Fifteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2023)

Market Structures for Service Providers in Advanced Air Mobility

Victor Qin, Geoffrey Ding, and Hamsa Balakrishnan
Department of Aeronautics and Astronautics
Massachusetts Institute of Technology
Cambridge, MA, USA
{victorqi, gding, hamsa}@mit.edu

***Abstract*—Proposed concepts of operations for advanced air mobility rely on private service providers being responsible for providing air traffic management services to uncrewed aircraft such as drones and autonomous air taxis. While such proposals are unprecedented in the aviation context, one can draw parallels to the Internet, and the role played by Internet Service Providers in managing web traffic. A study of the evolution of the Internet illustrates that, without clear rules for cooperation around a nascent market, private profit motives incentivize against service provider cooperation, especially for traffic flows that traverse multiple regions managed by different service providers. To address this problem, we propose a profit-sharing mechanism based on the Shapley value that incentivizes service providers to cooperate. We show that this mechanism: (i) ensures that service providers route flights along globally optimal routes, and (ii) encourages service providers to cooperate to alleviate congestion. Finally, we discuss some of the remaining challenges with having a federated network of private service providers supporting traffic management for advanced air mobility operations.**

***Keywords*—economics of advanced air mobility; traffic management service providers; profit-sharing mechanisms**

I. 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 [1]. Current estimates of the density, type, and number of these new flights [2] 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. These characteristics 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 [3]–[5].

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 [3], and urban air mobility (UAM) for operations of larger cargo- and passenger-carrying aircraft in “UAM corridors” [4]. 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 [5], [6]. The importance of USSs is reflected in the notional UTM architecture in Fig. 1; PSUs function in a similar role within UAM [4]. Throughout the remainder of this paper, 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 [4].
- 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 [4]. Such federated architectures–comprising of connected semi-autonomous components–were first proposed in the context of databases [7], [8],

This work was supported by NASA grant #80NSSC22K0794, but this article solely reflects the opinions and conclusions of its authors and not any NASA entity. Victor Qin was additionally supported by a Mathworks Fellowship. Geoffrey Ding was additionally supported by a National Science Foundation Graduate Research Fellowship.

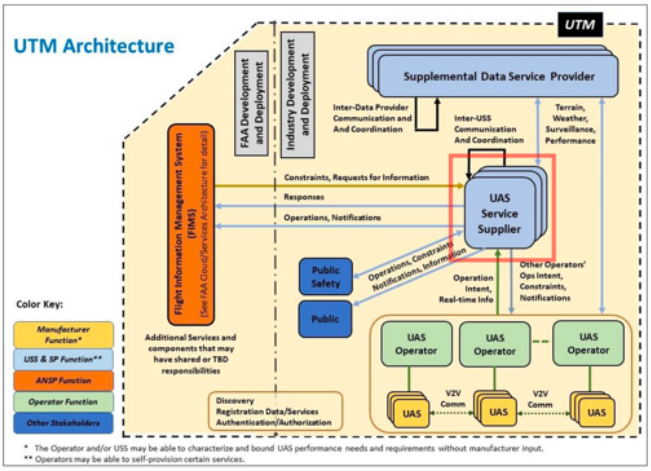

Figure 3. Notional UTM architecture

Figure 1: Proposed UTM system architecture (from [3]). 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).

and have since been considered in the context of the Internet [9].

- 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 in that sector.
- SPs are expected to be primarily private sector entities, although public sector SPs may also exist [10].
- Multiple SPs may provide services in the same geographical region [3].
- AAM flight operators may also be service providers, as long as they satisfy the relevant qualifications.

These envisioned characteristics are still quite loosely defined, and may be in conflict with each other when applied to real-world scenarios. In particular, if SPs are private entities, competition for customers among SPs may counteract the need for cooperation while moving flights between airspace regions managed by different 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 [11]. This method takes the total revenue earned for a flight operation, and divides it in a fair manner 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

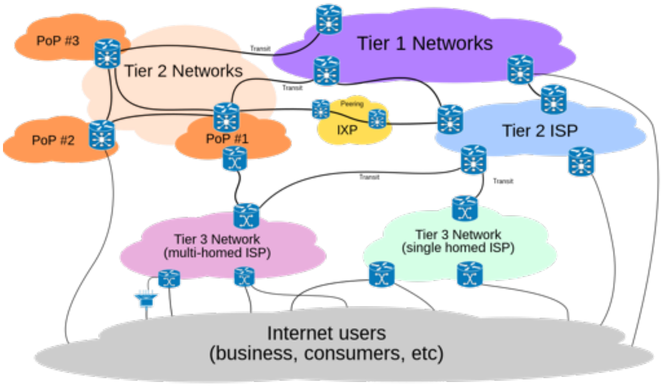

Figure 2: An illustration of how Internet Service Providers (ISPs) interact with each other (from [12]).

to support other SPs in the presence of congestion. This paper focuses on the first four of the desired characteristics listed above, leaving the last two (i.e., multiple SPs operating in the same region, and SPs also being flight operators) as topics for discussion and future research.

The rest of this paper is organized as follows. We first give a brief history of Internet traffic routing, and highlight parallels to current AAM traffic management concepts (Sec. II). 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 Sec. III, 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 Sec. IV, and analyze how the Shapley value would work in such a system (Sec. V). Sec. VI discusses possible challenges with using the Shapley value, as well as open questions that need further investigation. Finally, Sec. VII concludes this work.

II. 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.

A. 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 [3], 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. 2. 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 the Internet to understand how current conceptions of traffic management services for AAM might develop, and to preempt possible problems that may arise.

1) 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 [13].

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 [14].

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 [3]. 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 adapt to the pace of technological innovation, the increase in flight volumes, and the dynamic, on-demand requirements of the AAM operations.

2) 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 [15]. 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 [15], [16]. 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 [17]. 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.

B. Related work

The early history of Internet pricing and economics is well-covered in [13], which describes some of the basic properties of "sender-keeps-all" economics. [17] gives a deeper explanation of interconnection and Internet structures. Ma et al. outlines the concerns with "hot-potato routing" in [18], [19], 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.

III. The Shapley value

The Shapley value was first described by Lloyd Shapley in [11]. 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, $v(S) \rightarrow \mathbb{R}$ 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 [18].

A. 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.

B. 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

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

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.

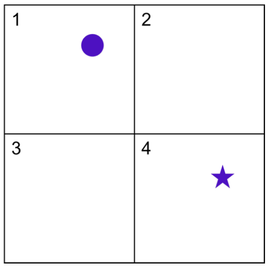

Figure 3: An example environment in which the value of a coalition depends on the presence of a contiguous path between the circle and the star entirely within the coalition. Then, coalitions $\{1, 2, 4\}$, $\{1, 3, 4\}$, and $\{1, 2, 3, 4\}$ have value 1; all others have value 0.

C. An illustrative example

We demonstrate the computation of the Shapley value using the example environment in Fig. 3. Here, we have $n = 4$ agents, given by $N = \{1, 2, 3, 4\}$. We define a coalition S to xhave 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 (ways to form a coalition), this only occurs in two cases: $(1, 4, 2, 3)$ and $(4, 1, 2, 3)$. For example, in the ordering $(1, 2, 4, 3)$, the marginal contribution of agent 2 is 0 because just the sectors $(1, 2)$ do not provide a contiguous path. In the ordering $(1, 4, 3, 2)$, the sectors $(1, 4, 3)$ already provide a contiguous path, and so agent 2 provides no marginal contribution.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

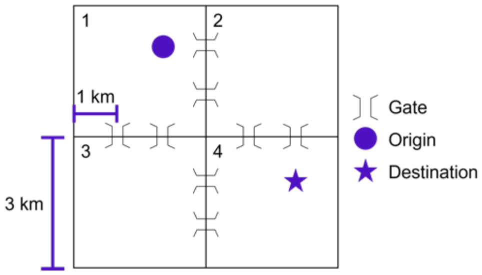

Figure 4: Our model of a small region of airspace consisting of four (numbered) sectors and gates between adjacent sectors that are the only locations where flights may cross a border. Origins and destinations may be arbitrarily located within a sector.

intuition suggested, this 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. III-A:

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.

IV. 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.

A. Airspace structure

We represent airspace as a grid of n sectors defined as bounded polygons, $G = \{P_1, \ldots, P_n\}$. 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. 4 shows this structure applied to a small region of airspace.

B. 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 P_i. For a given coalition of sectors $S \subseteq G$, $\mathcal{R}(f, S)$ 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 $R^*(f, G)$.

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.

C. 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 briefly discuss the case of sectors with multiple competing SPs in Sec. VI-B2.

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 $c_i(R) = \sum_{r_j \in P_i} \|r_j\|_2 \; \forall r_j \in R$. 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 $c(S) = \sum_{i, P_i \in S} c_i(R(S))$; then, the optimal cost using the best route possible is $c^*(S) = \sum_{i, P_i \in S} c_i(R^*(S))$.

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.

V. Analysis and Results

We now apply the Shapley value to the airspace model presented in Sec. IV. We provide and analyze an example scenario and demonstrate how, under our cost and revenue assumptions presented in Sec. IV-B and Sec. IV-C, 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.

A. 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 to include the cost structures from Sec. IV. 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 $c(S)$:

$$v(S) = u(S) - c(S). \tag{2}$$

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, $v(S)$ is also zero 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. III-C, 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.

1) 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. 5 are shown in Table I. An explicit list of all permutations of coalition formation and the respective marginal contributions of each SP may be found in Table III of the Appendix.

2) 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 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.

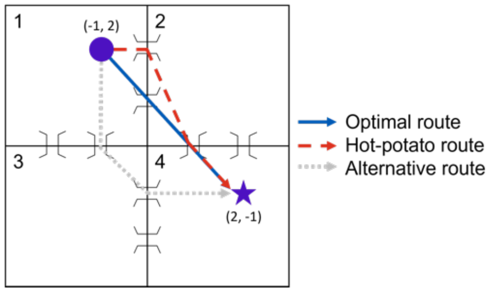

Figure 5: Optimal route in solid blue, hot-potato route in dashed red, and alternative route in dotted gray.

B. 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. 5 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 II, where the optimal routing strategy generates the most profit for every SP compared to other routing strategies, including the “hot-potato” routing strategy used in the Internet.

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.

C. 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

TABLE I. COMPUTATION OF SHAPLEY VALUES AND PROFIT SHARES FOR ONE ORIGIN-DESTINATION PAIR. AGENT 2 AND 3 EARN DIFFERENT PROFITS BECAUSE THE ADDITION OF AGENT 2 PROVIDES A SLIGHTLY SHORTER PATH BETWEEN THE ORIGIN AND DESTINATION UNDER THE DISTANCE BASED VALUE FUNCTIONS PRESENTED IN SEC. IV-B AND IV-C AND COMBINED IN (2)

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		

TABLE II. COSTS, PAYMENTS, AND PROFIT FOR DIFFERENT ROUTES TAKEN, WITH PROFIT SHARING BASED ON SHAPLEY VALUE FOR THE ORIGIN-DESTINATION PAIR IN FIG. 5. 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

in the scenario (a measure of social welfare, in kilometers). The airspace is structured as described in Fig. 4, 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. 4 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. 6. 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.

VI. 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.

A. Potential impacts of profit-sharing based on Shapley value

1) Truthful cost reporting: Since profits are computed taking as input the costs as reported by the service providers,

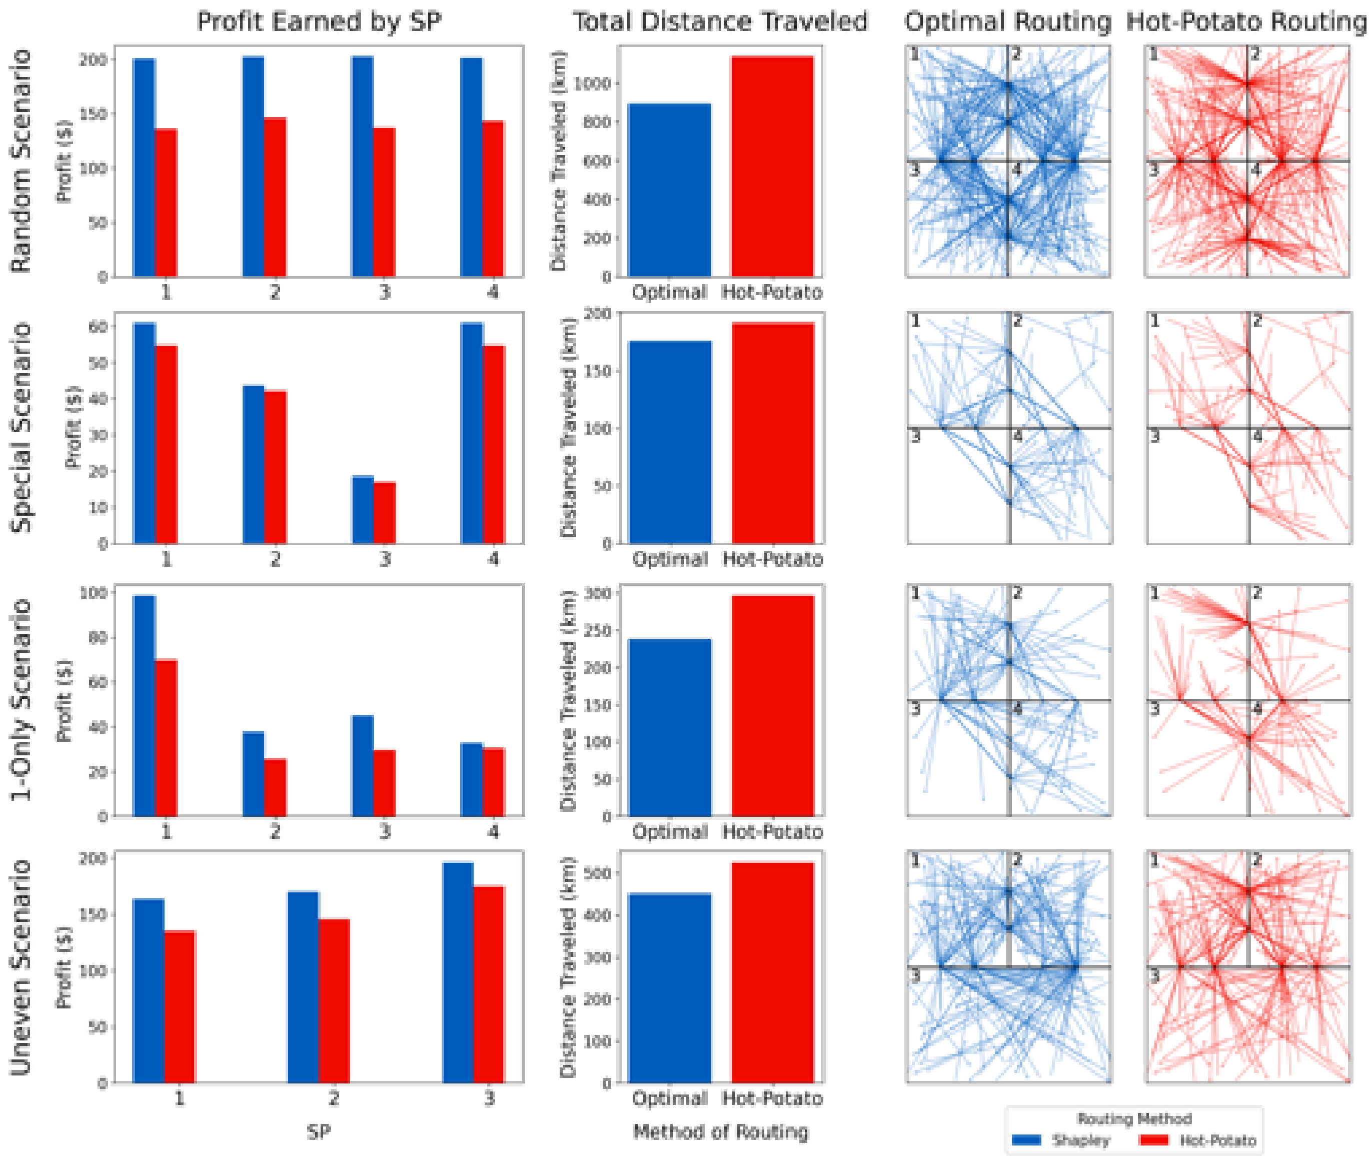

Figure 6: Profit per SP, total distance traveled, and routes traversed in each of the four traffic scenarios simulated.

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 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 [20]. 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.

2) 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.

3) 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 III-A.

B. Future Work

This paper represents an important first step in developing a market structure for AAM traffic management service providers. We now discuss several areas of future work needed to develop effective AAM traffic management techniques. In these discussions, we assume that the Shapley value can provide incentive-based profit-sharing solutions that encourage a baseline level of cooperation among different SPs. However, there remain unaddressed challenges that may require further study and even changes to the regulatory landscape for AAM.

1) Interaction with intra-SP traffic management: In this work, 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 or through centralized optimization [21]–[23]. 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.

2) Sectors with multiple competing service providers: A significant feature of the proposed concept of operations has been that SPs will not get exclusive rights to operate in any region, and that multiple SPs would be allowed to operate in the same airspace sector. This leads to the difficult problem of deconflicting flights working with different SPs, in addition to the issues of SP coordination. Each SP has an interest in seeing the flights for which it is responsible avoid delay, yet this could easily lead to a game of chicken (and raise safety concerns) as SPs refuse to divert flights. We offer some potential solutions to this problem:

- A "super-SP" entity, whether public or private, assumes ultimate responsibility for adjudicating conflicts between SPs. This could be a responsibility rotated between SPs on a regular interval, or assigned to the public SP.
- Using the Shapley value, interactions between SPs are regulated to encourage cooperation. Flights passing through a sector could split profit evenly to all SPs in that sector.
- A sector is further segmented into smaller regions, where each region is only controlled by one SP. SPs would swap control of different segments within the sectors on a regular interval.
- All flights use tactical deconfliction methods exclusively to coordinate avoidance maneuvers with other flights, without strategic intervention from SPs. This removes the need for SP coordination, but has an increased safety risk and possible loss in efficiency compared to strategic deconfliction methods.

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 [3]. 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 [10].

VII. 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, and promising future directions in the development of traffic management strategies for advanced air mobility systems.

Appendix

In Table III, we lay out the marginal contribution of every SP in every possible ordering, for the example flight illustrated in Fig. 4. Summation down each column, divided by the number of orderings, gives the Shapley value profit φ_i for SP i.

TABLE III. Marginal contributions for each coalition formation permutation.

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

References

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

[2] 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.

[3] Federal Aviation Administration, "Unmanned Aircraft System (UAS) Traffic Management (UTM) Concept of Operations Version 2.0," Federal Aviation Administration, Washington, DC, Tech. Rep., Mar 2020.

[4] ——, "Urban Air Mobility (UAM) Concept of Operations Version 1.0," Federal Aviation Administration, Washington, DC, Tech. Rep., Jun 2020.

[5] EUROCONTROL, "U-space ConOps (edition 3.10): U-space capabilities and services to enable Urban Air Mobility," EUROCONTROL, Tech. Rep., July 2022.

[6] Ministry of Economy, Trade and Industry (METI), Japan, "Advanced Air Mobility in Japan," METI, Tech. Rep., 2021.

[7] D. McLeod and D. Heimbigner, "A federated architecture for database systems," in *AFIPS '80*, 1980.

[8] M. Hammer and D. McLeod, "On database management system architecture," Massachusetts Institute of Technology, Tech. Rep., 1979.

[9] M. Serrano, S. Davy, M. Johnsson, W. Donnelly, and A. Galis, "Review and designs of federated management in future internet architectures," in *The Future Internet*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 51–66.

[10] Deloitte Consulting LLP, "UAM Vision Concept of Operations (ConOps) UAM Maturity Level (UML) 4 - v1.0," NASA, Tech. Rep., December 2020.

[11] L. S. Shapley, "Notes on the n-person game – ii: The value of an n-person game," The RAND Corporation, Tech. Rep., 1951.

[12] User:Ludovic.ferre. File:Internet Connectivity Distribution & Core.svg – Wikipedia. https://commons.wikimedia.org/wiki/File:Internet_Connectivity_Distribution_%26_Core.svg.

[13] L. W. McKnight and J. P. Bailey, *Internet economics*. MIT Press, 1997.

[14] R. Cohen-Almagor, "Internet history:," vol. 2, no. 2, pp. 45–64. [Online]. Available: https://services.igi-global.com/resolvedoi/resolve.aspx?doi=10.4018/jte.2011040104

[15] R. T. B. Ma, D.-M. Chiu, J. C. S. Lui, Misra, Vishal, and Rubenstein, Dan, "On cooperative settlement between content, transit and eyeball internet service providers," in *ACM CoNEXT 2008*, 2008.

[16] http://www.cio.com/article/252501/internet-level-3-cogent-spat-affects-15-percent-of-internet.html.

[17] 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.

[18] R. T. B. Ma, D.-M. Chiu, J. C. S. Lui, V. Misra, and D. Rubenstein, "Internet economics: The use of Shapley value for ISP settlement," *IEEE/ACM Transactions on Networking*, vol. 18, no. 3, pp. 775–787, 2010. [Online]. Available: http://ieeexplore.ieee.org/document/5464383/

[19] R. T. Ma, D.-M. Chiu, J. C. S. Lui, V. Misra, and D. 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*, ser. NetEcon '08. New York, NY, USA: Association for Computing Machinery, 2008, p. 61–66. [Online]. Available: https://doi.org/10.1145/1403027.1403041

[20] Federal Aviation Administration, "86 FR 4390: Remote Identification of Unmanned Aircraft," 2021, accessed: October 2021.

[21] V. Qin and H. Balakrishnan, "Cost-aware congestion management protocols for advanced air mobility," in *10th International Conference on Research in Air Transportation*, 2022.

[22] 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*, vol. 20, no. 9, 2021.

[23] C. Chin, K. Gopalakrishnan, H. Balakrishnan, M. Egorov, and A. Evans, "Protocol-based congestion management for advanced air mobility," in *Fourteenth USA/Europe Air Traffic Management Research and Development Seminar*, 2021.

Author Biographies

Victor Qin is a PhD student in the Department of Aeronautics and Astronautics at the Massachusetts Institute of Technology (MIT). His research interests include coordination and cooperation in multi-agent systems, as well as mechanism design for traffic management.

Geoffrey Ding is a PhD student in the Department of Aeronautics and Astronautics at MIT. He is interested in optimization of and interactions within multi-agent systems, especially for future air transportation systems.

Hamsa Balakrishnan is the William E. Leonhard (1940) Professor of Aeronautics and Astronautics at MIT.