Contents

Chapter 1

Current State of Affairs: Economic Impact

The FAA and NASA as well as all stakeholders in the development of Advanced Air Mobility (AAM) are developing a framework for the general organizational strategy for large-scale UAS traffic management (UTM) in urban areas (also called Urban Air Mobility or UAM). These stake holders include governmental bodies charged with overseeing air space utilization, as well as providers of UAM services (PSU), formerly called UAS Service Suppliers (USS). UAS operators such as Amazon, UPS, hospitals, etc. are anxiously awaiting operational UTMs which will enable package and drug delivery, as well as unmanned air taxi services. Companies like Airmap, Bell Helicopter, GE and others have expressed great interest in exploiting such a system. NASA has done market surveys which indicate that by 2030 there may be 750M air taxi flights and 500M package deliveries per year in 15 major cities. In addition, this work may allow efficient integration and synergy between ground and air vehicles. Finally, the existence of such a system will also enable the acquisition of a whole new source of big data (flight data, sensor data, communications data, weather data, etc.) which may form the basis for a wide variety of new services.

Current research and product development aim to catalyze the adoption cycle that underlies the nascent industry of Urban Air Mobility (UAM). In its 2020 forecast publication [1], the FAA acknowledges that “it is extremely difficult to put a floor on the growth of the commercial UAS sector due to its composition and the varying business opportunities and growth paths.” However in the same study they say, “if, for example, professional grade small UAS (sUAS) meet feasibility criteria of operations, safety, regulations, and satisfy economics and business principles and enter into the logistics chain via small package delivery, the growth in this sector will likely be phenomenal;” phenomenal, relative to the forecast of about one million non-model aircraft operating for commercial reasons in 2024, each registering multiple flights per day [1]. This fleet does not include the vehicles expected to deliver about one million express packages in that same year, according to a study conducted by NASA [4, 2]. The FAA also estimates between 12,000 and 23,000 passenger-carrying autonomous aircraft operating within urban environments by the year 2030. As the FAA suggested in their assessment however, these estimates rely on the assumption that UAM technology will be adopted and that efficient concepts of operations (CONOPS) can be developed.

Consulting reports and conversations with industry stakeholders indicate that most believe *regulation* to be the highest inhibiting factor to growth of the UAM industry. However, NASA's own funded study regarding the barriers to adoption indicate a much more complex landscape, including technical factors as well as market conditions (see Figure 1.1). Therefore, the more realistic view sees regulation as an outcome of progress in the technological development of this industry. The more realistic characterization is where conflicts exist between every pair of stakeholders, and it is the complexity of these relationships that inhibits growth.

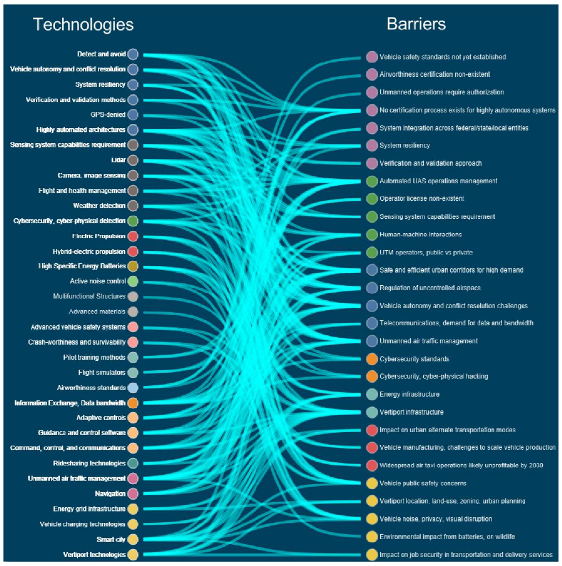

Figure 1.1: NASA's compilation of barriers for UAM is a complex landscape [4].

One of the authors, D. Sacharny, has developed the GeoRq platform which addresses these complexities by providing a collaborative integrated development environment with specialized system development tools, and by structuring the problem in terms of system-level policies and agent behaviors (see Figure 1.2) using the lane-based approach described throughout the book. Three organizational components form the platform: tools to create and store requirements (specifically geospatial-temporal requirements), tools to create impact and benchmark metrics, and tools to create real or simulated deployments. Both the lane-based approach and the platform are critical components because one provides the conceptual and computational framework for analysis, and the other provides a vehicle for collaborative engineering and commercialization.

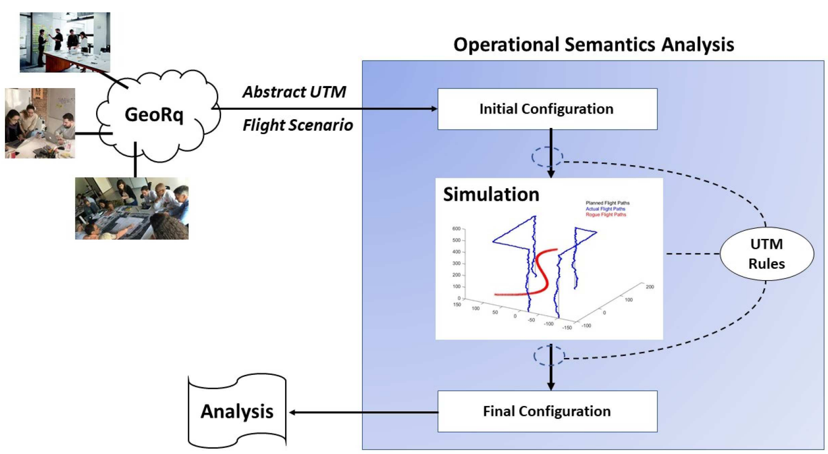

Figure 1.2: The Core Component of the GeoRq Platform.

Example Business Model

The main revenue streams for such a product include subscription to cloud-services (deployed and secured platform workspaces), access to APIs and microservices such as the Lane-Based UAS Management System, licensing and data-access fees. For example, the *GeoRq Workspace* is a cloud deployment consisting of multiple connected instances of virtual machines (VM), databases, and configurations. A GeoRq Workspace may feature an instance of a flight scheduling system, an instance of Eclipse Theia with GeoRq extensions, GeoServer to provide web-map services, two-instances of GeoRq's Provider of UAS Services (PSU), an OIDC security server, and 2TB of Google-backed storage. This setup supports designing, testing, and deploying large scale logistics operations: one PSU communicates with the region's UTM, while the other forms a digital-twin to simulate deployments, and the Eclipse Theia instance with GeoRq extensions serves both the end-user as an Air Traffic Operations Center (ATOC) and the developers as an integrated development environment. Workspace configurations can be updated dynamically with fine-grained resource pricing, and each workspace supports multiple users (contingent upon resource requirements).

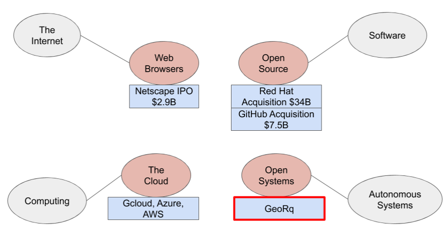

Figure 1.3: Similar Business Models Capitalizing on Trends in Large-Scale Collaborative Engineering.

In a nascent industry such as UAM, companies must replicate a similar structure of computational instances to conform to UAM system policies. However, the intense competition between current players to develop, and become the standard bearer of UTM software, has forced much of the common architecture into proprietary silos. The result is that non-recurring engineering (NRE) in this space, such as required by new-product development, is expensive and compounds with each new engineer that must climb the same hill.

Open source development, as with GeoRq, overcomes this problem by packaging up the common architecture, making it configurable, extensible and deployable, and by providing an integrated open-source systems development tool. Product developers can then repackage proprietary APIs, datasets, microservices, UIs, etc., and deploy the white-labeled GeoRq Workspace as a new product for their clients. Reducing NRE by building products using open-source and collaborative software enlarges the pool of qualified designers, engineers, and users, and it can have dramatic effects on the growth of industries (some examples shown in Figure 1.3). In the case of a minimal GeoRq Workspace, not including strategic deconfliction or PSU deployments, a standard software estimation tool applied to the current codebase estimates approximately 17 months and 8 engineers to complete this common architecture. The cost estimate of $\$1.6M$ assumes an average wage of $\$56,286$; however a higher average wage is likely due to the narrow expertise required.

After many discussions with potential subscribers, industry stakeholders, and government, our observation is that the drive to create products for the UAM industry exists across many disciplines. Table 1.1 shows a sample of the companies interested in this problem. For example, a company might acquire a patent for advanced trajectory generation. After integrating the capability into a web-based

API, they would spend considerable NRE developing visualizations using, for example, NASA's WorldWind libraries for marketing purposes. Given the chance to use a tool like the GeoRq workspace and the visualization capabilities available there, the API strategy might change considerably. The realization would be that packaging a company's technical capability within the a platform like GeoRq provides a powerful channel to market their product as part of a deployable system. Another example would be a developer engaged in the NASA AAM national campaign in order to commercialize communications research. This would require the development of a PSU for a valid simulation and the necessary infrastructure to deploy a production instance of their technology - this is a costly endeavour considering the NRE required. Access to a GeoRq-like system would accelerate their research. Integration of UAM infrastructure would allow product developers across industry to demonstrate the feasibility and potential for commercial investment. It would not be necessary to spend a considerable amount of NRE developing a web-based system for exploring and visualizing their data, including updates for changes in the AAM framework as this industry develops; systems like GeoRq are a cost-effective alternative.

A viable business model emerged through these discussions: offer product developers a configurable, cloud-deployable package containing the prerequisites for any UAM product. A basic set of features would be included, with additional cloud capabilities and deployments (such as large-scale publish/subscribe frameworks) available through fine-grained resource pricing. An open-source tier is provided to generate community engagement and sustainable commitment to the platform, allowing developers to customize the platform as the UAM industry evolves. An *Individual* tier addresses the needs of smaller firms, individual entrepreneurs and researchers. The *Enterprise* tier is for firms that plan to develop multiple products or to deliver the white-labeled platform as a product to downstream clients.

Table 1.1: Potential Subscribers and Active Discussions

Firm Name
Crown Castle
Crown Consulting
Skytelligence
SmartSky Networks Rockafellow
AiRXOS (GE)
UPS
Aerial Transportation Solutions (ATS) Essleman
AirMap
ANRA Technologies
University of North Texas Namuduri
Camel Works Design (Dubai Road Transit Authority) Khanjari
Anne Arundel Hospital System
Alakaí Technologies
Westinghouse Electric Company
Fortem Technologies
CogniTech Corporation
University of Utah Health

To estimate potential revenue given this pricing model, a sample list of potential subscribers was collected from pre-certified consulting firms for several state departments of transportation. The list was narrowed to consulting firms with the following capabilities, having a high likelihood of serving UAM requirements: surveying and mapping, geotechnical services, traffic operations design, traffic engineering and operations studies, and environmental studies. This compiled list included 1383 firms with an estimated median of 32 technical staff per firm. We expect that technical staff will be drivers, as well as end-users, for adopting GeoRq. As an example, the total number of technical staff present in one dataset (the most descriptive dataset) was 158,286 people. For this sample of the total addressable market, if 0.3% of the technical staff see potential in serving UAM requirements with their capabilities and each adopts a single enterprise tier package, then the **total annual revenue exceeds $28M**. This figure considers the first workspace adopted by these developers, and it becomes compounded as more products are developed, white-labeled, and adopted by downstream clients. Furthermore, this sample market represents a fraction of the developers that will enter this industry in the next few years. The total addressable market for a GeoRq like tool is likely orders of magnitude above this sample, especially if complementary markets (GIS, programming IDEs, cloud computing) are considered.

The margins on selling this type of NRE are large, the marginal cost to run the enterprise tier in the cloud runs annually about \$362. For a firm, or even an individual, deciding whether to venture into product development in this nascent UAM industry, the value proposition is dramatic: a GeoRq like product reduces the necessary investment by at least $\$2M$, and accelerates development by at least 1.5 years.

Complementary and competitive products exist, however they are either too broad, for example Google Cloud and IBM Rational, or too narrow to serve this industry well. Both Google and IBM possess the means to beat the NRE estimates given earlier and bring a competitive product to market quickly. However, the open-source components of a GeoRq tool is a mitigation tactic by diversifying the expertise capable of using the platform. A more recent trend for these large firms is to *acquire* companies such as GeoRq, for example Google and Android, IBM and Red Hat, and Microsoft and GitHub. Google may also have incentive to make an acquisition offer since Eclipse Theia is offered on Google Cloud for other purposes; acquiring GeoRq extensions is lower risk than developing its own version. For products that are developed atop GeoRq, such as UTMs and PSUs, direct competitors exist such as AiRXOS, ANRA, AirMap, and others. However, these products do not serve the same client base (product and technology developers) since their platforms are not open and reconfigurable as is a GeoRq. If a competitor decided to white-label the common architecture of their products, they would run into the issue of integration, which is what GeoRq's systems development tools address. Other complementary products exist, for example ArcGIS, which provides GIS capabilities, as well as some limited application development tools. For product developers, they provide a number of software development kits (SDKs) to interface with their offerings. With the tools that ArcGIS provide, an autonomous system product developer would still need to obtain and configure cloud computing resources, external IDEs, and expertise in deploying a product that interfaces with UAM systems - features that GeoRq provides as a package.

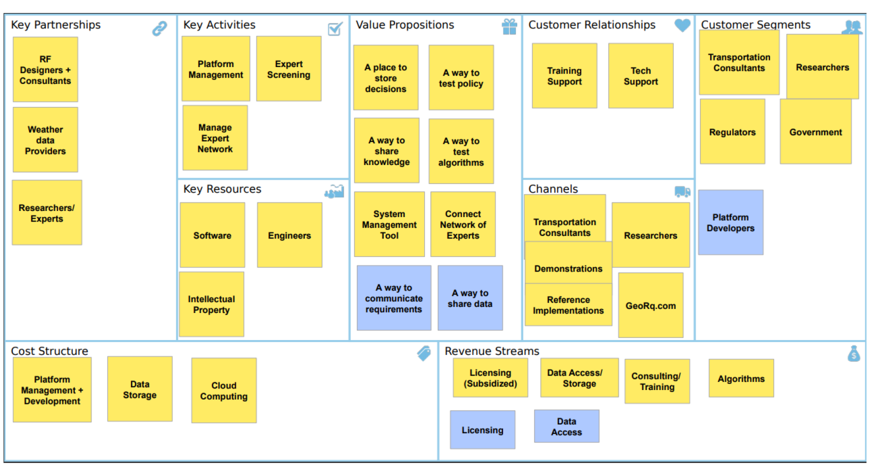

Figure 1.4: Business Model

Product Development

An example business model for a GeoRq system is given in Figure 1.4, and the product roadmap is shown in Figure 1.5.

[Describe business model.]

The roadmap is divided into sub-product tasks and encapsulated into representative elements called Epics. Each epic contains many software tasks, including design, implementation, and testing. Each epic contains tasks that improve the platform and provide open-source products, but the main thrust of each epic are revenue generating microservices and product specific tools.

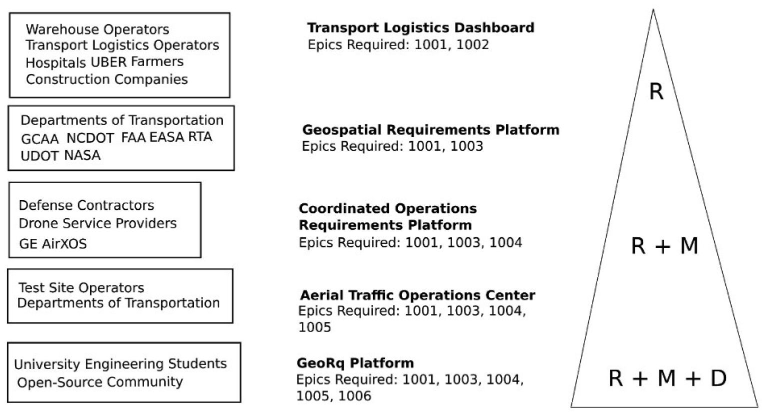

Figure 1.5: Product Road Map

Commercialization Approach

The GeoRq platform is an example vehicle for commercializing research. Research efforts produce software to perform simulations, record and validate benchmarks, and test assumptions. Source code can be delivered directly as part of a workspace configuration or wrapped in a microservice. Front-end code is engineered by programmers using GeoRq extensions, then included with individual or enterprise tiers. The commercial feasibility of each product is measured by the value (the marginal price of selecting this feature with a GeoRq workspace) over the cost of the computational resources required to run that feature in the cloud (e.g., required datasets, storage requirements, etc.) and the NRE required to produce it.

Developers of systems like GeoRq can apply for a variety of assistance from state and local entities to assist with portions of business development and commercialization. It is usually possible to work with the state agencies to identify, bid and win procurement opportunities with federal, state, and local government entities. Furthermore, it is possible to seek assistance from the appropriate Small Business Development Center (SBDC) to receive business counseling and assistance in business plan development.

Chapter 2

Introduction to UAS Traffic Management

2.1 FAA/NASA

Chapter 3

Lane Networks

3.1 LANE-BASED URBAN AIRWAY SYSTEMS

There are many reasons to fly UAS in an urban environment. Several expected high usage applications are general package delivery (e.g.,food, medical supplies, general goods), inspection (e.g., buildings, bridges, power infrastructure, etc.), and air taxi service. Major companies like Amazon, UPS, the Postal Service, etc., may deploy hundreds or thousands of UAS per day. Every one of these UAS will follow some trajectory accoding to a specified time schedule; this is a 1-dimensional curve in a 4-dimensional space. If every UAS creates an individual and arbitrary 4-D curve, then every pair of trajectories must be checked to ensure safety (i.e., minimal separation at all times). Moreover, given thousands of UAS in the air at one time, their safe operation is too complex to alllow human operators, and their safety monitoring is too complex for human air traffic controllers. This means that the flight of the UAS must be autonomous.

An alternative to the set of arbitrary trajectories is to create a pre-defined set of lanes through the air and to require that all UAS flights follow these lanes. Each flight must consist of a set of lanes; this starts with a launch lane that takes the UAS from the ground to the air, followed by a sequence of lanes through the air, and terminating wih a landing lane that takes the flight from the air to the ground. To ensure safety, a time slot through each lane (i.e., a lane entry time and a lane exit time) must be reserved for each flght so that at no time are any two flights too close. This is called strategic deconfliction, and an efficient method for lane-based networks is provided in the next chapter. To make such a method possible, some contraints are placed on the network:

- each lane is one-way (i.e., the network is a directed graph),
- the direction of an edge is related to compass heading,
- each lane is longer than some minimal length, and

  item roundabouts are used to allow multi-directional ingress and egress over a specific geographic ground location.

An easy way to obtain the basic layout of the airways over a given urban area is to define an undirected graph at ground level. For example, a simple grid may be affixed to ground locations, or the existing ground road network may be used wherein every road intersection or termination point is a vertex, and road segments between vertexes are edges. To achieve travel in both directions between air vertexes corresponding to to a ground vertex, the air lanes must be placed either side by side at the same altitude above ground level, or one above the other. The convention used here is that lanes with travel in opposite directions will be vertically separated; moreover, travel in directions $[0, \pi)$ will be at one altitude and in directions $[\pi, 2\pi)$ in the other.

To implement this, there are two levels of airways. In addition, a roundabout is created at each level above a ground vertex. Lanes to enter the airways from the ground are called launch lanes and connect a ground vertex to the lower altitude airway level. A landing lane connects the lower altitude airway level to a ground vertex. To achieve these connections, a vertex is placed in the roundabout to connect to a corresponding launch or lane location, respectively.

Figure 3.1 shows a 2x2 ground grid (an undirected graph with four nodes and four edges) and the corresponding air network (a directed graph with 40 nodes and 56 edges). Note that there are lanes connecting the the two roundabouts (separated by altitude) at a vertex location – one up and one down. This allows a flight to enter an air vertex at either altitude and exit along any lane leaving the vertex. Each ground vertex has an associated distinct ground location for its launch and landing lanes, if they exist. As can be seen in the figure, these launch and lane ground vertexes connect directly to corresponding vertexes in the low altitude roundabout above the ground vertex.

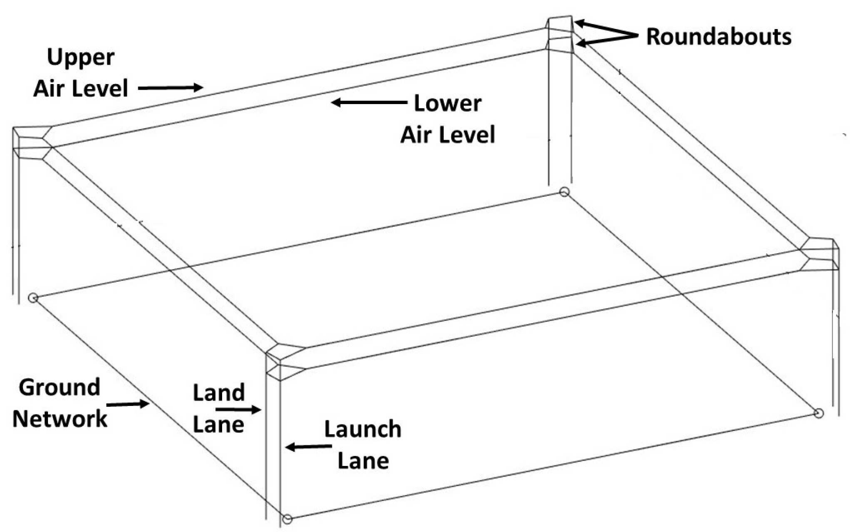

Figure 3.1: A Simple 2x2 Grid Network.

The lane-based approach defines a set of one-way lanes where each lane is defined by an entry point, an exit point, and a one-dimensional curve between the two (here we use straight line segments). UAS travel in three dimensions, and thus lanes are defined as 3D corridors (e.g., cylindrical-like tubes). The shape of corridors may change dynamically and should be constructed to account for the idiosyncrasies of the vehicles that they are meant to support; for example, smaller aircraft in windy environments may require a larger corridor radius than a heavier vehicle with better control dynamics. Further design constraints can be defined in terms of the headway – or safe separation distance – between UAS. The combination of headway and corridor design can support a range of vehicle trajectory constraints, while the directed graph (digraph) imposed on the airspace presents agents with a structured environment for computation (the lanes represent a complete model of the airspace under ideal conditions). Lanes may also have other associated properties (e.g., speed restrictions) specified by the UTM, enabling regulators to communicate requirements effectively to all agents in the system.

Lanes are connected so that every vertex has either in-degree or out-degree equal to one. This permits scheduling to be based on lanes as opposed to vertexes since all flights may be deconflicted based on one incoming or outgoing lane, and simplifies the analysis of congestion because various graph-based measures can be utilized to determine most likely high congestion parts of the network. This contrasts with zone-based deconfliction that presumes vehicles can enter and exit in any direction and the entire zone must be reserved (inefficient for large areas), and cell-based deconfliction that combines zone reservation with general motion planning within each cell (similar to the two-phase decoupled approach in [6]). The choice of the lane spatial layout is key to operational performance. As previously described, several alternatives exist:

1. airways modeled from ground road networks,
2. regular grid networks,
3. networks with specific properties (e.g., Delaunay networks).

We now give a more detailed account of the lane creation process.

Lane creation starts with a ground network defined as a graph, $G = (V, E)$, where V is a set of ground position vertexes, and E is a set of undirected edges between the vertexes. Figure 3.2 shows an example network for a small set of roads from San Francisco, CA. To create the two-level airways between vertexes, the ground network is duplicated as a set of airway lanes at two altitudes: one for travel in direction $[0, \pi)$, and the other in direction $[\pi, 2\pi)$. Since ground vertexes are road intersections, each is represented by two roundabouts in the air centered over the vertex; these ideas are demonstrated in Figure 3.3). A view of the airway network displayed over the San Francisco area is shown in Figure 3.4.

3.2 Spatial Network Measures

The behavior of requests and the strategy for scheduling can have a significant impact on the average density of lanes. Consider a single lane system of length x, with one entry and one exit. Further assume that vehicles consume a one-unit spatial interval within the lane, and requests arrive over time independently for a uniformly random unit interval. In the first scenario, assume that each vehicle either obtains the requested reservation or drops out, a "failure." This scenario mirrors a 1-dimensional sequential interval packing problem, also known as Renyi's parking problem [5]. Renyi showed that as the length of the lane approaches infinity, the mean filling density approaches 0.7476. This property also holds for the lane scheduling approach given here.

The layout of the lane system can also have significant effects on the behavior of the system. A common refrain among air mobility enthusiasts is that the ability to travel point-to-point in a straight line, Figure 3.5 for example, should be maintained and decreases the desirability of structured airspaces. However, a system of

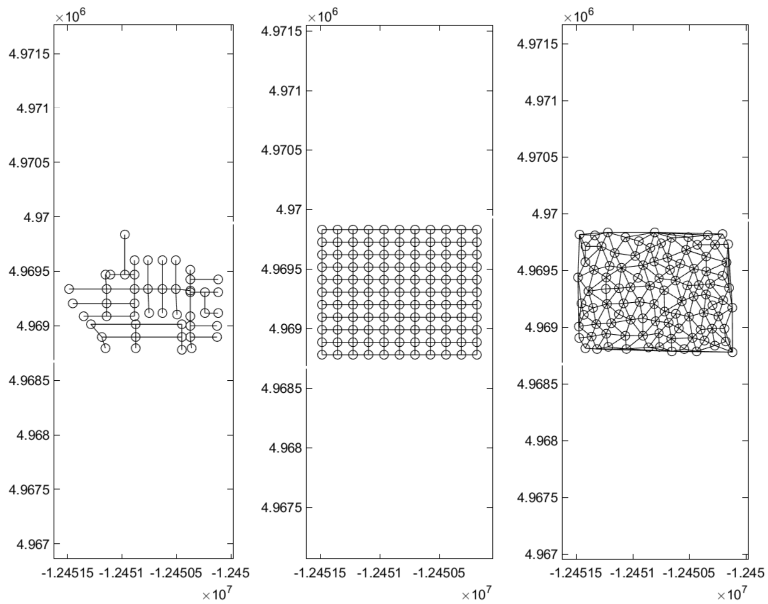

Figure 3.2: Three Types of Road Layouts over the Same Locale: Actual San Francisco Roads (left); Grid Layout (middle); Delaunay Trianglation (right).

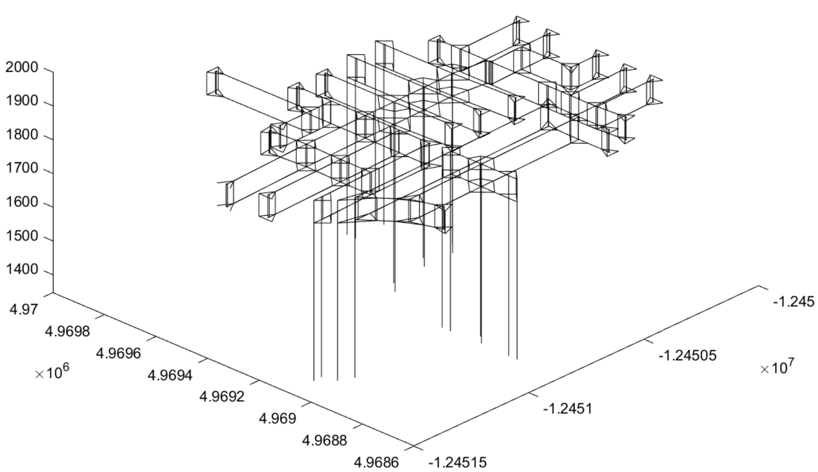

Figure 3.3: An Example Two-Level Grid Lane Layout of San Francisco Roads.

agents performing individually optimal trajectories in an unstructured airspace is unlikely to produce an efficient system. This is true in the case where agents can make decisions dynamically based on system-wide conditions, for example, Braess'

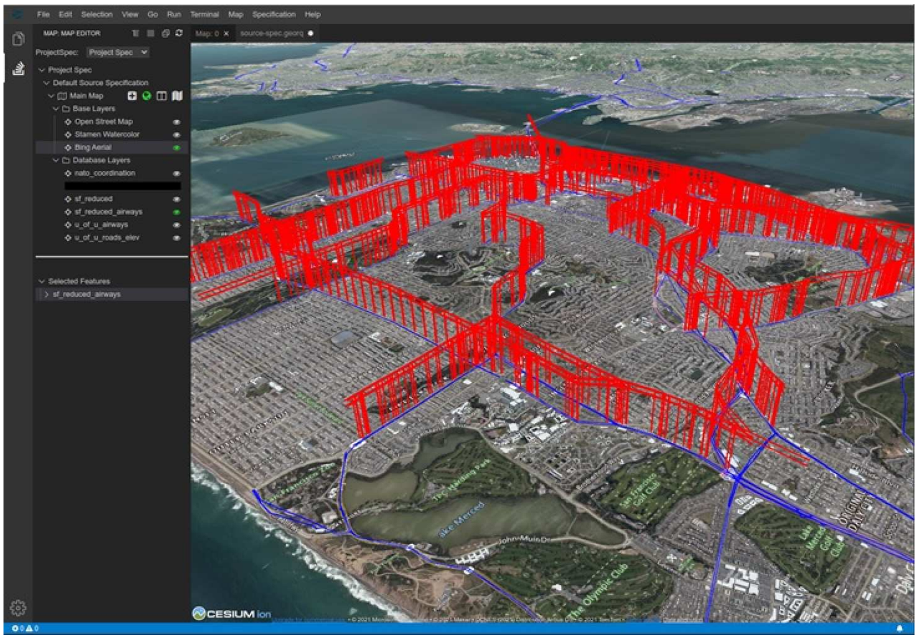

Figure 3.4: UTM Airways over the San Francisco Downtown Area.

paradox demonstrates where additional route options can result in an increase in travel time [3]. This also appears to be true when considering conflict counts for a simple cell-based deconfliction experiment (point-to-point flights deconflicted using ground-delay, we call the FAA-NASA approach). Figures 3.6 and 3.7 show histograms for cell traversals (how many times a flight crossed a cell) and intersections (how many flight paths intersect) for an experiment with 1000 UAS flying point-to-point in an unstructured airspace with uniformly distributed land and launch sites. These graphs show an increased density of conflicts focused in the center of the area of interest.

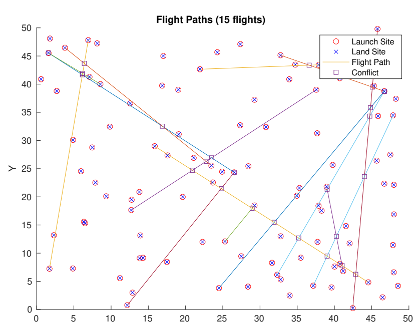

Figure 3.5: Sample Straight-Line Paths Between Launch and Land Vertices

This configuration of trajectories correlates to the structured regular-grid lane network, which exhibits the worst performance in the network comparison experiments described below. In Chapter 6, a simulation comparison between the point-

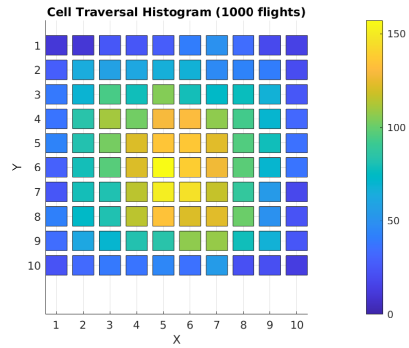

Figure 3.6: Cell Traversal Counts in an Unstructured Airspace with Point-to-Point Flights

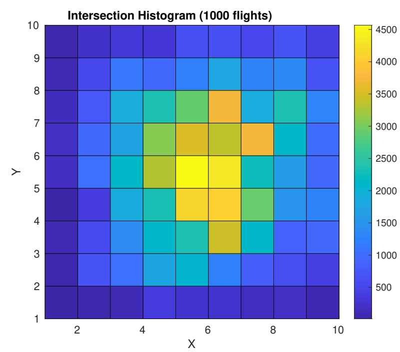

Figure 3.7: Intersection Counts in Unstructured Airspace (2-D Slice, Assuming Altitudes are within Minimum Separation)

to-point unstructured airspace and the lane-based approach is demonstrated.

Bibliography

[1] Federal Aviation Administration. FAA Aerospace Forecast Fiscal Years 2020-2040. Technical report, Federal Aviation Administration, 2020.

[2] Colleen Reiche and Rohit Goyal and Adam Cohen and Jacqueline Serrao and Shawn Kimmel and Chris Fernando and Susan Shaheen. Urban Air Mobility Market Study. Technical report, National Aeronautics and Space Administration (NASA), 2018.

[3] David Easley and Jon Kleinberg. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge University Press, Cambridge, 2010.

[4] National Aeronautics and Space Administration (NASA). Urban Air Mobility Market Study Executive Summary. Technical report, National Aeronautics and Space Administration (NASA), 2018.

[5] A Rényi. On a one-dimensional problem concerning random space filling. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences,*, 3:109–127, 1958.

[6] M. Saha and P. Isto. Multi-Robot Motion Planning by Incremental Coordination. In *IEEE International Conference on Intelligent Robots and Systems*, pages 5960–5963, 2006.