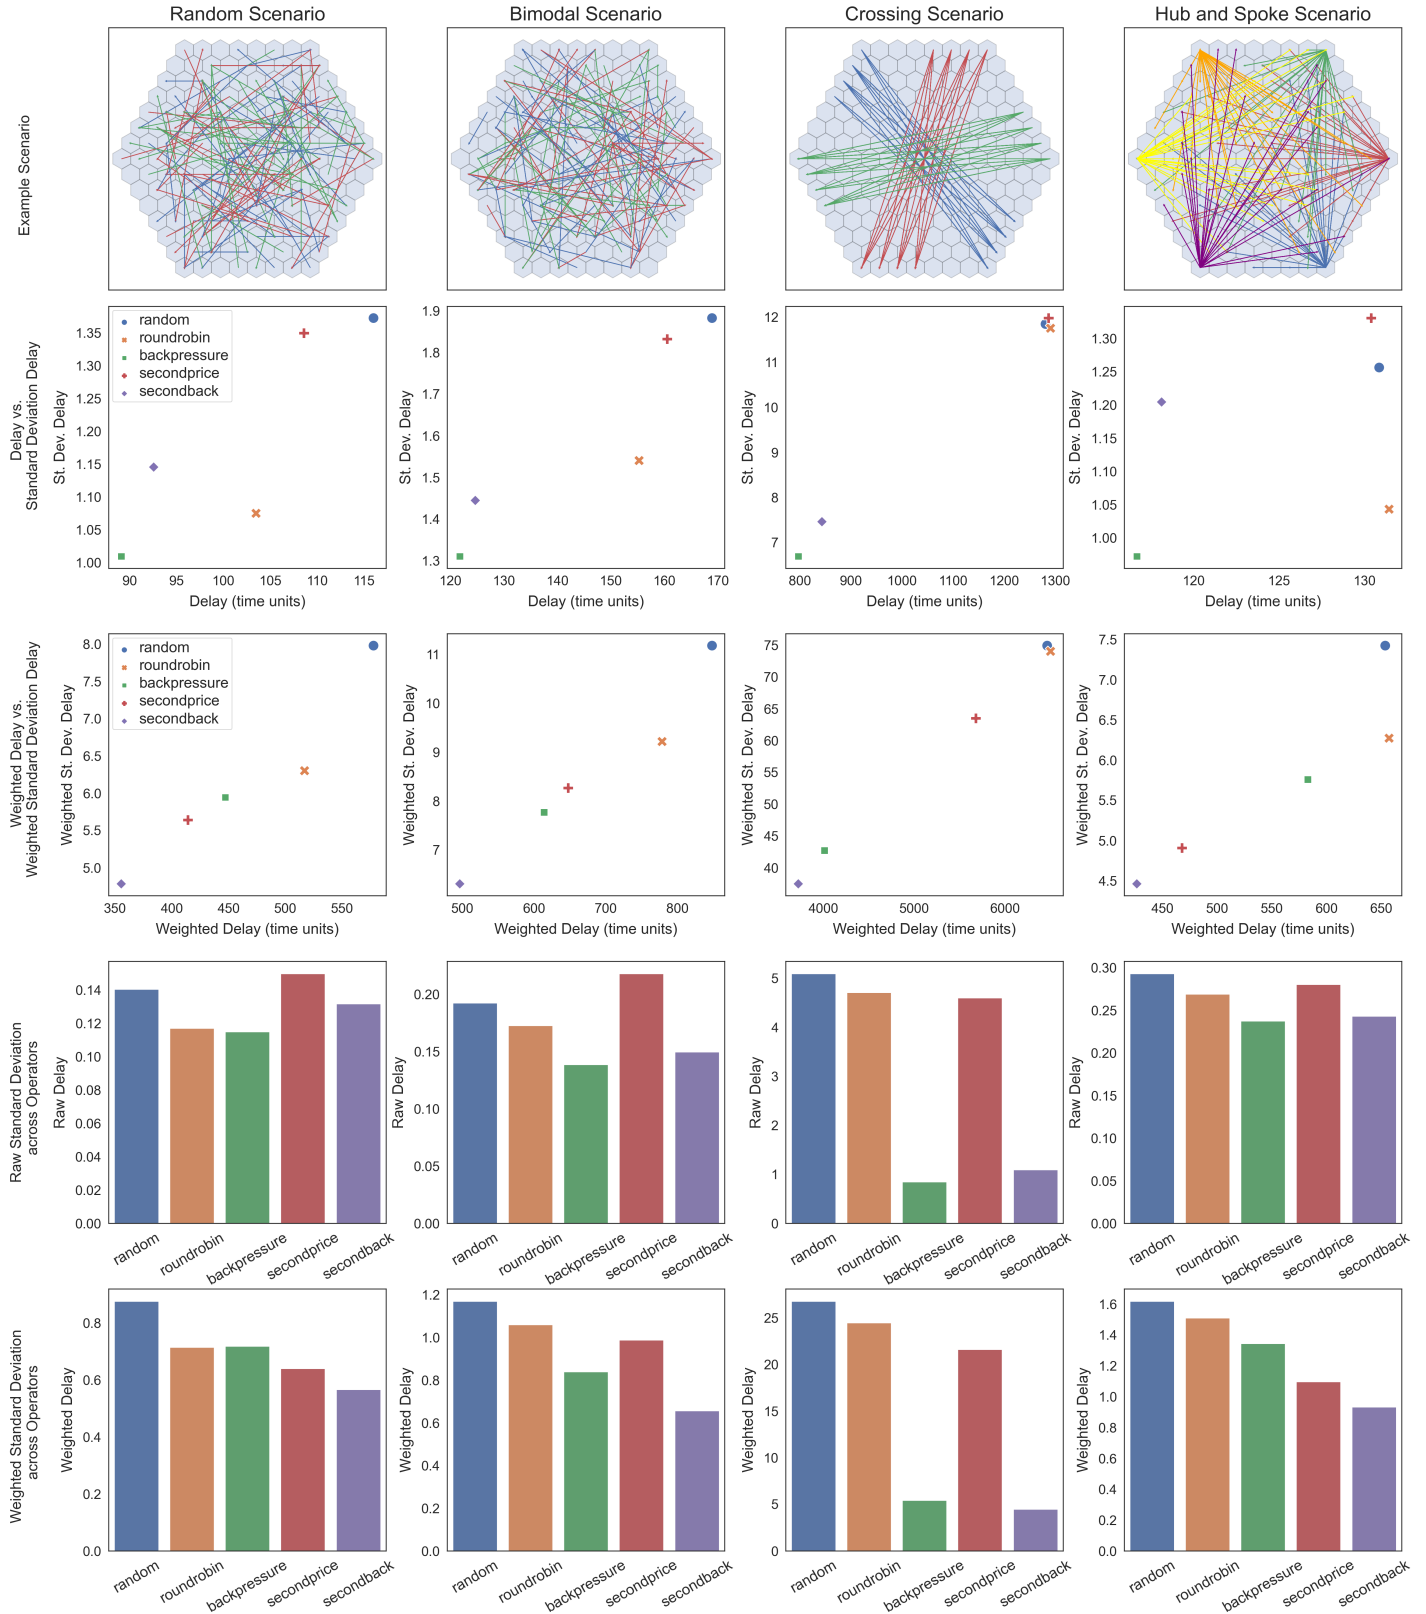


Fig. 2: Simulation results for all scenarios. Each scenario is in a column, with an example presented first. System-wide metrics are on rows 2 and 3 and operator metrics are on rows 4 and 5.