Market Mechanisms for Service Provider Operations in Advanced Air Mobility

by

Victor L Qin

S.B., Harvard University (2021)

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of

Masters of Science in Aeronautics and Astronautics

at the

Massachusetts Institute of Technology

May 2023

Authored by: Victor L Qin
Department of Aeronautics and Astronautics
May 23, 2023

Certified by: Hamsa Balakrishnan
William E. Leonhard (1940) Professor of Aeronautics and Astronautics
Thesis Supervisor

Accepted by: Jonathan P. How
R. C. Maclaurin Professor of Aeronautics and Astronautics
Chair, Graduate Program Committee

Market Mechanisms for Service Provider Operations in Advanced Air Mobility

by

Victor L Qin

Submitted to the Department of Aeronautics and Astronautics
on May 23, 2023, in partial fulfillment of the
requirements for the degree of
Masters of Science in Aeronautics and Astronautics

Abstract

The proliferation of advanced air mobility (AAM) flights in the form of vertical take-off and landing aircraft (VTOL) and uncrewed aircraft systems (UAS) in the near future will require a new air traffic management system adapted for on-demand flights flying at low altitudes and far from existing airports and aviation hubs. The FAA has proposed UAS traffic maangement (UTM) and urban air mobility (UAM) as two concepts of operations for AAM, where private service providers (SPs) will be responsible for managing these novel forms of air traffic alongside but independently from existing air traffic control services. The roles and characteristics of these new SPs are still not well defined today.

In this work, we propose possible methods that can fulfill the concepts proposed. First, we present cost-aware prioritization methods based on the second price auction for air traffic management protocols for use in an SP's internal operations. Next, we show a Shapley value profit-sharing mechanism to incentivize cooperation cooperation in efficiently routing flights between SPs. Finally, we extend the Shapley value framework to accommodate multiple SPs in the same region of airspace, and study how the combination of airspace structure, traffic demand, and sector allocation leads to differences in profit earned between SPs. We conclude with future directions for studying and building service providers in the AAM context.

Keywords: advanced air mobility, Shapley value, service providers, market mechanisms

Thesis Supervisor: Hamsa Balakrishnan
Title: William E. Leonhard (1940) Professor of Aeronautics and Astronautics

Acknowledgments

Many, many thanks to Prof. Hamsa Balakrishnan for being the best advisor one could ask for, and dedicating countless amounts of time night and day to teaching me. Thank you to Geoffrey Ding for your collaboration on so much of this work — so many of the ideas here have some imprint and inspiration from you. Thank you to Chris Chin for walking me through your work when I was just starting, and guiding conversations on where to go next in this journey. Many thanks to to Sydney Dolan and the rest of the DINaMo lab of Siddarth Nayak, Jasmine Aloor, Allan Shtofenmakher, Alissa Chavalithumrong, Kevin Zimmer, Jason Luo, Akila Saravanan, and Daniel Liu for the support throughout the years, from fun lab meetings and dinners to writing advice and countless ideas. Thank you to Ang Li, Lu Dai, Javid Mardanov, Josue Rivera Valdez, and Fanruiqi Zeng for the conversations that contributed to this thesis. Thank you additionally to Prof. Lina Li at Harvard for taking a chance on a sophomore new to the field of controls and robotics, and providing the opportunities that help get me to MIT.

To my friends that remind me to laugh, breathe and make everything feel alright — Hiromu Ryan Rose, Mohib Jafri, Evan Thompson, Kenneth Kim, Olivia Zhang, Liya Jin, Hannah McLaughlin, Lyna Kim, Karen Chan, Annie Lu, Shaan Jagani, Kaila Wong, Adriana Macieira, Kelly Shin, Katie Chiou, Paul Joo, I learn a little more about myself every time I talk to you. Thanks for being there along my journey, and for making my life a little brighter every time we talk. A very special thank you to Tingyu Li, for the best memories of my life, making me a better person and always reminding me that I'm a nerd — I'm finally catching up to your degrees!

Finally, to my sister, Lydia, thank you for listening when I needed to complain and get thoughts off my mind. And most especially, thank you to my parents Peter and Helen for supporting me all this way from the very beginning of my life. You have given countless hours of advice and comfort as I navigated school, and are the people I learn from every day.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The expected proliferation of advanced air mobility (AAM) in the near future will require the coordination of orders-of-magnitude more flights than ever before [5]. Current estimates of the density, type, and number of these new flights [6] have led to the conclusion that the existing air traffic management system is not equipped to effectively manage this emerging demand for airspace resources. While current air traffic management systems focus primarily on fixed-wing aircraft, scheduled flight operations, and airport infrastructure, AAM is expected to include novel vertical take-off and landing (VTOL) aircraft and uncrewed aircraft systems (UAS) flying without fixed schedules and on-demand, with origin and destination locations potentially far away from existing airports. While the focus of research and development efforts have been on the vehicles themselves, there has been limited attention on how large numbers of these vehicles will operate collectively, and how the increase in air traffic will be managed. The novel flight characteristics of these aircraft necessitate the development of novel air traffic management tools and strategies, built to support AAM aircraft and use cases, that will work in conjunction with existing air navigation service providers (ANSPs) to safely and efficiently realize new aerial transport opportunities [1, 7, 8].

The Federal Aviation Administration (FAA) in the United States has proposed two concepts of operations for AAM: UAS traffic management (UTM) for low-altitude aircraft operations [1], and urban air mobility (UAM) for operations of larger cargo- and passenger-

carrying aircraft in "UAM corridors" [7]. In these respective contexts, UAS service suppliers (USSs) and providers of services for UAM (PSUs) enable UAS and UAM operations, working alongside but independently from current air traffic control services. Similar constructs exist elsewhere as well, e.g., in Europe and Japan [8, 9].

The importance of USSs is reflected in the notional UTM architecture in Fig. 1-1; PSUs function in a similar role within UAM [7]. Throughout the remainder of this thesis, unless explicitly stated otherwise, we use "service provider" (SP) to refer to any AAM service provider, thereby encompassing both USSs and PSUs. In general, SPs are expected to support a wide range of aircraft operator needs, ranging from operational planning to communication to traffic management. This work focuses on the last of these services: similar to how the FAA currently provides traffic management services to crewed aircraft, we consider how AAM SPs will provide traffic management support for autonomous aircraft.

While there has been considerable—and justified—focus on the certification and operation of novel aircraft for AAM, the roles of an SP are only loosely defined today. The following list summarizes some envisioned characteristics and responsibilities of service providers:

- SPs will be responsible for the strategic deconfliction (preflight planning to account for anticipated traffic demand and capacity, and other traffic management functions) of AAM flights. By contrast, AAM operators, aided by SPs, will be responsible for the tactical deconfliction (collision avoidance) of flights [7].

- Service providers will support AAM operations through the exchange, analysis, and mediation of information among AAM flight operators, SPs, the FAA, and others. The proposed architecture is a federated network of service providers [7]. Such federated architectures–comprising of connected semi-autonomous components–were first proposed in the context of databases [10, 11], and have since been considered in the context of the Internet [12].

- The network of SPs will enable every AAM flight to traverse through the airspace sectors it needs to access, even if its directly-partnered SP does not manage airspace

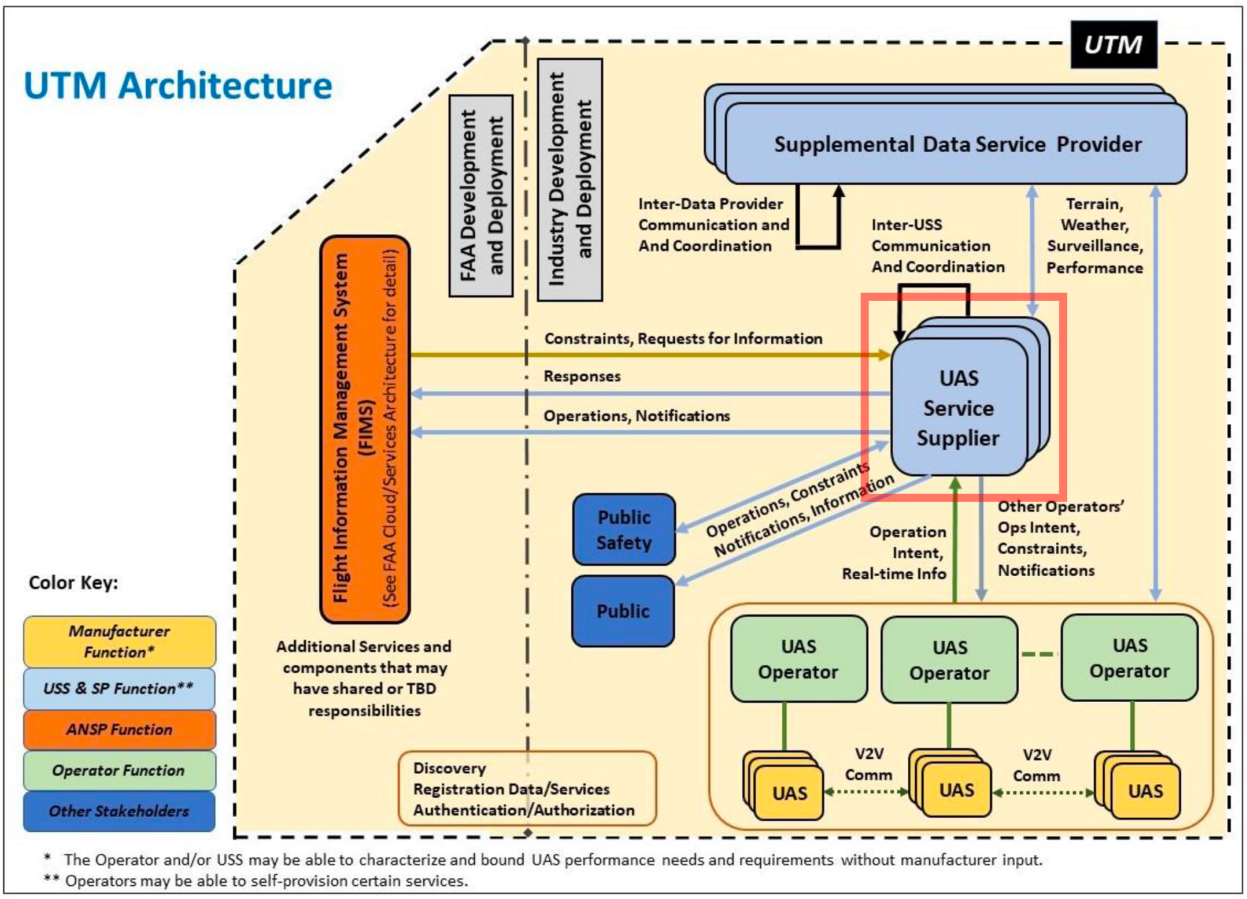

Figure 3. Notional UTM architecture

Figure 1-1: Proposed UTM system architecture (from [1]). Note the central role of the UAS Service Supplier (center right, red box), the SP in this context. It is responsible for coordination between public stakeholders, private operators, data service providers, and the flight information management system (FIMS) which exchanges data with the broader National Aviation System (NAS).

in that sector.

- SPs are expected to be primarily private sector entities, although public sector SPs may also exist [13].
- Multiple SPs may provide services in the same geographical region [1].
- AAM flight operators may also be service providers, as long as they satisfy the relevant qualifications.

These envisioned characteristics are not completely defined, and may be in conflict with each other when applied to real-world scenarios. Some of these characteristics are unprecedented in the context of large-scale physical transportation management. In this work, we identify three distinct problems with SPs as conceived of today, and analyze possible solution for these problems.

First, SPs must handle a large number of autonomous flights. Traditional centralized approaches to air traffic flow management assume that a central entity has complete knowledge of proposed flight trajectories in full. This assumption may not hold for AAM applications, where demand is expected to be more dynamic. While airline schedules are published months in advance and flight plans filed hours before scheduled departure, AAM demand may be known even to flight operators only minutes in advance. AAM flight operators may also hesitate to share trajectories in advance, as they may reveal information (e.g., areas of high demand for an air taxi company) that competitors can exploit. Even if they are willing to share information, operators may not be able to reliably share trajectory information, since it can change as new customers materialize and request flights, among other factors. Finally, mathematical formulations of centralized optimization approaches often scale poorly to large problem instances such as high-demand AAM scenarios [14].

Previous work has proposed the use of protocol-based, rules-of-the-road approaches to congestion management [15]. However, most work on congestion management protocols in the AAM context assume that all flights have equal delay costs, especially considering that private AAM operators would be unwilling to reveal information to competitors. We propose

using auctions as an effective method of eliciting information 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. This is based on our previous work published in [4, 16]. We present the use of auction-based prioritization methods in air traffic management protocols in Chapter 2.

Further issues of privacy and competition arise in the external relationship of an SP to other SPs. If SPs are private entities, the need for cooperation while moving flights between airspace regions managed by different SPs may be in tension with competition for customers among SPs. Even if regulatory frameworks require that SPs cooperate in the movement of AAM flights, SPs may be incentivized to route certain flights in inefficient or unfair ways. Such inefficient emergent behavior has been previously observed in other traffic management context, such as during the growth of the Internet.

To address these concerns and to incentivize collaboration among SPs as they fulfill their envisioned responsibilities, we propose a profit-sharing mechanism using the Shapley value [17]. This method divides the total revenue earned for a flight operation fairly among SPs based on their costs and contributions to the flight. We show that under this approach a SP maximizes its own profit if it routes flights along the globally shortest path, even if the SP incurs a higher individual cost. Furthermore, the SP is incentivized to support other SPs in the presence of congestion. We present our Shapley profit-sharing mechanism for use between SPs in Chapter 3.

The FAA has additionally argued that multiple SPs will operate in the same geographical region. It is unprecedented for two or more air traffic management entities to have control authority in the same region. We present a subdivision approach to airspace that, when combined with the use of our Shapley profit-sharing mechanism, satisfies the the FAA concept of multiple entities in one region while maintaining a clear delineation in traffic management responsibilities. We analyze how different sector geometries and demand profiles through the airspace, along with different allocations of sector responsibilities (and revenues) between

SP, lead to different profit distributions between multiple SPs. While there are few details available on the FAA's concept for how SPs will operate in the same region, we believe that our work offers a working model for the cooperative management of airspace. We show how airspace structure, demand, and allocations of SPs affects profits under the Shapley profit-sharing mechanism in Chapter 4.

Finally, we conclude and offer directions for future work in Chapter 5. This thesis presents a comprehensive conception for new internal and external methods for SPs to handle novel AAM flights. It addresses many of the desired operational properties of SPs proposed by the FAA and provides technical approaches to achieving these characteristics. As the nascent AAM industry grows and as the FAA continues to elucidate the role of SPs in UTM and UAM, we hope this work serves as a foundation for designing better ways to manage and guide flights in our skies.

Chapter 2

Cost-Aware Prioritization Mechanisms for Traffic Management Protocols

2.1 Introduction

Market analyses predict that the number of Advanced Air Mobility (AAM) operations will far exceed that of conventional aviation operations [5, 18–20]. 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 for SPs. Furthermore, centralized optimization methods for air traffic management will be harder to implement in a federated architecture of private service providers [7], where there is no central authority with the power to dictate commands to others.

Several characteristics of AAM operations motivate the use of protocol-based, or rules-of-the-road, approaches to congestion management [15]. Many AAM applications (e.g., urban air mobility, drone deliveries) tend to be on-demand in nature, making long-term planning ineffective. In addition, 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 assumes 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 details such as the destination or expected arrival time private. We propose the use of auctions as a prioritization mechanism for congested air traffic situations. Specifically, we extend the congestion management protocol proposed in [15] to incorporate auction-based prioritization schemes that account for aircraft operator valuations of their flights. This is based on our previous work published in [4, 16].

2.1.1 Related work

Market-based approaches have been studied for strategic demand management and tactical deconfliction in the aviation context, including airport slot auctions [21], slot trading during Ground Delay Programs [22], and mobility permits for airspace sector access [23]. More recently, there have proposals to consider auctions and other market-based mechanisms for AAM airspace use [24, 25]. Auctions for congestion management have been studied primarily for road networks, including for congestion pricing in a downtown area [26] and for managing autonomous traffic in an intersection [27]. 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 [28].

Congestion control protocols have been extensively studied in the context of communication networks [29] and road networks [30, 31]. We refer the reader to [15] for a discussion of other examples of congestion control protocols. In particular, [31] used the concept of back-pressure to formulate a control law across multiple intersections. Protocol-based methods

have also been used for aircraft trajectory deconfliction using heading and velocity changes [32]. Recently, [15] 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. We use the second-price auction (Vickrey-Clarke-Groves) to design this method.

2.1.2 Contributions

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 2.2 presents the problem setup. Section 2.3 discusses prioritization mechanisms, building from a basic auction into the full second-price backpressure (SPB) prioritization. Section 2.4 presents results on the performance of cost-aware congestion management protocols in four simulated AAM environments, and discusses the implications. We conclude in Section 2.5 with some promising directions for further investigation.

2.2 Problem Setup

We discretize space into a set of sectors $\mathcal{S} = \{s_1, s_2, \ldots, s_N\}$, represented by a hexagonal grid. We restrict capacity of each seactor to 1 to avoid the need for tactical deconfliction within a sector; future work could extend the protocol to scenarios with sector capacity greater than 1. We also discretize time into time-steps. In each time-step, aircraft can move to any adjacent sector (each hexagonal sector has up to six adjacent sectors). Alternatively, aircraft may stay in their current sector in the next time-step. An aircraft cannot be forced to leave a sector, and does not deviate from its planned trajectory of sectors. At each time-step, the protocol decides whether or not to allow an aircraft into a sector. We assume that aircraft comply with the protocol, and that aircraft move at the same speed of 1 hexagonal cell in a time-step. An aircraft can only occupy one sector at any time. Figure 2-1 illustrates an example of a grid of sectors. We use the following notation in this chapter to describe the scenario:

$\mathcal{T}$	=	Set of time periods $\{1, \ldots, t, \ldots, T\}$
$\mathcal{S}$	=	Set of sectors $\{1, \ldots, N_{sectors}\}$
V	=	Set of aircraft $\{1, \ldots, N_{aircraft}\}$
$\mathcal{V}_a$	=	Set of active aircraft, i.e., ready to depart or currently airborne
$x(i,t)$	=	Sector for aircraft $i \in \mathcal{V}$ at time $t \in \mathcal{T}$
$\widehat{x}(i,t)$	=	Intended sector at time $t+1$ for aircraft $i \in \mathcal{V}$ based on information at t
$\mathbf{x}(t)$	=	Sectors for all aircraft i at time t
$\widehat{\mathbf{x}}(t)$	=	Intended sectors for all aircraft i at time $t+1$ based on information at t
$\mathcal{G}$	=	Aircraft that can proceed to their next sector
$\mathcal{H}$	=	Aircraft that must hold in their current sector
$del(i)$	=	Total delay assigned to aircraft i
$del(i,t)$	=	Binary variable representing the delay assigned to aircraft i in time-step t

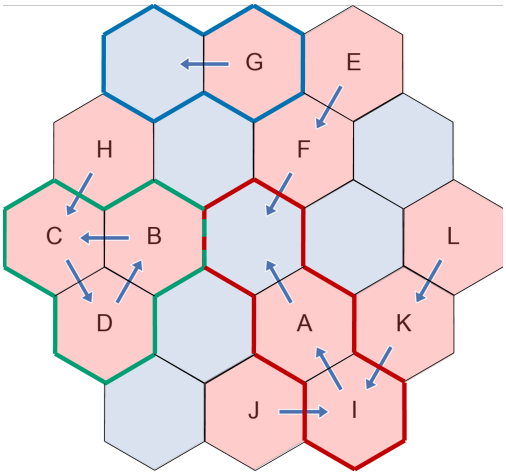

Figure 2-1: Grid setup for our protocol. Each red filled hexagon represents an occupied sector. The base of each arrow indicates the current sector occupied by an aircraft, while the head represents the desired next sector.

2.2.1 Protocol-based Congestion Management

The prioritization schemes discussed in this paper are implemented within the congestion management protocol given in [15]. The protocol first prioritized cycles of flights that gridlock the system, then resolves contested sectors by highest backpressure (see Sec. 2.3.2 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 [15], by only requiring that the bid amount for flight is passed between sectors.

2.2.2 Notation and Setup

To incorporate cost-aware methods, we introduce some more notation to describe *chains* of aircraft and auction mechanisms:

$$\begin{aligned} \hat{p}(i,t) &= \text{Bid for the next sector } \hat{x}(i,t) \text{ by aircraft } i \\ L &= \text{A chain of aircraft } \{k_1, k_2, \ldots, k_m\} \end{aligned}$$

V^s	=	The set of all aircraft involved with a contested sector s, a subset of V
$v_i(L^q)$	=	Valuation function for aircraft i for a certain outcome L^q
$v^l(L^q)$	=	Valuation function for chain l for a certain outcome L^q
$\chi(\mathbf{L})$	=	Choice mechanism that selects a winning aircraft or chain
$\rho_i(\mathbf{L})$	=	Payment mechanism that determines payment for every aircraft

In addition to each aircraft announcing a current sector $x(i,t)$ and a desired next sector $\hat{x}(i,t)$, it also announces a bid for the next sector $\hat{p}(i,t) \in \mathbb{R}$ for $i \in V$.

We define a chain of m aircraft as a set $L = \{k_1, k_2, \ldots, k_m\},\ k_j \in V$ centered on the contested sector s, where $\hat{x}(k_1,t) = s$, $\hat{x}(k_j,t) = x(k_{j-1},t)\ \forall\, j \in \{2,\ldots,m\}$. When there are r multiple chains at the contested sector, $\mathbf{L} = \{L^1, \ldots, L^r\}$, we will specify individual aircraft in a particular chain by $k_j^l \in L^l, l \in \{1,\ldots,r\}, j \in \{1,\ldots,m\}$. Chains are not restricted to all being of the same length. We will define $V^s = \bigcup_{l\in\{1,\ldots,r\}} L^l$ as the set of aircraft involved in the resolution of the contested sector s—that is, the union of aircraft in all chains of $\mathbf{L}$.

We define a valuation function $v_i : L \to \mathbb{R}$ for aircraft i as its valuation of a certain outcome. We modify it with a slight abuse of notation for $v^l, l \in \{1,\ldots,r\}$ to represent the valuation of a certain outcome for chain l. We assume that aircraft bid truthfully, so that $\hat{p}(i,t)$ reflects v_i, and we assume that the bid remains constant for all times t.

$$
\begin{aligned}
v_i(L^q) &= \begin{cases} \hat{p}(i,t) & i \in L \\ 0 & o.w. \end{cases} \\
v^l(L^q) &= \begin{cases} \sum\limits_{k\in L^q} \hat{p}(k,t) & l = q \\ 0 & o.w \end{cases}
\end{aligned}
\tag{2.1}
$$

Each prioritization method will be treated as a mechanism with two parts: a choice mechanism $\chi(\mathbf{L}) : \mathbf{L} \to L$ that determines the chain of aircraft that will proceed, and a payment mechanism $\rho(\mathbf{L}) : \mathbf{L} \to \mathbb{R}_+^{V^s}$ that returns the payment each aircraft has to make. We additionally define an exclusion operator /: given a set of aircraft X and an aircraft k,

X/k is the set X without the aircraft k; if k isn't in the set, $X/k = X$. We overload notation: for the set of chains $\mathbf{L}$, $\mathbf{L}/k = \{L^1/k, \ldots, L^r/k\}$ removes aircraft k from every chain in $\mathbf{L}$.

For example, in Fig. 2-2, there are 4 chains of different color: $L^{orange} = \{A\}$, $L^{purple} = \{C, G\}$, $L^{brown} = \{B, D\}$, and $L^{blue} = \{B, E, F\}$. The set $\mathbf{L} = \{L^{orange}, L^{purple}, L^{brown}, L^{blue}\}$ is the set of all chains; $\mathbf{L}/B$ would remove the aircraft B from both the brown and blue chains.

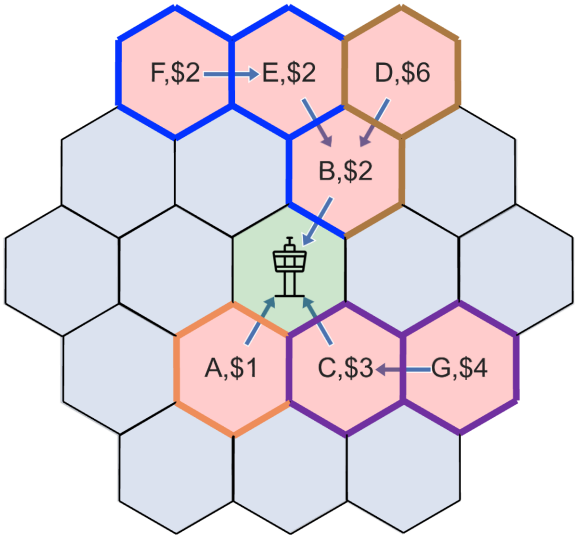

Figure 2-2: Example scenario for cost-aware prioritization. The central sector of conflict is the green sector, with the highest backpressure, contested by 3 aircraft (A, B, and C) and 4 chains of aircraft (orange, purple, blue, and brown).

2.3 Cost-Aware Prioritization Schemes

In this section, we consider cost-aware prioritization schemes for aircraft. Aircraft have varying levels of need and urgency, and a corresponding cost of delays. These delay costs vary due to many different factors, such as environmental concerns, societal factors, operating costs, or others. For example, an aircraft carrying passengers may have a higher value of time than a single delivery drone, while a medical delivery could be more urgent than both. However, these factors cannot be expressed by agents using previously-developed

prioritization methods (for example those in [15]).

To give agents more flexibility in expressing delay costs, we implement cost-aware prioritization methods. There are several properties we desire of the resulting methods:

1. *Economic Efficiency:* The sum of delay costs should be minimized and weighted throughput of aircraft should be maximized throughout the system.

2. *Ex post rationality:* Aircraft should rationally want to participate in the system, and the mechanism should never make an aircraft worse off (i.e., the operator should not pay more than their value of a timestep of delay).

3. *Fairness:* Costs of delay incurred should be evenly spread across aircraft in the system. In this section, we aim for the fair distribution of costs, instead of solely delay.

We use modified versions of the second price auction, which satisfies the first two properties [33] and gives a basis for exploring the third. The second price auction, also known as the Vickrey-Clarke-Groves mechanism, is a well-known result from game theory that distributes a good (in our case, an airspace sector) in the most efficient way. Second price mechanisms select the winner of the auction as the player with the highest bid, but the price paid by the winner is the second-highest bid. This ensures truthful reporting of player valuations in the bid and mechanism efficiency (the highest valuation actually wins), while also being budget-balanced.

To extend auctions to consider aircraft across multiple intersections, we also introduce the concept of proportional payment within the second price auction [27, 28]. While previous work focused only on considering a single intersection, we combine the auction mechanism with backpressure to develop efficient prioritization methods that address flight delay costs for all aircraft in a subproblem.

2.3.1 Second Price

We first consider the simple case of the second-price mechanism that ignores backpressure, where we only consider the aircraft adjacent to the contested sector. We pass to χ and ρ

truncated chains of length one. For contested sector s, let $X = \{i|\hat{x}(i,t) = s, i \in V\}$. Then:

$$\begin{aligned} \chi(X) &= \underset{i \in X}{\operatorname{argmax}}\, \hat{p}(i,t) \\ \rho_k(X) &= \sum_{j \neq k} v_j(\chi(X/k)) - \sum_{j \neq k} v_j(\chi(X)) \end{aligned} \tag{2.2}$$

The choice function selects the winning chain by examining the bids from the first aircraft in each chain (the aircraft adjacent to the contested sector) and choosing the highest bid as the winner. The payment function then selects the second highest price as the winner's payment, while all other aircraft pay nothing (for losing aircraft k, the outcome from $\chi(X/k)$ and $\chi(X)$ is the same). This can also be seen as a method of prioritizing between chains of length 1, ignoring any backpressure or bids beyond the first aircraft. Note that the mechanism is considering X, instead of the set of chains $\mathbf{L}$, similar to how round-robin or random prioritization methods might operate. Only one aircraft will advance when we run this mechanism, and later conflicts will be resolved by using the mechanism for lower backpressure sectors.

In Fig. 2-3, the set of agents considered would be $X = \{A, B, C\}$. Aircraft C has the highest bid, so it advances. During the same time-step, aircraft G would advance as the protocol proceeded to deconflict lower backpressure sectors. Aircraft C pays the price of the second highest bid, which is \$2 from aircraft B.

The above mechanism is straightforward, and maintains many of the positive traits of the VCG mechanism (including efficiency and truthfulness, among others). However, this ignores potentially serious delays incurred by aircraft not adjacent to the contested sector. In our previous example for instance, not selecting aircraft B also delays aircraft D, E, F, which together incur a very large delay cost. This motivates the following prioritization mechanism, which accounts for weighted backpressure and proportionally distributes costs along the winning chain.

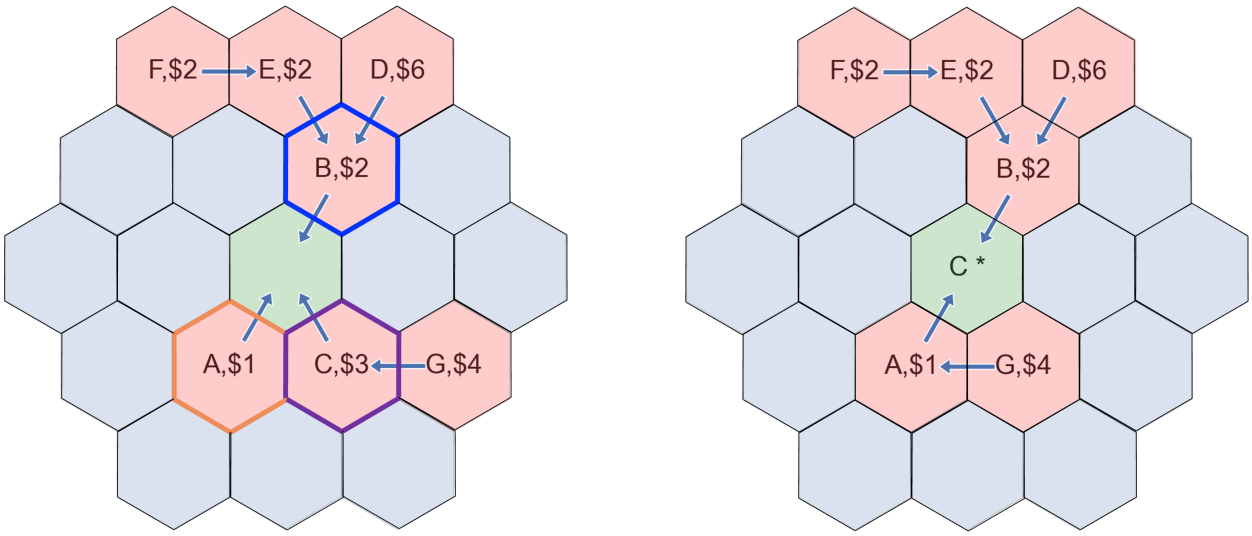

Figure 2-3: Before (left) and after (right) using the second price prioritization method, for the green sector marked. Aircraft C proceeds and has a payment of $2.

2.3.2 Second Backpressure

When we have multiple contests in a subproblem, the resolution of the most important sector has important implications for which subcontests must be resolved next. If we resolve sectors in a sequential manner, this could lead to inefficiencies as less efficient chains of aircraft are moved forward. For example in Fig. 2-3, if we resolve the central conflict at the green sector first, we may make a suboptimal decision eg. moving aircraft C when there are higher cost chains of aircraft. A more efficient mechanism would account for information from a broader set of aircraft to find the best solution for the entire subproblem.

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 on the left (before flight prioritization) in Fig. 2-3, Flight A has a backpressure of 1 because no flights are requesting its sector. Flight C has a backpressure of 2, because Flight G is requesting its current location, $x(C,t) = \hat{x}(G,t)$. Flight B has a backpressure of 3, because the longest chain directed back to flight B is the chain formed by Flights B, E and F.

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 Fig. 2-3 Flight C would have a weighted backpressure of 7.

In this section, we develop the second backpressure prioritization method, which accounts for the preferences of aircraft not immediately adjacent to the contested sector. We do this by adapting the backpressure prioritization method for agent bids—instead of the length of the chain, we consider the total sum of bids from aircraft along that chain. Second backpressure is one-step optimal with respect to agent bids. To handle chains of aircraft together, we use proportional payment from [28]. Proportional payment divides the cost charged to a group of agents proportionally to each agent based the fraction of their bid to the total bid of the group. For example, if an agent bid \$3, and the total group bid was \$8, then that agent would pay 3/8 of the cost to the group. This method allows chains to collectively bid together and then distribute the cost among its agents. We illustrate how subconflicts are also simultaneously resolved using the second backpressure method by creating distinct chains that may share aircraft.

We implement the mechanism as follows: for conflict s, let the list of chains originating at s be $\mathbf{L} = \{L^1, \ldots, L^r\}$. Then:

$$\begin{aligned} \chi(\mathbf{L}) &= \operatorname*{argmax}_{i \in r} \sum_{q \in L^i} \hat{p}(q,t) \\ \rho^i(\mathbf{L}) &= \sum_{l \neq i} v^l(\chi(\mathbf{L}/i)) - \sum_{l \neq i} v^l(\chi(\mathbf{L})) \\ \rho_k(\mathbf{L}) &= \sum_{L^i, k \in L^i} \frac{\hat{p}(k,t)}{\sum\limits_{q \in L^i} \hat{p}(q,t)} \rho^i(\mathbf{L}) \end{aligned} \tag{2.3}$$

There are three parts to this mechanism. First, our choice mechanism is $\chi(\mathbf{L})$, which selects the chain with the highest sum of bids (what we will call the *chain bid*). Next, the payment that the chain as a whole ($\rho^i(\mathbf{L})$) must make is the second highest chain bid. Finally, the payment is divided among aircraft in the winning chain using proportional payment ($\rho_k(\mathbf{L})$, losing chains i have $\rho^i(\mathbf{L}) = 0$). Proportional payment ensures that aircraft that bid more (and thus indicated greater need) are responsible for a greater share of the

payment needed. In practice, the method does the following:

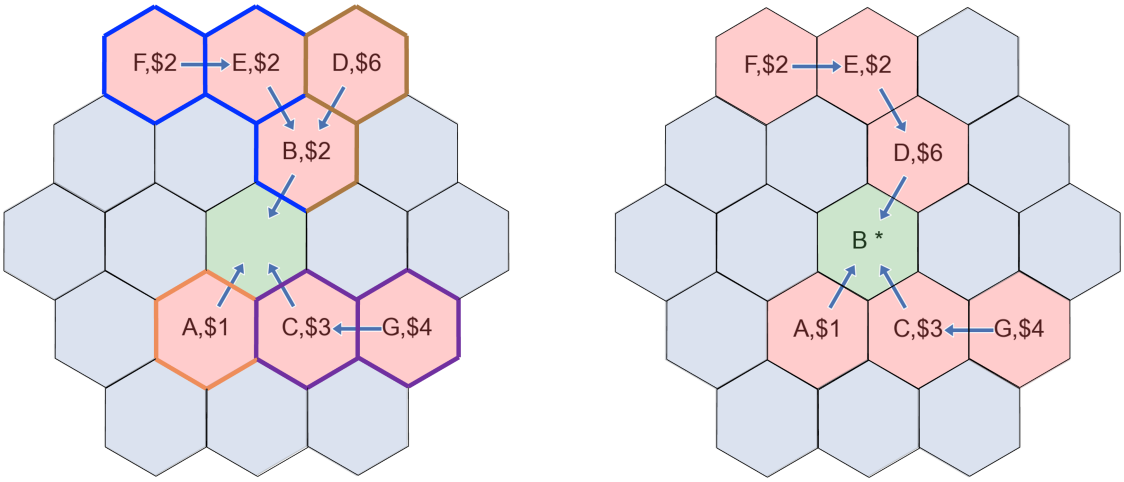

Figure 2-4: Before (left) and after (right) using the second backpressure prioritization method, for the green sector marked. We define 4 chains - orange, purple, brown and blue. The highest chain bid comes from brown, and aircraft B and D move forward with a payment of 7/4 and 21/8 respectively.

1. The central sector under conflict collects a list of all chains $\mathbf{L}$ and their bids $\hat{p}(q,t)\ \forall q \in L^i$, which can be done recursively in a manner similar to the backpressure calculation.

2. The central sector returns $\chi(\mathbf{L})$ and $\rho_k(\mathbf{L})$, which is then disseminated back down the winning chain. In a real implementation, payments can be done through reporting to a centralized third-party handling payment transfers. We for now abstract away the information-sharing constraints.

3. Aircraft in the winning chain advance at the next time-step.

In Fig. 2-4, we consider the sum of bids from the four chains in our example. The brown chain of $\{B, D\}$ has the highest sum total bid of 8, with the second highest price of 7 coming from the purple chain. Thus, aircraft B pays a total of $\rho_B = 7 * 2/8 = 7/4$, and aircraft D pays $\rho_D = 7 * 6/8 = 21/8$.

The second backpressure method is guaranteed to select the most efficient chain of aircraft to proceed in a conflict [4], when flights are weighted by bids. A proof sketch is provided in Appendix A.1; for a more complete proof see [4].

However, it is not incentive compatible. Aircraft are able to achieve better outcomes for themselves by bidding untruthfully at values v_i' different from their actual valuation v_i. Future work using game theory and mechanism design can mitigate these issues and eventually remove the truthful bid assumption we have made here.

2.4 Results

In this section, we demonstrate our protocol and prioritization methods on four simulated traffic scenarios. We compared the SECONDPRICE and SECONDBACK prioritization methods presented in Section 2.3 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 [15] for a comprehensive explanation.

The backpressure method has been shown to be Pareto-efficient in the one-step optimization by [15], and is our main point of comparison for these results. We compare the different prioritization methods on two main metrics: average delay μ (with the goal of minimizing delay) and standard deviation of delay σ across aircraft (with the goal of fairly distributing delay), in both unweighted forms $\mu^{\mathcal{V}}, \sigma^{\mathcal{V}}$ and weighted forms $\mu_w^{\mathcal{V}}, \sigma_w^{\mathcal{V}}$ (with respect to variable cost of delay). We show that backpressure-based methods perform well along these metrics, which are formally defined below:

$$\mu^{\mathcal{V}} = \frac{1}{|\mathcal{V}|}\sum_i del(i) \qquad \sigma^{\mathcal{V}} = \sqrt{\frac{\sum_{i\in\mathcal{V}}(del(i)-\mu^{\mathcal{V}})}{|\,\mathcal{V}\,|}} \tag{2.4a}$$

$$\mu_w^{\mathcal{V}} = \frac{1}{|\mathcal{V}|}\sum_i \hat{p}(i,t)del(i) \qquad \sigma_w^{\mathcal{V}} = \sqrt{\frac{\sum_{i\in\mathcal{V}}(\hat{p}(i,t)del(i)-\mu_w^{\mathcal{V}})}{|\,\mathcal{V}\,|}} \tag{2.4b}$$

2.4.1 Scenarios

In each scenario, we utilize the hex grid setup shown in our examples. Each hex cell s in simulation consists of two sectors: a ground sector that aircraft depart from or land to and an airspace sector connected to all airspace sectors around it. We assume that there is one layer of airborne sectors (i.e., no vertical separation possibilities), but multiple layers could be explored in the future. We use a 7-radius (169 sector) hex grid for simulation. At time t, the protocol accepts requests bids from aircraft for sectors, then determines and gives approval to winners to enter their requested sector at time $t + 1$. Aircraft 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. Aircraft "finish" their trajectory at the end of time-step t_f, which is when they enter their destination sector (a ground sector). At $t_f + 1$, the destination sector is available for other aircraft to use.

Trajectories are assumed to be the shortest path between origin and destination, determined by the hex cells intercepted by a straight line from the origin to destination sector. Each aircraft takes one time-step to traverse one sector. Aircraft are initialized with a random cost of travel $p(x)$ between the integers [1, 10) in every scenario, used for cost-aware prioritization.

We demonstrate our protocol and simulation on 4 scenarios, with varying characteristics in the numbers of aircraft, the origin/destination locations, and the flight schedules:

1. **Random:** 124 aircraft travel across a radius 7 grid (169 cells). Origin and destination points are randomly and uniformly drawn across all hex cells in the grid, and departure times are uniformly drawn from between 0 to 50. This serves as a baseline example of the protocol and prioritization methods in action.

2. **Bimodal:** 126 aircraft travel across the same radius 7 grid. Origin and destination sectors are determined by assigning every sector a probability in $[0, 1)$, with all probabilities summing to 1. Aircraft departure times were drawn between $t = [0, 50]$, with

$p(t) = \mathcal{N}(40, 5) + \mathcal{N}(20, 8)$ normalized so that $\sum_{t=0}^{50} p(t) = 1$, where $\mathcal{N}$ is the normal distribution. This simulates traffic demand over a day, where there may be peak demand times and variation in the popularity of origin/destination locations.

3. **Crossing:** The cross-flow scenario studies how the protocol and prioritization methods handle a heavy amount of traffic through central sectors. 3 “operators" each send aircraft from 4 origins along the top of the grid to 4 possible destinations on the opposite side. This creates a large amount of traffic in the central sectors, where many aircraft intersect. Departure times were generated using the above method from the bimodal scenario. 30, 40, and 30 aircraft were sent respectively from the top left, top, and top right sides towards the opposite sides of the grid.

4. **Hub and Spoke:** This scenario represents a package delivery system, where aircraft originate on the outskirts of the grid and move to destinations across the whole grid. Six “warehouses" on each corner of the grid send out 25 flights each. The start times are determined by a discrete Poisson process with $\mu = 2$, and destinations are randomly and uniformly drawn from all hex cells in the grid.

The results are presented in Fig. 2-5, averaged over 100 samples for each scenario. To measure efficiency and fairness for prioritization methods, we plot total delay (where less is more efficient) and standard deviation of delay across flights in that scenario (where less is more fair). These metrics can be measured in an unweighted form, where each aircraft is treated the same, or a weighted form, where the delay for each aircraft is normalized by the aircraft's variable cost. We expect that cost-aware prioritization methods should perform less efficiently on unweighted metrics, and more efficiently for weighted metrics.

2.4.2 Discussion

We first note that scenarios with more interactions between aircraft have higher delays and standard deviations of delay. This is most pronounced with the crossing scenario, which has the highest delays due to the congestion in the central hex grids. Among the cost-

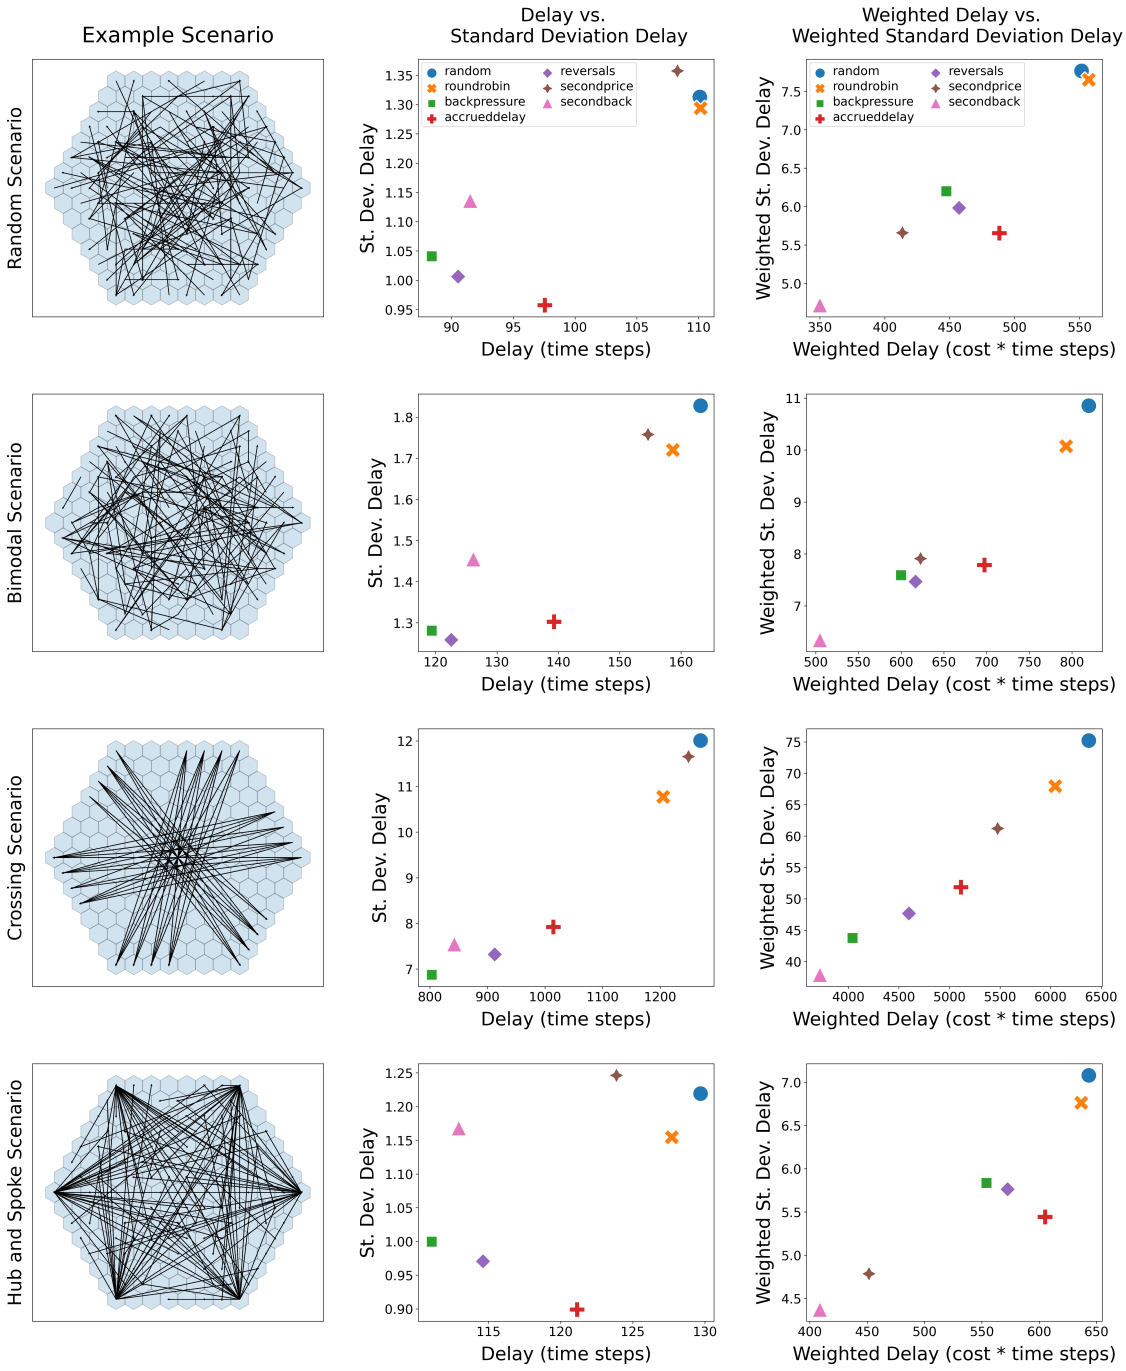

Figure 2-5: Simulation results of the protocol with all prioritization methods. Each scenario is in one row, with system metrics in columns 2 and 3.

agnostic prioritization methods, RANDOM performs the worst, in terms of both delay and standard deviation. ROUNDROBIN performs slightly better than RANDOM in most cases. The three other cost-agnostic prioritization methods (BACKPRESSURE, ACCRUEDDELAY, and REVERSALS) have between 7-35% lower delay and 15-37% lower standard deviation than these naive baselines. BACKPRESSURE consistently has the lowest delays, which makes sense given that it leads to minimal delay in one time-step. BACKPRESSURE also performs well in terms of standard deviation. REVERSALS and ACCRUEDDELAY result in higher delay than BACKPRESSURE, but sometimes have lower standard deviation (e.g., random and hub and spoke scenarios).

For the cost-aware prioritization methods, 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 aircraft. This makes sense because BACKPRESSURE has been shown to be optimal in the unweighted case in [15], but its cost-agnostic approach leads it to suffer after weighting by variable costs. SECONDPRICE is clustered with the other protocols methods as it ignores backpressure, but it outperforms the ROUNDROBIN and RANDOM protocols after weighting. This shows that adding second-price considerations to the prioritization protocol improves economic efficiency. Notably, SECONDBACK and SECONDPRICE 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.

Additional sensitivity analysis was done to study the sensitivity of our protocol and prioritization methods to different congestion levels, by varying the ratio of aircraft in a scenario to sectors. These results are presented in Appendix A.2.

2.5 Conclusion

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. Additional protocol functions (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.

Chapter 3

A Shapley Value Profit-Sharing Mechanism for Service Providers

3.1 Introduction

The FAA has proposed that third-party, likely private SPs will be responsible for managing UTM and UAM traffic and routing flights safely from origin to destination. The private nature of the entities will lead to competition in search of profit even as SPs are required to cooperate to safely route flights through a congested airspace. A SP could refuse to work with another SP if it was financially beneficial. Even if regulation required SPs to cooperate in moving AAM flights, SPs may have incentives to inefficiently route flights to minimize their own costs, or hide information from other SPs to exploit information asymmetry for business purposes. Such inefficient emergent behavior has been previously observed in other traffic management context, such as during the growth of the Internet.

To encourage more cooperation among SPs, we propose a profit-sharing mechanism based on the Shapley value, a game theoretic method for the fair assignment of profit. We develop an understanding of the Shapley value for a continuous space, then show how it can be used with a distance based value function to share profit earned for a flight among SPs responsible for managing various sectors in an airspace.

This chapter is organized as follows. We first give a brief history of Internet traffic routing, and highlight parallels to current AAM traffic management concepts in Section 3.2. We then discuss the risks of allowing AAM traffic management services to evolve in a laissez-faire manner, similar to the evolution of the Internet. In Section 3.3, we describe the Shapley value, a well-established method of fair division of rewards among a coalition of agents. We present a simple airspace system model with SPs in Section 3.4, and analyze how the Shapley value would work in such a system in Section 3.5. Section 3.6 discusses possible challenges with using the Shapley value, as well as open questions that need further investigation. Finally, Section 3.7 concludes this chapter.

3.2 Background

A recent example of traffic service providers in action can be found in the provision of Internet traffic, a federated, decentralized routing system run mostly by private companies. In this section, we draw parallels between the Internet and AAM traffic management, and note some of the challenges and key differences that prevent the direct application of market structures used in the Internet to the AAM context.

3.2.1 The Internet as a model for advanced air mobility

One prominent example of a networked infrastructure that evolved from centralized to decentralized management, and from a public to private service providers, is the Internet. Over the past 25 years, the Internet has grown into one of the most vibrant and innovative parts of society, and a mainstay of our everyday existence. Similar to the proposed AAM architectures [1], the Internet is a collection of federated and decentralized services, with private Internet Service Providers (ISPs) managing different local and regional routes for data packets to travel, as shown in Fig. 3-1. ISPs are independent entities that transport information from many different Internet users and other ISPs. AAM service providers are envisioned to serve a similar role in the airspace context. It is therefore worth tracing the development of

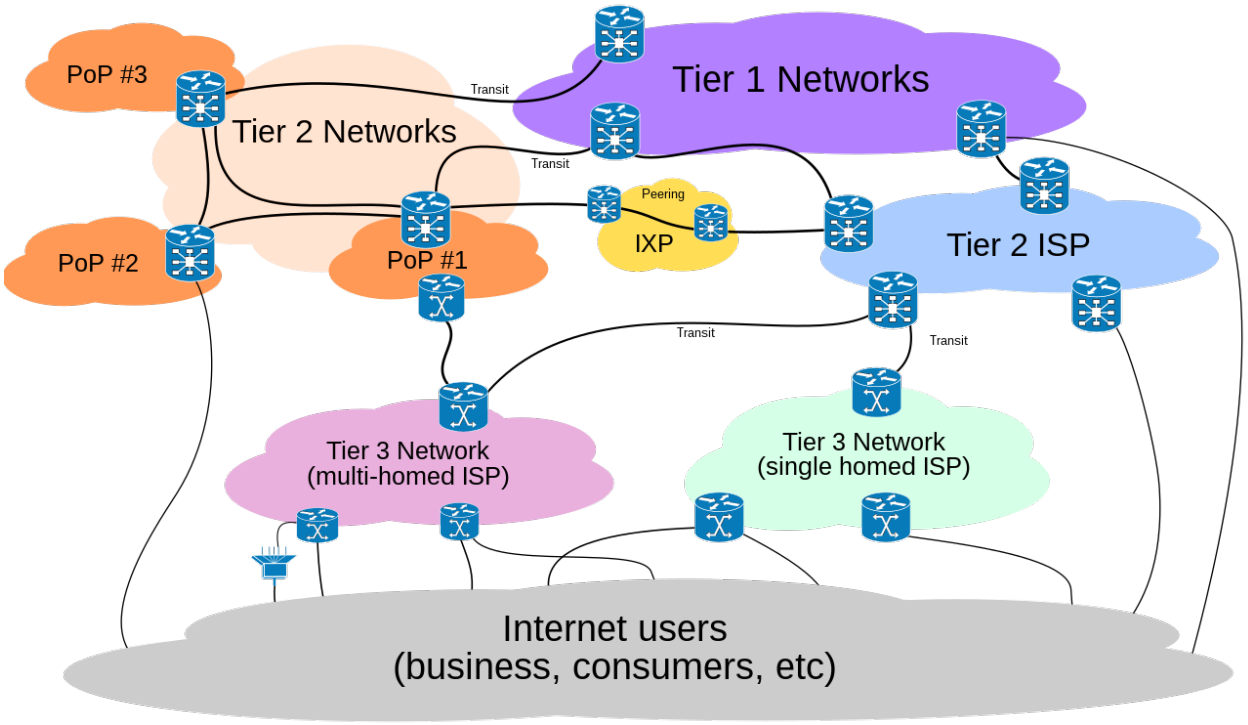

Figure 3-1: An illustration of how Internet Service Providers (ISPs) interact with each other (from [2]).

the Internet to understand how current conceptions of traffic management services for AAM might develop, and to preempt possible problems that may arise.

Parallels to the Internet

The proposed vision for AAM mirrors the development of the Internet, where the responsibility of routing and managing traffic has transitioned from public to private entities. The Internet in the U.S. began with a series of government-funded efforts, culminating in NSFNET, a cross-country Internet backbone supported and operated by the National Science Foundation (NSF) to connect its supercomputers to various research and education networks. Participation in NSFNET came at no cost to institutions, but eventually the use cases and traffic volume of the Internet ballooned to a degree unsustainable through purely government-backed support and growth [34].

In the mid-1990s, companies began developing private fiber-optic networks to carry a

growing volume of commercial Internet traffic, forming the first ISPs. These ISPs grew rapidly, driven by high demand for their services, along with legislation allowing commercial network connections to NSFNET. New private networks, driven by profit motives, expanded rapidly and eventually became the Internet that we know today [35].

There are clear parallels between the growth of the Internet and the forthcoming wave of AAM. At present, air traffic controllers employed by largely-public ANSPs (e.g., the FAA) are responsible for all traffic flow management, much like NSFNET initially formed the backbone of the Internet. This cannot continue, and as the volume and variety of AAM operations increase, conventional ANSPs will not be able to manage all airborne operations. It is envisioned that multiple private entities will step in to form a distributed network of federated service providers to perform routing and other services [1]. Competition between these SPs will result in better quality of service to the AAM aircraft operators who will contract with them. The hope is that privatized SPs will be better able to better adapt to the pace of technological innovation, the increase in flight volumes, and the dynamic, on-demand requirements of the AAM operations.

Challenges with applying the Internet model to AAM service providers

While the Internet has been remarkably successful in connecting the world, a number of key differences between the Internet and aviation contexts means we cannot directly adapt ISP operating paradigms to AAM service providers.

Internet traffic is routed via ISPs and follows a settlement-free peering model. This is a "sender-keeps-all" system in which each ISP only profits from its own customers [36]. The effectiveness of this model hinges on one of two conditions: (i) Traffic in both directions must be approximately equal, or (ii) secret bilateral deals between ISPs must compensate for imbalanced traffic flows. However, both of these conditions are far from guaranteed in the AAM context. While certain types of traffic demand (e.g., commutes) may be approximately symmetrical, traffic from other applications such as drone package delivery is far more likely to be directional (e.g., from a warehouse to customers). If the service provider covering the

vertiport near the warehouse kept all revenue from the drone operator, there would be no incentive for other service providers to cooperate to route flights through the airspaces that they manage. On the other hand, secret bilateral deals between service providers pose a safety concern, as the lack of transparency could create a culture of competition and distrust in inter-SP relations and obfuscate critical SP operations from the FAA. Even for ISPs, these deals have been an occasional source of dramatic disagreements, leading to the complete severing of parts of the Internet: For example, a dispute between Level 3 and Cogent severed 15% of the Internet for three days in 2005 [36, 37]. Such breakdowns of operations would be undesirable for emerging AAM applications.

Once revenue streams are solidified, there are operational concerns with directly using the ISP model for AAM service providers. In the Internet, TCP/IP deals with congestion through the graceful handling of dropped packets: If part of the network is congested, packets are dropped and then retransmitted to improve reliability. In the airspace context, dropping—literally—a flight is a major safety issue and unacceptable in any proposed approach. Instead, SPs will need to manage congestion by cooperating to reroute and delay flights entering and exiting their region, instead of dropping their "buffer."

Furthermore, the sender-keeps-all revenue structure of the Internet incentivizes "hot-potato" routing, in which an ISP passes data along the path of least cost to itself, even if that path may then result in a reduced quality of service for the customer [38]. Among SPs, such routing would lead to inefficiencies such as longer delays and routes traveled, and greater energy consumption. While this may be tolerable in the Internet context, given the very low cost per packet transmitted and the general lack of safety concerns around increased Internet congestion, inefficient routing of aircraft can waste fuel, lead to flight delays, and decrease system safety.

The gradual evolution of the Internet make it more subject to industry inertia and established market structures, and poses a challenge to significant change. By contrast, the forthcoming transformation of the airspace system to support AAM operations presents an unprecedented opportunity for *clean-slate design*, i.e., to implement a novel market structure

determined by the AAM concept of operations and expected behaviors of the emerging demand. By doing so, we can offer innovative solutions that circumvent some of the problems experienced by the Internet, as well others that are unique to the AAM context.

3.2.2 Related work

The early history of Internet pricing and economics is well-covered in [34], which describes some of the basic properties of "sender-keeps-all" economics. [38] gives a deeper explanation of interconnection and Internet structures. Ma et al. outlines the concerns with "hot-potato routing" in [39, 40], with an accompanying solution of profit-sharing based on the Shapley value. We consider how these concepts can be adapted to the context of advanced air mobility.

3.3 The Shapley value

The Shapley value was first described by Lloyd Shapley in [17]. It is a concept from cooperative game theory that provides a way to allocate the value obtained by a collection of agents. Suppose we have a set of agents, N, with $n = |N|$. A subset of agents, $S \subseteq N$, is also called a *coalition*. For every coalition S, $v(S) \to \mathbb{R}$ is the *value* of the coalition; typically, v is determined by the model, game, or environment. We wish to distribute the value accrued by the agents collectively, i.e., $v(N)$, among the agents in some manner, such that each agent i earns a profit share of $\varphi_i(N, v)$. Next, we discuss some desirable properties of profit-sharing mechanisms, and then present the Shapley value, the only mechanism which satisfies all these properties [39].

3.3.1 Desirable properties of a profit-sharing mechanism

We list some desirable properties for a profit-sharing mechanism, centered on

Property 1 (Efficiency)**.** $\sum_{i \in N} \varphi_i(N, v) = v(N)$.

The sum of the values of individual agents equals the total value of all agents. Efficiency ensures that the system does not distribute out more value than it receives, similar to a *budget-balance* property in other fields of mechanism design.

Property 2 (Symmetry)**.** *If* $v(S \cup \{i\}) = v(S \cup \{j\})\ \forall S \in N \backslash \{i, j\}$*, then* $\varphi_i(N, v) = \varphi_j(N, v)$*.*

If the marginal contributions of agent i and agent j to all subsets of agents not including either agent i or agent j is identical, then the shares of profits awarded to the two agents are identical. Symmetry ensures that all agents are treated equally: If the contribution of two agents to a coalition are equal, then the values they each receive are equal.

Property 3 (Additivity)**.** *For two systems* (N, v) *and* (N, w)*, if* $(N, v + w)$ *has the worth function* $(v + w)(S) = v(S) + w(S)$*, then* $\varphi_i(N, v + w) = \varphi_i(N, v) + \varphi_i(N, w)$*.*

The sum of the profits allocated to an agent across two systems equals the profit allocated to the agent in the combined system. In other words, we can calculate the total profit-share allocated to an agent by calculating the profits corresponding to each individual service it provides, and summing them.

Property 4 (Dummy agent)**.** *If agent* i *is a dummy, where* $v(S \cup \{i\}) - v(S) = 0\, \forall S \subseteq N \backslash i$*, then* $\varphi_i(N, v) = 0$*.*

An agent that does not add any value to any coalition is allocated a profit-share of zero. This property ensures that if an agent does not contribute to the system, it does not receive anything from the profit-sharing mechanism.

3.3.2 Computation of the Shapley value

The Shapley value represents the average marginal contribution of an agent to a set of agents and is computed as follows. Let Π be the set of all permutations of N, representing all possible orderings of coalition formation for the agents in N; as such, $|\Pi| = n!$. When required, a specific permutation $\pi \in \Pi$, i.e., a specific ordering of coalition formation, will be written in

parentheses (e.g., (A, C, B, D)). As we wish to find the average *marginal* contribution of an agent i, let p_π^i be the set of agents that strictly precede agent i in permutation π. Note that $i \notin p_\pi^i \ \forall \pi$ and $p_\pi^i = \emptyset$ if and only if i is the first agent in π. Then, the Shapley value of agent i in a group of agents N under value function v, $\varphi_i(N, v)$, is given by

$$\varphi_i(N, v) = \frac{1}{n!} \sum_{\pi \in \Pi} \left[v(p_\pi^i \cup \{i\}) - v(p_\pi^i) \right]. \tag{3.1}$$

In situations where N and v are clear from context, we drop the arguments and denote the Shapley value of agent i as φ_i for brevity. The Shapley value has the aforementioned desirable properties; these properties make it effective as a method of attributing agent contributions to various coalitional games, especially in cooperative network settings.

3.3.3 An illustrative example

We demonstrate the computation of the Shapley value using the example environment in Fig. 3-2. Here, we have $n = 4$ agents, each one representing a sector, given by $N = \{1, 2, 3, 4\}$. We define a coalition S to have value 1 if and only if there is a contiguous path from sector 1 to sector 4 passing only via edges (i.e., not along corners). As such, $v(\{1, 2, 4\}) = v(\{1, 3, 4\}) = v(\{1, 2, 3, 4\}) = 1$, and $v(S) = 0$ for all other coalitions.

We now reason through the computation of the Shapley value for each agent. The marginal contribution of agent 2 in a permutation is 1 if and only if it enters after agents 1 and 4 and before agent 3. Of the $4! = 24$ permutations, this only occurs in two cases: $(1, 4, 2, 3)$ and $(4, 1, 2, 3)$. Thus, the Shapley value of agent 2 is $\varphi_2 = \frac{2}{24} = \frac{1}{12}$. By symmetry, $\varphi_3 = \varphi_2 = \frac{1}{12}$.

Because agent 1 is an endpoint, we might expected it to have a higher Shapley value than agents 2 and 3; indeed, this is the case. Its marginal contribution in a permutation is 1 if and only if it enters after agent 4 and at least one of agent 2 or 3. This occurs in ten cases: $3! = 6$ when it is the last agent to enter and $\binom{2}{1} \times 2! = 4$ when it is the third agent to enter, with $\binom{2}{1}$ ways to choose one element from $\{2, 3\}$ and 2! ways to order that element and 4 as the first two entrants to the coalition. Then, $\varphi_1 = \frac{10}{24} = \frac{5}{12}$. As our intuition suggested, this

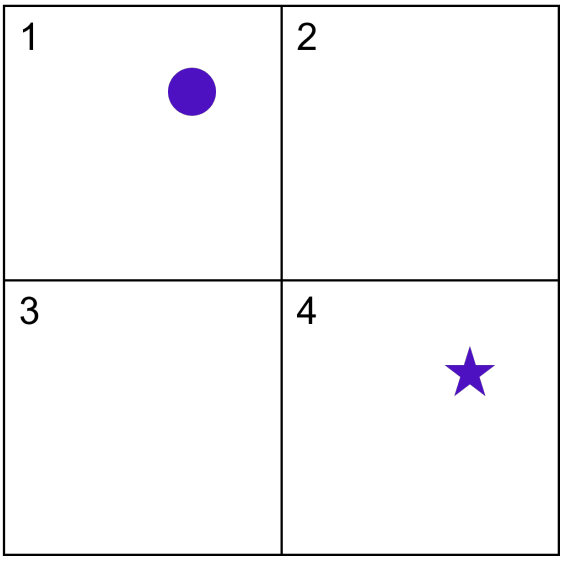

Figure 3-2: An example environment in which the value of a coalition depends on the presence of a contiguous path between the circle and the star entirely within the coalition. Then, coalitions $\{1, 2, 4\}$, $\{1, 3, 4\}$, and $\{1, 2, 3, 4\}$ have value 1; all others have value 0.

is a significantly higher value than φ_2. Again by symmetry, $\varphi_4 = \varphi_1 = \frac{5}{12}$.

We can verify that the Shapley value satisfies all of the desirable properties of a profit-sharing mechanism listed in Sec. 3.3.1:

1. *Efficiency:* $\sum_{i=1}^{n} \varphi_i = \frac{5}{12} + \frac{1}{12} + \frac{1}{12} + \frac{5}{12} = 1 = v(N)$.

2. *Symmetry:* We used symmetry to argue $\varphi_1 = \varphi_4$ and $\varphi_2 = \varphi_3$; a corollary of that argument is satisfaction of the symmetry property.

3. *Additivity:* With only one value function, additivity is not relevant.

4. *Dummy agent:* Every agent provides a marginal contribution to some coalition, so the dummy property is trivially true.

3.4 System model

Our simplified AAM traffic management system model consists of three components: a two-dimensional airspace structure, aircraft operators, and service providers. We briefly describe

the structure of each in this section.

3.4.1 Airspace structure

We represent airspace as a grid of n sectors defined as bounded polygons, $\mathcal{S} = \{F_1, \ldots, F_n\}$. We assume, without loss of generality, that there are two "gates" spaced evenly along the border between every pair of neighboring sectors; gates are the only locations where a flight can transit a border. These gates simplify the calculations that allow us to illustrate the impact of profit sharing on SP routing decisions. Fig. 3-3 shows this structure applied to a small region of airspace.

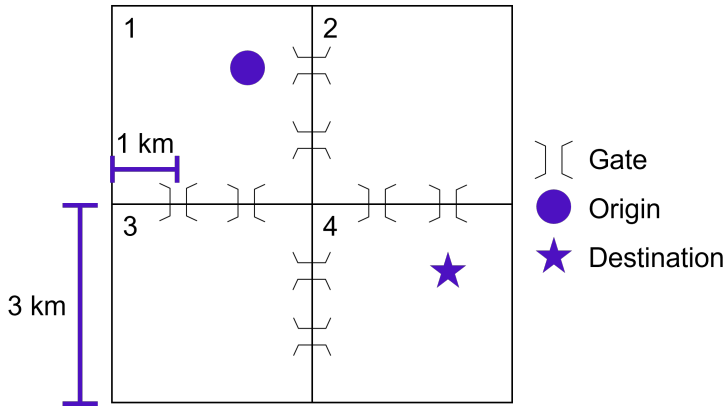

Figure 3-3: Our model of a small region of airspace consisting of four (numbered) sectors and gates between adjacent sectors that are the only locations where flights may cross a border. Origins and destinations may be arbitrarily located within a sector.

3.4.2 AAM aircraft operators

An aircraft operator is an entity that is interested in directing a flight from one location, its origin $o \in \mathbb{R}^2$, to another location, its destination $d \in \mathbb{R}^2$. While simple, this abstraction encompasses a wide variety of airspace applications ranging from package delivery to emergency services to aerial surveillance (which can be viewed as a series of origins and destinations in close proximity to one another).

We define a route that a flight f can take as a collection of m vectors $R = \{\vec{r}_1, \vec{r}_2, \ldots, \vec{r}_m\}$,

where $\vec{r}_j \in \mathbb{R}^2$ in our model. $\vec{r}_1$ originates at o, $\vec{r}_n$ terminates at d, and the endpoint of $\vec{r}_j$ is the starting point of $\vec{r}_{j+1}$ for $1 \leq j < m$. We assume without loss of generality that each vector $\vec{r}_j$ is fully contained within a sector F_i. For a given coalition of sectors $S \subseteq \mathcal{S}$, $\mathcal{R}(f, S)$ is the set of valid routes where all segments are contained within the sectors S. We define the shortest possible route for a coalition as $R^*(f, S)$, and the shortest route overall is $R^*(f, \mathcal{S})$.

In our system model, we assume that an aircraft operator is willing to pay an amount equal to twice the Euclidean distance between a flight's origin and destination if a valid route exists. We define the revenue function for a flight f and coalition S as $u(f, S) = 2\|o - d\|_2$ if there exists a path from origin to destination contained entirely within S, i.e., $|\mathcal{R}(f, S)| > 0$, and 0 otherwise. For simplicity of notation, we drop f in the following discussions that follow, and assume that these routes are being discussed for a given flight.

3.4.3 Traffic management service providers

A service provider (SP) offers traffic management services within a subset of sectors and is responsible for safely routing flights within the airspace under its jurisdiction. This entails strategic deconfliction of flights under its authority, as well as coordination with other SPs to manage flights with origins or destinations outside its service region. In this paper, we focus on the latter problem of incentivizing competing SPs to collaboratively route a flight. To do so, we assume that there is a unique SP responsible for each sector; we will discuss how we analyze multiple competing SPs in Ch. 4.

We use a simple cost structure in which the cost incurred by an SP in routing a flight is equal to the Euclidean distance routed within the SP. Mathematically, we define the cost for SP i of carrying a flight along a route $R \in \mathcal{R}(S)$ as $c_i(R) = \sum_{r_j \in F_i} \|r_j\|_2 \ \forall r_j \in R$. This is equivalent to a model in which the costs of routing are entirely variable; it can be viewed as the cost of a flight occupying airspace. We ignore for now SP decision-making under congestion and deconfliction, and consider how SPs would route flights at the highest strategic planning level. For a given coalition S, we define the overall cost using route $R(S) \in \mathcal{R}(S)$ as $c(S) = \sum_{i, F_i \in S} c_i(R(S))$; then, the optimal cost using the best route possible

is $c^*(S) = \sum_{i,F_i \in S} c_i(R^*(S))$.

Combined with the flight revenue model, this implies that the maximum profit for a group of SPs occurs when a flight is routed exactly along its shortest path and that the profit is equal to the length of the shortest path. We assume SPs are rational agents and seek to maximize profit for the set of flights being managed.

We distinguish between optimal and hot-potato routing for SPs. Under optimal routing, SPs direct flights along the globally optimal path from o to d, while under hot-potato routing SPs will direct flights to the nearest SP and minimize its own costs $c_i(R)$, regardless of the globally optimal path.

3.5 Analysis and results

We now apply the Shapley value to the airspace model presented in Sec. 3.4. We provide and analyze an example scenario and demonstrate how, when we apply our cost and revenue assumptions presented in Sec. 3.4, profit sharing based on the Shapley value incentivizes SPs to route flights along the globally optimal solution, not just the best path for an individual SP. Then, we show that randomly generated flights in our airspace model are routed more efficiently under the Shapley value framework. We conclude with some discussions of possible issues and solutions with using the Shapley value to determine profit share.

3.5.1 Profit sharing with the Shapley value

We propose using the Shapley value as the means by which to divide profits for routing flights among service providers. We extend the example in Fig. 3-2 to include the cost structures from Sec. 3.4. Rather than $v(S) = 1$ if there is a contiguous path from origin to destination and 0 otherwise, $v(S)$ now represents the *profit* obtained by a coalition S for a flight, defined as the revenue $u(S)$ minus the cost $c(S)$:

$$v(S) = u(S) - c(S). \tag{3.2}$$

If the profit from a coalition is negative, i.e., the route it forms is more than twice as long as the Euclidean distance between origin and destination, then the flight is not served and $v(S) = 0$. As before, $v(S) = 0$ also if there is no path from origin to destination. Note that here, each origin-destination pair has its own value function; in the illustrative example in Sec. 3.3.3, the value function depended only on the sectors in which the origin and destination were found.

To determine the distribution of revenue for each SP, we proceed in two steps. First, we calculate the Shapley values and profit shares for each SP before a flight is routed. We then use these fractions along with actual routing costs to distribute revenue after the flight is routed.

Pre-flight

We compute the Shapley value for each agent based on its marginal contributions under optimal routing, i.e., along the globally shortest path from origin to destination within the coalition. Then, these Shapley values are used to determine profit share. If $\varphi_i(N, v)$ is the Shapley value of SP $i \in N$ for value function v, the profit share of agent i, $\rho_i(N, v)$ is proportional to its Shapley value as a share of total value of the coalition, or $\rho_i(N, v) = \varphi_i(N, v) / \sum_{k \in N} \varphi_k(N, v)$. For brevity, we will refer to this as ρ_i when N and v are clear from context. Marginal contributions, Shapley values, and profit shares for each SP for the origin-destination pair in Fig. 3-4 are shown in Table 3.1. An explicit list of all permutations of coalition formation and the respective marginal contributions of each SP may be found in Table A.2 of Appendix A.3.

Post-flight

The profit share for an origin-destination pair is computed *before* any actual routing occurs, as it is based on Shapley values from optimal routes. Then, compensation for routing proceeds as a reimbursement of true costs and a share of overall profit. Suppose SP i incurs an actual cost of c_i in the course of routing a flight, for a total cost of $c_{real} = \sum_{i \in N} c_i$ across the entire

Table 3.1: Computation of Shapley values and profit shares for one origin-destination pair.

SP	Marginal contribution	Frequency	Shapley value φ	Profit share ρ
1	$6\sqrt{2} - 3\sqrt{2} = 3\sqrt{2}$	8	1.670	0.394
	$6\sqrt{2} - (4 + \sqrt{2}) = 5\sqrt{2} - 4$	2		
2	$6\sqrt{2} - 3\sqrt{2} = 3\sqrt{2}$	2	0.646	0.152
	$3\sqrt{2} - (5\sqrt{2} - 4) = 4 - 2\sqrt{2}$	6		
3	$6\sqrt{2} - (4 + \sqrt{2}) = 5\sqrt{2} - 4$	2	0.256	0.060
4	$6\sqrt{2} - 3\sqrt{2} = 3\sqrt{2}$	8	1.670	0.394
	$6\sqrt{2} - (4 + \sqrt{2}) = 5\sqrt{2} - 4$	2		

route. We would like SP i to receive a ρ_i share of profit, so the total payment to the SP is $c_i + \rho_i(u_{real} - c_{real})$, where u_{real} is the total revenue from the flight.

3.5.2 Example and analysis

We now examine how routing may change in the presence of multiple service providers across sectors. To begin, we assume that each sector has a unique service provider associated with it that offers routing services to aircraft operators. As such, for the remainder of this work, we will use sector number and service provider number interchangeably. We consider three cases: optimal routing, hot-potato routing, and alternative routing. These cases are shown in Fig. 3-4 as a solid blue line, a dashed red line, and a dotted gray line, respectively.

We briefly observe that some degree of profit sharing is required; a sender-keeps-all model, as the Internet has, will not work for AAM. Any SP that is not the "sender" would incur nonnegative cost and have zero profit, so there would be no incentive to cooperatively route flights. And, as previously mentioned, the argument of symmetrical traffic flow does not hold because many AAM applications, most notably package deliveries, will be directional.

Under the Shapley value, SPs will be incentivized to carry flights along the optimal route, as any deviation will decrease the total profit earned by all flights. Because the Shapley profit-sharing framework ensures that each SP i ultimately earns $\rho_i(u - c)$ in profit, any deviation from the optimal route will increase c and cause the SP to profit less, even if the cost to the individual SP decreases. An example of this this is shown in Table 3.2, where the optimal

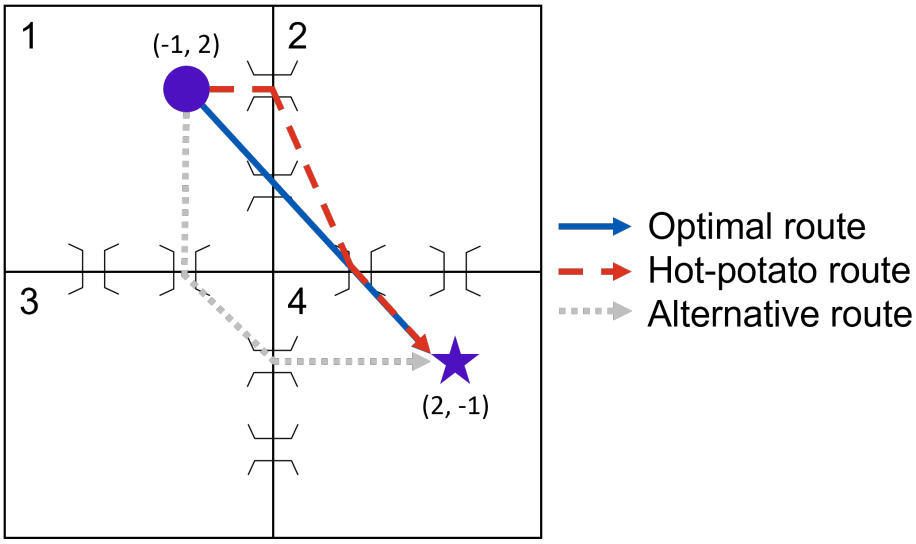

Figure 3-4: Optimal route in solid blue, hot-potato route in dashed red, and alternative route in dotted gray.

routing strategy generates the most profit for every SP compared to other routing strategies, including the "hot-potato" routing strategy used in the Internet.

Table 3.2: Costs, payments, and profit for different routes taken, with profit sharing based on Shapley value for the origin-destination pair in Fig. 3-4. The payment is $c_i + \rho_i(u_{real} - c_{real})$, while the profit is $\rho_i(u_{real} - c_{real})$.

SP	Optimal routing			Hot-potato routing			Alternative routing		
	Cost	Payment	Profit	Cost	Payment	Profit	Cost	Payment	Profit
1	$\sqrt{2}$	3.084	**1.670**	1	2.510	1.510	2	3.209	1.209
2	$\sqrt{2}$	2.061	**0.645**	$\sqrt{5}$	2.820	0.584	0	0.468	0.468
3	0	0.256	**0.256**	0	0.231	0.231	$\sqrt{2}$	1.600	0.185
4	$\sqrt{2}$	3.084	**1.670**	$\sqrt{2}$	2.924	1.510	2	3.209	1.209

We also observe that under congestion or flight rerouting, all SPs are incentivized to keep delays to a minimum in order to maximize profit. For example, if SP 2 faces delays, SP 3 has an incentive to provide an alternative route, as it is compensated for doing so. The most important factor is that, regardless of route, all profits are positive; thus, entering the profit-sharing arrangement has a positive return for participants.

3.5.3 Simulation results

In this section, we present simulation results where the Shapley value is used to divide profit when SPs use globally optimal and hot-potato routing. We measured profit earned per SP (in dollars) and total distance traveled by all flights in the scenario (a measure of social welfare, in kilometers). The airspace is structured as described in Fig. 3-3, with four 3 km-by-3 km square sectors with connecting gates separated by 1-km arrayed in a grid. This is done over four different simulation scenarios, with varying characteristics:

1. *Random traffic scenario*: Each SP sends 20 flights to every other SP. A total of $12 \times 20 = 240$ flights are sent. This serves as a benchmark scenario, where the average effects of the Shapley value and routing decisions can be studied.

2. *Special traffic scenario*: SP 1 sends 20 flights to destinations in SP 4, and vice-versa. SP 2 sends 10 flights to destinations within its sector, while SP 3 sends and receives no flights, receiving profit only through participation in the system. A total of $2 \times 20 + 10 = 50$ flights are sent.

3. *1-Only traffic scenario*: SP 1 sends 20 flights to destinations in SPs 2, 3, and 4. No other flights are sent or received. A total of $3 \times 20 = 60$ flights are sent.

4. *Uneven traffic scenario*: SP 3 controls the merged bottom two sectors (with the border separating SPs 3 and 4 in Fig. 3-3 removed). SPs 1, 2, and 3 each send 20 flights to destinations in every other SP, and 10 flights to destinations within itself. A total of $6 \times 20 + 3 \times 10 = 150$ flights are sent.

The results are presented in Fig. 3-5. The first column shows profit by SP, the second column compares total flight distance under optimal and hot-potato routing respectively, and the third and fourth columns visualize the routes taken under optimal and hot-potato routing respectively. In every scenario, SPs earn more by taking the optimal route compared to using greedy hot-potato routing, and the overall distance traveled by flights is shortened. In the *Random* traffic scenario, we see that optimal routing improves profit by approximately

43% and decreases distance traveled by 27%. Hot-potato routing sometimes forces flights to take long detours to minimize cost to the SP—flights from SP 1 to 3 originating close to the boundary between SP 1 and 2 are routed through SPs 2 and 4 to reach SP 3.

Sending or receiving more flights generates more profit, as demonstrated by SPs 1 and 4 in the *Special* traffic scenario—the dominant position when calculating the Shapley value is at the origin or destination, so having more flights originate or terminate in a sector improves the profit of that sector. The difference in profit from optimal to hot-potato routing becomes smaller if the routes by these methods must follow the same path, shown by the *Special* traffic scenario where flights from SP 1 to SP 4 dominate, and must pass through many of the same gates in both optimal and hot-potato routing.

Serving a larger area could also result in slightly more profit, as shown by SP 3 in the *Uneven* traffic scenario. However, this is likely because of the gates used in our airspace model, as flights originating in the lower right that might have had to pass through gates between SPs 3 and 4 can now take a direct and shorter path to SP 1 or 2 without the gate.

3.6 Discussion

We begin with a discussion on possible impacts of implementing profit-sharing based on the Shapley value among AAM service providers. We then discuss challenges in the emerging field of AAM traffic management, and how the Shapley value might help address or otherwise impact these problems.

3.6.1 Potential impacts of profit-sharing based on Shapley value

Truthful cost reporting

Since profits are computed taking as input the costs as reported by the service providers, a reasonable question involves the incentives for the truthful (or not) reporting of incurred routing costs. Suppose an SP is compensated with $c_i + \rho_i(u - c)$ in accordance with our scheme for a profit of $\rho_i(u-c)$. Now, consider a situation in which the SP misreports its cost

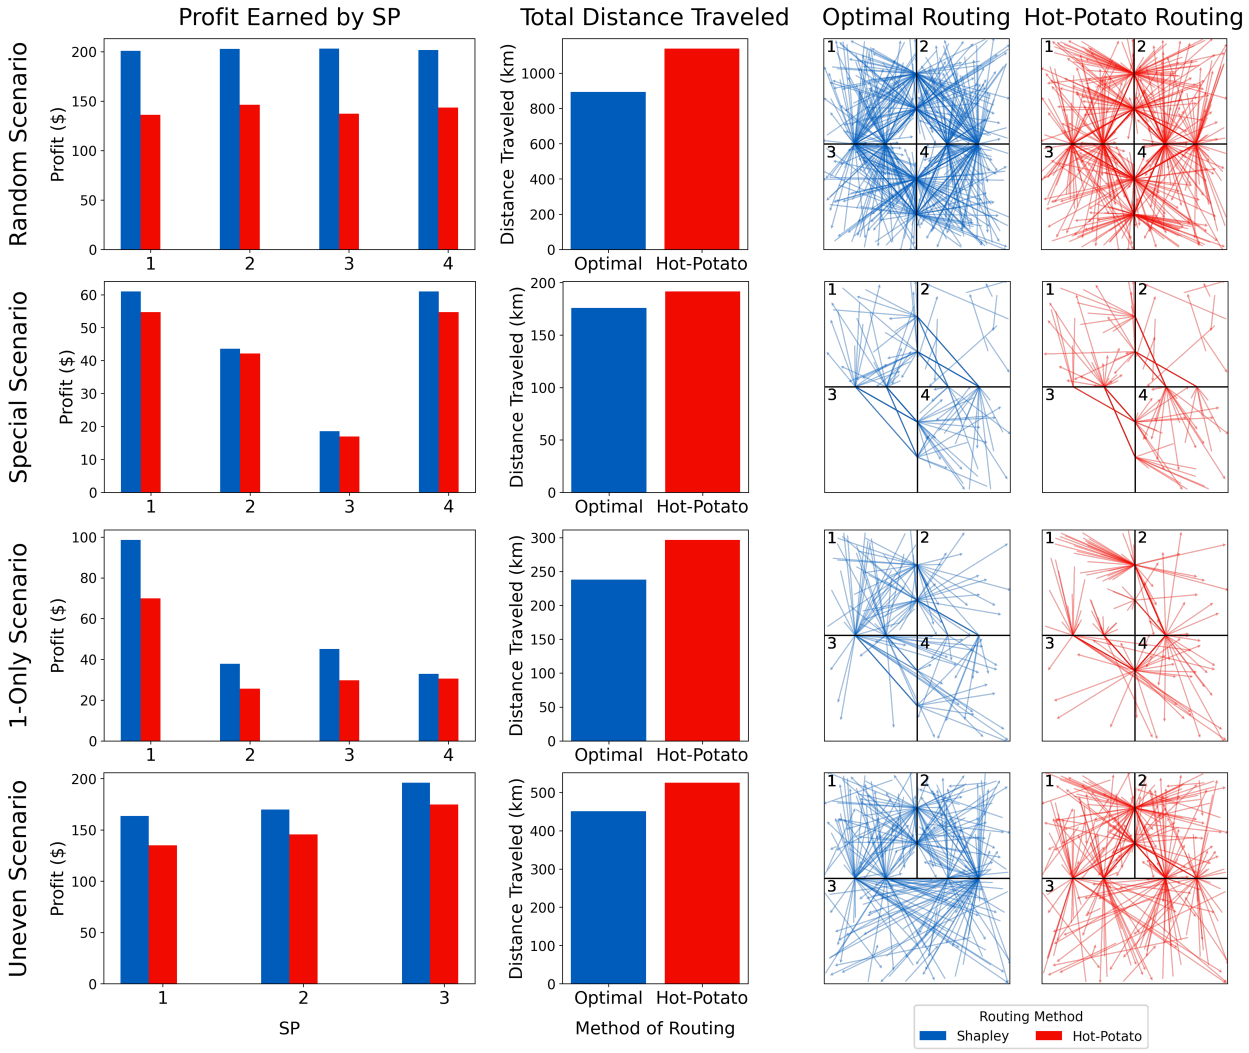

Figure 3-5: Profit per SP, total distance traveled, and routes traversed in each of the four traffic scenarios simulated.

as δ_i more than its true routing cost. Its profit would then be $c_i + \delta_i + \rho_i(u - (c + \delta_i) - c_i = \delta_i + \rho_i(u - c - \delta_i) = \rho_i(u - c) + \delta_i(1 - \rho_i)$, or an increase of $\delta_i(1 - \rho_i)$ over the original profit. Clearly, this is unacceptable—if all SPs attempted to game our system in this way, it could seem impossible to route a flight profitably!

Indeed, untruthful reporting of costs would be a problem if we were to implement this cost reimbursement and profit share compensation structure in the Internet setting, where true routing costs are not transparent. However, in the aviation context, it is straightforward to track the actual route (and thus distance) traveled. This is due to location broadcast requirements such as ADS-B Out and remote identification for drones [41]. Then, one way to ensure truthful cost reporting would simply be for the regulatory authority to assign a fixed cost per unit distance routed, periodically updated based on changes in technology, economic conditions, or policy.

Profit share determination

While we have argued for use of the Shapley value in determining profit share, it is not the only valid distribution. In fact, *any* profit-sharing mechanism with a positive allocation to all agents (i.e., service providers) along the route will incentivize optimal routing. This desirable property is inherent to any form of *profit* sharing; if we had used a revenue sharing model instead, the guarantee would not hold.

To see why, we consider the common economic pie metaphor. Under a profit-sharing mechanism with fixed positive allocations, suboptimal routing will decrease the size of the pie. Therefore, all agents are incentivized to route optimally and, if optimal routing is impossible (e.g., due to congestion), to minimize any additional cost incurred. On the other hand, under revenue sharing with fixed positive allocations, an agent will try to minimize its own cost to maximize its profit, i.e., by using hot-potato routing, because it will always receive the same revenue regardless of the routes flown by aircraft.

It should be noted that, while any profit-sharing mechanism can work, the selection of *which* participants have a nonzero allocation must be considered. If a participant that

could provide an alternative route is not given an allocation, that participant will have no incentive to cooperate. This is particularly important when the system is congested and such alternative routes can relieve the congestion, which is enabled by the Shapley value. On the other hand, if a participant without any practical value is given an allocation, that participant becomes a free rider, benefiting without having to make any contribution. One approach to determine which participants are allocated a share leverages the concept of *spatial locality*, which we discuss next.

Spatial locality

One potential concern with using the Shapley value for computing profit share involves geographic proximity. It is possible for a service provider extremely far removed from the actual area of service to nevertheless receive a small share of profit. In particular, one can show that a poor choice of value function to compute the Shapley values, such as a binary function that values every coalition that creates a path as 1 (a revenue-sharing method), can result in such counter-intuitive allocations. However, this type of behavior disappears when we use the value function based on the *profit* accrued by a coalition of SPs, which we have done in this work. At a certain point, routing through an SP far away from the shortest path generates a negative marginal contribution (negative profit), which turns it into a dummy agent that receives no share of the profit by Property 4 from Section 3.3.1.

3.7 Conclusions

With the vast emerging market for advanced air mobility, private third-part service providers are expected to provide traffic management services. Drawing lessons from Internet Service Providers, we propose a profit-sharing mechanism based on the Shapley value. The proposed mechanism encourages cooperation among service providers by routing flights on the globally optimal paths regardless of individual costs. In addition to optimal routing, it incentivizes AAM traffic management service providers to cooperatively manage congestion. We also

discuss some limitations of the proposed approach.

Chapter 4

The Effects of Sector Allocations to SPs

One characteristic of SPs that was not discussed in the previous chapter was that the FAA expects multiple SPs to operate and be responsible for flights in a single region or sector [1, 7]. This is unprecedented in air traffic management. Traditionally, only one controller or control entity (the control tower at an airport, or an air route traffic center) is responsible as there must be a clear chain of accountability for commands issued or actions taken. We believe that the FAA crafted this policy for reasons of antitrust, and intends this policy to prevent local monopolies of traffic control over a region or city.

So far, the FAA has not clarified how multiple SPs will operate in the same region. We propose a concept of multi-SP operations in this chapter, by segmenting a larger region into smaller (sub)sectors, each exclusively served by a single SP. When combined with the incentives of the Shapley value, which encourage SPs to cooperate and increase global profit, this will allow for one large airspace (eg. New York City) to have multiple SPs operating but still delineated SP control responsibilities clearly in geographical subsectors (eg. Upper West Side). The smaller subsectors discussed here could be regions of airspace where aircraft can freely fly, or AAM corridors moving flights along dedicated airspace lanes [7]. Even if SPs refused to cooperate, a checkered allocation of airspace that spreads out sectors controlled

by a SP would allow for more flexible rerouting of flights around non-cooperative sectors. Given the guarantees of the Shapley value, this should not be necessary.

In this chapter, we analyzes the difference in profit between SPs given a sectorization of a region of airspace, a demand profile of the flow for every pair of origin-destination sectors, and an allocation of sectors to SPs. Of these three factors, only the allocation is a flexible economic and policy tool for the FAA to affect change on the airspace system. The difference in profit is important because it indicates the fairness of a particular allocation to SPs — any differences could create allegations of unfair treatment by the FAA, or cause distortions to the economic incentives of SPs. The goal of the framework we present and data provided will inform future study of economic mechanisms focused on the allocation of SPs, the sectorization of sectors, and other policy considerations necessary to plan out future AAM airspace.

We first give background information and show our overall vision of how a multi-SP system could work inf Sec. 4.1, and explain the different modeling considerations that went into our analysis in 4.2. We demonstrated a fast Shapley value calculation method based on [3] that allows us to use the Shapley value on large spaces in Sec. A.4, show the results of our analysis in Sec. 4.4, and conclude in Sec. 4.5.

4.1 Method

The FAA has indicated that they intend for multiple SPs to operate in the same airspace as each other [1]. This could be to mitigate monopolies in sector markets, where operators can only contract with one SP for traffic services. A monopoly SP, especially in a high density area, could stymie economic development of UAS operations and might not effectively innovate in air traffic control methods. This effect on innovation is, we suspect, one of the main reasons why the FAA has proposed that private SPs should be responsible for managing the proliferation of AAM aircraft instead of doing it through governmental bodies.

The goal of this section is to show how given a sectorization (airspace graph), demand profile, and allocation of sectors to SPs leads to a different amounts of profit are earned by

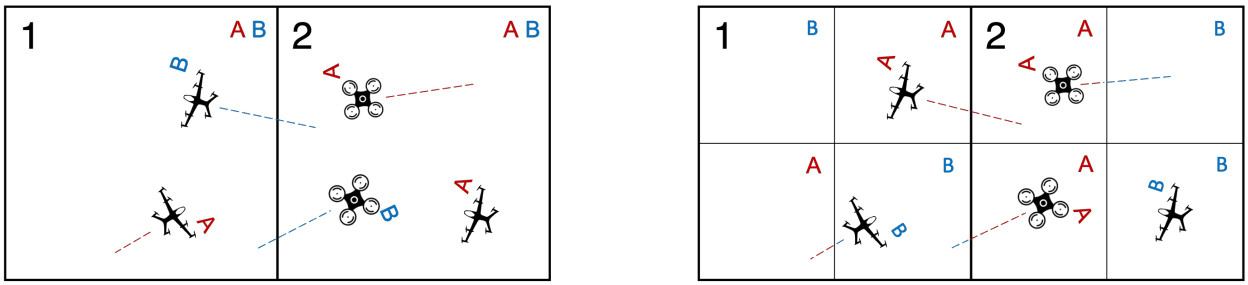

Figure 4-1: A comparison of superposition and subdivision. In superposition (left) both SP A and B can operate in both sector 1 and 2, with flights from both SPs flying in the same sector. In subdivision (right) responsibility for flights is handed off to the SP in charge of the subsector.

SPs. We particularly care about the difference in profit because we believe that the multi-SP stance was taken for economic policy reasons and can be treated as an unfair advantage granted to some parties by the FAA. While we do not propose solutions to dealing with (probably inevitable) differences in profits between SPs, we believe this is a major metric of analysis for future work.

Of the three factors given (sectorization, demand profiles, and allocations), allocation is likely a responsive one. Sector boundaries likely will not change constantly from tradition and regulatory inertia, and can be based on existing reasonable boundaries that exist. Demand profiles for origin/destination pairs is a function of the market and not something that SPs or the FAA can directly influence. Allocation of different sectors to SPs can be a controlled process that the FAA uses to manage the other two factors.

We considered several solutions to having multiple SPs operate in a single sector. The two most important ones are:

1. *Superposition:* Multiple SPs have operational rights in a single sector. Flight 1 under Operator 1 is contracted under SP A, while Flight 2 under Operator 2 is contracted under SP B. SP A and B are responsible for deconflicting Flights 1 and 2 (from a strategic and possibly also tactical perspective). While this ensures that each SP gets to maintain control of the flights they were originally contracted for, it is hard to identify between SP A and B who has priority or ultimate decision-making authority.

Unless a SP is identified as the superior partner, or the FAA steps in as a super-SP, this ambiguity of authority in the sector is unsafe and dangerous while each SP tries to minimize delay.

2. *Subdivision:* A single sector is divided into subsectors, with each SP having exclusive control of a subsector. Flights passing through to a new subsector are handed off to the responsible SP of that subsector (this can be done either through information routed through their SP of original contracting, or directly through the responsible SP). While this resolves the issue of ultimate authority, it potentially creates many small subsectors, which exponentially increases the runtime of calculating the Shapley value. Ultimately, the subdivisioning method attacks the problem of the monopoly SP by splitting regions into small enough sectors and allocating these sectors to SPs in a diverse way, such that each SP will ultimately receive approximately equal benefit.

See Fig. 4-1 for an example of the two methods. The safety issues from the superposition method are serious - a responsible authority must be found for such air traffic management systems. The FAA has indicated that it does not want to directly step into resolving strategic air traffic management. Other solutions would lead to the privileging of one SP above others, and creates unwanted competition between SPs to gain that extra benefit. For example, one SP could be designated as the super-SP that has final power in arbitrating conflicts between SPs, or SPs could be ranked in priority order of flights.

On the other hand, subdivisioning magnifies the effect of SPs being uncooperative. If SPs choose not to cooperate with their neighboring SPs, they act as inaccessible sectors for flights navigating the airspace. But if the subsectors created by subdivisioning are small enough, the delays incurred from a lack of cooperation should be minimized, as the chance of a similar path of sectors that works should be higher in a highly discretized space. In addition, the Shapley value incentivizes cooperation for SPs and a Shapley value framework would prevent negative interactions between SPs. For a large number of sectors it takes a long time to calculate the Shapley value. This can be mitigated by running the Shapley value as an offline calculation (see below for justification), and by using heuristic or faster

exact methods of calculating the Shapley value.

I chose to proceed with developing the subdivisioning method. We continue to use the Shapley profit-sharing method described in Chapter 3, and carry over the same guarantees of efficient routing from SPs. We envision the FAA process of portioning out subsectors to SPs like so:

1. Multiple SPs receive slots or rights to operate in a region (eg. NYC metro region). These could be distributed by auction like spectrum licences.

2. The FAA divides that region into sectors. This could be based on existing boundaries (eg. NYC city districts), census data, data on AAM traffic patterns, regulatory considerations, and other data sources. This sets the airspace system.

3. The Shapley values for each pair of *(sector, origin-destination)* is calculated. This implements a Shapley profit-sharing mechanism.

4. The FAA allocates sectors to different SPs, through optimization, random techniques, auctions, or other options. This could be updated on a timely basis (eg. anually or semi-annually) to ensure that imbalances in profit are averaged out.

5. SPs manage flights that pass through the sectors they are responsible for. This process creates costs for the SP, while generating revenue and profit that is distributed through the Shapley profit-sharing mechanism. Sectors are treated independently regardless of which SP is responsible in the Shapley mechanism. The total profit an SP earns is equal to the sum of the Shapley profits from the sectors it is responsible for. This is the UAM demand profile.

In this section, we model the last four steps in our envisioned FAA process — building an airspace graph, calculating the Shapley values for all origin-destination pairs and all sectors on that graph, deciding on an allocation of sectors to SPs, and considering different traffic demand patterns. The output from these four factors is a profit for each SP. We focus in particular on the difference in profit between SPs — while the previous chapter established

the *efficiency* of the Shapley value (by showing the incentives for SPs to move flights), in this chapter we consider the *fairness* of allocating multiple sectors to certain SPs and the disparities that may arise in given situations.

4.2 System Model

First, we mode the sectorization of a region as a graph $G = (\mathcal{S}, E)$. Each node in the graph $s \in \mathcal{S}$ represents a sector, while an edge connecting two nodes (s_i, s_j) represent that two sectors share a border. Future work can loosen this constraint. Nodes will have a location in $l(s_i) \in \mathbb{R}^2$, and edge lengths are the distance between nodes $||l(s_i) - l(s_j)||$. A coalition of sectors is defined as $S = \{s_{a_1}, s_{a_2} \dots s_{a_k}\}$. In this work, we do not consider flights as moving in free space, and instead limit flights to traveling on the graph G, unlike in Chapter 3. Any flight within in a sector will be assumed to be at the node, and the distance it travels from sector to sector will be given by the edge between those sectors. Future work will address different flight paths through a sector.

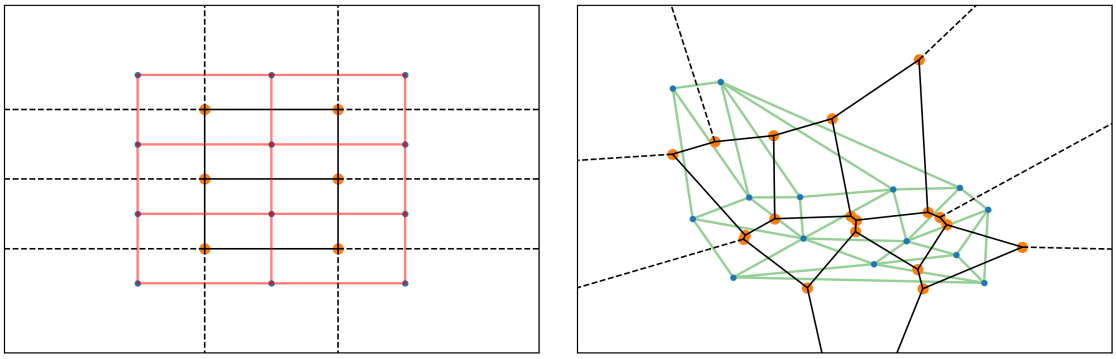

(a) (b)

Figure 4-2: A demonstration of how we generate graphs based on existing sectors. Each existing sector is represented by a node, and neighboring sectors are connected by edges. A 3x4 grid is shown in (a), and an example random graph is shown in (b). The random graph contains a few links that are created by sectors intersecting with each other outside the scope of the figure.

We can then find the Shapley value $\varphi_i^{(o,d)}$ for each node (each sector) across this graph

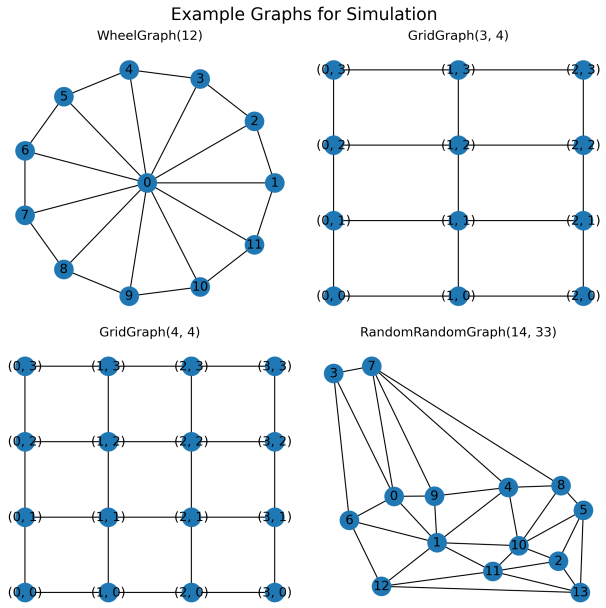

Figure 4-3: Graphs used for testing different SP allocations

structure. We need to evaluate the Shapley value for every origin-destination pair, because every pair involves nodes in the graph differently. In the previous chapter this was done online, where the Shapley value calculation was run for every flight. This is a very expensive calculation to do exactly, in addition to the large computational overhead of calculating the Shapley value (see the factorial term in Eqn. 3.1). However, the large computation burden is not a major concern in this setting for two reasons:

1. We use an improved method first demonstrated in [3] to calculate the Shapley value for graphs. This improves the runtime of the Shapley value by at least 100x.

2. We can design the FAA process described in Sec. 4.1 to make the Shapley value calculation to be an offline calculation. If an airspace sectorization and its corresponding graph is created months in advance, supercomputing resources can be dedicated to solving even very massive Shapley value calculations in that period of time and return the exact Shapley value. The flexibility in timeline for our airspace process allows for offline computations of the Shapley value for large graphs. When combined with the performance of our improved Shapley value calculation, the runtime scales appropri-

ately slowly for exact calculation to be feasible.

The Shapley profit sharing method requires us to define a profit function for origin-destination pairs (o, d), $v(S) = u(S) - c(S) \;\; S \subset \mathcal{S}$. We use two types of value functions. The first is the same value function described in Sec. 3.4: if there is a route from o to d contained within a coalition S then the value is twice the Euclidean distance; otherwise the value is 0.

$$v_{dist}(S, (o, d)) = \begin{cases} 2 * ||l(o) - l(d)|| - \sum_{i, P_i \in S} c_i(R(S)), & \text{if } |\mathcal{R}(f, S)| > 0 \\ 0, & \text{otherwise} \end{cases} \tag{4.1}$$

The second function is a path-based value function that only checks for a valid path from the origin to the destination, used previously in Sec. 3.3.3. Here, I have just directly defined $v(S)$.

$$v_{path}(S, (o, d)) = \begin{cases} 1 & \text{if } |\mathcal{R}(f, S)| > 0 \\ 0, & \text{otherwise} \end{cases} \tag{4.2}$$

Using the Shapley allocation rule defined in Sec. 3.5, we find the Shapley value $\varphi_i^{(o,d)}$ for each sector s_i and origin-destination pair (o, d).

Next, we model the demand of flights. For each origin-destination pair, we assign a traffic demand for that pair. Formally, we will define demand $D(o, d) \in \mathbb{R}$ as the demand for flights between the origin o and destination d. We do not make any constraints such as flow conservation at this step; instead, we provide a simple value for how many flights will be traveling from the given origin to the given destination. This allows us to create different traffic profiles. In this work, we study three traffic profiles - UniformProfile, DegreeProfile, and RandomProfile. The first is a uniform traffic profile, where every origin-destination pair has the same amount of traffic, the second models demand as proportional to the sum of the degrees of the origin and destination nodes, and the third assigns a uniform random amount of traffic (drawn from [0, 1]) between origin-destination pairs. All

profiles are normalized between [0, 1].

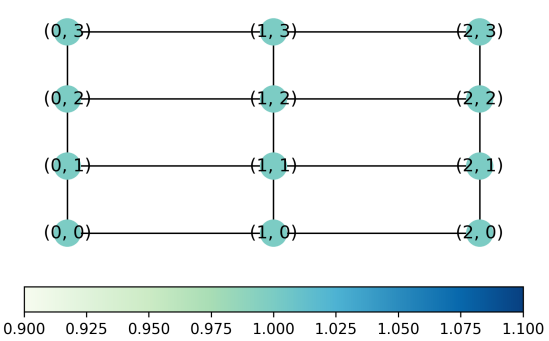

(a) UniformProfile

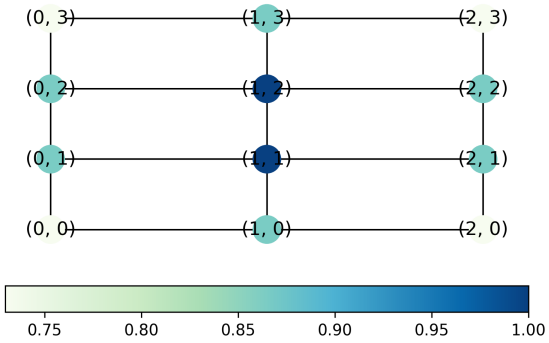

(b) DegreeProfile

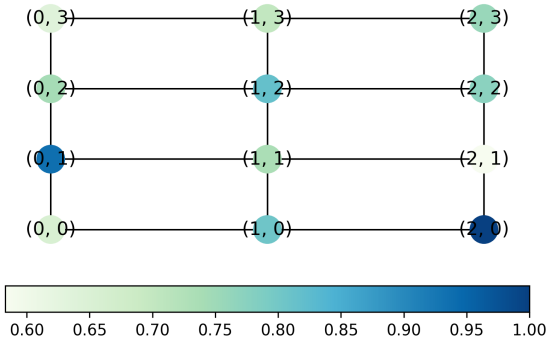

(c) RandomProfile (seed 10)

Figure 4-4: The three different demand profiles used. For each origin-destination pair, the uniform profile assigns each par an equal amount of demand; the degree profile assigns each pair a demand based on the sum of the degrees of each node; the random profile randomly assigns each pair a demand. This figure shows the total sum of all origin and destination demand at a node under a profile, normalized to the largest total.

Finally, we model the allocation of sectors to SPs. In this work we study only one allocation rule, EvenAllocation, which distributes an equal number of sectors to each SP. For each graph, we generate all possible allocations of sectors under this rule. Formally, given a SP in the set of all SPs $p \in P$, we define the function $Q(p) = \{s_{a_1}, s_{a_2}, \ldots s_{a_k}\}$ $s_{a_i} \in \mathcal{S}$ as assigning a set of sectors to SP p. Then, the EvenAllocation rule is the set of assignments $\mathbf{Q_{even}}$ such that $|Q(p)| = \frac{|\mathcal{S}|}{|P|}$ $\forall Q \in \mathbf{Q_{even}}$. This is a fair assignment rule for a certain type of fairness, in that each SP gets the same number of sectors.

The final profit for an SP is as follows: given an airspace structure $G = (\mathcal{S}, E)$ with all

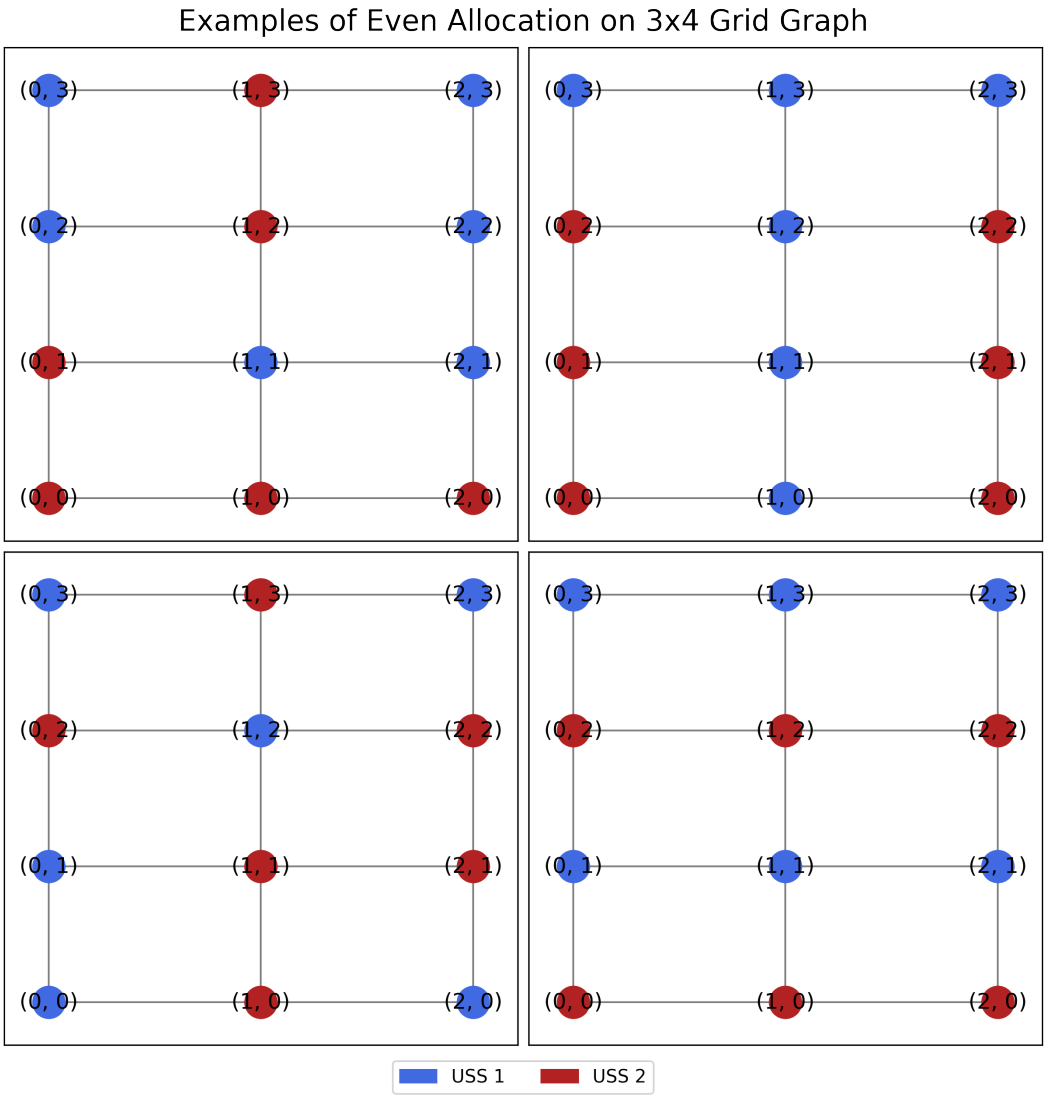

Figure 4-5: Example of allocations between two SPs on the 3x4 Grid Graph

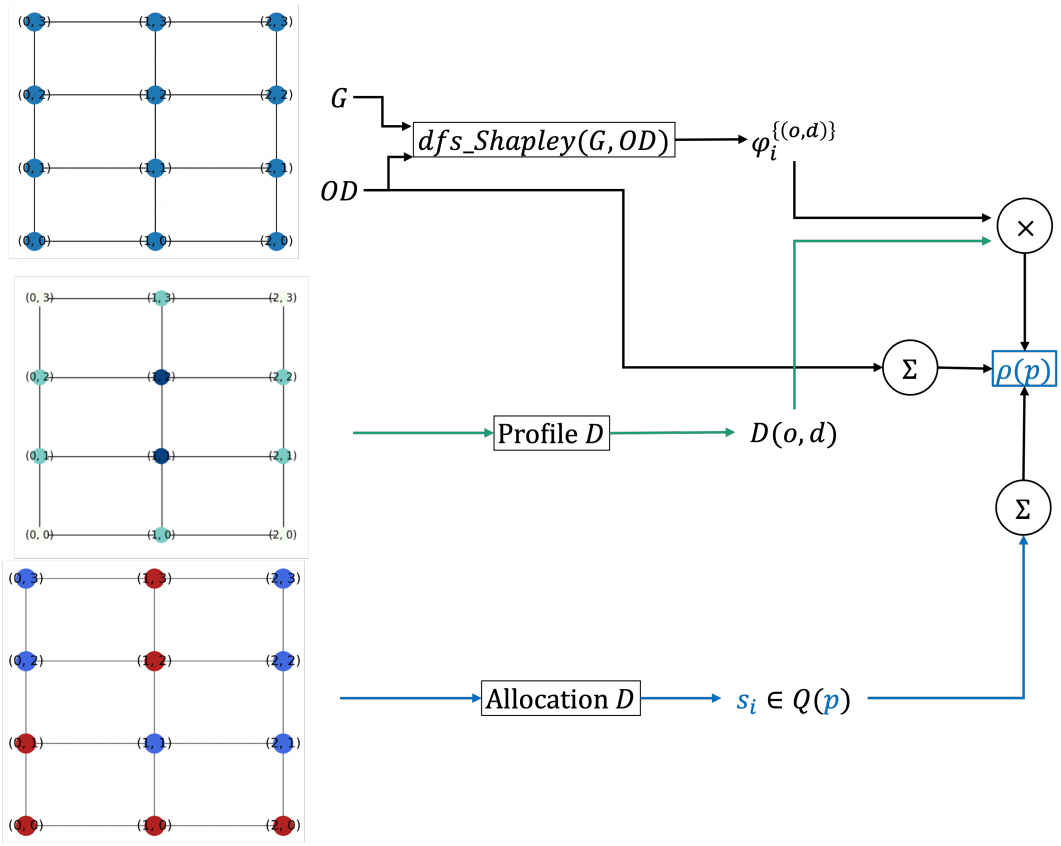

Figure 4-6: The calculation of the total profit for a SP based on Eqn. 4.3, in this case the blue SP

orgin-destination pairs OD, a demand profile D, and an allocation of sectors to SPs Q, we find that for SP p the final profit is:

$$\rho(p) = \sum_{s_i \in Q(p)} \sum_{(o,d) \in OD} \varphi_i^{(o,d)} D(o,d) \tag{4.3}$$

In this final profit equation, the runtime limiting step is the calculation of the Shapley value $\varphi_i^{(o,d)}$. This calculation is independent from the definition of the demand profile D or the allocation of SPs Q, which allows us to find φ for a certain graph offline and quickly iterate through different profiles or allocations.

4.3 Algorithm

A major concern of the Shapley value is the computational runtime on large number of agents - in our case, the nonlinear increase in runtime needed as the number of sectors grows scales by $O(N!)$. The Shapley value is also unique for each OD pair, which scales by $O(N^2)$. However, in the SP setting the Shapley value can be run offline and infrequently because of the fixed nature of airspace sectorization.

In addition, we build a fast Shapley value algorithm, based on [3] by Skibski et al. This method enumerates all connected induced subgraphs (which enumerates connected coalitions of sectors), and then calculates the Shapley value for every subgraph. We modify the algorithm to account for the value of all coalitions, because the focus on Skibski et al. is on calculating the value of connected coalitions. Non-connected coalitions, while having no value in [3], could have value in our setting because a path from the origin to destination could still exist among some connected nodes.

We first explain the algorithm presented in [3], then describe the modifications made for the SP setting. Skibski et al. in [3] focuses on finding the Shapley value for value functions dependent on the connectivity of the graph. Let $\mathcal{C}$ be the set of all connected coalitions, where every node in the coalition is connected to every other node. Then, Skibski et al. considers value functions of the form:

$$f_G^{\mathcal{C}}(S) = \begin{cases} v_G(S) & \text{if } S \in \mathcal{C} \\ 0 & \text{otherwise} \end{cases} \tag{4.4}$$

The algorithm centers around a depth-first search method to enumerating connected induced subgraphs. Non-connected graphs have no value, and so do not need to be enumerated. For each subgraph, the Shapley value for all vertices are updated by the value function given. We based our algorithm on the *DFS-Shapley* algorithm in [3].

The value function given by Eqn. 4.4 differs slightly from the value functions we use given in Eqn 4.1 and 4.2, in that our value functions could return a nonzero value if $S \notin \mathcal{C}$. The

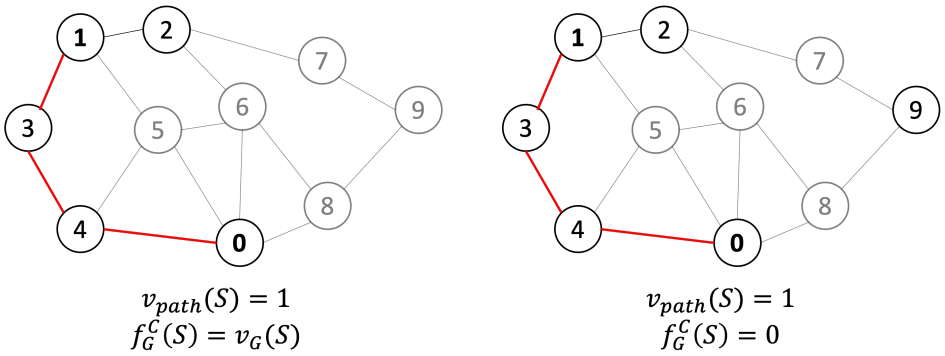

Figure 4-7: Comparison of our value function v_{path} versus the value function for a connected subgraph f_G^C in [3].

disconnect of an arbitrary node s_i from the origin o does not affect the value of a coalition for v_{dist} and v_{path} as long as there is path between the origin and destination o, d (see Fig. 4-7). We ignore the tracking of cut vertices provided in [3] and modified how the Shapley Value was calculated for every node s_i in a coalition S, in Alg. 1 on Line 29, to account for coalitions where disconnected nodes are included.

For every node in the coalition on Line 29 of Alg. 1, we find the difference in value with the addition of the node s_i to the coalition based on the value function v_G (either v_{path} or v_{dist}). The coefficient $\lambda_{S'}$ is a fraction of how many ways a coalition can be constructed with s_i in this situation, divided by the total number of ways to construct a coalition $|\mathcal{S}|!$. Because our value function does not necessarily depend on the full connectivity of the graph, we do not look for cut vertices as explained in Alg. 4 of [3].

For each node, we split all the other nodes in three, then identify how they slot into an order. The three are: 1) the other coalition members which have to come before the node in order (to form the coalition); 2) the neighbors of the node not in S which have to come after (they will make a larger connected subgraph with this node and so are already enumerated); and 3) the disconnected far nodes (which can go anywhere in the order because they don't form a connected subgraph). Inspired by the stars-and-bars problems, where the faraway nodes are the bars, we determine that there are $\binom{|\mathcal{S}|}{|Far|}$ locations to place the faraway nodes and $(|Far|)!$ ways to order the nodes. Thus $\lambda_{S'}$, illustrated in Fig. 4-8, is:

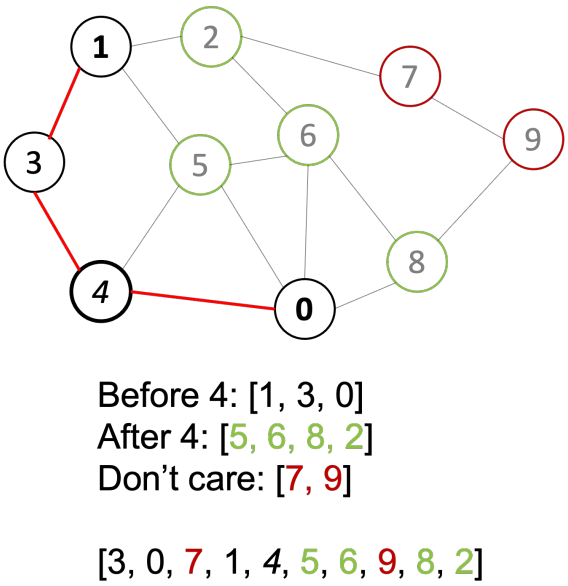

Figure 4-8: $\lambda_{S'}$, used for all other nodes.

$$\lambda_{S'} = \frac{(|Neigh|)!(|S|-1)!(|Far|)!\binom{|\mathcal{S}|}{|Far|}}{|\mathcal{S}|!} \tag{4.5}$$

The marginal contribution of a node s_i to a coalition S for an origin-destination pair (o, d) is thus:

$$\lambda_{S'} * (v_G(S,(o,d)) - v_G(S \backslash s_i,(o,d))) \quad s_i \in S \tag{4.6}$$

A node has the same coefficients $\lambda_{s'}$ for a given coalition, regardless of the value function used or origin-destination pair. At each enumerated subgraph, we can iterate through all origin-destination pairs and find how each combination of (subgraph, pair) contributes to the Shapley value φ_i for a node s_i. The complete algorithm is shown in Alg. 1 in Appendix A.4. Eqn. 4.6 can be found on Line 29 of Alg. 1.

4.4 Results

We study the combination of our four graph types and three demand profiles under the EvenAllocation rule for two SPs. We simulated all allocations allowed under EvenAllocation, and produced general summary statistics showing the difference in profit earned between SPs under each combination of graph, demand, and allocations in Sec. 4.4.1. We also study the outlier results for different scenarios to find the factors that may produce higher differences in profit. In Sec. 4.4.2, we recreate these factors by studying extreme demand and allocation scenarios, to magnify the difference in profit and demonstrate the effects of these factors.

4.4.1 General Summary Results

In Fig. 4-9, we show summary statistics for the difference in earnings between two SPs over all allocations. For each graph, the x-axis is the difference in earnings between two SPs normalized by the total earnings, while the y-axis is the percentage of allocations that have a particular difference value. We compare different graphs in each row and different demand profiles in each column. The values generated from both v_{path} and v_{dist} are compared. Generally the difference in profit does not become greater than 20% of the total profit earned by both SPs, and the standard deviation is relatively low. The distance based measure has a wider distribution of difference in profit for the random graph, which is likely from the larger difference in revenue for a path under v_{dist}.

We look at the outliers from the 3x4 grid to look for what leads to high differences in profit between SPs. In general, we see that controlling the nodes that have a high degree of connection and are more towards the center of the graph is important, as shown in Fig. 4-10a and 4-10b. This makes sense from the perspective of the Shapley value, because the closer and more well connected you are in the graph the more likely a path between an OD pair runs through you. The structure of the largest outlier is different for the random graph, where the centrality and degree of a node seems to be important. This can be partly explained by the nature of the Shapley value as a measure for centrality.

4.4.2 Specific Analysis

Based on the analysis of the last section, we test specific allocations and demand patterns to show the importance of certain sectors. In this section, we focus on studying distinguishing between nodes involved in the transit of flights and nodes as the origin and destination. In Fig. 4-11, we set certain nodes in the 3x4 grid graph as origin-destination nodes and compare the profits. Fig. 4-11a sends traffic between the top and bottom layers of nodes, while Fig. 4-11b and 4-11c sends traffic between the middle layers. In Fig. 4-11a, we can see that the transit nodes for this graph have less earnings than the origin and destination node generally, because multiple paths exist through the transit nodes (in blue). If the origin and destination sectors are in a limited area, other sectors outside that space are unlikely to receive any profit, unless they can provide paths that are profitable (for example, under the v_{path} instead of v_{dist} value function in Fig. 4-11c).

Transit sectors are able to make profit if they hold an important transiting space. In Fig. 4-12, traffic is sent and received at all nodes except the central node, yet the central node receives more profit than other nodes (10.5 vs 9.1). This is because for most origin-destination pairs, the central node is the only node in the middle of the shortest path. This centralized power compared to the fragmented sectors on longer paths along the edge of the graph gives the central node a high Shapley value.

Finally, we test a variation of the Special scenario in Fig. 2-5 in Fig. 4-13, modified to fit our airspace graph. We again send demand between 4 nodes in SP 1 and SP 4 respectively (for 16 origin-destination pairs), and send flights internally in SP 3 (for 6 origin-destination pairs). We see that SP 1 and 4 generate a large amount of profit, SP 3 earns some from internal traffic, while SP 2 is benefiting a little as a transit SP. The relative traffic between

4.5 Conclusion

In this chapter we consider how multiple SPs can responsibly manage flights in the same region of airspace. We do so by further subdividing the airspace and using the Shapley

value to incentivize SPs to cooperatively move flights. In light of the heavy computational required to calculate the Shapley value, we design a procedure for airspace segmentation that gives plenty of time to complete the Shapley value calculation offline, and modify faster Shapley value algorithms for our purposes. We then analyze how the combination of airspace sectorization, origin-destination demand, and allocation of sectors for SPs affects the difference in profits between SPs, and identify notable factors that increase this difference.

The results of this work are necessarily future-looking, as (at the time of this publication) the FAA has not yet announced more concrete proposals for how to achieve their vision for UTM and UAM. We hope that this work sets the foundation for grappling with the challenging problems of developing a novel air traffic management architecture for a novel set of requirements. Future work could consider this problem in free flight trajectories instead of on a graph, and analyze how SPs might be incentivized against cooperation and incorporating those considerations in our design for a multi-SP management system.

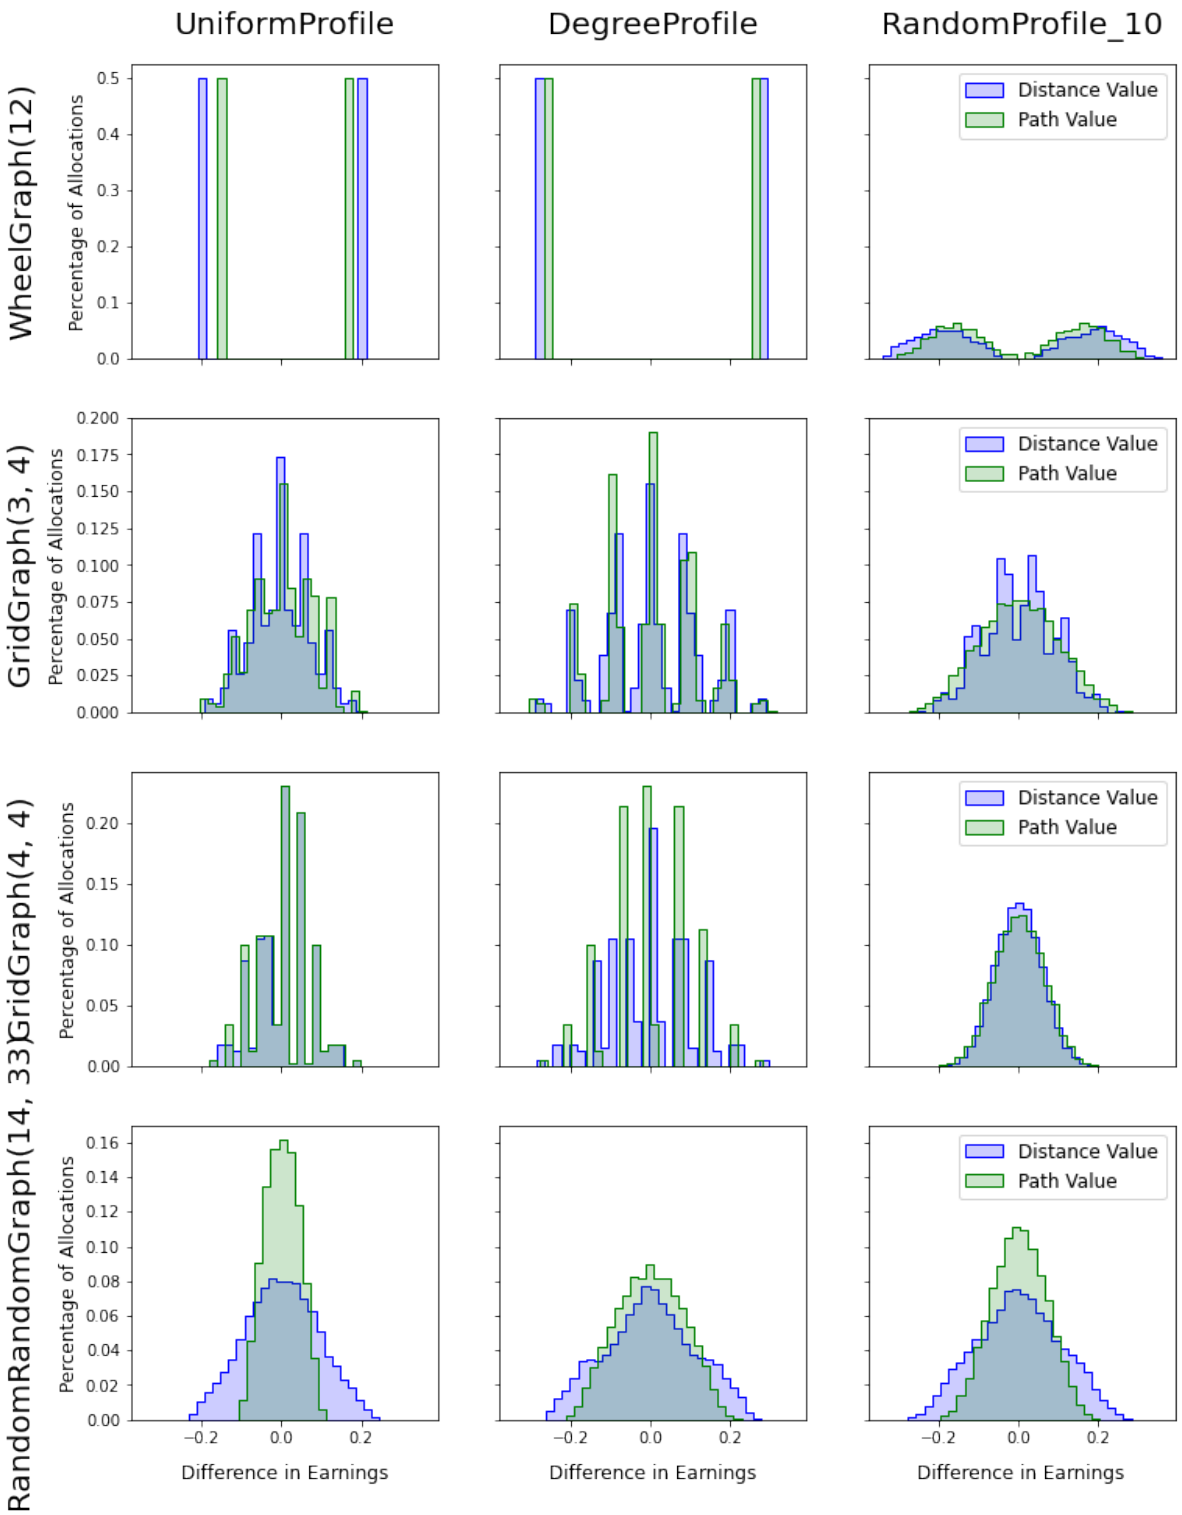

Figure 4-9: Difference in profit for 2 SPs as a fraction of total profit under $v_{distance}$ and v_{path} for combinations of the 3 demand profiles (columns) and 4 graphs (rows).

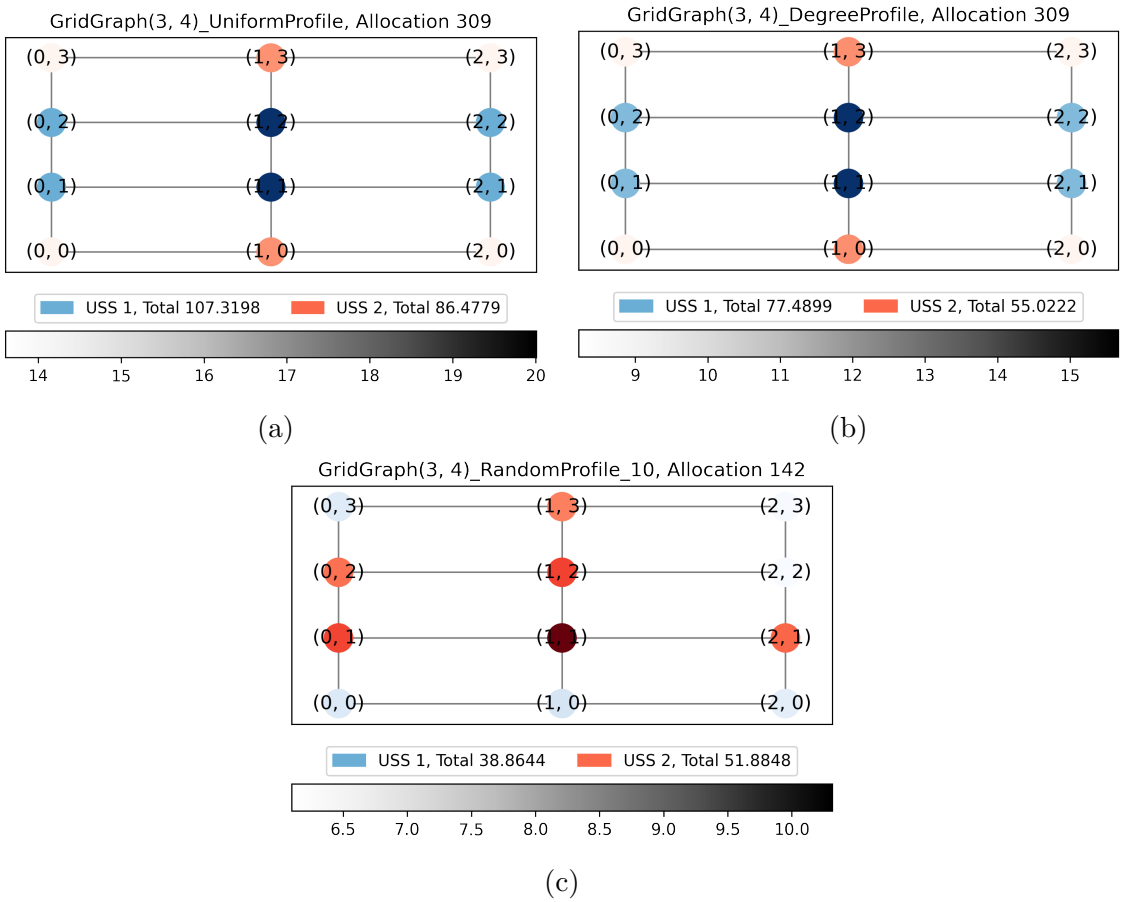

Figure 4-10: Comparison the largest outlier of the 3x4 Grid on the 3 different profiles.

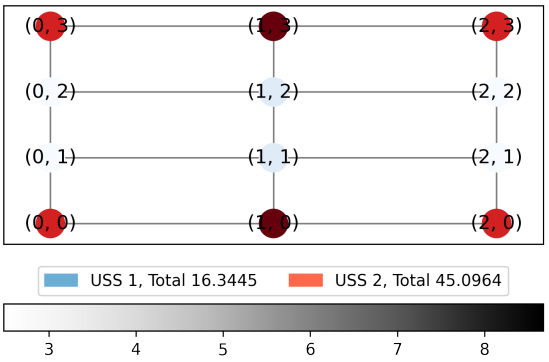

(a) Between top and bottom nodes, v_{dist}

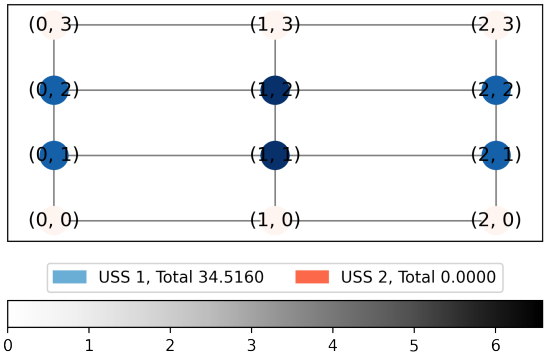

(b) Between middle nodes, v_{dist}

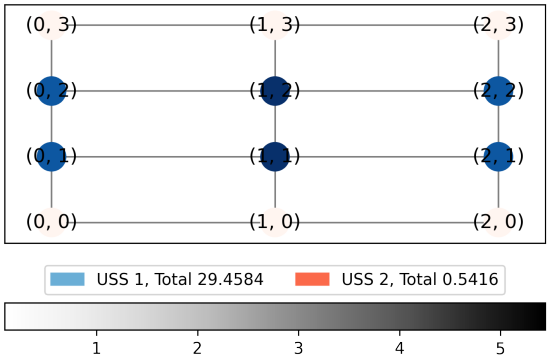

(c) Between middle nodes, v_{path}

Figure 4-11: Comparison of transit scenarios on the 3x4 grid graph.

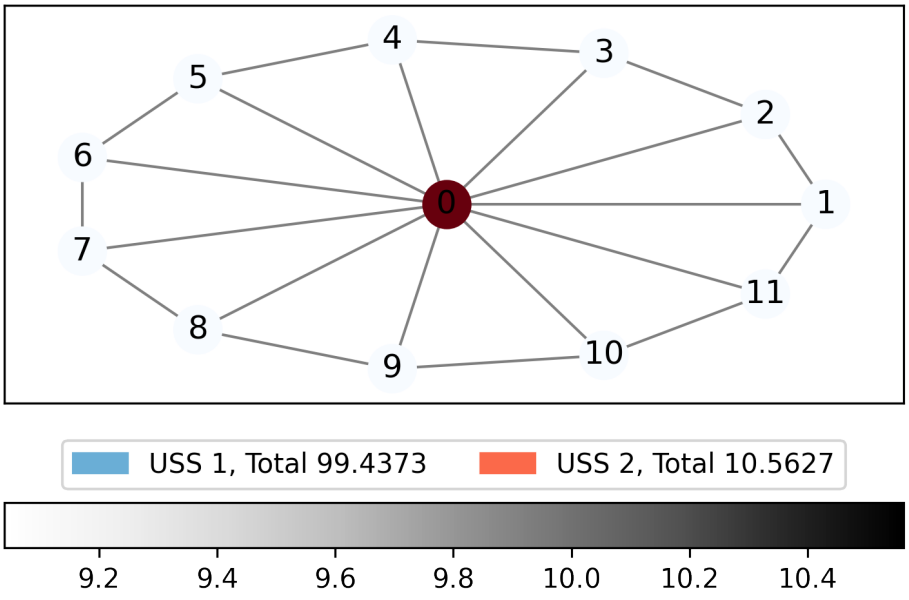

Figure 4-12: Wheel graph, with origin-destination pairs on the outside of the graph.

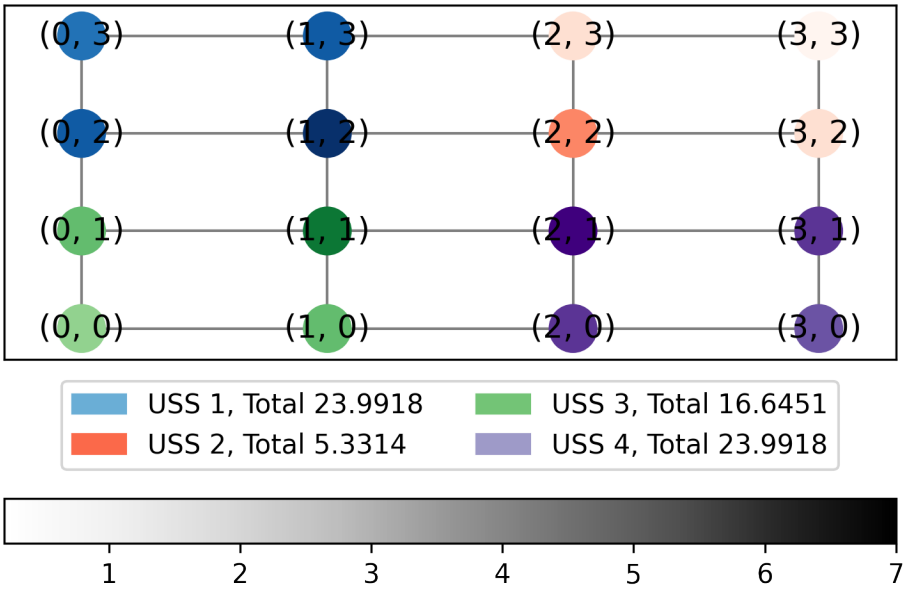

Figure 4-13: Replication of the Special Scenario in Ch. 3.

Chapter 5

Conclusion

This work represents an important first step in developing a market structure for AAM traffic management service providers. In this work, we considered how future service providers responsible for managing air traffic for AAM and UAM flights might operate. We presented cost aware priotization methods and protocols that an SP might use for strategic deconfliction within its area of responsibility, used the Shapley value to incentivize cooperation among SPs, and showed how large scale use of the Shapley value can allow multiple SPs to operate in a region of airspace. These solutions represent a market-based approach towards mitigating some of the concerns with creating private air traffic management for future air mobility. They allow flights and entities to reveal their preferences without sharing complete information, thus maintaining privacy while also empowering limited coordination within the system. When implemented in conjunction with regulations on safety and unified procedures, they can mitigate negative competitive pressures that might arise in the privation of airspace.

At the time of this writing, many of the details on how the FAA's proposed concept for AAM traffic management have not been specified. There remain many unaddressed challenges that may require further study and even changes to the regulatory landscape for AAM. This leave open many opportunities for future work, which I will list below.

5.1 Future Work

5.1.1 Game theoretic considerations in traffic management protocols

The protocol and prioritization methods presented in Ch. 2 have many game-theoretic implications that remain unexplored. For example, the second backpressure method relies on an assumption of truthfulness because of its use of proportional payment; a truthfulness guarantee could be obtained with modifications to the method. Second price auctions are generally susceptible to coalitions, such as an operator who owns several flights, and it is important to analyze how these prioritization methods might respond or be made resilient to these problems. Additionally a real implementation of the protocol should consider flights moving at heterogenous speeds several timesteps in the future. Holding auctions for airspace asynchronously will affect the incentives and behavior of participating flights.

5.1.2 Interaction with intra-SP traffic management

In Ch. 3 and 4, we abstracted away congestion management within a SP, and assumed that flights do not conflict. In a real airspace system, SPs may treat flights transiting their sector differently depending on the fraction of profit they earn, and inter-SP coordination could be affected by internal SP traffic management methods, whether protocol-based (like in Ch. 2) or through centralized optimization [4, 15, 16, 42]. Preliminary study suggests that using the Shapley value for congestion management in an airspace system does not significantly affect performance, regardless of the traffic management methods used within a sector. This deprioritization can be tolerable because, without the Shapley value, SPs that do not receive any benefit from a flight would have no incentive to ever carry that flight, forcing a regulatory solution that would have to specify complex rules around SPs assisting each other instead of a more flexible incentive-driven solution.

5.1.3 Net neutrality type challenges in AAM

In the context of the Internet, net neutrality refers to the notion that Internet service providers should treat all content equally, without favoring one content creator or type of content over another. This has become a complex ethical and economic question in Internet policy, with companies picking sides in the debate based on their business models and affiliations. In many instances, ISPs that are also content creators will give preference to data from an affiliated content creator, rather than a competing content generator. An analogous situation in the AAM context would be when a service provider is also an aircraft operator.

The FAA has explicitly stated that entities that are aircraft operators may also be service providers, rather than relying on third-party SPs [1]. Consequently, a single entity may serve as both an operator of flights and a service provider for other aircraft operators. This dual role as both aircraft operator and service provider is analogous to an ISP also being a content creator, with the physical airspace being analogous to capacity-constrained bandwidth. It remains to be studied how different regulatory policies (e.g., similar to ones that try to ensure equal treatment of all aircraft operators) might affect AAM traffic operations. These are some of the open questions that need to be resolved before the full potential of advanced aerial mobility can be realized in practice [13].

Appendix A

Appendix

A.1 Efficiency of Second Backpressure Prioritization

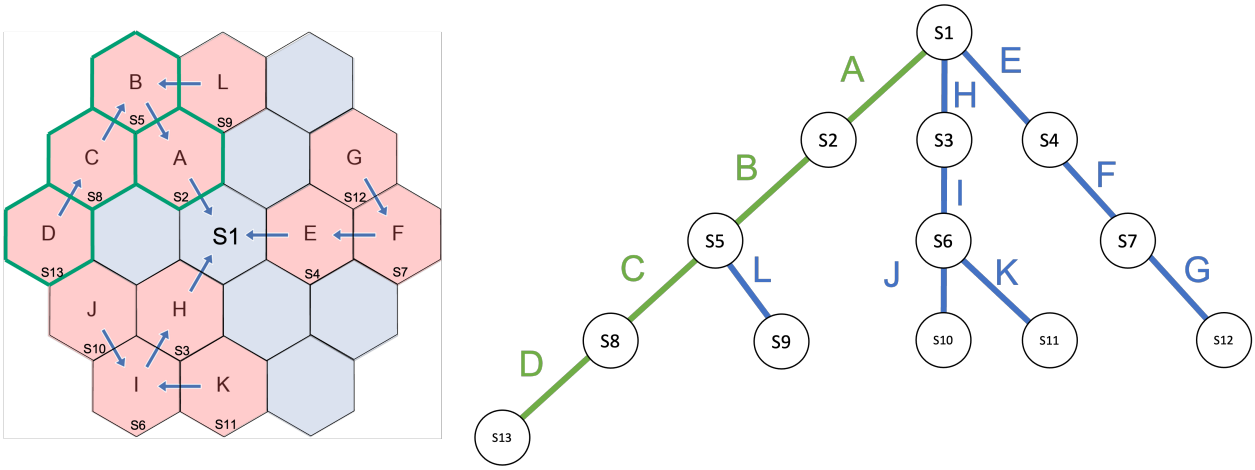

Figure A-1: Abstraction of a set of aircraft in conflict from the protocol (left) into a tree structure (right). Backpressure prioritization solves for the deepest path of the tree, which is also the optimal one-step solution [4].

The second backpressure method is optimal with respect to weighted delay, where each aircraft's flight delay is weighted by its variable cost. We can prove this by mapping the second backpressure prioritization method to the backpressure prioritization method, which is optimal with respect to raw delay. The proof for backpressure prioritization is given in [4]. To do this, we simply duplicate the number of nodes for a given sector by the variable cost,

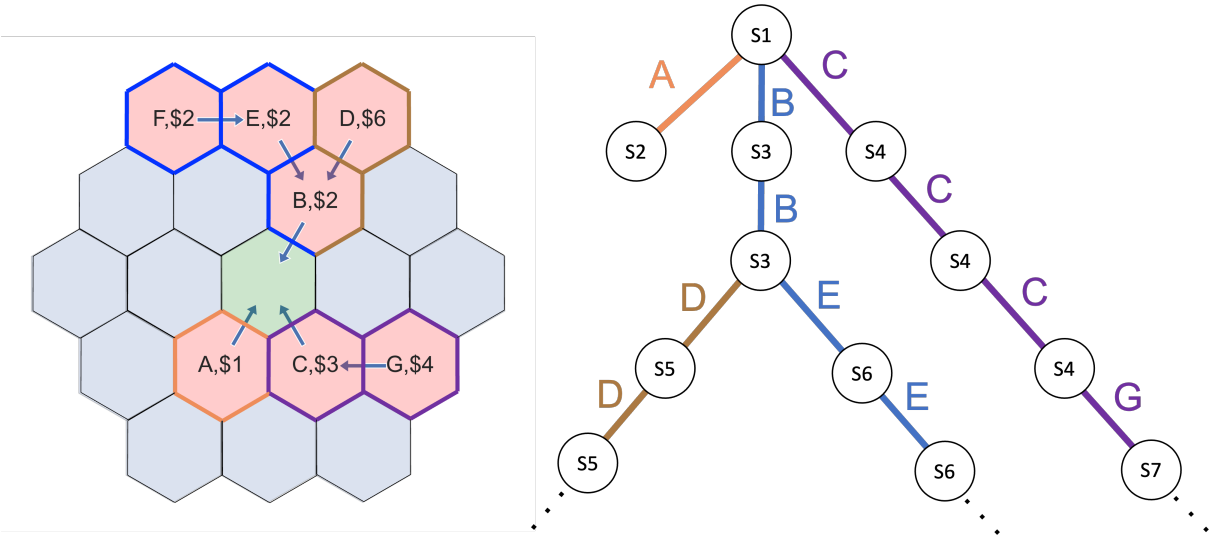

Figure A-2: Abstraction of a set of aircraft in conflict from the protocol (left) into a tree structure (right), to prove second backpressure prioritization optimality. Note that nodes are repeated by the variable cost of the aircraft in the sector.

which turns it into a backpressure prioritization. More specifically, for a given set of aircraft $i \in V'$ and their bids $\hat{p}(i,t) \in \mathbb{R}$, we take the following steps, illustrated in Fig. A-2.

Proposition 1. *Any weighted prioritization problem can be mapped to a unweighted prioritization problem presented in Fig. A-1.*

1. We normalize all $\hat{p}(i,t)$ to be integers $\hat{p}'(i,t) \in \mathbb{Z}$. This is possible because for all fractions $\hat{p}(i,t) = \frac{r}{q}$ $\forall i \in V'$, the denominators share a least common multiple $lcm = qw$ for some $w \in \mathbb{R}^+$. We can then define $\hat{p}'(i,t) = \hat{p}(i,t)\, lcm = rw$ $\forall i \in V'$.

2. We construct the tree structure provided in Fig. A-1, except we repeat nodes. For a given sector s, where $\hat{x}(i,t) = s$, we repeat the node-edge creation process for sector s by $\hat{p}'(i,t)$ times. In-neighbors to node s are connected to the lowest node of s. This effectively means that for each aircraft in a branch, instead of the edge having a length of 1 it has a length of $\hat{p}'(i,t)$.

We then follow the same proof as backpressure optimization to show that the second backpressure method provides a solution that minimizes one-step weighted delay: $\min \sum_{i \in V_a} \text{del}(i,t)\hat{p}(i,t)$.

Scenario	$N_{aircraft}/N_{sectors} = 0.4$	$N_{aircraft}/N_{sectors} = 0.6$	$N_{aircraft}/N_{sectors} = 0.9$
Random	68	100	152
Bimodal	68	100	152
Crossing	20, 28, 20	30, 40, 30	46, 60, 46
Hub and Spoke	12 x 6	17 x 6	25 x 6

Table A.1: Number of aircraft per scenario for different values of $N_{aircraft}/N_{sectors}$, on the 7-radius grid.

See [4] for more details of the proof of optimality for both backpressure and second backpressure.

A.2 Sensitivity Analysis

In this section, we show that the results presented in Fig. 2-5 hold across different congestion levels, at different ratios of aircraft to sectors $N_{aircraft}/N_{sectors}$. The parameters used for testing different ratios of congestion are presented in Table A.1. We varied $N_{aircraft}$ while holding $N_{sectors}$ constant by maintaing the 7-radius grid. Results for unweighted delay are shown in Fig. A-3 and for weighted delay are shown in Fig. A-4. Most figures show the expected result - Backpressure has the lowest unweighted delay, Second Backpressure has the lowest weighted delay, and methods like Reversals and Accrued Delay sometimes outperform Backpressure in fairness. This is sometimes contradicted at lower levels of congestion (where the ratio is smaller), which could be occurring because there are less deconfliction events and thus the probability of deviation from expected results is higher.

A.3 Table of Marginal Contributions

In Table A.2, we lay out the marginal contribution of every SP in every possible ordering, for the example flight illustrated in Fig. 3-3. Summation down each column, divided by the number of orderings, gives the Shapley value profit φ_i for SP i.

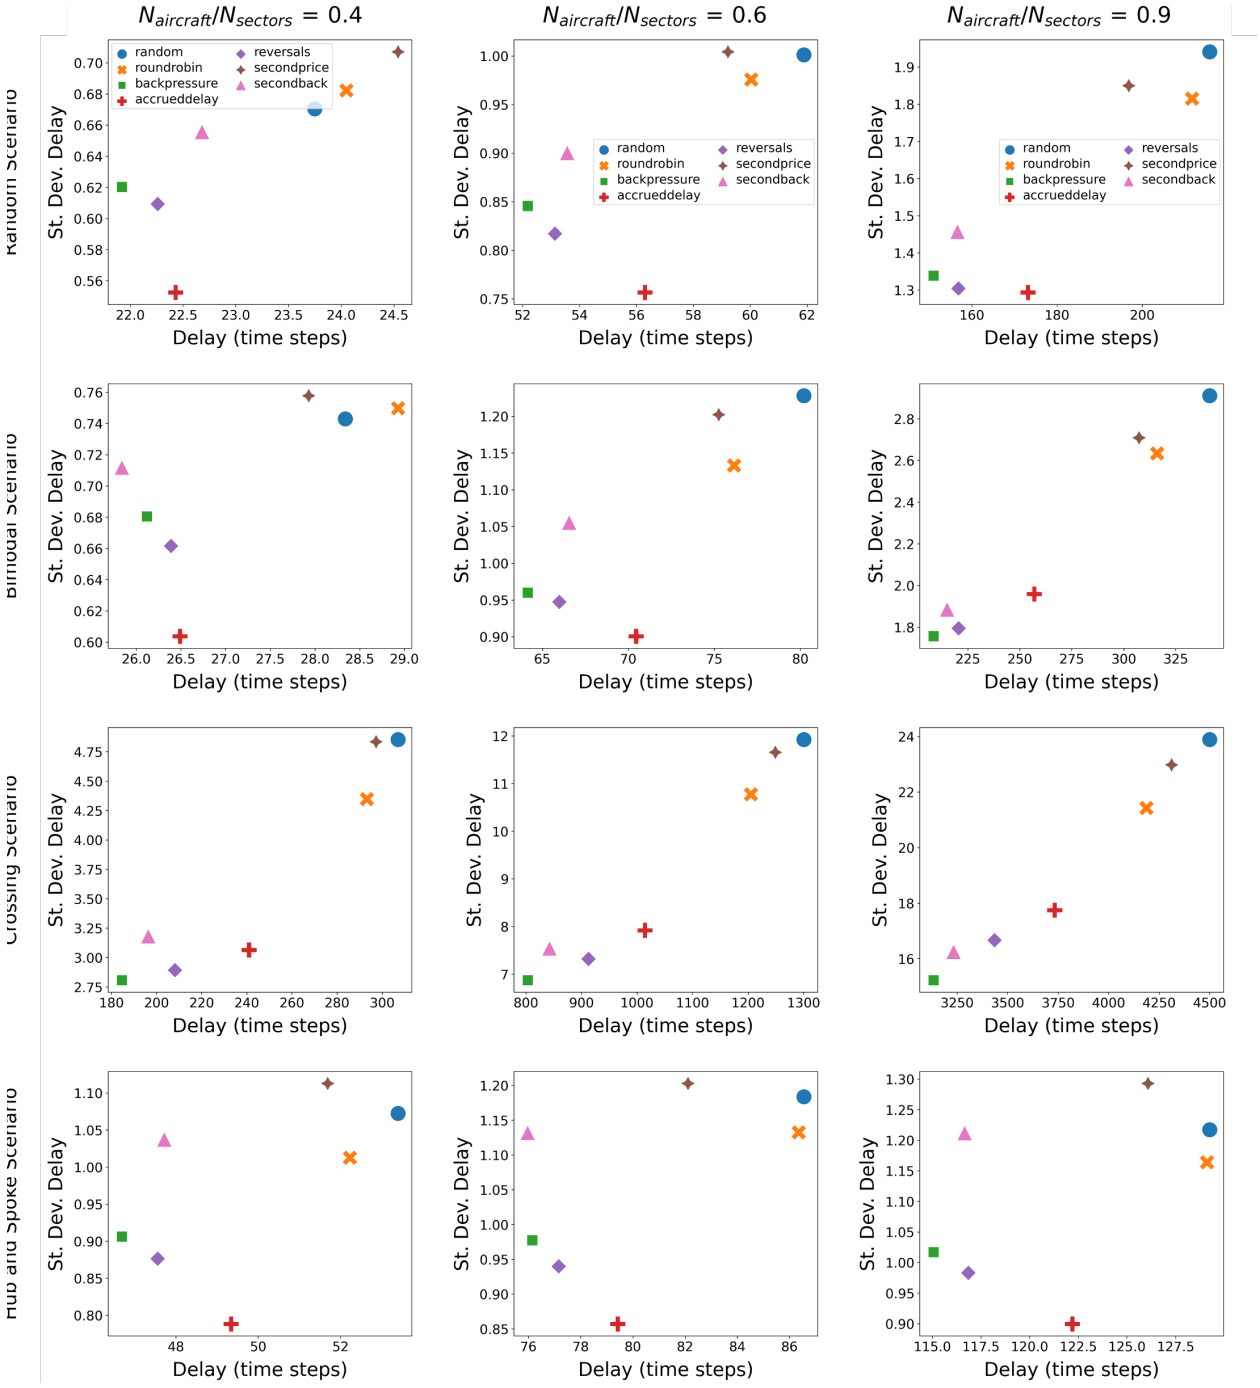

Figure A-3: Sensitivity results of the protocol to different levels of congestion with respect to unweighted delay metrics. Each column demonstrates a different ratio of aircraft to sectors.

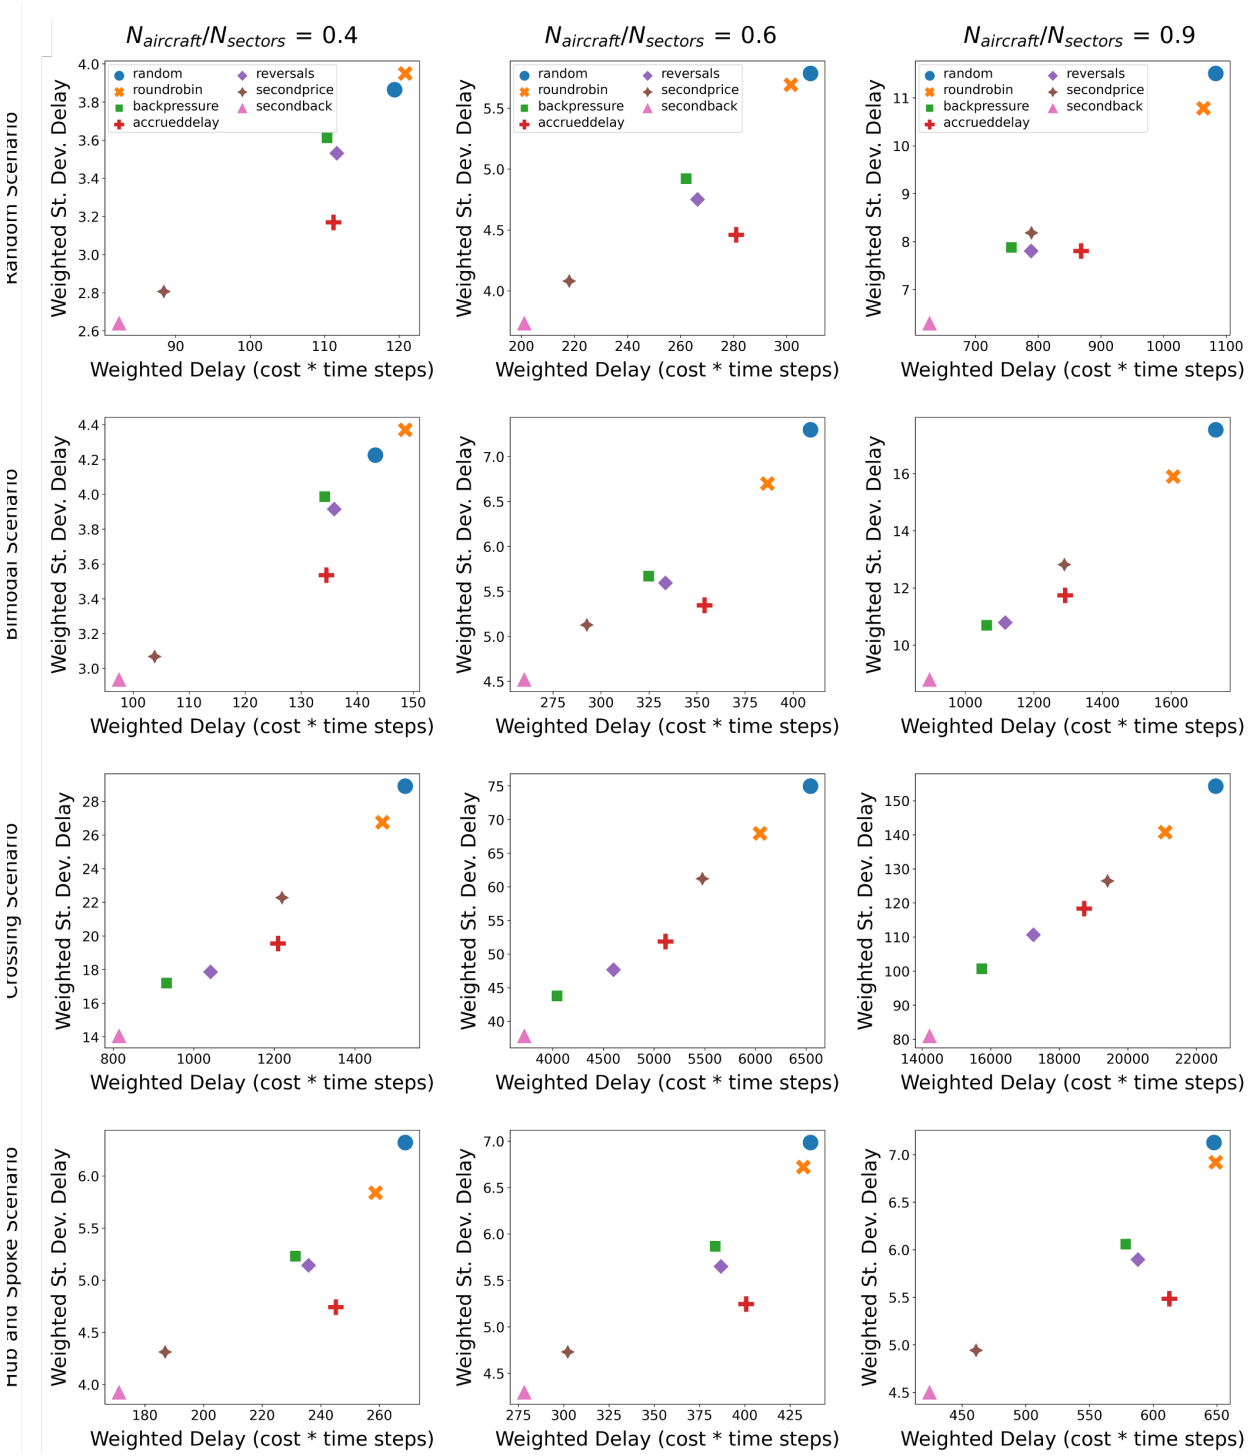

Figure A-4: Sensitivity results of the protocol to different levels of congestion with respect to weighted delay metrics. Each column demonstrates a different ratio of aircraft to sectors.

Ordering	Marginal contribution			
	SP 1	SP 2	SP 3	SP 4
(1,2,3,4)	0	0	0	$3\sqrt{2}$
(1,2,4,3)	0	0	0	$3\sqrt{2}$
(1,3,2,4)	0	0	0	$3\sqrt{2}$
(1,3,4,2)	0	$4-2\sqrt{2}$	0	$5\sqrt{2}-4$
(1,4,2,3)	0	$3\sqrt{2}$	0	0
(1,4,3,2)	0	$4-2\sqrt{2}$	$5\sqrt{2}-4$	0
(2,1,3,4)	0	0	0	$3\sqrt{2}$
(2,1,4,3)	0	0	0	$3\sqrt{2}$
(2,3,1,4)	0	0	0	$3\sqrt{2}$
(2,3,4,1)	$3\sqrt{2}$	0	0	0
(2,4,1,3)	$3\sqrt{2}$	0	0	0
(2,4,3,1)	$3\sqrt{2}$	0	0	0
(3,1,2,4)	0	0	0	$3\sqrt{2}$
(3,1,4,2)	0	$4-2\sqrt{2}$	0	$5\sqrt{2}-4$
(3,2,1,4)	0	0	0	$3\sqrt{2}$
(3,2,4,1)	$3\sqrt{2}$	0	0	0
(3,4,1,2)	$5\sqrt{2}-4$	$4-2\sqrt{2}$	0	0
(3,4,2,1)	$3\sqrt{2}$	0	0	0
(4,1,2,3)	0	$3\sqrt{2}$	0	0
(4,1,3,2)	0	$4-2\sqrt{2}$	$5\sqrt{2}-4$	0
(4,2,1,3)	$3\sqrt{2}$	0	0	0
(4,2,3,1)	$3\sqrt{2}$	0	0	0
(4,3,1,2)	$5\sqrt{2}-4$	$4-2\sqrt{2}$	0	0
(4,3,2,1)	$3\sqrt{2}$	0	0	0

Table A.2: Marginal contributions for each coalition formation permutation.

A.4 Fast Shapley Value Algorithm

Algorithm 1 is the modified DFS-Shapley algorithm adapted from [3] used to calculate the Shapley value for path-based value functions and all origin-destination pairs in Sec. A.4.

Algorithm 1 Shapley Value Calculation

Input: Graph $G = (\mathcal{S}, E)$, list of origin-destination pairs OD, value function v_G
Output: Shapley value $\varphi_i^{(o,d)}$ for every $s_i \in \mathcal{S}$

function DFS-Shapley(G, OD)
  Sort $\mathcal{S}$ by degree descending
  for $s_i \in \mathcal{S}$ **do**
    for $(o,d) \in OD$ **then** $\varphi_i^{(o,d)} = 0$ **end for**
    Sort $Nb(s_i)$ by degree descending, and assign to $M(s_i)$
  end for
  for $s_i \in \mathcal{S}$ **do**
    $S = \{s_i\}$
    Expand-Subgraph($G, (s_i), S, \{s_1 \dots s_{i-1}\}, 1, \emptyset, (\infty), \emptyset, OD$)
  end for
end function
function Expand-Subgraph ($G, path, S, Forbidden, indexOfFirstNeighbor, Neighbors, low, SC, OD$)
  $s = path.last()$; $l = low.last()$
  for $indexOfCurrentNeighbor \in indexOfFirstNeighbor, \dots, |M(s)|$ **do**
    $u = M(s).get(indexOfCurrentNeighbor)$
    if $(u \notin S) \wedge (u \notin Forbidden)$ **then**
      Expand-Subgraph($G, path \cup \{u\}, S \cup \{u\}, Forbidden, 1, Neighbors, low \cup \{\infty\}, SC, OD$)
      $Forbidden = Forbidden \cup \{u\}$; $Neighbors = Neighbors \cup \{u\}$
    else if $u \in Forbidden$ **then** $Neighbors = Neighbors \cup \{u\}$
    end if
  end for
  $path.removeLast()$; $low.removeLast()$
  if $path.length() > 0$ **then**
    $w = path.last()$
    $indexOfFirstNeighbor = M(w).getIndex(s) + 1$
    Expand-Subgraph($G, path, S, Forbidden, indexOfFirstNeighbor, Neighbors, low, SC, OD$)
  else
    for $(o,d) \in OD$ **do**
      for $s_i \in S$ **then** $\varphi_i^{(o,d)} = \varphi_i^{(o,d)} + \lambda_{S'} * (v_G(S,(o,d)) - v_G(S \backslash s_i, (o,d)))$ **end for**
    end for
  end if
end function

Bibliography

[1] Federal Aviation Administration. Unmanned Aircraft System (UAS) Traffic Management (UTM) Concept of Operations Version 2.0. Technical report, Federal Aviation Administration, Washington, DC, Mar 2020.

[2] User:Ludovic.ferre. File:internet connectivity distribution & core.svg - wikipedia.

[3] Oskar Skibski, Talal Rahwan, Tomasz P. Michalak, and Michael Wooldridge. Enumerating connected subgraphs and computing the myerson and shapley values in graph-restricted games. 10(2):1–25.

[4] Christopher Chin, Victor Qin, Karthik Gopalakrishnan, and Hamsa Balakrishnan. Traffic management protocols for advanced air mobility. *Frontiers in Aerospace Engineering*, 2, 2023.

[5] K. Balakrishnan, J. Polastre, J. Mooberry, R. Golding, and P. Sachs. Blueprint for the Sky: The roadmap for the safe integration of autonomous aircraft. Technical report, Airbus UTM, 2018.

[6] McKinsey & Company. The future of air mobility: Electric aircraft and flying taxis. `https://www.mckinsey.com/featured-insights/the-next-normal/air-taxis`, November 2021. Accessed: 2022-12-02.

[7] Federal Aviation Administration. Urban Air Mobility (UAM) Concept of Operations Version 1.0. Technical report, Federal Aviation Administration, Washington, DC, Jun 2020.

[8] EUROCONTROL. U-space ConOps (edition 3.10): U-space capabilities and services to enable Urban Air Mobility. Technical report, EUROCONTROL, July 2022.

[9] Ministry of Economy, Trade and Industry (METI), Japan. Advanced Air Mobility in Japan. Technical report, METI, 2021.

[10] Dennis McLeod and Dennis Heimbigner. A federated architecture for database systems. In *AFIPS '80*, 1980.

[11] Michael Hammer and Dennis McLeod. On database management system architecture. Technical report, Massachusetts Institute of Technology, 1979.

[12] Martín Serrano, Steven Davy, Martin Johnsson, Willie Donnelly, and Alex Galis. Review and designs of federated management in future internet architectures. In *The Future Internet*, pages 51–66, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[13] Deloitte Consulting LLP. Uam vision concept of operations (conops) uam maturity level (uml) 4 - v1.0. Technical report, NASA, December 2020.

[14] Dimitris Bertsimas and Sarah Stock Patterson. The air traffic flow management problem with enroute capacities. *Operations research*, 46(3):406–422, 1998.

[15] 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.

[16] Victor Qin and Hamsa Balakrishnan. Cost-aware congestion management protocols for advanced air mobility. In *10th International Conference on Research in Air Transportation*, 2022.

[17] Lloyd S Shapley. Notes on the n-person game – ii: The value of an n-person game. Technical report, The RAND Corporation, 1951.

[18] Booz Allen Hamilton. Urban Air Mobility (UAM) Market Survey. Technical report, NASA, November 2018. https://ntrs.nasa.gov/citations/20190001472.

[19] Crown Consulting. Urban Air Mobility (UAM) Market Survey. Technical report, NASA, November 2018. https://ntrs.nasa.gov/citations/20190002046.

[20] Parker D Vascik, Hamsa Balakrishnan, and R John Hansman. Assessment of air traffic control for urban air mobility and unmanned systems. page 9.

[21] Michael O Ball, George L Donohue, and Karla Hoffman. Auctions for the safe, efficient, and equitable allocation of airspace system resources. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, pages 507–538. The MIT Press.

[22] Thomas W. M. Vossen and Michael O. Ball. Slot trading opportunities in collaborative ground delay programs. 40(1):29–43.

[23] Luca Corolli, Tatjana Bolić, Lorenzo Castelli, and Desirée Rigonat. Tradable mobility permits for the strategic allocation of air traffic. page 8, 2014.

[24] Brent Skorup. Auctioning airspace. 21:79–109, 2019.

[25] Sven Seuken, Paul Friedrich, and Ludwig Dierks. Market design for drone traffic management.

[26] Dušan Teodorović, Konstantinos Triantis, Praveen Edara, Yueqin Zhao, and Snežana Mladenović. Auction-based congestion pricing. 31(4):399–416.

[27] M. Vasirani and S. Ossowski. A market-inspired approach for intersection management in urban road traffic networks. 43:621–659.

[28] Dustin Carlino, Stephen D. Boyles, and Peter Stone. Auction-based autonomous intersection management. In *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*, pages 529–534. IEEE.

[29] Steven H Low, Fernando Paganini, and John C Doyle. Internet congestion control. *IEEE control systems magazine*, 22(1):28–43, 2002.

[30] Ayesha Atta, Sagheer Abbas, M Adnan Khan, Gulzar Ahmed, and Umer Farooq. An adaptive approach: Smart traffic congestion control system. *Journal of King Saud University-Computer and Information Sciences*, 2018.

[31] Jean Gregoire, Xiangjun Qian, Emilio Frazzoli, Arnaud de La Fortelle, and Tichakorn Wongpiromsarn. Capacity-aware back-pressure traffic signal control.

[32] Inseok Hwang, Jegyom Kim, and Claire Tomlin. Protocol-based conflict resolution for air traffic control. 15(1):1–34.

[33] Yoav Shoham and Kevin Leyton-Brown. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, 2008.

[34] Lee W. McKnight and Joseph P. Bailey. *Internet economics*. MIT Press, 1997.

[35] Raphael Cohen-Almagor. Internet history. 2(2):45–64.

[36] Richard T B Ma, Dah-Ming Chiu, John C S Lui, Misra, Vishal, and Rubenstein, Dan. On cooperative settlement between content, transit and eyeball internet service providers. In *ACM CoNEXT 2008*, 2008.

[37] Dionne Searcey. Level 3, Cogent Affect Web Users In Internet Fight. `https://www.wsj.com/articles/SB112865268189462541`, Oct 2005. Accessed: May 2023.

[38] P Faratin, D Clark, P Gilmore, S Bauer, A Berger, and W Lehr. Complexity of internet interconnections: Technology, incentives and implications for policy. In *The 35th Research Conference on Communication, Information and Internet Policy (TPRC)*, 2007.

[39] Richard T. B. Ma, Dah-Ming Chiu, John C. S. Lui, Vishal Misra, and Dan Rubenstein. Internet economics: The use of shapley value for ISP settlement. *IEEE/ACM Transactions on Networking*, 18(3):775–787, 2010.

[40] Richard T.B. Ma, Dah-Ming Chiu, John C. S. Lui, Vishal Misra, and Dan Rubenstein. Interconnecting eyeballs to content: A shapley value perspective on isp peering and settlement. In *Proceedings of the 3rd International Workshop on Economics of Networked Systems*, NetEcon '08, page 61–66, New York, NY, USA, 2008. Association for Computing Machinery.

[41] Federal Aviation Administration. 86 FR 4390: Remote Identification of Unmanned Aircraft, 2021. Accessed: October 2021.

[42] C. Chin, K. Gopalakrishnan, M. Egorov, A. Evans, and H. Balakrishnan. Efficiency and fairness in unmanned air traffic flow management. *IEEE Transactions on Intelligent Transportation Systems*, 20(9), 2021.