Evaluating Fairness in UTM Architecture and Operations

Airbus UTM

Version 1.1
February 2020
TR-010

Peter Sachs, Airbus UTM
Antony Evans, Airbus UTM
Maxim Egorov, Airbus UTM
Robert Hoffman, Metron Aviation
Bert Hackney, Metron Aviation

Table of Contents

1. Introduction

UAS (Unmanned Aircraft System) Traffic Management (UTM)[1] operations are rapidly advancing in complexity around the world, as technological capabilities move from simulation to real-world operations and as standards take shape. Numerous demonstrations and trial projects have begun to validate the capabilities of UTM providers to manage multiple aircraft, including urban air mobility (UAM). In the United States, several UAS operators have received their Part 135 air carrier certificates, a key regulatory milestone that enables payload-carrying commercial operations.

Now that traffic numbers are increasing and complex operations are becoming a reality, we have a unique opportunity to refine UTM architectures to ensure that our system designs, implementations, and operating rules remain fair for all users, rather than benefiting first-movers and the largest operators or UTM Service Suppliers (USSs).[2]

Most industry participants are outwardly collaborative and purport to be “good actors.” But history shows us that, faced with the competitive pressure of open markets, participants will seek ways to gain an advantage, in some cases unfairly. This can have severe consequences for the ultimate goal of providing access to our airspace and enabling a wide variety of future missions.

This document fills an important gap by articulating what fairness is in a UTM context, and why it’s important to start addressing it now. Ultimately, “fair airspace access” is a function that will be enacted through policies and requirements on services such as demand management and deconfliction, and through rules set for negotiating a resolution to a conflict between aircraft. Fairness is closely related to efficiency of airspace management and operations, as well: In general, when multiple stakeholders choose to cooperate, all users realize the benefits of more efficient operations. We believe we can provide a useful framework for identifying, describing and quantifying fairness-related problems, while remaining agnostic as to how best to approach their solutions.

Quantifying fairness is inherently difficult, because up until now it has been tied closely with other metrics, such as those related to airspace efficiency or auction practices. Ensuring we do it correctly within UTM will be even harder. We recommend that industry and academia conduct a series of near-term studies to better quantify the effects of unfair operations. The goal should be a path forward that keeps the UTM ecosystem competitive and open to innovation, while ensuring fair airspace access for UAM and conventional aviation. To further balance the unique needs of commercial spaceflight launches and recoveries, this work should be as broadly applicable as possible. We also see the need to define and collect a series of metrics related to fair and unfair operations, as this data will help industry and regulators make informed operational and policy decisions.

---

[1] For consistency, we use the term UTM in this document. UTM in Europe is referred to as U-Space.

[2] For consistency, we use the term USS in this document. A U-Space Service Provider, or USP, is an equivalent term within the European U-Space concept.

2. What is fairness and why is it important?

2.1. Fairness as a guiding principle in UTM

The Airbus UTM architecture is built upon nine guiding principles (see sidebar) (Balakrishnan *et al.*, 2018). These criteria are intended to ensure that overall system design considerations around safety, security, reliability and interoperability (among others) are not compromised by a desire to deploy UTM systems as quickly or inexpensively as possible. The ninth guiding principle is fairness: that the UTM ecosystem ensures that availability and access to airspace takes into account the needs of all stakeholders. This mirrors one of the principles in ICAO's UTM Framework (ICAO, 2019), that "access to the airspace should remain equitable." This view on fairness is also supported by the FAA UTM Conops (FAA, 2018), where UTM must "maintain fair and equitable access to airspace," and the European U-Space Conops (Hately *et al.*, 2019), where UTM must "guarantee equitable and fair access to airspace for all users."

Nine Guiding Principles for UTM Systems
from the Airbus UTM Architecture

1. **Safe.** The UTM system must ensure safe operations at all times.
2. **Scalable.** The UTM ecosystem can support high numbers, varieties, and densities of operations and service providers.
3. **Interoperable and Compatible.** The UTM ecosystem must be conducive to operations involving multiple stakeholders, which includes providers, implementations and jurisdictions.
4. **Reliable.** The UTM ecosystem must be sufficiently reliable and available for safe operation at scale.
5. **Secure.** The UTM ecosystem is sufficiently secure for safe operation, including appropriate authorization, authentication, and defense-in-depth mechanisms..
6. **Open architecture.** The UTM ecosystem's design meets normal criteria for architectural "openness."
7. **Future-proof.** The UTM ecosystem must support vehicles, missions, and systems in all environments and airspaces — including future, unforeseen uses.
8. **Risk-Aware.** It must be possible to know and manage safety and failure risks in the UTM ecosystem.
9. **Fair.** The UTM ecosystem ensures that availability and access to airspace takes into account the needs of all stakeholders.

The current UTM environment lacks a clear and agreed upon set of mechanisms, protocols, or services for resolving competition or conflicts for resources, such as airspace, vertiports, or air traffic services. This has not been an operational problem yet, because the number of operations has not warranted it. But the inability to resolve conflicting requests for resources will soon be an impediment to integration of new entrants. Of particular concern is establishing a UTM environment in which flight operators and other stakeholders view their access to airspace and related resources as fair.

Preliminary industry standards work related to fairness is underway today. But these efforts are related to collecting metrics to better understand the problem, and research and testing will be required to find solutions that best meet the needs of operators and other UTM stakeholders. Regulators and competent authorities may have a variety of roles in helping to ensure fairness.

Sometimes this may take the form of specific policies and regulations. Other times, it may be simply to nudge the industry toward consensus decisions when collaborative processes aren't otherwise yielding positive results.

2.2. Defining fairness and equity

In everyday use, the terms fairness and equity are often used interchangeably, even within the Air Traffic Management (ATM) community. For the purposes of this paper,

> Fairness is the state in which each stakeholder's welfare is increased to the extent possible, given limited resources, after taking proper account of disparate claims and individual circumstances.

This definition has been adopted in similar form by other treatments of the subject in traditional ATM, such as (del Pozo de Poza, 2012) and (Metron Aviation, 2003). Thus, "equity" refers to the special case of fairness when stakeholders with similar characteristics are treated the same. This is a key element of fairness, and often its primary focus. However, given this distinction between terms, in a UTM system with a variety of stakeholders — which may have different characteristics, claims and circumstances — similar participants should be treated equitably. At the same time, implementations should also treat disparate participants fairly, but not necessarily equitably.

There are three reasons why we need fairness in allocation of UTM resources, which are based on obligation and the need for cooperation and efficiency.

Obligation: Unlike today's ATM environment, there may not be a central entity in the UTM environment to perform and oversee resource allocation; it may be accomplished by decentralized or federated systems. Our present systems of centralized air traffic management oblige air navigation service providers (ANSPs) to act fairly. In a decentralized UTM architecture, the various participants can similarly be obliged to be fair by regulation, policy, or "rules of the road," laid down by the civil aviation authority. These obligations must, however, be enforced by a central entity, such as the ANSP or civil aviation authority. Under any of these schema, the decisions we make today in system design and architecture should ensure that fairness is addressed.

Cooperation: When participants feel that they are being treated unfairly, they are more likely to take action on their own to assure better outcomes for themselves. In aggregate, this results in a breakdown of cooperation more widely, and puts stress on the entire system. Fair allocation improves customer satisfaction; operators are more likely to participate in the decision-making process; trust grows; and enhanced understanding leads to sharing of improved data and user intent information. Such cooperation leads to efficiency improvements for all participants.

Efficiency: The third argument for fairness is that it can enhance efficient use of resources, both indirectly through increased cooperation, and directly by incorporating utility (preferences) of the participants into the resource allocation scheme. Making efficient use of resources is always of interest to those managing or operating within the system, and may be of broader interest from a public acceptance and sustainability standpoint. This is very much in line with the view of modern economists, who focus on whether scarce resources are awarded to those entities that can make most effective use of them. However, whether an entity can make effective use of the resources they receive depends very much on their business preferences and the utility for those resources, which are often proprietary. The efficiency and fairness of a system may therefore be enhanced by capturing and using these preferences without exposing them to other users and operators. Efficiency can also be enhanced by encouraging competition, and ensuring that resources are not monopolized. The latter is a further key outcome of a fair allocation.

2.3. Related definitions

This paper uses a number of terms that have specific meanings, described here.

Centralized and federated: The implementation of any UTM system falls along a spectrum from fully centralized at one end, to fully federated, or decentralized, at the other. Within a fully centralized system, a single entity provides all data and services and operators have no choice of who to sign up with for their missions. This is similar to how most ANSPs operate today, and is sometimes referred to as a monopolistic system. In a fully federated system - sometimes called a competitive system - all services can be provided by any number of competing data and service providers, which differentiate their offerings by features, prices or other attributes. Between these extremes, different jurisdictions may desire to delegate some functions, such as airspace information, surveillance data or filing operational intents, to a central entity. Federated services may provide supplemental data, deconfliction services or more advanced functions. Different architectures may be appropriate for different use cases and in different regions. While fairness must be considered in both centralized and federated architectures, the more distributed the services, the greater the challenge to ensure fairness.

Operational intent: There is an ongoing debate within industry and standards bodies about the most appropriate terminology to describe the messages that comprise an operator's intent, as well as the approved or deconflicted routing that the vehicle is expected to fly. We use "operational intent" to refer to a four-dimensional trajectory (4DT) or a four-dimensional volume request. The term is distinct from, but analogous to, an ICAO-formatted flight plan used by conventional aircraft.

Strategic and tactical: In the UTM context, we refer to "strategic" actions as pre-departure, and "tactical" actions as in-flight. We recognize that this is, however, context dependent, as some traffic flow management actions in traditional ATM are considered "strategic."

2.4. Why must fairness be considered now?

2.4.1. Fairness is a concern even at low traffic volumes

Among industry and standards groups, conversations sometimes include arguments that traffic levels will not be sufficiently high to impact fairness well into the future. Even if conflicts for resources are rare, they should be handled fairly. Moreover, distribution of delay across operators, which is a common metric for fairness, is driven by traffic density. Golding (2018) has shown that, in UTM, the absolute number of vehicles in a region of airspace is not, by itself, a good metric for indicating density. In fact, UTM traffic can be dense at low traffic volumes.[3] This suggests that there may be issues with fairness at lower traffic volumes than expected.

2.4.2. Incentives must be established for fair behavior

It is tempting to hope that as we reinvent our airspace, all UTM operators will act in a way that betters both themselves and the common good. While this desire is certainly commendable, history suggests that organizations placed in a competitive environment will naturally focus on their own betterment, and may even work to undermine their competitors. In fact, game theory (and much of economics) is predicated on the assumption that players make rational decisions to maximize their utility. This can lead to a “tragedy of the commons,” where system level utility is sacrificed for individual utility. This does not mean that UTM operators will break rules, but they will find creative ways within the confines of the rules to better their situation, even if it is at the expense of system level efficiency. This suggests that rules for fair allocation have to be carefully designed and agreed upon, and have proper incentives for behavior that is fair and efficient at the system level. Initially, this could include logging of unfair behavior followed by appropriate penalties.

2.4.3. Fairness issues must be addressed proactively, not reactively

Among industry participants and observers, some have proposed that fairness issues need not be resolved until they manifest, reducing the need to consider fairness in the early stages of development. However, the architecture of the UTM system and rules governing its operation must be developed from the beginning in such a way as to accommodate future traffic volumes and types of operations. Architecture, and especially rules governing deconfliction, can have a significant impact on fairness, as discussed in Section 4. It is therefore important to consider fairness throughout the development of the UTM architecture and standards that govern its operation, when adjustment is easy. Failing to account for future needs contributes to technical debt as well, which adds time, cost and programmatic delays to all industry participants in UTM systems, and hurts the long-term viability of the UTM ecosystem. While we can’t anticipate every eventuality, there is an opportunity now to avoid later rewrites of standards and requirements by leaving room for fairness considerations. However, it is also important not to discourage early adopters by overly onerous rules to enforce fairness. Rules must therefore be developed carefully and in close collaboration with industry.

3. UTM use cases

Ensuring fair airspace access to airspace in a UTM context will require more than simply balancing the requests of two vehicles in flight that are in conflict with each other. Some considerations will most certainly center around interactions between UAS and other aircraft and, just like today’s air traffic management strategies, may be resolved either before takeoff or during flight. Access will also be affected by other constraints, such as noise, other

[3] This can happen in a variety of situations, but of note here are regions with opposite-direction or converging routes that must be deconflicted; or vehicles operating with large separation minima (or within large airspace volumes) because of the uncertainty associated with communication, navigation or surveillance requirements.

environmental constraints, security, and privacy, which will require resource allocation in a fair way.

To help frame the types of problems we envision, this section describes four very broad use cases in which fairness issues will arise. We focus on the direct effects that a vehicle, operator or service provider might experience. We therefore specifically consider fairness from the perspective of the operator and service provider. Indirect effects to consumers, users or other stakeholders are intentionally left out of the scope of this initial analysis. To properly evaluate and implement fair responses, these use cases will need further subdivision and refinement.

3.1. Airspace interactions between vehicles

There are currently two paradigms for representing how an operator would share its intent to use the airspace. One is through a series of airspace volume reservations: three-dimensional shapes that adjoin one another and cover all flight phases, including regions where a vehicle might hover, loiter or survey. Each volume has beginning and ending times, and the vehicle is expected to stay within that volume. However, the vehicle need not enter at one end and follow the most direct path to exit at the opposite end: it could perform any set of maneuvers, as long as it leaves the volume before the reservation expires. This concept has been advanced through the NASA TCL demonstrations (Rios *et al*., 2018-1, 2018-2); the NASA UTM Architecture (FAA, 2018); ongoing industry standards development such as DSS (Discovery and Synchronization Service); and related implementations, such as InterUSS (InterUSS, 2019). The second paradigm represents a four-dimensional trajectory through a series of waypoints, each defined by latitude, longitude, altitude and crossing time.

Regardless of how operational intent is represented, there are important issues of fairness that arise when trying to avoid a conflict. Deconfliction maneuvers always have some cost, which may in the best case be negligibly small. These measures either increase flight time or distance traveled, which may increase battery usage and decrease overall safety margins. They may also delay departure, impacting the business efficiency of the operator and demand for their service from customers. Some vehicles, operators or missions may have the tolerance to absorb a greater amount of deviation from the original or desired trajectory. Other operators may be better able to handle a change in speed or departure time than a change in lateral routing. Finally, many vehicles may need to have their desires balanced in relation to a prioritization scheme. In these cases, “fairness” isn’t an explicit, standalone service, but rather a function that is enacted through demand management and deconfliction tools, and through rules set for the resolution of an identified conflict.

The above issues may initially be easier to resolve when airspace demand is relatively low and most encounters are pairwise. However, as demand increases for package delivery, inspection, public safety and other missions, the solutions we come up with may not scale. Just as in today’s air traffic management context, operators may want a predefined pool of departure times, slot constraints or other allocated resources so that they can adjust their schedules across a large number of flights while remaining within those constraints.

3.2. Airspace structures with conditional usage

Existing research shows that adding any level of airspace structure — even just a handful of compulsory flyover waypoints — increases safety by reducing the rate of collisions and losses of separation (Bruggeman, 2015). As demand for UAS traffic increases, we expect that there will be a desire to design and implement a variety of airspace structure concepts to improve the safety and efficiency of operations, but at the expense of free-flight opportunities as a greater proportion of vehicles are directed into predefined route segments. The exact form and definition of these airspace structures, as well as the interactions between them, may have profound impacts on fair airspace access for users. For example, consider a network of corridors across cities, predefined in a playbook based on operational conditions such as winds and airport runway configurations. If access to such a dense network were limited to only the operators for which it was requested, other users could have difficulty finding an available route that avoids all reserved routes. The outcome is the same if a route network is nominally available to all operators, but only connects to takeoff and landing points with restricted access, or if the performance requirements for using a given route effectively bar all but a single operator from using it. Thus, the well-intentioned structure implemented to improve efficiency and safety may result in reduced airspace availability for others. This suggests the need for processes that create routes based on the demands of many users, not just the largest operators or the first mover in a region.

3.3. Landing sites

There are a number of fairness questions surrounding allocation of constrained resources, specifically those used for takeoff and landing. Exclusive use of landing sites by one operator could have significant implications for fair access to the market, especially in high demand areas where it is difficult to build landing sites, such as city centers. Other landing sites will be shared - especially those leveraging government grants and matching funds. The allocation of arrival and departure slots at these landing sites is most comparable to the issues faced by slot-constrained airports today, where high demand, known runway occupancy time and a limited number of runways (physically, as in the case of London Heathrow, or operationally on days with poor weather at many large airports) combine to require a coordinated slot management program to maintain fairness and efficiency.

The first generations of UAM vertiports may have a limited number of pads that can be used for takeoff and landing, and vehicles will need to remain on the pad or be moved to an adjacent parking spot for several minutes before taking off again. In the best case, operators will swap batteries, but if a vehicle must recharge, it will need to remain in a parking spot for significantly longer. Operators will have to coordinate their schedules and account for contingencies (no fully charged batteries, or a late outbound passenger). Vertiports may not have room to park additional vehicles, nor will vehicles have the ability to hold in the air for more than a couple of minutes while waiting for a parking spot to open up.

The pad allocation problem risks several types of negative outcomes. Schedule too tightly, and there is a safety and/or diversion risk, especially as a problem early in the day cascades through the day. Leave too much slack in the schedule to allow for unforeseen delays, and pads remain empty, even as operators have demand to otherwise use them.[4] Neither scenario can be considered efficient, though setting quantitative diversion and unallocated-pad rates may help in evaluating models and solutions to this problem. Efficiency will be of particular importance to business cases that are sensitive to delay, such as medical package transport and UAM.

Fairness will also need to be considered for other resource allocation problems. For example, future regulations designed to limit noise impacts might limit the total number of flights from all operators over a particular community based on their noise, creating fairness of access issues.

3.4. Interoperability and compatibility between multiple USSs

At a broader level, the mechanisms we create to implement multiple-USS concepts may create unfair conditions as a UTM system grows. Consider an early system with only one or two USSs that coordinate a relatively small number of flights. The first and second USSs may conform to industry standards for service provider interoperability, such as DSS. Their implementations may involve very large airspace volume reservations (much bigger than what the vehicle will actually occupy), but because density is low, there are relatively few conflicts. These two USSs may also enter into direct discussions and side agreements so that one's operations don't impinge on the other: shifting operations to different times of the day, or perhaps keeping most routes over certain parts of a city.

Depending on how access to resources is granted, this could create compatibility issues and first-mover advantages when the system grows and additional USSs join that UTM system. Those USSs may find that the ways of working preclude many routes, and the first one or two USSs may not have any incentive to negotiate conflict resolutions beyond the bespoke agreements they already have. New entrants may be disadvantaged in the negotiation process.

We consider this type of scenario to be interoperable (USSs all conform to industry standards and data protocols) but incompatible. Therefore, incompatible operations tend to be a flag of unfairness in the system, which may be captured through metrics related to efficiency.[5] Since capacity is artificially suppressed in ways that favor some operators, the airspace resources are unfairly allocated. These issues present themselves in jurisdictions with early UTM adoption such as Australia, where concerns over a first-entrant USS controlling drone access to the airspace are already at the forefront of UTM discussion (Evans, 2019).

[4] This scenario can have damaging consequences for the economic viability of UAM operations. Excess demand will increase prices, making UAM less accessible to the general population. Since UAM purports to enable improved mobility and address social equity issues, this necessitates optimizing pad utilization to ensure adequate supply.

[5] Inefficient operations in air traffic are characterized by, e.g., long delay times, denied slot allocations, and very long reroutes; similar effects may manifest in UTM under this scenario as well.

4. Implications for fairness in UTM concepts

This section outlines the fairness implications for a number of concepts proposed for UTM. While we do not claim that these approaches are the solution to airspace fairness, we outline the pros and cons associated with each concept to provide a broad and comprehensive survey of the potential solution space.

4.1. Exclusive allocation of airspace blocks may be inequitable

Skorup (2018) proposes an approach whereby some resources, such as airspace blocks demarcating aerial travel corridors, could be allocated to users in advance, through an auction. This would allow exclusive licenses to be auctioned to operators for use of the corridors, similar to the way regulators auction radio spectrum licenses and offshore wind energy sites. This approach has some benefits, in that exclusive rights to corridors would "allow transfer and sale to more efficient operators and would also give operators the certainty they need to finance the substantial capital investments" (Skorup, 2018).

However, depending on how the airspace blocks are defined, such exclusive rights may lead to significant inequities. This is because the access to the airspace is based entirely on the criteria for winning the auction. Only one operator gains that right, to the exclusion of all other operators. If there is no requirement for operators to use the airspace they win, supply could be artificially constrained and access starved for other operators who would want to use some of that airspace. This would be very limiting for high-demand blocks of airspace, and if those operators were forced out of the market, we could be left with a monopoly. While such consequences are highly dependent on how the auction is designed, the negative consequences of poor designs can be significant, as demonstrated in the California electricity sector (e.g., Borenstein, 2001; Morey, 2001). The ability for an operator to sell the license would further limit access to only the operator with the greatest financial resources. Time constraints on the allocated resources, similar to the definition of airport slots, described in Section 5.1.3, could increase fairness, with rules ensuring equitable distribution of slots across users. However, this could be difficult for many UAS operations, such as those that are not scheduled far in advance. Allocation of the resources would also be significantly more complex than today's airport slot allocation, because multiple resources would have to be allocated simultaneously.

Another mechanism by which an operator could gain exclusive rights to a region is through certification of its proprietary UTM system for operating UAS in a region. In such a case, the UTM system operator would be able to limit access to the airspace by limiting which other companies can use its UTM services or communicate with its UTM system. This has already been raised as a concern in Australia (Evans, 2019).

4.2. The need for strategic conflict resolution

In the NASA UTM concept, conflicts identified by strategic deconfliction services can be resolved by peer-to-peer negotiation. While there is a requirement that known conflicts be resolved prior to departure (Rios, 2018), no rules have yet been agreed upon governing the negotiation process itself, how the conflicts should ultimately be resolved, or how the negotiation should ensure fair access to airspace. The FAA UTM ConOps states that “equity of airspace access for UTM operations is ensured through appropriate performance authorizations and operation orchestration/Operator negotiation to optimize airspace use among the participants,” but the mechanisms by which that is done have not been agreed upon.

It has also been proposed by some members of industry that while negotiation may be required, no formal requirement for strategic conflict resolution is necessary — only conflict detection. There may be an assumption that the cost of deconfliction would be sufficiently low, and the value of ensuring safe operations sufficiently high, that USSs would always choose to resolve conflicts. However, as traffic grows, delays may no longer be negligible, especially for applications where delivery time is critical to the business model, or where vehicles have extremely limited extra battery capacity to absorb airborne delays.

In such cases, it is likely that some conflicts will not be resolved by unregulated peer-to-peer negotiation, particularly when they arise between USSs or operators that are in competition on high-demand routes, at high-demand times. In such cases, the conflict may not be resolved strategically, leaving it for tactical deconfliction or even detect-and-avoid systems to resolve. Regulators are likely to be resistant to certifying such an approach, especially for electric vehicles which have limited battery reserves for unplanned airborne maneuvering. The FAA UTM ConOps states that “FAA right of way rules are imposed when collaborative de-confliction cannot successfully resolve demand issues.” This would provide a degree of fairness in that access to airspace would not be restricted, and FAA right of way rules are generally accepted in the community. However, it is unclear whether these rules would be applied to resolve conflicts strategically as well as tactically, and how they would be applied, especially in multi-vehicle or successive conflicts.

Some governmental bodies have suggested that a centralized approach to conflict resolution, in the form of an “honest broker”, could be more effective at resolving all conflicts in a fair way. Such a centralized function would also be able to balance multiple considerations, including safety, efficiency, reliability, environmental compliance and fair access. Such an architecture may, however, have some limitations in scalability, and would require that the single USS (or ANSP) providing the centralized deconfliction service be available for all operators and vehicles. Any outage in such a service would have a significant impact on all vehicle operations. Some of these limitations could be mitigated by a distributed implementation of a system that is still managed centrally.

4.3. Potential prioritization rules

As described in Section 5.1.1, the first-come, first-served approach to allocating resources in traditional ATM is widely accepted, and generally considered fair for tactical situations. For strategic (pre-departure) allocation of airspace in UTM, this becomes a first-requested, first-served approach if allocation is based on when operators request, or file, operational intent. This means that operators that are able to file early get an advantage in that they are less likely to receive strategic delay. Simulation results (Evans *et al.*, 2020) show that this advantage could lead to significant inequalities, even at relatively small differences in file-ahead time in the order of minutes. This is of particular concern in UTM because of the diversity of types of operations envisioned. While many operations may have the predictability to file early, on-demand operators may be highly constrained in their ability to do so.

One approach to constrain operational intent request times is to define a Reasonable Time To Act (RTTA), proposed by Hately *et al.* (2019). This is a time before flight operation after which it becomes difficult to change the requested operational intent. Under this concept, all flights would be considered to have equal priority before their RTTAs, with the exception of specific high priority operations (e.g., public service operators, emergencies). Any conflict resolution at this point would not be governed by first-come, first-served or first-requested, first-served rules. After the RTTA, however, previously filed operational intents would have priority over later filed operational intents. UTM would "protect that flight from any further change in all but the most extreme situations." (Hately *et al.*, 2019) This concept is more consistent with approaches to resource allocation in traditional ATM. Ramifications of this approach, however, need to be studied carefully.

Size constraints on the operational intent volumes may also be required to prevent overreach of resource requests. It has been suggested that service providers are already incentivized to request the smallest volumes possible, so as to limit the number of other service providers with which they must negotiate strategic deconfliction. However, larger volumes would provide less risk of an operation falling out of conformance, and the increased negotiation associated with the larger volume may provide some information about competitors' operations. Furthermore, under a first-requested, first-served approach, the reservation of large volumes of airspace could be used to limit access by competitors. These challenges can be addressed at least in part if volume reservations are defined based on communications, navigation and surveillance performance requirements, rather than the operator's desire for arbitrarily smaller or larger volumes.

It has also been suggested that prioritisation could be used to incentivize desired policy outcomes, such as minimizing safety risks or reducing negative environmental impacts. This could be achieved by giving higher priority to safer, quieter or greener vehicles. The goal of fairness may therefore need to be balanced with other desired policy outcomes.

4.4. Should resource allocations be tied to specific flights?

One proposed approach is that flight corridors could be filed independent of specific operations, which operators then have the discretion to use when they wish. Under a first-requested, first-served approach, such corridor reservation could be highly constraining, and could have significant negative consequences for fairness, with the possibility of anti-cooperative behavior to reserve high-demand airspace for long periods of time, during which access by competitors could be limited. There are also more practical concerns with reservations becoming stale, but still requiring operators to negotiate access to. For this reason, it has been suggested that operational intent must be tied to specific flight operations — as in today's air traffic control (ATC) system. However, one consequence may be that a monitoring service is then required to confirm that flights are being operated after being requested, and that the system is not being gamed. In such an approach, there must be a mechanism to deal with cancellations in a fair way.

5. Applicability of existing fairness mechanisms to UTM

Fairness in the allocation of resources has been studied by economists and mathematicians for centuries (e.g., O'Neill, 1982; Aumann and Maschler, 1985), across a diverse range of domains. In this section we identify lessons learned from a diverse set of domains, with a particular focus on ATM, and discuss expectations and limitations to their applicability in UTM.

5.1. Air Traffic Management

In controlled airspace, the ANSP has strict control over the use of public resources, such as airspace and runways. This means that the processes for allocation are centrally controlled by the ANSP, airport authority or civil aviation authority (CAA). This standard is in contrast to proposals for a more federated architecture in UTM. In traditional ATM, control is distributed hierarchically, typically including a flow control or command center level, with regional ATM services below that. In the U.S., these are air traffic control facilities and airport control towers. ATM also almost exclusively allocates resources for trajectory based operations. This is not true for UTM, where many operations may be area based, requiring resources to be allocated that allow a vehicle to loiter in an area of airspace for periods of time.

The manner in which resources are allocated within the ATM hierarchy depends on the planning time horizon, which can be broken down into tactical resource allocation, strategic resource allocation, and long-term airport slot allocation. These are each discussed below.

5.1.1. Tactical allocation

In the tactical time frame (less than two hours, and often minutes) resources are allocated by ATC in such a way as to maintain safety, with a particular emphasis on maintaining separation between aircraft. Resources such as airspace sectors, jet routes, arrival/departure routes, and airport runways are generally awarded on a first-come, first-served (FCFS) basis.[6] This applies to ATC, where use of resources is decided by controllers, and tactical traffic flow management using tools such as the FAA's Time-Based Flow Management (TBFM) or EUROCONTROL's Enhanced Tactical Flow Management System (ETFMS). TBFM allocates arrival times at a network of metering points in order to maintain appropriate aircraft spacing given operating conditions. Allocated times are frozen after flights cross specified freeze horizons, allowing controllers to use speed control and vectoring to ensure aircraft meet these times. This type of demand management has particular applicability to UTM, but would need to be adapted if a more federated architecture is adopted for UTM, with service providers self-separating their vehicles from other vehicles.

[6] Practically speaking, controllers responding to an impending conflict may not consider the broader operational preferences of the aircraft in question (if they're even known) and will execute a plan that has the highest likelihood of quickly resolving the conflict. Controllers also have broad discretion to change the arrival or departure sequence, regardless of which aircraft was first, if there is an operational advantage for the controller.

In a similar way, ETFMS arranges the flow of traffic through a regulated sector with the aim of maintaining the order of the flights. The resulting delays are allocated to the flights by delaying their departure times. In much the same way, TBFM allows local departures to an airport to be scheduled into an overhead stream of airborne arrival traffic by delaying their departure time. Recent work on improving the fairness of how this is done (e.g., Smith *et al.*, 2016) may provide a useful model for pre-departure deconfliction in UTM for air taxi operations and pre-scheduled package deliveries.

FCFS in tactical allocation has long-standing acceptance in the aviation community, and generally makes sense for situations in which there may potentially be a physical queue for services. In ATC, physical queues for services do form, such as approaching a sector boundary; on arrival routes into an airport; or taxiing to depart a runway. It would also be impractical (if not hazardous) to reorder aircraft at the last minute. Such queuing is also relevant to many UTM operations.

FCFS can, however, have unintended inefficiencies. For example, during periods of high departure demand, long queues of aircraft build up near the runway with engines idling, when their delays would better have been served at the gate. Different definitions of what constitutes "first-come" are also likely to have very different implications for fairness, e.g., depending on whether "first-come" is defined by filing, pushing back, taking off, crossing a specified boundary, arriving, etc. Fairness issues associated with the timing of resource allocation are of particular relevance to UTM more broadly because of the proposed federated architecture. An approach similar to the freeze horizon used in TBFM may provide a solution, limiting how early resources can be requested. An example of such an approach is the RTTA proposed by Hately *et al.* (2019) and described in Section 4.3.

5.1.2. Strategic allocation

In the strategic time frame - considered hours in advance of operation for traditional ATM - airspace and airport resources are allocated by air traffic flow and capacity management (ATFM).[7] Resources are allocated in order to manage demand levels at those resources in such a way as to enable efficient separation management at the tactical level by controllers or tactical traffic management tools. Separation is not managed strategically because of the high uncertainties in conditions in the strategic time frame. In UTM, most operations will occur in the tactical time frame for traditional ATM, meaning that separation management may be possible pre-departure, with potentially reduced need for demand management. This will have to be confirmed for different operation types.

Both in the U.S. and in Europe there is heavy reliance on using scheduled times of operation as the basis for a claim to resources strategically. In structured ATFM programs,[8] the basic paradigm is to create a virtual queue of airport or airspace 'slots'. In the U.S., allocation is based on the ration-by-schedule (RBS) algorithm (Wambsganss, 2001), with earlier slots

---

[7] Particularly in Europe, this process is known as air traffic flow and capacity management (ATFCM). The key differentiator is that while capacity constraints in the United States often relate to weather events, in Europe the issues are generally a factor of flight scheduling causing sector overload. Thus, ATFCM puts special emphasis on dynamic resectorization as a means to optimize available capacity. We use the term ATFM here since it is the more common term, but the reader should understand these concepts to also apply to ATFCM where implemented.
[8] In the U.S., these are ground delay programs, airspace flow programs, and collaborative trajectory options programs.

allocated to flights that were scheduled earlier. While most UTM operations will not be scheduled, a virtual queue based on some rationing scheme will likely have relevance to many types of operations.

RBS processes three nested queues of flights: scheduled flights; exempt flights; and flights that have received a prior allocation. In ATFM in the U.S., it is common for half of all flights to be exempt. In UTM, exemptions may be a mechanism to accommodate high-priority operations, such as public safety or emergency operations. In RBS, ties between exempt flights are broken by estimated time of arrival. In UTM, other criteria may be more appropriate.

RBS is executed in batch mode, and therefore also accounts for flights that have already received slots under a prior execution of RBS. This allows RBS to honor its prior allocations, rather than unfairly revoke previously awarded resources. This has similarities to the freezing of metering times by TBFM, which prevents changing resource allocation beyond a specified time. This could also be relevant to UTM, preventing the need for replanning close to departure.

In the U.S. ATFM system, RBS is supplemented by a dynamic compression algorithm that moves flights up in the arrival hierarchy to fill slots vacated by canceled flights, making more efficient use of arrival resources. Compression gives as much compensation as possible to the operator vacating the slot. Similar algorithms may be needed to accommodate canceled flights in UTM, which may be more common because of the on-demand nature of many UTM operations.

ATFM allocation practices and procedures have been heavily shaped in the U.S. by the joint FAA-industry venture known as Collaborative Decision Making (CDM) (Wambsganss, 2001). CDM targeted incremental progress, with no attempt made to solve generalized ATFM resource allocation problems or abstractions. Rather, they began with some initial processes, identified flaws, and proposed specific solutions. Stakeholder consensus was also reached at every level of development. Such a model could also be effective for UTM. CDM also led to the creation of incentives for operators to submit timely and accurate information. Prior to CDM, airport arrival slot allocation was performed based purely on estimated time of arrival, which led operators to stop disclosing flight delays and cancellations to the FAA. CDM allowed operators to retain slots for canceled or delayed flights and use them for other operations. Similar mechanisms may be necessary in UTM to ensure truthful behavior.

Europe has taken a more authoritative stance on ATFM than the U.S. because of differences in the predictability of capacity constraints (impacted primarily by the lower incidence of convective weather). Airspace and airport allocation are applied days to weeks in advance, and allow little or no overscheduling. Airport and airspace resources are allocated on a more proactive, continual basis than in the U.S., using a heuristic version of a large-scale optimization algorithm, in contrast to the as-needed, batching approach used in the U.S. Such optimizations may be appropriate for a more centralized approach to UTM, or within single service providers in more federated architectures. In Europe, notification of assigned slots is not sent until 3 hours prior to expected gate pushback time. In this sense, it is more of a first-planned, first-served paradigm, which may be more representative of proposed UTM architectures.

5.1.3. Long-term allocation

In the U.S. and Europe, airport authorities, facilitated by trade bodies, impose long-term (i.e., months or years in advance) airport slot allocation on chronically congested airports. The processes rely heavily on slots that have been allocated historically, meaning preference is given to carriers that have operated for a qualifying duration. The allocation process takes place over an extended period of time, in which interactions and reviews take place with flight operators. Slots may be transferred or swapped between airlines, or used as part of a shared operation. Slots may only be transferred to another airline that is serving or planning to serve the same airport.

The process is highly centralized and well established, but contains significant subjective review and determinations by the coordinator. Efficient use of the resource is achieved simply by allocating all of the slots. The allocation of slots is based on a prioritization rule, heavily skewed toward incumbent operators and those operating a published schedule. This process necessarily assumes a lengthy petitioning, grievance, and allocation process (several months). To prevent awarded airport slots from going unused, utilization minima are set (usually 80%). Another key feature of the airport slot allocation process is that it guards against speculation, whereby slots are acquired merely to sell or manipulate pricing.

This type of highly procedural process will not work on the smaller time scales of real-time UTM (minutes or hours), nor is it clear what percent of operations at vertiports will be scheduled in advance. However, it has been proposed that some resources — even airspace blocks — could be allocated in advance without specifics of exactly how they will be used (Skorup, 2018). This would allow the lengthy petitioning and grievance process present in airport slot allocation. As discussed in Section 4.1, however, it would be very difficult to ensure fair access to resources for all users in such a scheme.

5.2. Non-aviation domains

Fairness in the allocation of resources has also been considered across a wide range of domains outside aviation, such as:

- Wireless networks, where radio spectrum is allocated in real time (Kelly, 1997; Eryilmaz and Srikant, 2005);
- The reallocation of wireless spectrum by the Federal Communications Commission (FCC) using two auctions designed by Paul Milgrom that have Pareto optimal[9] strategies that are truthful, thus incentivising truthful behavior from players (Milgrom *et al.* 2012);
- Automated trial decisions, where bias must be avoided in judicial decision making (Corbett-Davies *et al.*, 2017);
- Free museum ticket allocation, which allocates free timed-entry passes at high-demand times;
- Legislative seat assignments in the U.S. Congress;
- Estate inheritance;

[9] Pareto optimality is a state of resource allocation from which it is impossible to reallocate resources in such a way as to make any one individual better off without making at least one individual worse off.

- Supplying organs to transplant patients, which uses a points-based schema, based on weighting a number of different factors (Young, 1994);
- Army discharges at the end of WWII, which also used a points-based schema (Young, 1994).

There are several valuable lessons to learn from approaches to resource allocation in these various fields. One lesson learned is that a structured, points-based schema (first rights go to the claimant with the most points accumulated) can be an effective way to prioritize for services, and this is usually more palatable than relative ranking (claimant *A* always has priority over claimant *B*). However, it is difficult to prevent anti-competitive behavior using points-based systems. The definition of values for different actions/missions, the protocol for transferring points, and any initial allocation of points between parties is difficult to do in a fair way. Auditing such a system that is integrated in a distributed manner would in itself also be difficult.

Prioritization in the cellular services industry has demonstrated that structured procedures can perform rapidly and are capable of balancing efficiency with fairness, although these procedures may not translate well to UTM. Information packets are mostly uniform, and the consequences of dropping or rerouting them are less severe. Statistical inference techniques (machine learning) can be applied to track performance over long periods of time and correct for bias.

Another lesson learned is that market-based mechanisms can be an effective way to alleviate all or most of the need for centralized allocation, although some regulation is often still required. Market-based mechanisms are desirable because they are often the best way to ensure that each stakeholder's welfare is represented and maximized to the extent possible. Market-based solutions do, however, optimize efficiency rather than fairness, so have to be designed with fairness in mind. Auctions are a common way to implement market-based mechanisms. Recent technologies, such as online bidding by advertisers wanting to post ads to online viewers, has shown that electronic auctions can be conducted at high speed without the need for human intervention. Auctions can also be designed in ways to incentivize truthful behavior. Before applying this to a UTM environment, however, one would have to consider the technological sophistication required of all users. Congestion pricing is a common alternative to auctions, but it suffers from the defect that there is usually no basis for setting the appropriate prices.

Perhaps the most important lesson learned is that air transportation is orders of magnitude more complex than the commonly cited resource allocation situations. It has technical, stochastic, political, legal, dynamic, and safety aspects that must be considered. Moreover, resource allocation techniques necessarily vary by the nature of the resources and the claimants. For both these reasons, allocation paradigms ported over from other domains may require significant modification.

6. Metrics for evaluating fairness

Defining metrics for fairness are particularly challenging because fairness may be perceived differently by different stakeholders. Many metrics proposed to quantify fairness are based on concrete notions of claims and entitlement associated with resources. However, because neither the resources nor the claims to them have been clearly established in UTM, a different approach is required which measures the degree to which different participants have been denied free access to the airspace. For example, for traditional ATM, del Pozo de Poza *et al.* (2009, 2012) propose a metric for whether aircraft are receiving even treatment of ATC separation services, which is strongly tied to delay costs. The distribution of costs across operators can then be evaluated to quantify fairness. For similar operators, a uniform distribution can be considered fair. However, for disparate operators, other distributions may be more appropriate. Costs accounting for operator utility can be considered if such data is available, or they can be represented by metrics for each operator that are known to the system, such as:

- Average, total and maximum allocated or incurred delay minutes. This may include delay accrued throughout the operation, and even because of previous operations (e.g., causing a delayed departure), as suggested by Idris *et al.* (2019).
- Number of negotiations, and negotiations with a successful outcome. In a federated UTM architecture, any time two or more operational intents are in conflict, the respective USSs for those operations are expected to negotiate a resolution. Tracking high-level metrics about negotiations, without recording specific mission details, may be a useful metric in evaluating overall airspace density and congestion.
- Average additional delay per resolved negotiation. When two operations are in conflict, a negotiation process may attempt to resolve the conflict by adjusting the operational intents, which could result in one or more operators incurring a delay.
- Number of delayed, excessively delayed, or canceled operations. These metrics are recorded by each airline today, and generally broken down by specific flights, routes and other aspects. This lets stakeholders view system wide impacts in a variety of ways.

Two additional metrics, which are more complex and are based on the assumption that a first-come, first-served allocation is fair, are:

- Number of reversals, referring to changes in sequence relative to a first-come, first-served allocation (Bertsimas and Gupta, 2011);
- Delay deviation relative to the maximum expected delay under a first-come, first-served allocation (Barnhart *et al.*, 2012).

Key to implementation of these metrics is a definition of the first-come, first-served allocation. The traditional definition based on scheduled arrival times may not be relevant to the on-demand operations supported by UTM, while basing it on when operational intent is requested has been shown to have significant negative implications for fairness (Evans *et al.*, 2020).

Each of these metrics will vary by operator. Properties of the distributions of these metrics across operators can then be used as metrics to describe fairness, such as:

- Standard deviation across operators, used by Idris *et al.* (2019);
- Ratio of the geometric mean to the arithmetic mean across operators, providing a normalized fairness metric between 0 and 1, used by del Pozo de Poza *et al.* (2009, 2012);
- Gini coefficient (also between 0 and 1), which is a statistical measure of inequality used particularly to quantify wealth inequality;
- Product across all operators, which quantifies proportional fairness;
- Maximum across all operators, used by Rodionova *et al.* (2016) and Evans *et al.* (2014), which quantifies max-min fairness.

Under max-min fairness (Bertsimas *et al.*, 2011), users seeking a small amount of a resource receive their allocations first. Those with increasing demands on the resource follow, with the heaviest users dividing up equally what remains of the resource. Ultimately, this maximizes the smallest allocated resource. Achieving max-min fairness is considered more desirable than alternative allocations because, "any other allocation can only benefit the rich at the expense of the poor (in terms of utility)." (Bertsimas *et al.*, 2011) Max-min fairness has been used to define fairness in ATM (e.g., Rodionova *et al.,* 2016; Evans *et al.*, 2014) and communication networks.

It is important to note that there are also tradeoffs between fairness and other metrics, including efficiency, predictability, flexibility, and safety. Further metrics can then be calculated, including Pareto optimality, which assumes we can quantitatively compare operator utilities, and the price of fairness (Bertsimas *et al.*, 2011).

The price of fairness is significant because it quantifies the degradation of other metrics (e.g., safety, efficiency) by imposing fairness. For example, the most efficient allocation, from a system perspective, would maximize the sum of player utilities. A fair allocation of resources, however, would likely see a lower sum of player utilities, even if they were more fairly distributed. This reduction in sum of player utilities is the price of fairness. There are qualitative metrics to consider as well, which are discussed by Metron Aviation (2003).

Other metrics may also need to be recorded to inform auditing functions. Such metrics may be needed, e.g., to quantify conformance between filed operational intent and trajectories actually flown. This may include metrics such as the percentage utilization of operational intent.

7. How to measure and obtain fairness

7.1. Proposed Six-Step Evaluation Process

Based on our review and analysis of fair resource allocation across a range of domains, we propose the following six-step framework and process that should be used for evaluating fairness within a UTM setting and in UTM standards development (Young 1994, Metron Aviation 2003). In Section 3, we discussed a range of use cases in which fairness may need to be considered, from interoperability and compatibility policy to in-flight tactical decisions.

Because these uses are so different, the following steps are kept intentionally general: they must adapt to the specific use case.

1. **Define the domain.** What are the resources to be allocated, who uses them, and in what time frame? Each use case discussed in Section 3 should be treated as a distinct domain, so that the rest of this process repeats for each case. Indeed, some of the high-level use cases may need to be further refined and subdivided to make the rest of the evaluation process more valuable, especially when decomposing functional requirements into a given set of service requirements.

2. **Determine operator utility and value criteria.** What do different operators approximately value? This may vary dynamically. By identifying these values, one can begin to measure operational tradeoffs between operators for a given solution. This step will require soliciting input from stakeholders, who may be leery to divulge their preferences initially, or may be blind to the process (e.g., the general public). Therefore, we identify the need for further work, perhaps in collaboration with a small set of willing stakeholders, to develop reasonable ways to capture entitlement values. These could be represented as a range of acceptable bounds on added route lengths or delay times, so that individual flight preferences are not exposed. The weighted form of these values may be combined to create a generalized utility term which can later be used for evaluation – again, without regard to individual operator and/or flight constraints. These utilities may also change over time, and this temporal component along with the uncertainties associated with it will need to be considered.

3. **Determine allocation method.** This could include universal prioritization schemes (e.g,. first-come, first-served; best-equipped, first-served; or random lottery), as well as more complex schemes, such as auctions, that allow operators to weigh their internal flight priorities against prices for access to the resource. The most appropriate allocation methods will vary by application and operational environment. We identify a need for further research to better understand the tradeoffs and applicability of various allocation tools to given use cases, as a precondition to being able to apply this step to most use cases.

4. **Consider behavior incentives.** By accounting for how participants might react to a given allocation method, one can attempt to close as many loopholes as possible that would allow "gaming" the system. The "incentives," therefore, may be a combination of regulatory or policy constraints as well as economic ones. Similar to step 3, we identify a need for further research to identify the most effective incentives for a given use case and allocation method. This step also implies that loopholes within a given mechanism are first identified and understood.

5. **Measure and analyze fairness before and after allocation.** These metrics require establishing a baseline for comparison, as well as definitions of the metrics themselves. We identify a number of potential metrics in Section 6, but these may need to be modified or added to based on a given set of circumstances being evaluated. Our recommendation is that during initial phases of study and implementation, we collect a greater range of underlying metrics, so that the UTM community can identify those that are ultimately most critical to effective evaluation of fairness in a given use case. This is

not a one-time step; recurring analysis of fairness should be done as UTM evolves.

6. **Iterate Steps 1-5, as needed.** Based on the quantified fairness metric values (and particularly their distribution, trend over time, or other appropriate statistical analysis), modifications or alternative allocation methods should be considered until an allocation method is identified that ensures fairness. While some of this process will necessarily be driven by the analysis of large amounts of data, we recommend that affected stakeholders be involved in this evaluation process as well. This qualitative aspect is important in capturing effects that might not be measured yet (or measurable at all) by the quantitative metrics, and in collaboratively defining solutions and best practices.

This process should be applied to the various use cases in the development of common standards for UTM.

7.2. Recommendations for further research

Feedback is needed from the UAS operator, UTM service provider and UAM communities to validate the fairness considerations highlighted in this document and to refine the six-step approach described in Section 7.1. This six-step approach should then be tested by applying it to quantify the fairness implications of existing proposed approaches to resource allocation in UTM, to develop prioritization schemes and resource allocation algorithms that improve fairness, and to evaluate the impact of using these schemes and algorithms in different UTM architectures, at scale. These studies should target answering many of the open questions presented in this document about how to maintain fairness in UTM. Simulation will be invaluable in this process, and can be used to capture how operators could game the system in order to gain an advantage over competitors, using e.g., agent-based modeling. These simulations should quantify the relative impact of such behavior on fairness as well as efficiency and predictability.

There is also a need to explore the trade space of methods by which this resource allocation can be performed, and explore how these will work in a federated or centralized architecture. Simulation work could be followed by interoperability demonstrations that focus on the allocation of resources, and on negotiation.

8. Conclusions

There is a clear need to study and evaluate the implications of the UTM architecture on fairness, even at this relatively early point where complex implementations don't yet exist. This paper has established a common set of terms and theory around fairness, as well as several broad use cases in which fairness may apply in a UTM setting. Each one of those use cases can be studied, through analysis, simulation and demonstration, to further refine assumptions about the effects of fair or unfair practices, and to develop the best ways to ensure fair operations. These studies may all benefit from following the six-step process we explain in Section 7, as it provides a straightforward rubric to decompose multi-combinatorial problems into ones that can be evaluated piecewise through a series of smaller studies.

Having established a general framework for understanding the implications of fairness, we also see an immediate need to begin collecting and saving operational metrics, especially those around delays and route adjustments in areas with more than one operator or USS. Building up this real-world data now will be invaluable both in understanding the right kinds of metrics to collect long-term and in validating proposed solutions in simulation.

9. References

Aumann, R.J., and M. Maschler, "Game Theoretic Analysis of a Bankruptcy Problem from the Talmud," Journal of Economic Theory 36:195–213, 1985.

Balakrishnan, K., Polastre, J., Mooberry, J., Golding, R., Sachs, P., "Blueprint for the sky: The roadmap for the safe integration of autonomous aircraft," A^3 by Airbus, 2018. Available at https://storage.googleapis.com/blueprint/Airbus_UTM_Blueprint.pdf.

Barnhart, C., Bertsimas, D., Caramanis, C. and Fearing, D., 2012, "Equitable and efficient coordination in traffic flow management," Transportation science, 46(2), pp.262-280.

Bertsimas D., Farias V.F., and Trichakis N., "The Price of Fairness", Operations Research, vol. 59, No. 1, pp. 17-31, 2011.

Bertsimas, D. and Gupta, S., 2011, "A proposal for network air traffic flow management incorporating fairness and airline collaboration," Operations Research, pp.32-45.

Borenstein, S., 2001. "The trouble with electricity markets and California electricity restructuring disaster," POWER Paper, PWP-081, University of California Energy Institute. https://pdfs.semanticscholar.org/94af/df6a646bf33034d31bffd0a42f09779689bd.pdf

Bruggeman, A., VanCasteren, R., Feith, R., Gooijaers, L., El Sioufy, K., Naorniakowski, A., Xu, X., Udluft, H., "Simulation of mid-air encounters depending on air traffic density using the OSSAR simulator," Proceedings 5th Air Transport Operations Symposium, 2015.

Corbett-Davies, S., E. Pierson, A. Feller, S. Goel, and A. Huq, "Algorithmic decision making and the cost of fairness," in Proceedings of KDD '17, Halifax, NS, Canada, August 13-17, 2017. doi: 10.1145/3097983.3098095

del Pozo de Poza, I., "Assessment of fairness and equity in trajectory based air traffic management," Ph.D. thesis, University of Glasgow, 2012. http://theses.gla.ac.uk/3237/

del Pozo de Poza, I., M.A. Vilaplana Ruiz, and C. Goodchild, "Assessing Fairness and Equity in Trajectory Based Operations," 9th AIAA Aviation, Technology, Integration, and Operations Conference, Hilton Head, SC, September 21-23, 2009.

Eryilmaz, A., and R. Srikant, "Fair Resource Allocation in Wireless Networks using Queue-length-based Scheduling and Congestion Control," Proceedings of IEEE Infocom, 2005.

Evans, A., Egorov, M. and Munn, S., 2020. “Fairness in Strategic Deconfliction in UTM,” AIAA Science and Technology Forum and Exposition, Orlando, FL, 2020.

Evans, A., Vaze, V. and Barnhart, C., 2014. Airline-driven performance-based air traffic management: game theoretic models and multicriteria evaluation. Transportation Science, 50(1), pp.180-203.

Evans, J., *Google drones could have a monopoly over Canberra's airspace as Wing takes flight*, ABC News, 10 Jun 2019, 5:37pm.
https://www.abc.net.au/news/2019-06-07/google-drones-could-have-a-monopoly-over-canberra-airspace/11188062 (Cited September 27, 2019)

Federal Aviation Administration, 2018. *Unmanned Aircraft System (UAS) Traffic Management (UTM) Concept of Operations v1.0*, FAA NextGen Office.

Golding, R. 2018. *Metrics to characterize dense airspace traffic*, Airbus Altiscope Report TR-4, 7 June, 2018.

Hately, A., *et al.* 2019. *CORUS U-Space Concept of Operations*, 2016 SESAR 2020 RPAS Exploratory Research Call, EUROCONTROL

Idris, H., Chin, C., and Evans, A.D., “Accrued Delay Application in Trajectory Based Operations,” ATM R&D Seminar, 2019.

ICAO, 2019. “Unmanned Aircraft Systems Traffic Management (UTM) – A Common Framework with Core Principles for Global Harmonization,” available at
https://www.icao.int/safety/UA/Documents/UTM-Framework.en.alltext.pdf

InterUSS, URL: https://interuss.org (Accessed Oct 25, 2019)

Kelly, F.P., “Charging and rate control for elastic traffic,” European Transactions on Telecommunications, Vol. 8, pp. 33–37, January 1997.

Metron Aviation, “Equitable Allocation of Limited Resources (EALR): Defining, Measuring, and Implementing Equity,” Technical Report to FAA, November 2003.

Milgrom, P., Ausubel, L., Levin, J. and Segal, I., “Incentive auction rules option and discussion,” Report for Federal Communications Commission, September, 12, 2012.

Morey, M., 2001. “Power market auction design: Rules and lessons in market based control for the new electricity industry,” Prepared for Edison Electric Institute.
http://web.mit.edu/esd.126/www/MktsAuctions/EEI.pdf

Moulin, H., and W. Thomson, “Can Everyone Benefit From Growth?: Two Difficulties,” Journal of Mathematical Economics 17:339-45, 1988.

O’Neill, B., “A Problem of Rights Arbitration from the Talmud,” Mathematical Social Sciences 2:345–371, 1982.

Rios, J., 2018. “Strategic Deconfliction: System Requirements”, NASA UAS Traffic Management (UTM) Project, July, 2018.

Rios, J., I.S. Smith, P. Venkatesen, D.R. Smith, V. Baskaran, S. Jurcak, R. Strauss, S. Iyer, P. Verma, 2018. “UTM UAS Service Supplier Development, Sprint 1 Toward Technical Capability Level 4,” NASA Ames Technical Memorandum, NASA/TM-2018-220024. URL: https://utm.arc.nasa.gov/docs/UTM_UAS_TCL4_Sprint1_Report.pdf (Accessed Oct 25, 2019)

Rios, J., I.S. Smith, P. Venkatesen, D.R. Smith, V. Baskaran, S. Jurcak, S. Iyer, P. Verma, 2018. “UTM UAS Service Supplier Development, Sprint 2 Toward Technical Capability Level 4,” NASA Ames Technical Memorandum, URL: https://utm.arc.nasa.gov/docs/2018-UTM_UAS_TCL4_Sprint2_Report_v2.pdf (Accessed Oct 25, 2019)

Rodionova, O., Arneson, H., Sridhar, B. and Evans, A., 2017, September. Efficient trajectory options allocation for the collaborative trajectory options program. In 2017 IEEE/AIAA 36th Digital Avionics Systems Conference (DASC) (pp. 1-10). IEEE.

Skorup, B., “Auctioning Airspace,” Working Paper, Mercatus Center at George Mason University, Arlington, Va., November 2018.

Smith, N.M., Brasil, C., Lee, P.U., Buckley, N., Gabriel, C., Mohlenbrink, C.P., Omar, F., Parke, B., Speridakos, C. and Yoo, H.S., 2016. Integrated demand management: Coordinating strategic and tactical flow scheduling operations. In *16th AIAA Aviation Technology, Integration, and Operations Conference* (p. 4221).

Wambsganss M.C. (2001) Collaborative Decision Making in Air Traffic Management. In: Bianco L., Dell’Olmo P., Odoni A.R. (eds) New Concepts and Methods in Air Traffic Management.

Young, H.P., Equity in Theory and Practice, Princeton University Press, 1994.