NASA/TM–20230012622

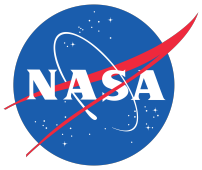

Vertiport Dynamic Density

Lilly Spirkovska
Ames Research Center, Moffett Field, California

September 2023

NASA STI Program . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI Program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI Program provides access to the NASA Aeronautics and Space Database and its public interface, the NASA Technical Report Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers, but having less stringent limitations on manuscript length and extent of graphic presentations.
- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.
- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.
- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.
- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.
- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include creating custom thesauri, building customized databases, and organizing and publishing research results.

For more information about the NASA STI Program, see the following:

- Access the NASA STI program home page at *http://www.sti.nasa.gov*
- E-mail your question to *help@sti.nasa.gov*
- Fax your question to the NASA STI Information Desk at 443-757-5803
- Phone the NASA STI Information Desk at 443-757-5802
- Write to:
  STI Information Desk
  NASA Center for AeroSpace Information
  7115 Standard Drive
  Hanover, MD 21076–1320

NASA/TM–20230012622

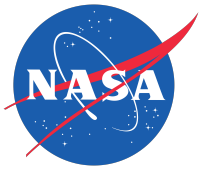

Vertiport Dynamic Density

Lilly Spirkovska
Ames Research Center, Moffett Field, California

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California 94035

September 2023

Acknowledgments

Funding for the effort was provided by the System Wide Safety (SWS) project of the Aviation Operations and Safety Program (AOSP), Aeronautics Research Mission Directorate (ARMD), NASA.

Available from:

NASA Center for AeroSpace Information
7115 Standard Drive
Hanover, MD 21076-1320

National Technical Information Service
5301 Shawnee Road
Alexandria, VA 22312

Available electronically at http://www.sti.nasa.gov

Vertiport Dynamic Density

Lilly Spirkovska
National Aeronautics and Space Administration
Ames Research Center
Moffett Field, California 94035

Summary

Advanced Air Mobility (AAM) is envisioned to be another spoke in a region's transportation system, supplementing ground-based travel with air-based travel. Eventually, air taxis will be pervasive, convenient, and affordable. As demand and operations tempo increase, congestion may arise. Several characteristics of AAM reduce the applicability of techniques used in today's National Airspace System (NAS) to manage congestion. AAM will include both scheduled and non-scheduled, on-demand operations, challenging strategic deconfliction algorithms. Flights will be fairly short, traversing an urban area, not hundreds or thousands of miles, with corresponding low energy reserves, prohibiting excessive delays. Flight operators will need flexibility in operations, scheduling a flight only minutes before departure, or diverting to an alternate vertiport if it becomes advantageous, further adding to trajectory uncertainty that would challenge strategic deconfliction. Finally, operators will desire privacy to protect sensitive information or preserve competitive advantage, limiting early access to intent information. In this paper, we present an approach for managing congestion at vertiports by providing insight into the traffic situation to support operationally-advantageous and safe land or divert decisions. We propose a metric designed with usefulness and usability in AAM operations in mind. The metric uses the sociology concept of dynamic density (DD) that takes into consideration not only number of flights, but also the interaction of those flights with the vertiport's limited resources, namely the landing pads and the parking spots. DD supports a Pilot in Command (PIC) with decisions about whether to proceed, expedite, delay, or divert; and supports air traffic control (ATC) and vertiport operators in airspace management and vertiport usage. We demonstrate the metric on a notional vertiport scenario. We also show that DD provides better insight into congestion and resulting flight delays than an aircraft count metric used in traditional air traffic management.

Nomenclature

Acronyms

AAM	Advanced Air Mobility
AF	arrival fix
DCB	demand capacity balancing
DD	dynamic density
ETA	estimated time of arrival
ETD	estimated time of departure
eVTOL	electric vertical take-off and landing
HITL	human-in-the-loop
NAS	National Airspace System
SWS	System Wide Safety
UAM	Urban Air Mobility

Variables

$Delay_F$	delay for flight F
ETA_{CF}	conflict-free ETA
ETA_P	proposed ETA
H_F	freeze horizon (min)
H_S	intent sharing horizon (min)

1 Introduction

Advanced Air Mobility (AAM) is evolving from a few flights using existing infrastructure, procedures, and regulations to an ecosystem resembling ground transportation (Ref. 1). Under the FAA's Innovate28 implementation plan*, AAM is starting with a few pre-scheduled flights commanded by an onboard pilot, operating from a few existing airports and heliports, flying under fairly benign visual conditions along charted routes, abiding by visual flight rules, coordinating with air traffic controllers, and generally not adding much stress to the existing National Airspace System (NAS). It is expected to evolve to include hundreds of simultaneous flights within an urban area; operating from a variety of large vertihubs, medium-size vertiports, and small vertistops (Ref. 2, 3); relying on automation and a federated air traffic control system; flown by remote pilots or autonomous systems; pre-scheduled and on-demand; and abiding by community-based flight rules. AAM is envisioned to be another spoke in a region's transportation system, supplementing ground-based travel with air-based travel.

Congestion is not expected to be an issue in the *crawl* stage of the FAA's *crawl-walk-run* AAM development strategy. The low demand, low tempo operations can be readily integrated into the NAS. On the contrary, congestion can become a significant issue as demand and operations tempo increase, and mitigating action will be required to maintain airspace safety. One mitigation strategy is strategic deconfliction, such as demand capacity balancing (DCB). Much like the metering lights that slow traffic entering a highway, DCB strategies can reduce congestion at busy crossing points or vertiports. A challenge with DCB is that both demand and capacity are uncertain. A number of events, such as delayed passenger boarding, adverse en route winds, or excessive tactical deconfliction, could decrease demand for a landing slot, while other events, such as unexpected, unscheduled operations, can significantly increase demand (Ref. 4). Similarly, an event such as a problem on the landing pad, faster than anticipated weather deterioration or improvement, or faster or slower charging can change capacity. To accommodate such unforeseen events and operational uncertainties, capacity can be preemptively reduced to provide a buffer. If everything proceeds as planned, however, that buffer may result in unnecessary delays, significantly increasing both the number of delayed flights and the mean delay (Ref. 5).

We propose an approach that reduces the need for operational buffers, potentially boosting capacity and increasing resource utilization, while supporting an acceptable level of safety. One of the problems of excess demand for a shared resource like a vertiport is the likelihood of increased delays while flights wait their turn. AAM flights will predominately use electric aircraft, for noise and emissions reasons. Flight duration will be limited and energy reserves will be minimal until battery technology improves considerably. An extended delay waiting to land greatly increases risk of energy exhaustion. To maintain safety while scheduling flights to actual vertiport capacity with no buffers, pilots (or autonomous systems) will need adequate insight about expected delays so they can make informed wait or divert decisions while diversion is still an option. Dynamic density has been studied for related tasks such as predicting controller workload (Ref. 6), predicting urban air mobility (UAM) corridor congestion (Ref. 7), and predicting airspace complexity in upper class E† (Ref. 8). We apply the concept of dynamic density to predicting vertiport congestion.

Dynamic density (DD) refers to the combination of population density and the amount of

*`https://www.faa.gov/air-taxis/implementation-plan`

†Upper class E is the airspace above 60,000 ft above mean sea level.

interaction within that population[‡]. It was suggested in sociology by Emile Durkheim's theory of modernization as an explanation for how society transformed from more primitive to more modern, impelled by limited resources. For air traffic control, DD was investigated as a metric to more effectively predict controller workload compared with predictions based solely on number of aircraft (Ref. 6). It was posited that controlling an airspace with few intersecting airways with flows of similar aircraft types requires less effort than controlling an airspace with many intersecting flows with widely differing performance, capabilities, and equipage. A variety of airspace structural characteristics, such as number of coordination points or volume available for maneuvering, also affect workload. In previous work, we developed a DD metric to assess congestion in AAM (also referred to as Urban Air Mobility (UAM)) *corridors* (Ref. 9) to provide awareness of the likelihood of delays for the en route phase of flight (Ref. 7).

In this work, we investigate DD as a measure of vertiport congestion. We provide two values to pilots (or autonomous systems) of aircraft in the vicinity of a vertiport. First, we provide expected delay for flights that are within the estimated time of arrival (ETA) freeze horizon (defined in Subsection 2.1). And second, we provide a categorical value representing anticipated congestion for flights that are within the intent sharing horizon (again, defined in Subsection 2.1). This information can inform decisions about whether to proceed inbound as planned, speed up or slow down to arrive ahead of or after a predicted congested period, or divert to a suitable alternate.

The remainder of the report is organized as follows: Section 2 will outline assumptions and set the context for the research. Following that, the algorithm will be described. Section 3 will describe how test data is generated and presents an example schedule of flight operations at a single vertiport. Section 4 will show the results of the algorithm for the example schedule and will compare DD results with predictions based solely on aircraft count. Finally, conclusions and ideas for future work will be presented in Section 5.

2 Dynamic Density Metric

The objective of the DD metric is to support safe landing or facilitate timely diversion for high tempo vertiport operations. DD is meant to forewarn of potentially unsafe landing conditions caused by vertiport congestion, particularly, energy exhaustion due to arrival delays arising from too many flights competing for limited shared resources.

One approach to reducing contention for resources is to limit the number of flights through, for example, DCB. Fewer flights, in general, would lead to fewer resource conflicts. But, excessive metering of traffic is not the most efficient way to operate, somewhat negating AAM's promise that urban air travel can save time[§]. Operational uncertainties could result in flights missing the landing window assigned by DCB, in effect wasting resources that could have been used by another flight. Delayed flights could bunch up, causing the very congestion DCB aims to control. A common mitigation strategy is to leave empty slots in the vertiport schedule to accommodate delayed flights. However, if flights are not delayed, those empty slots unnecessarily decrease vertiport throughput. Another approach is to increase en route separation. That has a similar

[‡]`https://en.wikipedia.org/wiki/Dynamic_density`

[§]Although the same flight time may be required with and without metering, delaying a flight's departure time would cause the customer to arrive later, perhaps as late or later than they could using a different transportation modality.

impact on throughput.

In our approach, we endeavor to maximize utilization of the vertiport and airspace, while ensuring safety, without operational buffers to absorb delayed flights and without extra en route separation to reduce demand. Congestion that may arise due to delayed or unexpected[¶] flights will require operational flexibility. Rather than accepting delays in flight, the best available option may be to divert to a nearby vertiport, especially if extra ground travel time required from an alternate vertiport to the customer's final destination is less than the expected arrival delay. Also, diverting to an alternate may be the only option if expected delays exceed the flight's energy reserves. It is critical that the Pilot in Command (PIC) is made aware of that situation while mitigation options are available.

In this section, we provide the details of the DD metric including airspace overview and definitions of relevant parameters (subsection 2.1); assumptions and simplifications (subsection 2.2); the intended users (subsection 2.3); and, finally, the DD algorithm (subsection 2.4).

2.1 Definitions

Before proceeding to the details of the algorithm, we define and illustrate several notional features of the airspace, as shown in Figure 1. For our work, we assume a vertiport has one or more arrival fixes [‖] and one or more departure fixes. The arrival paths and departure paths are procedurally separated, so that arrivals and departures compete only for parking spots and takeoff/landing pads.

Arrivals become known to DD only upon reaching the *intent sharing horizon*. Known arrivals are separated into two groups. The separation point is known as the *freeze horizon* and flights in the first group are referred to as *frozen flights*. Frozen flights are scheduled for conflict-free pad arrival time and are guaranteed a parking spot on arrival. Remaining flights, also known as *unfrozen flights*, maintain schedule flexibility until reaching the freeze horizon. We define the two horizons, shown in Figure 1, as follows:

Intent sharing horizon Lookahead time, H_S, from a vertiport at which a flight's ETA is shared and contributes to the DD calculation. This allows for equal treatment of scheduled and unscheduled flights. For simplicity, we assume an intent sharing fix rather than an arc, with a specific fix for each arrival stream. We also assume the same approach performance for all flights so that H_S (and H_F, defined next) will be at the same distance for all flights.

Freeze horizon Time H_F from a vertiport at which a flight commits to its ETA. The frozen ETA may differ from the proposed ETA used for DCB at departure and even its ETA when it arrives at the intent sharing horizon, H_S, due to real-world uncertainties. It is expected that a flight's actual time of arrival will rarely vary from its frozen ETA. For simplicity, we assume a freeze fix rather than an arc, with a specific fix for each arrival stream.

2.2 Assumptions and Simplifications

AAM operations are being formulated and in flux. We therefore make several assumptions about the conduct of operations and the layout of vertiports and surrounding airspace. Assumed

[¶]Unexpected flights include, e.g., unscheduled on-demand flights, or flights that need to divert due to aircraft system failures.

[‖]A fix is a point in the airspace, given a name for easy reference in procedures and to facilitate pilot/controller communications.

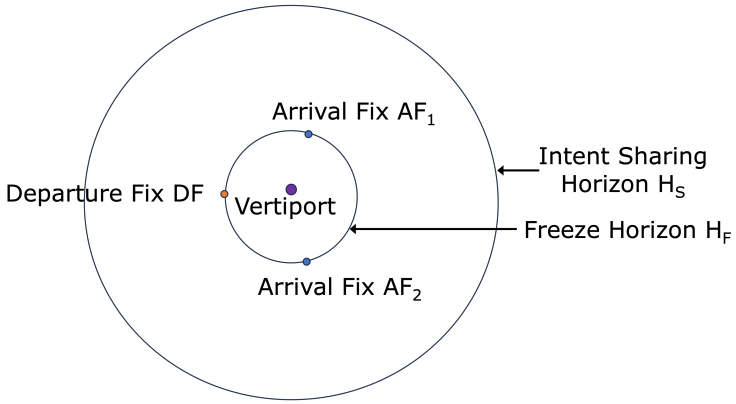

Figure 1.—Notional diagram of relevant airspace features. Not to scale.

AAM operations follow traditional air traffic management (ATM) operations, though with modified separation intervals to reflect automation-supported capabilities. Airspace layout assumptions intend to simplify the narrative and facilitate communication of the proposed solution by eliminating the complexities concerning multiple vertiports, arrival and departure streams, or takeoff/landing pad slots.

We assume the following:

1. Strategic deconfliction, such as DCB, may be in use. However, capacity buffers are not preemptively imposed for contingencies or uncertainties.

2. Arrivals may be strategically deconflicted pre-departure based on proposed 4D trajectories and may be delayed on the ground as a result of DCB. Arrivals may also be delayed or expedited en route due to environment or airspace events, including winds not as forecast, vehicle faults, priority emergency operations, operational needs, or vertiport closures. Unscheduled flights may also *pop up* at the intent sharing horizon.

3. Similarly, departures may be delayed, expedited, or pop up.

4. Arrivals have priority over departures. If parking is available, arrivals are cleared to land. Any conflicting departures get delayed until there is a break in arrivals or parking is full. (This mimics current right-of-way rules.**)

5. Arrivals maintain at least 30 sec en route separation without air traffic control (ATC) assistance. They arrive at the intent sharing horizon separated from others by no less than 30 sec.

6. Arrivals and departures occur at 30 sec slots. That is, an ETA may be at, e.g., 3:00 or 12:30, but cannot be at, e.g., 3:15 or 12:43. This detail simplifies the narrative and does not affect the algorithm's generality.

7. Vertiport occupancy time is 60 sec. This allows time for an arrival to land and taxi or be towed to parking, and for a departure to get into position and depart. It does not include time for passenger ingress/egress, charging, etc. Those activities occur away from the pad.

**Federal Aviation Regulations (FAR) 91.113(g).

8. Vertiport has one takeoff/landing pad and five parking spots, a number consistent with expectations for vertiports. (Vertihubs may have ten or more including overnight parking, and vertistops may have fewer than two (Ref. 2).)

9. Flights share the takeoff/landing pad and must have a parking spot, except during takeoff or landing. There is no temporary holding area that can accommodate extra aircraft on the vertiport surface.

10. Vertiport has N arrival streams, each with its own arrival fix. For the initial prototype, we assume two arrival streams.

11. Vertiport has one departure stream and the departure path is clear of arrivals. Departing flights are not considered a collision factor for arrivals, but are considered a competitor for resources, namely the takeoff/landing pad and parking spots.

2.3 Intended Users

The target users for DD are the PIC, air traffic controllers, and vertiport operators.

Federal Aviation Regulations[††] require each PIC to be familiar with all available information concerning a planned flight, including weather, fuel requirements, available alternatives, landing zone conditions, and any known traffic delays. Similar pre-flight regulations will likely apply for AAM, even if the PIC may be a remote operator simultaneously responsible for several flights or an autonomous system.

DD is most pertinent to informing the PIC about traffic delays. Pre-flight, insight into traffic congestion is critical to deciding whether the flight can be conducted safely, including determining how much fuel (or energy) will be required. Once airborne and until the aircraft is parked, the PIC continually reassesses risk to ensure that evolving conditions will not result in energy exhaustion.

DD can also provide situational awareness to air traffic controllers and vertiport operators. For example, an air traffic controller responsible for surrounding airspace may decide to vector flights away from a congested vertiport. Or, a vertiport operator may delay a vertiport configuration change, i.e., change landing direction to better conform to changing wind direction, to follow an arrival push when fewer flights would be affected.

DD metric is independently computed for all vertiports, and users can request updates for any vertiport of interest. DD for the planned arrival vertiport and vertiports in its vicinity will likely be of predominate interest under nominal operations, whereas in off-nominal operations, such as early diversion, DD for en route vertiports can support vertiport selection.

2.4 Dynamic Density Algorithm

DD metric must provide adequate insight into the vertiport traffic situation to support operationally-advantageous and safe decisions. To be useful in practice, it must be easy to comprehend, easy to corroborate, predictable, not overly sensitive to small changes in the traffic situation, and stable. To be usable, especially during the busy final arrival phase, it must sufficiently encapsulate the situation without providing extraneous details that require much time and effort to interpret. Finally, to protect privacy and business decision making, the algorithm should not be too invasive (Ref. 10). That is, it should require minimal intent sharing as late as

[††]For example, FAR Section 91.103, Section 121.647, and Section 135.617.

possible. In this section, we describe the algorithm designed with these usefulness, usability, and intent sharing characteristics in mind.

Subsection 2.1 introduced the H_S intent sharing and H_F freeze horizons. In describing the algorithm, the H_F duration is referred to as a *chunk*. The H_S duration is split into H_F-sized chunks, and DD is computed individually for each chunk. For example, if H_F is 3 minutes from the arrival fix (AF) and H_S is 10 minutes from the AF, DD is computed for four chunks: 0-3 min, 3-6 min, 6-9 min, 9-10 min. Although computed individually, the chunks are not independent; delays that originate in one chunk may propagate to flights in the next chunk.

The algorithm needs the following data:

1. Arrivals' proposed ETA, as computed within H_S and updated as necessary until H_F
2. Departures' proposed estimated time of departure (ETD), updated if flights are delayed
3. Number of empty parking spots

The flow of the algorithm is shown in Figure 2 and the outline of the chunk computation subroutine in Figure 3. Although *Compute chunk DD* computes the same information for all flight chunks, the output to the user varies according to whether the flights are inside H_F, succinctly called the *frozen chunk*, or flights are within H_S but not yet inside H_F, or *unfrozen chunks*. For the frozen chunk, the user is provided:

1. Categorical DD value (negligible, low, moderate, or high)
2. A delay severity score indicating risk of energy exhaustion
3. Residual delay for the frozen chunk that will propagate to the next chunk
4. Revised cleared-to-land conflict-free ETA for each flight in the frozen chunk

The rationale for providing the cleared-to-land ETA is that flights are fairly close to landing and typical operational uncertainties can be effectively managed to comply with a required arrival time. Improbable events that require landing priority or a diversion are expected to be rare and are outside the scope of the current work. The residual delay of the chunk will inform future flights how much delay needs to be absorbed by flights in the following chunk.

For unfrozen chunks, the user is provided:

1. Categorical DD value (negligible, low, moderate, or high)
2. Residual delay for the chunk that is predicted to propagate to the next chunk

Expected DD, operational needs, or environment uncertainties may influence a decision to proceed, delay, expedite, or divert. To maximize flexibility, flights remain unrestrained. The categorical DD value encapsulates the dynamic traffic situation without misrepresenting the inherent uncertainty or burdening the PIC to meet an impractical arrival time. Like for the frozen chunk, the residual delay is intended to indicate how much delay needs to be absorbed by flights in the next chunk.

The algorithm to compute DD for a chunk is outlined in Figure 3 with additional details provided in Algorithm 1. It begins with the proposed arrival or departure times for all flights in

the chunk, the number of parking spaces available at the vertiport, and the conflict-free schedule for the previous chunk. In ascending order of ETA_P, flights are rescheduled to a conflict-free ETA_{CF} that is no earlier than the proposed ETA and provides adequate temporal separation between flights at the pad. The temporal separation is not just for flights in the chunk being processed; rather, it also considers the pad occupancy time of the last flight in the previous chunk. The parking constraints are then enforced. If parking will not be available, the ETA_{CF} is moved forward to permit a flight to depart. The ETA_{CF} is then set to the departure's time plus the required temporal separation between flights at the pad. The process continues for the remaining arrivals. Note that arrivals have precedence over departures. Arrivals are scheduled for the first landing slot that satisfies the constraints just described, i.e., no earlier than proposed, temporal separation from previous pad occupant, and available parking spot. Departures are held at their parking spots until a sufficient gap exists in the arrival flow or no parking spots remain. Flights are then permitted to depart, but only until the arrival constraints are again satisfied.

The next step in the algorithm is to determine the resulting delay for each arrival, computed as the difference between its proposed and conflict-free ETAs. The risk of energy exhaustion is not linear with amount of delay. Even with the limited energy reserves that will be required for 15–20 min flights, a 30 sec delay can likely easily be absorbed and pose little risk to energy exhaustion. Similarly, a five-min delay is likely manageable, but the increased risk is starting to get the PIC's attention. In contrast, a ten-min delay is significant and a safe landing may be a concern. The DD metric aims to capture that exponentially increasing risk by computing a weighted average of the delays expected for flights in a chunk. The weighted delay signifies the severity and is computed as 1.25^d where d is delay in minutes. Weighted severity for all flights in a chunk are then averaged and mapped to a categorical DD value as shown by Table 1. This dilutes the contribution of any one flight to the chunk's score. Other approaches for mapping delay to DD were also considered, including un-weighted average and sum of delays. An *un-weighted* average does not capture the situation well. As an example, it is difficult to distinguish between one flight being delayed 10 min and four flights being delayed 2:30 min each. The first case is more concerning than the second, especially for the delayed flight's PIC. The sum of delays has a similar issue and is even less informative. An alternative approach that may be superior and will be tested in future work is to present both the average and standard deviation of the chunk's severity scores. We will investigate whether the decision value of the extra data outweighs the interpretation burden.

Finally, as shown in the flow chart, all of the results are returned to the calling routine.

Table 1.—Mapping between weighted average delay severity and DD. The intervals for each category are configurable for each vertiport and should reflect the risk of delays given the availability of alternate vertiports in the vicinity and the minimum energy reserves on landing required by regulation[‡‡].

Chunk Delay Severity Score	Dynamic Density
> 5.25	HIGH
$(2.5, 5.25]$	MODERATE
$(1.25, 2.5]$	LOW
$<= 1.25$	NEGLIGIBLE

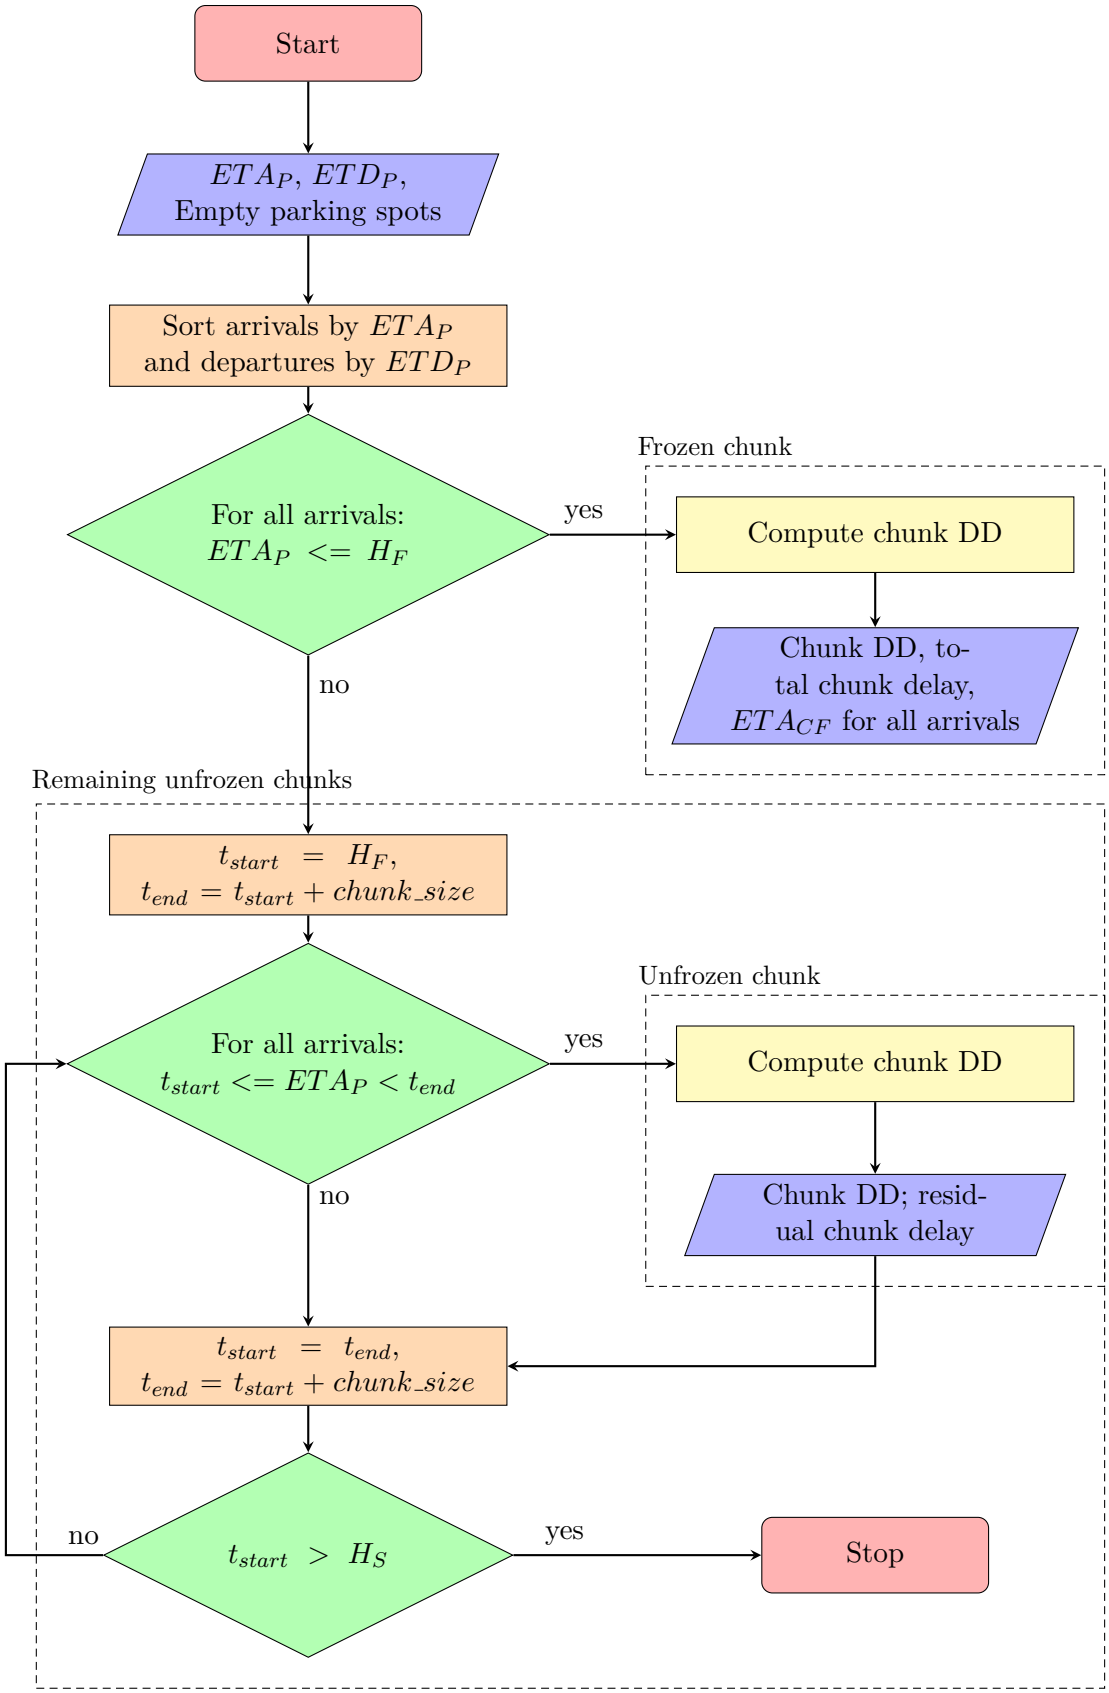

Figure 2.—Flow chart for dynamic density computation for vertiports. Element *Compute chunk DD* is a subroutine, outlined in Figure 3.

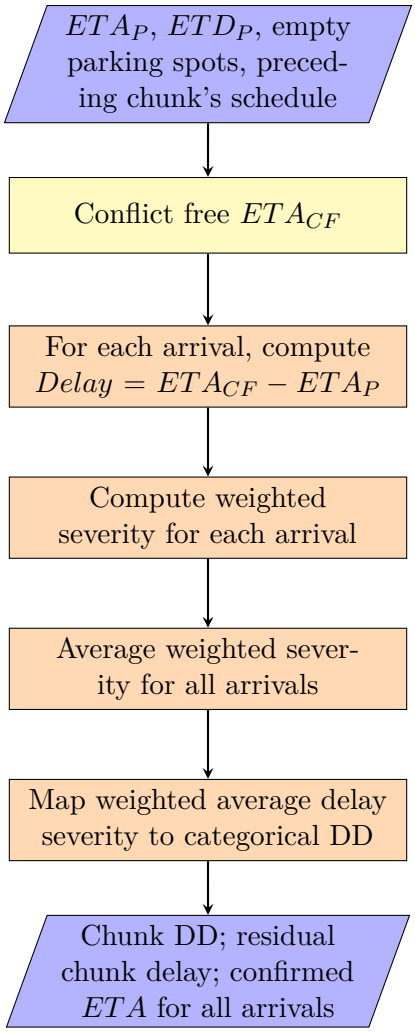

Figure 3.—Flow chart for subroutine *Compute chunk DD* referenced in Figure 2.

3 Simulation

For development and test purposes, we randomly generate arrival and departure flights. Only ETA and ETD are generated, not full trajectories. The inputs to the simulator are number of flights to generate, number of minutes lookahead (i.e., the intent sharing horizon, H_S), the minimum number of seconds between arrival and departure slots (e.g., 30 sec, as stated in Subsection 2.2), and the number of arrival streams. A single departure stream is assumed for simplicity. ETD between now and H_S is randomly assigned for the specified number of departures. For arrivals, the desired number of flights is first randomly split into flights per stream, and then ETA between now and H_S is randomly assigned for each flight.

To simulate operations being perturbed due to environmental uncertainties and unanticipated events, ETA and ETD are updated at the start of each chunk's computation for all remaining flights. That is, using the example of chunking from Subsection 2.4, when computing chunk DD for duration 3-6 min, ETA or ETD for any flight proposed to land or depart, respectively, between 3 minutes and H_S or 10 minutes (from now) can be revised to simulate changed

Algorithm 1 Algorithm to compute conflict free ETA for a chunk of flights.

```
If chunk is first to be scheduled:
    Conflict free ETA for first flight is its proposed ETA
Else treat first flight same as remaining flights in the for-loop below

For each remaining flight F in chunk:
  If there's parking available:
    Conflict free ETA for F is later of its proposed ETA or
      previous flight's conflict free ETA plus pad separation time
  Else if there's no parking available:
    Schedule next departing flight for later of its proposed ETD or
      previous arr flight conflict free ETA plus pad separation time or
      previous dep flight plus pad separation time
    Conflict free ETA for F is prev departing flight plus pad separation time
```

environmental conditions.

An example vertiport schedule is shown in Table 2 and the simulation parameters used to generate it are shown in Table 3.

Table 2.—Example vertiport schedule, generated with the input parameters shown in Table 3. The schedule is divided into *chunks* of flights processed as a group, delineated by dashed lines.

Slot	Time	Departures	Arrival Stream 1
1	0:30		UAM_a_0_1
2	1:00	UAM_d_2	
3	1:30		UAM_a_0_3
4	2:00		UAM_a_0_4
5	2:30		
6	3:00	UAM_d_6	UAM_a_0_6
7	3:30	UAM_d_7	UAM_a_0_7
8	4:00	UAM_d_8	
9	4:30		UAM_a_0_9
10	5:00		
11	5:30	UAM_d_11	UAM_a_0_11
12	6:00		
13	6:30	UAM_d_13	
14	7:00		UAM_a_0_14
15	7:30	UAM_d_15	UAM_a_0_15
16	8:00	UAM_d_16	
17	8:30	UAM_d_17	UAM_a_0_17
18	9:00		
19	9:30	UAM_d_19	
20	10:00		

Table 3.—Parameter values used to randomly generate the schedule in Table 2.

Parameter	Value
Intent sharing duration (H_S)	10 min
Chunk duration (H_F)	3 min
Slot interval	30 sec
Number arrival streams	1
Number arrival flights	10
Number departure flights	10
Perturbation probability	10%

4 Results

In this section, we show results of the DD algorithm, first for the example scenario shown in Table 2 and then three scenarios shown in Figure 5 used to compare DD to a simple aircraft count metric.

We begin with an example scenario to demonstrate the algorithm. DD is computed for each three-minute chunk of the simulated schedule, per the simulation input parameters in Table 3. The flights in the first chunk, i.e., those with ETAs between 0:00 and 2:30, are close enough to landing to accommodate a firm arrival time. Thus, the DD algorithm computes a conflict-free revised ETA for each flight; and DD, delay severity, and residual delay for the chunk. For the remaining unfrozen chunks, t =[3:00,6:00), t =[6:00,9:00), and t =[9:00,10:00), covering the remainder of the 10-minute simulated schedule, predicted DD and residual delays are provided to the user. The (potentially) revised schedule is hidden due to the continued uncertainty of the situation. Figure 4 shows the output values.

In practical use, DD would be updated periodically to cover the next H_S period. To simulate time passing during development, DD is updated every H_F min, but only for the initially scheduled flights, without including additional flights beyond the initial H_S period. Thus, DD is rerun starting at simulated time t =3:00 and considering all flights until t =10:00, with the frozen chunk including flights with proposed ETA between 3:00 and 5:30. (In practical use, DD would be rerun at t =3:00 and would consider flights proposing to land/depart from t =3:00 to t =13:00.) The proposed times for all flights are updated to simulate flights being delayed or expedited due to operational uncertainties. An example is shown in Figure 4b where flight UAM_a_0_11 was delayed from its previously proposed ETA of 5:30 to a new proposed ETA of 6:00. Figure 4b shows the DD, chunk delay severity, revised ETAs, and residual delay for the frozen chunk, and the DD and residual delay for the remaining two chunks. Note that the delay has pushed UAM_a_0_11 into the next chunk. It still gets delayed due to congestion and lack of parking, but because of the 30-sec operational delay, the congestion-imposed delay is shorter, lessening its severity score, and, thereby, the DD category for the chunk. Also note that the responsibility for tracking total borne and predicted delays remains with the operator. Although congestion has eased by the operational delay of UAM_a_0_11, perhaps improving the situation for others, its reserves will still be impacted. To conclude the example, Figure 4c and Figure 4d show the results after again simulating time passing to 6:00 and then to 9:00, each time possibly perturbing the proposed ETA or ETD to imitate effects of operational uncertainties.

We now compare results of the DD algorithm with an aircraft count metric often utilized in traditional air traffic management. To ensure air traffic controller workload remains manageable and does not impact the efficiency of the NAS, a limit is placed on the number of flights that can

```
Chunk: [00:30-02:30]  DD: NEG   WtAvg Del Sev: 0.7
Schedule:
       Flight Prop ETA Cleared ETA Delay
0  UAM_a_0_1     0:30        0:30  0:00
1  UAM_a_0_3     1:30        1:30  0:00
2  UAM_a_0_4     2:00        2:30  0:30
Residual delay:  00:30
-------
Prediction for remaining 3 chunks:
Chunk           DD      Residual-delay
[03:00-05:30]   NEG       03:30
[06:00-08:30]   MOD       06:30
[09:00-11:30]   NEG       00:30
```

(a) Output at t =0:00.

```
Updated ETA: UAM_a_0_11 was 5:30        now 6:00
-----------------------------------
Chunk: [03:00-05:30]  DD: NEG   WtAvg Del Sev: 1.0
Schedule:
       Flight Prop ETA Cleared ETA Delay
3  UAM_a_0_6     3:00        3:30  0:30
4  UAM_a_0_7     3:30        4:30  1:00
5  UAM_a_0_9     4:30        6:30  2:00
Residual delay:  02:00
-------
Prediction for remaining 2 chunks:
Chunk           DD      Residual-delay
[06:00-08:30]   LOW       06:30
[09:00-11:30]   NEG       00:30
```

(b) Output at t =3:00.

```
Chunk: [06:00-08:30]  DD: LOW   WtAvg Del Sev: 2.4
Schedule:
        Flight Prop ETA Cleared ETA Delay
6  UAM_a_0_11     6:00        8:30  2:30
7  UAM_a_0_14     7:00       10:30  3:30
8  UAM_a_0_15     7:30       12:30  5:00
9  UAM_a_0_17     8:30       14:30  6:00
Residual delay:  06:00
-------
Prediction for remaining 1 chunks:
Chunk           DD      Residual-delay
[09:00-11:30]   NEG       00:30
```

(c) Output at t =6:00.

```
-----------------------------------
Chunk: [09:00-11:30]  DD: NEG   WtAvg Del Sev: 0.0
-------
No remaining chunks.
```

(d) Output at t =9:00.

Figure 4.—DD algorithm output for example schedule shown in Table 2: (4a) Output for frozen chunk includes DD, weighted average delay severity score, revised ETA for the frozen flights, and residual delay that will propagate to next chunk. Output for unfrozen chunks is predicted DD and residual delay. (4b) Process is repeated at t =3:00 minutes to simulate time passing. Figure 4b shows an arrival that has been delayed due to operational uncertainties. It still gets delayed due to congestion and lack of parking, but because of the 30-sec operational delay, the congestion-imposed delay was shorter, thus moving the predicted DD category from *MOD* shown in Figure4a to *LOW* shown in Figure4b. The process is repeated with simulated time passing to (4c) t =6:00 and (4d) t =9:00.

be in an airspace (a controller's *sector* or airport) during a time frame, say 15 minutes .

Figure 5 shows three scenarios that could occur under DCB with a specified capacity limit of at least four flights within a three-minute period. Flights are approaching the vertiport from two independent directions via *Stream A* and *Stream B*. The landing pad is shared by both streams. In each scenario, the flights are well separated from each other, in accordance with the en route separation assumption. In the tables in Figure 5, the first column shows the proposed ETA, the second and third columns show flights arriving on each stream, and the final column shows the DD-computed ETA that enforces the required pad separation assumed in Subsection 2.2 and the amount of delay for that flight. Note that for illustration purposes and to keep the example simple while still communicating the gist of the message, we switch the delay-to-DD mapping to use a simpler sum of delays per chunk mapping and split the categories as NEG if less than 1 min

For more information on the Monitor Alert Parameter (MAP) and how the count value is set, see FAA 7210.3, Chapter 18, Section 9, at https://www.faa.gov/air_traffic/publications/atpubs/foa_html/chap18_section_9.html

delay, LOW if between 1 but less than 3 min delay, MOD if between 3 but less than 5 min delay, and HIGH if 5 min delay or greater.

The aircraft count for each scenario is four. None of the scenarios would trigger an alert from the aircraft count metric if the threshold is set to at least four.

In contrast, because DD considers not just capacity constraints but also shared-resource constraints, DD varies from no delay and negligible DD to a total delay of 5 min and a high DD. Even in the scenario shown in Figure 5b, where the flights on each stream are separated by the pad-spacing interval, sharing with another stream results in moderate DD.

The longest delay experienced by a flight in the given scenarios is 2.5 min, an amount likely within required reserves and not likely to alter a pilot's flight plan. (And is an example of why we use the weighted average described in Subsection 2.4.) However, flights in succeeding chunks could quickly exhaust reserves due to propagated residual delays from previous chunks if their groups are similarly scheduled. For example, consider the chunks in Figure 6. Four aircraft are scheduled in each chunk, satisfying DCB restrictions. However, because they are all unfavorably scheduled at the end or beginning of a chunk, delays grow quickly. This is one scenario where the additional information provided by DD over just a simple aircraft count metric facilitates planning that may lead to safer decisions.

5 Conclusion

This paper presents a metric that supports safe landing or facilitates timely diversion for high tempo Advanced Air Mobility (AAM) vertiport operations. The metric uses the sociology concept of dynamic density (DD) that takes into consideration not only number of flights, but also the interaction of those flights with the vertiport's limited resources, namely the landing pads and the parking spots. We describe a simple algorithm that supports landing decisions while maintaining operational flexibility as long as practical. Arriving flights are separated into two groups: those that are very close to the vertiport and can accommodate a fixed arrival time without undue stress, and those that may still require schedule flexibility. The separation point is known as the *freeze horizon* and flights in the first group are referred to as *frozen flights*. Frozen flights are scheduled for conflict-free pad arrival time and are guaranteed a parking spot on arrival. Remaining flights, also known as *unfrozen flights*, maintain schedule flexibility until reaching the freeze horizon. For these flights, DD and expected delays support Pilot in Command (PIC) decisions about whether to proceed, expedite, delay, or divert. DD also supports air traffic control (ATC) and vertiport operators in airspace and vertiport usage.

In follow-on work, we plan to include multiple vertiports and multiple landing pads per vertiport, additional operational uncertainties, and off-nominal events. Scheduling of flights will begin pre-departure using strategic deconfliction, such as demand capacity balancing (DCB), though pop-up flights, such as unscheduled on-demand flights, will remain possible. Scheduling inside the intent sharing horizon will utilize multiple landing pads per vertiport. Departures and arrivals will be linked more tightly, taking into account required turn-around time between flights. Flights will be impacted by additional real-world uncertainties, such as arriving earlier than anticipated and diverting to a different vertiport due to, e.g., aircraft system failures or imminent fuel exhaustion. Further, events will be incorporated that simultaneously affect multiple flights, such as closing a vertiport. These extensions will contribute to a more comprehensive test of the efficacy of the DD metric.

ETA_P	Stream A	Stream B	ETA, delay
0:00	A1		A1, 0
0:30			
1:00		B3	B3, 0
1:30			
2:00	A5		A5, 0
2:30			
3:00		B7	B7, 0

Total Aircraft = 4 Total Delay = 0 min
DD = NEG

(a) Interaction of flights results in negligible DD.

ETA_P	Stream A	Stream B	ETA, delay
0:00	A1		A1, 0
0:30		B2	
1:00	A3		B2, 0:30
1:30		B4	
2:00			A3, 1:00
2:30			
3:00			B4, 1:30

Total Aircraft = 4 Total Delay = 3 min
DD = MOD

(b) Interaction of flights results in moderate DD.

ETA_P	Stream A	Stream B	ETA, delay
0:00	A1	B1	A1, 0
0:30	A2	B2	
1:00			B1, 1:00
1:30			
2:00			A2, 1:30
2:30			
3:00			B2, 2:30

Total Aircraft = 4 Total Delay = 5 min
DD = HIGH

(c) Interaction of flights results in high DD.

Figure 5.—Comparison between dynamic density metric and aircraft count metric. Chunks with the same number of aircraft can have significantly different dynamic density, depending on the distribution of the flights' proposed ETA. Interaction between flights and the shared resources significantly changes the DD metric, from negligible (5a) to moderate (5b) to high (5c).

To fully vet the metric, we plan to conduct a human-in-the-loop (HITL) interactive simulation on a variety of representative scenarios. Because *how* the information is provided to the user impacts its usefulness, the HITL study can compare the situational awareness and decision support benefits of several visualization options, as well as inform the design of future visualization alternatives and data requirements.

References

1. Uber Elevate: Fast-forwarding to a future of on-demand urban air transportation. `https://evtol.news/__media/PDFs/UberElevateWhitePaperOct2016.pdf`, 2016. Accessed: 2023-08-16.

2. Mallela, J.; Wheeler, P.; Le Bris, G.; and Nguyen, L.-G.: *Urban Air Mobility: An Airport Perspective*. No. ACRP Project 03-50A, 2023.

ETA_p	Stream A	Stream B	ETA, delay
0:00			
0:30			
1:00			
1:30			
2:00			
2:30	A1	B1	A1, 0
3:00	A2	B2	

(a) First chunk of flights scheduled for end of chunk.

ETA_p	Stream A	Stream B	ETA, delay
3:30	A3	B3	B1, 1:00
4:00	A4	B4	
4:30			A2, 1:30
5:00			
5:30			B2, 2:30
6:00			
6:30			A3, 3:00

(b) Second chunk of flights similarly scheduled.

ETA_p	Stream A	Stream B	ETA, delay
7:00	A7	B7	
7:30	A8	B8	B3, 4:00
8:00			
8:30			A4, 4:30
9:00			
9:30			B4, 5:00
10:00			

(c) Third chunk, also similarly scheduled.

Figure 6.—Unfavorable scheduling of flights within a chunk quickly propagates delays and increases risk to flights scheduled for later chunks. Note that each chunk has the same number of aircraft, as allowed by DCB, but delays grow significantly in later chunks.

3. Doole, M.; Ellerbroek, J.; and Hoekstra, J.: Drone delivery: Urban airspace traffic density estimation. *8th SESAR Innovation Days, 2018*, 2018.

4. Chen, S.; Evans, A.; Brittain, M.; and Wei, P.: Integrated Conflict Management for UAM with Strategic Demand Capacity Balancing and Learning-based Tactical Deconfliction. *arXiv preprint arXiv:2305.10556*, 2023.

5. Lee, H.; Moolchandani, K. A.; and Arneson, H.: Demand capacity balancing at vertiports for initial strategic conflict management of urban air mobility operations. *2022 IEEE/AIAA 41st Digital Avionics Systems Conference (DASC)*, IEEE, 2022, pp. 1–10.

6. Roychoudhury, I.; Spirkovska, L.; Oconnor, M.; and Kulkarni, C.: Survey of Methods to Predict Controller Workload for Real-Time Monitoring of Airspace Safety. , 2018.

7. Spirkovska, L.; Kulkarni, C. S.; Watkins, J.; and Martin, L.: Urban Air Mobility Airspace Dynamic Density. *AIAA AVIATION 2022 Forum*, 2022, p. 3403.

8. Pradeep, P.; Xue, M.; Lee, P. U.; Sridhar, B.; Li, J.; and Huynh, P. T.: Study of Pairwise Deconfliction Metrics to Analyze Air Traffic Complexity in Upper Class E Airspace. *AIAA AVIATION 2023 Forum*, 2023, p. 4103.

9. Dulchinos, V.; Wood, R. D.; Farrahi, A.; Mogford, R.; Shyr, M.; and Ghatas, R.: Design and analysis of corridors for UAM operations. *2022 IEEE/AIAA 41st Digital Avionics Systems Conference (DASC)*, IEEE, 2022, pp. 1–10.

10. Chin, C.; Gopalakrishnan, K.; Balakrishnan, H.; Egorov, M.; and Evans, A.: Protocol-Based Congestion Management for Advanced Air Mobility. *Journal of Air Transportation*, vol. 31, no. 1, 2023, pp. 35–44.