A FRAMEWORK FOR DIMENSIONING VDL-2 AIR-GROUND NETWORKS

Leila Z. Ribeiro, Leone C. Monticone, Richard E. Snow, Frank Box, The MITRE Corporation
Rafael Apaza, Steven Bretmersky, NASA Glenn Research Center

Abstract

This paper describes a framework developed at MITRE for dimensioning a Very High Frequency (VHF) Digital Link Mode 2 (VDL-2) Air-to-Ground network. This framework was developed to support the FAA's Data Communications (Data Comm) program by providing estimates of expected capacity required for the air-ground network services that will support Controller-Pilot-Data-Link Communications (CPDLC), as well as the spectrum needed to operate the system at required levels of performance. The Data Comm program is part of the FAA's NextGen initiative to implement advanced communication capabilities in the National Airspace System (NAS).

The first component of the framework is the radio-frequency (RF) coverage design for the network ground stations. Then we proceed to describe the approach used to assess the aircraft geographical distribution and the data traffic demand expected in the network. The next step is the resource allocation utilizing optimization algorithms developed in MITRE's Spectrum Prospector^{TM} tool to propose frequency assignment solutions, and a NASA-developed VDL-2 tool to perform simulations and determine whether a proposed plan meets the desired performance requirements.

The framework presented is capable of providing quantitative estimates of multiple variables related to the air-ground network, in order to satisfy established coverage, capacity and latency performance requirements. Outputs include: coverage provided at different altitudes; data capacity required in the network, aggregated or on a per ground station basis; spectrum (pool of frequencies) needed for the system to meet a target performance; optimized frequency plan for a given scenario; expected performance given spectrum available; and, estimates of throughput distributions for a given scenario.

We conclude with a discussion aimed at providing insight into the tradeoffs and challenges identified with respect to radio resource management for VDL-2 air-ground networks.

Background

A 2010 MITRE study [1] first introduced the framework presented here. That framework utilized an earlier version of the MITRE-developed Spectrum Prospector tool, and the NASA-developed VDL-2 simulator. In 2012, as a follow-on of that study [2], [3], MITRE developed algorithms in Spectrum Prospector to optimize frequency plans given a limited set of spectrum. That study also utilized an improved version of the NASA VDL-2 simulation tool, capable of simulating nationwide scenarios.

In 2013 MITRE performed an update to the study, described in [4], in which updates in inputs were considered, including more recent estimates of aircraft equipage profiles, traffic model inputs and other assumptions. Also in 2013, another parallel MITRE study [5] explored algorithms for load-balancing.

VDL-2 Air-Ground Networks

The VDL-2 protocol encompasses Data Link Layer functions including Aviation Link Control (AVLC) for the data link and *p*-CSMA (p-persistent Carrier Sense Multiple Access) for the Media Access Control (MAC) sub-layer, and Physical Layer. The main characteristics of each layer are briefly described below:

- The Physical Layer uses a differentially encoded 8-phase shift keying (D8PSK) modulation, operating at a channel data rate of 31.5 kbps over a 25 kHz channel.

- The Data Link Layer provides the following sublayer functions:

- The MAC sub-layer implements a p-persistent CSMA protocol to access the shared medium.

- The Data Link Service (DLS) sublayer provides connection oriented, half-duplex air-ground links using the AVLC protocol, which is derived from the High-Level Data Link Control (HDLC) protocol. The DLS

supports point-to-point and broadcast links.

- The VHF Data Link (VDL) Management Entity (VME) establishes and maintains Data Link Entities (DLEs) between the aircraft radio and ground stations using the Link Management Entities (LME). The LME procedures include the handoff initiation and management to maintain connection across multiple service volumes.

The p-persistent CSMA Method

Carrier-Sense Multiple Access (CSMA) is a method of **random access** to the RF medium (as opposed to other non-random methods such as reservation-based, or polling); i.e., nodes access the medium without a centralized control.

In **CSMA**, the node senses the medium before transmitting. If it senses the channel as busy (i.e., if it can receive a signal above a defined threshold), it refrains from transmitting and tries again later.

In **p-persistent CSMA**, if the medium is sensed as idle, the node transmits with probability p or defers with probability 1-p. After deferring, the node waits a short amount of time before attempting to transmit again with probability p or deferring with probability 1-p. This process repeats until a successful transmission or the medium becomes busy. If the medium is busy, it keeps sensing the medium until it becomes idle.

In the Data Communications scenario, the medium is the VHF channel, and the nodes are the aircraft and ground station radios for a given service volume that are currently assigned to that channel.

Random access methods have to deal with the **collision** problem, which occurs when two nodes' transmissions overlap or "collide" in a way that affects the intended reception.

Although the carrier-sensing performed in CSMA methods reduces the probability of collision compared to purely random access methods (such as ALOHA), collisions may still occur due to the **hidden node problem**, and the carrier-sensing may also cause unnecessary transmission postponements due to the **exposed node problem**. Those two effects are illustrated in Figure 1 in the context of an air-

ground scenario. In the illustration, consider that ground station GS-A and ground station GS-B are both utilizing the same frequency and have an overlap in their area of coverage, as illustrated by the intersecting blue shades above each. Assume also that they cannot sense each other due to the horizon or some other obstacle along the path. Their intended receiving aircraft in the example are AC-A and AC-B, respectively.
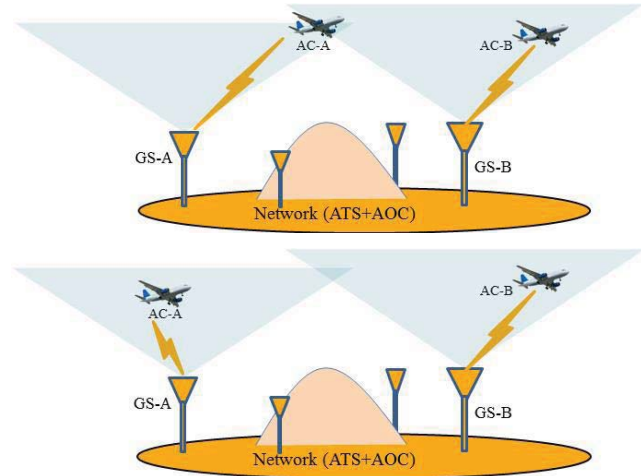
Figure 1. Illustrating the Hidden Node (top) and Exposed Node Problems (bottom)

- The **hidden-node problem** illustrated at the top of Figure 1 would occur when: GS-B starts transmitting to AC-B while GS-A is transmitting to AC-A. A collision would occur in AC-A, because AC-A is in the overlap area where signals arrive from both GS-A and GS-B. This event results in reduced throughput (due to retransmissions) and additional delay at the network. From the point of view of an aircraft receiving in AC-A position, the probability of collision is proportional to the probability of GS-B transmitting, which may be high since GS-B is a ground station.

Hidden-node is a severe problem in the en route environment and is typically related to areas where there is overlap of air-coverage without corresponding ground-station visibility. Note that there could exist other GSs between GS-A and GS-B that would not be relevant to the discussion if they were assigned to different frequencies (as in a cellular type of architecture). There is no need to assume that GS-B and GS-A are adjacent

neighbors for the examples above to be valid, only that they do have some coverage overlap and cannot sense each other.

- The **exposed-node problem** on the air-air interface is illustrated at the bottom of Figure 1. It would occur when AC-B has something to transmit to GS-B but refrains from doing so because it can sense that AC-A is transmitting to GS-A at the time. This would occur if AC-A and AC-B have line of sight to each other, and would result in an unnecessary delay (waiting for AC-A to finish), because the two transmissions (AC-A to GS-A and AC-B to GS-B) could occur simultaneously without a collision. From the point of view of AC-B, the probability of unnecessary delay is proportional to the combined probability of aircraft in its visible air-area (from the air point of view) transmitting to GS-A while outside the coverage overlap area.

The exposed-node is a less severe problem than the hidden node in en route domain, due to the lower probability of AC transmission (compared to ground station) and because its effect (unnecessary delay) has a lesser impact than a collision, which is followed by retransmission. Furthermore, it affects latency mainly in the downlink, which is not typically as critical for latency in the en route domain as the uplink.
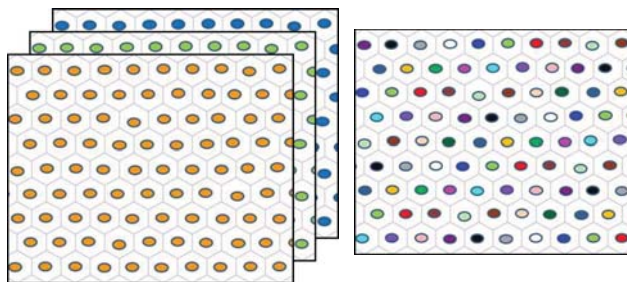
Architecture Approaches

In the en route domain, which is the focus of this discussion, two different approaches can be considered for deploying VDL-2 networks. We will refer to them as "reuse-1" and "cellular" approaches.

Figure 2 illustrates those two approaches. Each dot represents a ground station and the colors represent frequencies.

In a **reuse-1** approach, frequencies are reused in neighboring cells. The term "reuse-1" means that a single frequency could potentially be reused in every cell in the system. Reuse-1 could be used also in the context of multiple frequencies being utilized in the system, in which case it refers to each of those frequencies being reused in neighboring sites. This is represented by the multiple layers of reuse-1 on the

left side of Figure 2 (to indicate that all ground stations have the same three frequencies).

In a cellular architecture, different frequencies are assigned to cells that are adjacent or within a defined distance (or RF signal ratio) from each other. In the right side of Figure 2, this is illustrated by the reuse of frequencies (colors in the picture) in non-adjacent cells (reuse >1).

Figure 2. Notional example of different frequency reuse approaches for en route VDL-2. Left: reuse-1 multi-frequency; Right: cellular architecture

Reuse-1 allows hidden transmitter pairs to occur often, as most of the time ground stations do not have line-of-sight to each other in the en route domain. This has a high impact on the channel performance and latency (due to retransmissions). To keep latency low and meet performance requirements, the probability of collisions (and therefore retransmissions) should be kept low by limiting the load per channel. This can be achieved by limiting the maximum traffic allowed in the channel, which translates to reducing the effective capacity allowed per frequency.

On the other hand, reuse-1 has the advantage of being a simpler approach for deployment when traffic levels are low, as there is no need for developing complex frequency plans and managing multiple frequency handoffs as an aircraft moves through multiple ground-stations' coverage areas.

To increase capacity as traffic grows in a reuse-1 system, additional frequencies may be added in a reuse-1 arrangement as well (multiplying the capacity by the number of frequencies), or instead the system can be migrated to a cellular-like architecture.

A **cellular architecture** has the potential to avoid hidden node pair occurrences by providing spatial isolation between reusing cells. Since frequencies are reused geographically far apart from each other they become more protected against

interference, i.e., it is less likely that a node will detect two transmitters on the same frequency at reasonable reception levels. The interference protection, which translates also into collision protection, makes the capacity per frequency higher than what would be achieved with the reuse-1 type of approach, for the same levels of delay requirements.

One of the challenges of achieving that higher capacity per frequency, however, is that increasing the reuse distance requires a careful frequency-assignment task. In practical scenarios where network topologies hardly resemble those of hexagonal grids used in the notional examples, the problem of frequency assignment becomes a complex one, often too complex to be adequately solved manually, and therefore it becomes necessary to use automated optimization algorithms.

The effect of overlap between coverage areas affects the probability of collision in both reuse-1 and cellular scenarios (at different proportions), so the overlap allowed in the network during coverage design has an important role in the dimensioning of its required spectrum resources later on. While some overlap is desired for coverage, reliability, and to ensure continuous (and redundant, if desired) coverage over the target geographical area, beyond that point overlap is an effect to be avoided in order to increase spectral efficiency.

Regional mixes of reuse-1 and cellular approaches are also possibilities, with multi-frequency playing a role only in high traffic density areas, for instance.

Data Comm Network Service (DCNS)

The infrastructure that will support data communications for Air Traffic Services (ATS) in the continental US is the Data Comm Network Services (DCNS). The same Communication Service Providers (CSPs) that will provide the ATS services are also expected to continue to provide Aeronautical Operational Control (AOC) services over the same air-ground network. Furthermore, since no additional aircraft radios are being planned to support CPDLC communications, ATS and AOC traffic are to be supported on shared links. In other words, from the point of view of one aircraft, communicating with a given ground station, AOC and ATS messages will flow on the same frequency. This scenario requires only one VDL-2 radio to handle both flows at the

aircraft and ground-station nodes. The concept of shared ATS + AOC is illustrated in Figure 3.
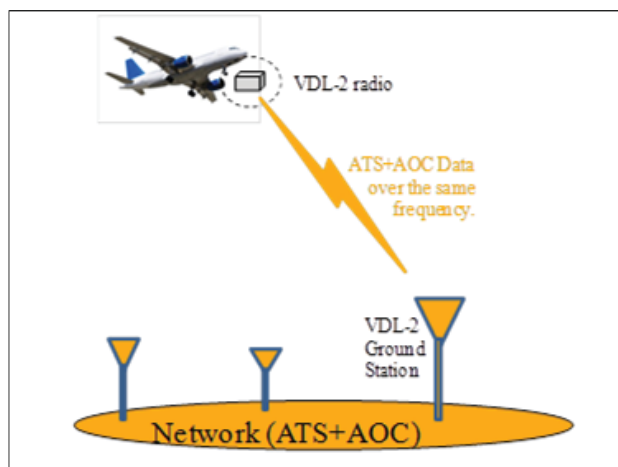
Figure 3. Concept of Shared ATS+AOC Network

Methodology Overview

The framework described in this paper comprises a combination of interacting processes, specialized software tools and various input data sets and models. These components are designed to take into account multiple aspects and dimensions of the air-ground network, and to include interdependencies and metrics that allow for solution optimization. Figure 4 shows the block diagram of its main components.

Inputs are highlighted in yellow. Moving clockwise from the top, the first set of orange blocks in the diagram comprises the tasks related to RF coverage design. These include using baseline network data (existing and candidate radio site locations), terrain database and link budget assumptions to develop an initial propagation assessment, iterating successively to add new sites as needed to fill coverage gaps and excluding sites as much as possible to minimize excessive overlap.

Following the RF coverage design, the next step is the evaluation of traffic demand per ground-station (GS), which is done using service-volume (SV) polygons, as well as aircraft distribution for a set of selected timeframes of interest, in order to evaluate the maximum count of aircraft in each polygon at any given time. The geographical airspace analysis was performed using MITRE's sectorEvaluator tool. Assumptions are used to estimate the correlation of carrier, model and airframe type to VDL-2 equipage

projections. Forecasts of growth are applied to obtain the total number of equipped aircraft per SV, per provider, per year.
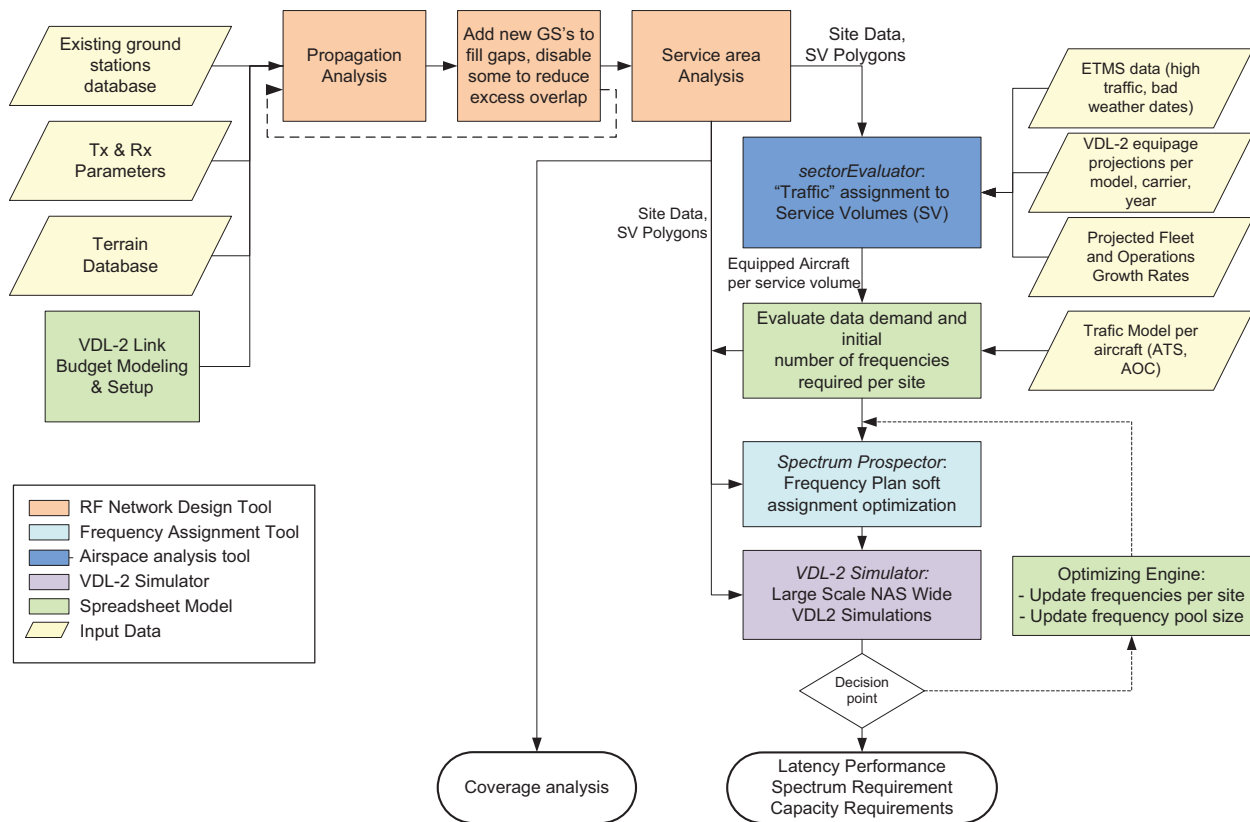
Figure 4. Block Diagram of MITRE's Network Dimensioning and Optimization Framework

Once the information on equipped aircraft per SV is available, the next task is to estimate the number of frequencies required per SV. We start this iterative process with initial estimates of expected maximum capacity per VDL-2 channel.

With the initial estimated number of frequencies needed per SV, we then use Spectrum Prospector to generate a frequency plan that most closely satisfies those requirements given a pool of VHF frequencies.

Once a frequency plan is generated, we simulate that particular solution using the NASA VDL-2 simulator. The output of the VDL-2 simulator is the latency (and related statistics) per SV, which allows an evaluation of whether all SVs have met the performance requirement, and by how much margin.

Next an iterative loop is executed with the generation of new plans by Spectrum Prospector, in an effort to reduce the pool of frequencies if overall system performance is satisfied with a specified margin, or to increase the pool of resources if overall system performance is not met. The iterative process converges when a frequency plan is reached that satisfies the latency requirements and no other plan with a smaller pool of resources has been found to satisfy the overall system performance.

The set of results that can be obtained with the framework include:

- Coverage analysis, including the capability to generate availability plots and reports per site or groups of sites.
- Assessment of the capacity needed in the network on a per GS basis, groups of GSs or system-wide.
- An estimate of the minimum spectrum required to meet a target performance threshold or (alternatively) the expected performance given an assumption of spectrum available.
- A suggested frequency plan that minimizes interference for the assumed spectrum pool

size. The frequency pool size can be given as input, or the result of the previous step.

- An estimate of throughput distribution (mean and other percentiles) achieved by the system in the frequency plan of choice.

RF Coverage Design

The RF design for coverage and overlap was performed using a commercial RF design tool, CelPlanner Suite.

The target of the RF coverage design is to postulate network solutions that satisfy the coverage requirements established for en route Data Comm operation. These requirements are to provide solid and reliable VDL-2 coverage for the full CONUS within the en route domain altitude ranges defined for Data Comm [6]. In our study, we assumed those altitudes to be between 16,000 feet and 60,000 feet Above Mean Sea Level (AMSL). Link-budget assumptions were considered as in [7].

For air-ground coverage the limiting scenario is at lower altitudes, due to horizon restrictions and the closer proximity to obstacles such as mountains, i.e., if coverage is achieved for low altitude, it is generally also available at the higher altitudes at the same location. Therefore, the coverage target can be translated into achieving solid redundant coverage within CONUS for the lowest en route altitudes.

A secondary objective in the coverage design, not derived from the system requirements but from the need to increase spectral efficiency, is the minimization of excessive overlap among cells. Excessive overlap increases the chance of hidden node transmissions in a CSMA system, resulting in collisions and therefore reducing spectral efficiency.

In order to achieve these objectives, we start by estimating the expected baseline coverage for each provider, using the information we have available about their current sites as potential initial candidates. Then we proceed by successively planning for new sites as needed to fill coverage gaps at the lowest en route altitudes, and excluding sites as much as possible to minimize excessive overlap.

Once this iterative process converges to a design that meets the requirements with as little overlap as acceptable, we proceed to generating outputs that will be used in the site data and Service Volume (SV) polygon definitions.

The SV polygons are defined based on the best-server analysis of each station, which identifies, at each location, the ground station that provides the strongest signal level. The approximately contiguous polygons corresponding to each station's best-server map at 60,000 ft. altitude (AMSL) were used as the basis for defining the horizontal cross section of the station's corresponding SV. Figure 5 illustrates those polygons for one hypothetical (postulated) network solution. In our studies, a ground station's SV was defined as the vertical projection from 60,000 ft. to 16,000 ft. of its 60,000 ft. best-server polygon.

Figure 5. Example Postulated Design: Service Volumes Polygons at 60,000 Feet Altitude
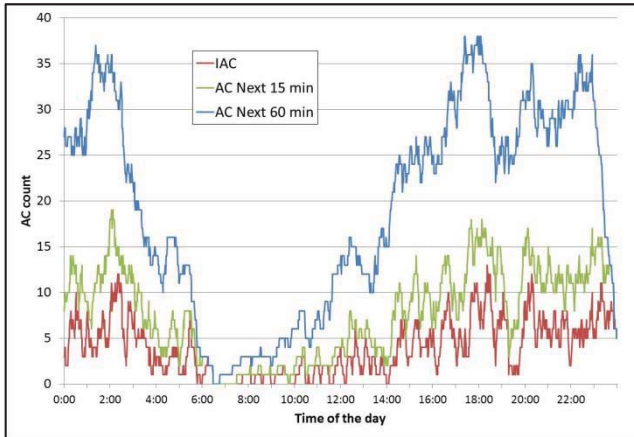
Aircraft Traffic Analysis

The next step in the framework is the generation of the aircraft geographical distributions relevant to the postulated VDL-2 network, leading to the estimation of the number of aircraft to be served by each SV on a per year basis.

The evaluation of aircraft traffic quantities per SV is performed using the SV polygons as input, resulting from the RF design. The analysis uses aircraft distributions for a set of selected dates of interest, to estimate the maximum count of AC occurring at each polygon at selected time ranges. The maximum count is used to ensure that the system is dimensioned with sufficient capacity for a worst-case traffic scenario.

The geographical airspace analysis was performed using the sectorEvaluator tool, developed by MITRE to perform airspace analysis over the NAS en route sector boundaries or user-defined polygons (such as SV polygons). The tool is capable of analyzing en route air traffic at selected dates, and

of keeping count of the number of AC located at any instant within each polygon (Instantaneous Aircraft Count, or IAC). Other outputs can also be generated, such as the temporal distributions of the number of AC per service volume, and accumulated counts within larger sliding time windows (e.g. 15-minute, 60-minute). Figure 6 illustrates the concept of instant, 15-min, and 60-min counts for a given SV. For this study, the IAC was considered as the relevant output.

Figure 6. Example AC Count Results for One Service Volume as an Output from the Airspace Analysis Tool (sectorEvaluator)

As input for the sectorEvaluator tool for the NAS traffic analysis we used data from FAA's Enhanced Traffic Management System (ETMS) dataset, including aircraft situation data and weather data. We selected a set of key dates for analysis, representing scenarios of heavy traffic and bad weather situations.

For each of the selected dates, data was analyzed from 0600 UTC on the specific day until 0559 UTC on the next day to generally cover full days based on local time. For each SV, the maximum IAC (Peak IAC, or PIAC) was obtained within that date. Then, the highest PIAC was selected for each SV among all dates of interest (high traffic, bad weather). That selected value was used as the AC count for that SV.

Growth in fleet was then applied to those dates to project traffic into the future.

With respect to the forecast of future VDL-2 equipage evolution, MITRE developed avionics equipage assumptions. To determine the fleet forecast, MITRE leveraged current fleet data from FAA WebOPPS and Registration Databases and combined these current fleets with the FAA Fleet

Forecast. MITRE then generated a fleet retirement profile to identify aircraft that will depart the fleet during the analysis period. Once the fleet of interest was determined, equipage profiles were created from MITRE's CAASD Avionics Information Repository for NextGen (CAIRN) aircraft capability database.

The sum of the peak estimates for all SVs in a scenario/year does not correspond to the expected aircraft volume at any given time, since each SV is estimated based on its own PIAC, which may be uncorrelated in time with the surrounding SVs. In other words, each SV is investigated at its worst case scenario (maximum) for a conservative approach.

The thematic map in Figure 7 is a notional illustration of the estimated VDL-2 equipped aircraft for a postulated network for a given year.
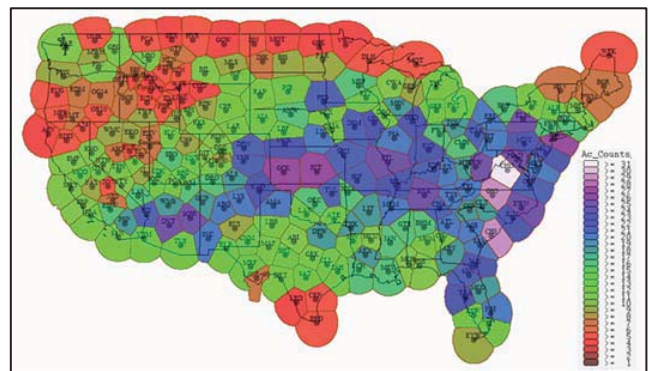
Figure 7. Estimated VDL-2 Equipped Aircraft (PIAC) per Postulated SV for a Given Year (Notional, Hypothetical Network)

Traffic Model

The message profile per aircraft was developed and updated considering three categories of traffic: Air Traffic Service (ATS), Airline Operations Center (AOC) and network traffic. The following is a definition of each type:

- **ATS Traffic:** Data Comm messages between controller and pilot including all the uplink and downlink messages to/from the aircraft.

- **AOC Traffic:** Operational messages between the AOC and the aircraft including engine performance and flight status, fuel status, gate and connecting flights status, NOTAMS, position and weather reports, and real time maintenance information.

- **Network Traffic:** Includes network layer connection establishment and keep alive, link establishment and link handoff associated with the VDL-2 sub-network.

The DCNS is designed to initially support Future Air Navigation (FANS) version 1/A+ and at a later date the Aeronautical Telecommunications Network (ATN) as well. For FANS, the Aircraft Communications Addressing and Reporting System (ACARS) over AVLC (AOA) is the service that encapsulates ACARS messages within an AVLC frame for transmission over the VDL-2 system. For ATN this function is provided by the Mobile Subnetwork Dependent Convergence Function (Mobile SNDCF).

For the scope of this discussion, we focus on traffic models for FANS scenarios as they correspond to the bulk of the traffic in the Segment 1 phase of the Data Comm program. The framework developed is also capable of modeling ATN traffic, which we did consider for evaluation of mixed traffic scenarios in some of our studies.

For **ATS traffic**, we generated traffic profiles for the uplink and downlink messages per aircraft to model services to be supported in each timeframe, guided by the services roadmap plan from the Data Comm program, and by use-case inputs that allowed us to model messages per transaction for each type of service. A summary of the traffic profile per aircraft for ATS services utilized in our analysis is shown in Table 1. This estimate is subject to changes as the program service offerings evolve and should not be considered in anyway as an official FAA estimate of traffic.

Table 1. Summary of the ATS Traffic Profile for FANS per Aircraft

Summary ATS Traffic Profile (per Aircraft)	UL	DL
Mean Message Arrival Rate (msg/sec)	0.0073	0.0087
Mean Message Size (bits/msg)	1351	350
Mean bits/second (bps)	9.81	3.06

The **AOC traffic** per SV was based on the Communications Operating Concepts and Requirements (COCR) v2.0 model [8]. A summary of the traffic profile considered in the simulations per aircraft for AOC is shown in Table 2.

Table 2. Summary of the AOC Traffic Profile for FANS per Aircraft

Summary AOC Traffic Profile (per Aircraft)	UL	DL
Mean Message Arrival Rate (msg/sec)	0.0047	0.0044
Mean Message Size (bits/msg)	5508	2288
Mean bits/second (bps)	25.94	9.98

In addition to the ATS and AOC traffic, network traffic was part of the traffic profiles used in the simulations. This category includes estimates of overhead traffic such as network layer connection and keep-alives. The total overhead traffic modeled is presented in Table 3, at the row "Network", along with a summary of the other traffic categories (ATS and AOC) mentioned above.

Table 3. Combined Traffic Profile for FANS per Aircraft (ATS, AOC, Network)

Traffic Category	Mean Demand per Aircraft (bps)		
	UL	DL	UL + DL
ATS Traffic	9.81	3.06	12.87
AOC Traffic	25.96	9.98	35.94
Network	0.62	0.45	1.07
Total	36.39	13.49	49.89

Frequency Plans

Once the peak number of VDL-2 equipped aircraft per SV and the expected data traffic per aircraft are estimated, we proceed to the next step in the framework, which is the frequency assignment.

We start by assuming that one frequency will be sufficient per ground station and generate frequency plans for multiple frequency pool sizes in order to verify iteratively by simulation which pool will suffice to meet the system latency requirements.

In order to generate a candidate frequency plan, the Spectrum Prospector tool takes into account the specific frequency requirements per SV (initially one for all SVs), the shape of the SVs, the aircraft altitude distribution associated with the SV, and the

requirements for interference protection desired for the system.

Spectrum Prospector generates frequency plans using "soft-assignment-rule" algorithms, which were created and incorporated into the tool [3] to support this framework and allow the optimization of frequency-assignment to VDL-2 systems. Given a set of performance requirements and a pool of resources, the soft-assignment algorithms construct a frequency plan that minimizes the impact of interference in the system, while not necessarily meeting strict signal-to-interference thresholds (which would characterize the hard-assignment). This is done because the VDL-2 systems can tolerate some interference at the expense of additional retransmissions and increased latency. Therefore, the algorithms work to minimize occurrence of co-channel interference, and optionally adjacent-channel interference if desired, with heavier penalties for areas with smaller Carrier-to-Interference (C/I) levels. Traffic distributions per altitude are also allowed to have different weights as input in the tool.

Three different meta-heuristic algorithms are supported by the soft-assignment capability of Spectrum Prospector: Simulated Annealing, Tabu Search, and genetic algorithms. From analyses performed for multiple scenarios and reported in [3], we found that Simulated Annealing and Tabu Search algorithms consistently provided better scores than genetic algorithms and were considered our preferred methods.

Using the preferred soft-assignment algorithms mentioned above, we generated frequency plans for the postulated network designs for different frequency pool sizes.

The plans with the best (smallest) scores provided by the tool for each pool size were used as input for the simulations to verify expected latency performance, as described in the next section.

Table 4 shows an example of results from Spectrum Prospector of different plans for each scenario which include: the system-wide score, the channel reuse factor (i.e., average number of times each channel was assigned), and the number of pairs of neighboring co-channel SVs.

Table 4. Frequency Plan Results from Spectrum Prospector for One Postulated CONUS Network Solutions, Assigning One Frequency per Ground Station

Frequency Pool Size	Best plan for pool size		
	Reuse Factor	Number of Neighbor SV Co-channel Pairs	Spectrum Prospector Score
2	115.0	251	10,653,554
3	76.7	143	5,925,427
4	57.5	86	3,639,975
5	46.0	39	2,330,331
6	38.3	19	1,467,095
7	32.9	10	930,879
8	28.8	6	582,705
9	25.6	5	368,917
10	23.0	1	227,121
12	19.2	0	79,881
15	15.3	0	11,274

As an example, the thematic map shown in Figure 8 depicts a 9-channel frequency plan for one of the postulated networks. Polygons painted the same color correspond to SVs reusing the same frequency.
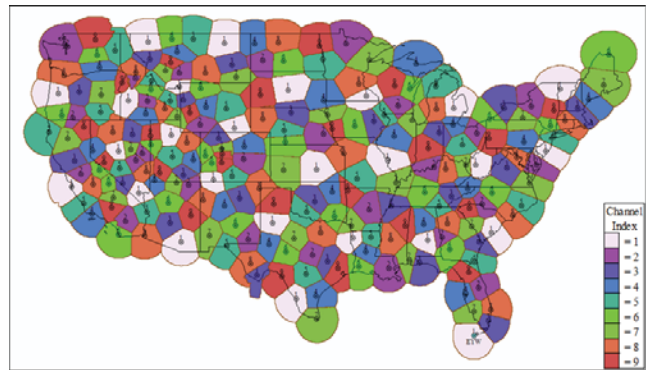

Figure 8. Example 9-Channel Plan

As part of the assessment of the different frequency plans, and to further our insight of acceptable C/I levels that would be likely to meet latency performance requirements, we also performed ground-to-air C/I predictions corresponding to the channel plans selected from Spectrum Prospector as candidates for a given scenario simulation.

An example of such prediction is shown in Figure 9 for different altitudes for the same frequency

plan shown above. The plot at the top shows the C/I for 18,000 feet, while the plot at the bottom shows the C/I for 50,000 feet AMSL. Both plots show some areas with C/I ≤ 5 dB, that correspond to neighboring ground stations having co-channel assignments. Those areas may still perform with acceptable latency in regions of low traffic.
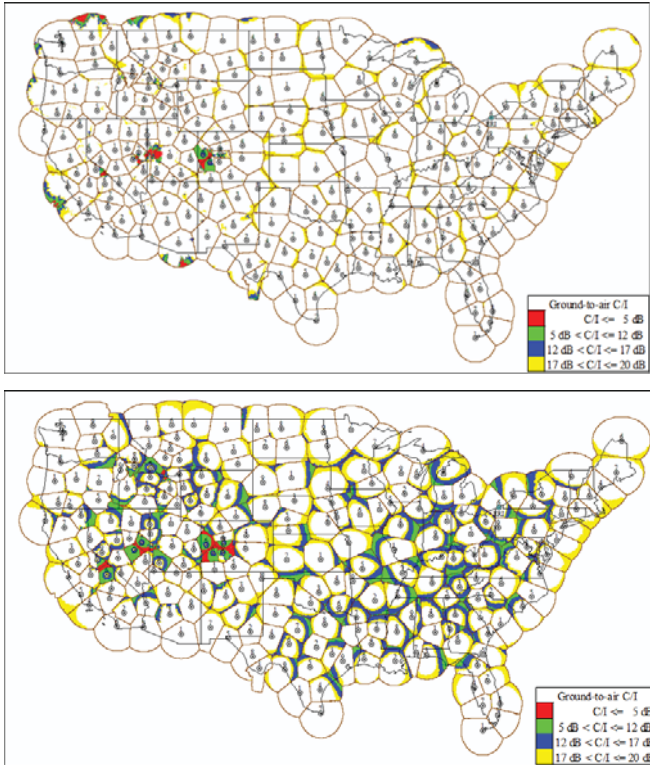
Figure 9. C/I map for 9-Channel Plan at FL180 (Top) and at FL500 (Bottom)

The following observations are made from the interference analysis presented in this section:

- The service volumes are susceptible to more interference with the increase in altitude as expected. This is due to the increased coverage overlap at higher altitudes for the same network topology (ground-station locations). If we observe any two maps for the same network with the same channel plan, the C/I results for the lower altitude map (at FL180) are generally much better than the C/I for the higher altitude (FL500). There are only a few isolated areas where C/I results may be better for higher altitudes than for lower altitudes. They occur at the edge of coverage of a few service volumes and are due to

terrain effects that impact the signal strength (C) at lower altitudes.
- For the same altitude, if we compare the C/I performance for any two frequency-plans for the same network, we observe better results for the channel plan with more resources (i.e., larger pool).

Latency Simulation

The best frequency-plans generated for each spectrum pool-size were simulated to identify the smallest pool-size that meets the latency requirements.

Although the C/I plots presented in the previous section provide a good insight on how C/I ratios are spread geographically, they do not show the complete picture of the expected performance of a system, because traffic is not taken into account. While a SV may have large areas with poor C/I (low values) in one scenario, those areas may be of very little significance in the SV performance due to low traffic, whereas another SV with a higher C/I ratio may end up having a poorer performance due to its location in a high traffic area. This occurs because VDL-2 can tolerate more interference at lower traffic where the probability of collision is reduced even in hidden-node prone conditions.

We utilized the VDL-2 simulation model developed by NASA that uses the OPNET platform. The VDL-2 simulator includes the detailed protocols (packets, algorithms, and timers) and the physical layer modeling (link budget, signal-to-noise ratio, bit error rate, line-of-sight analysis, and interference analysis) for the VDL-2 network. The simulator was improved in 2012 to support this framework and allow the execution of continental-US-wide scenarios.

Figure 10 shows the VDL-2 simulation process with the inputs and outputs used in the simulator. This process begins by using the frequency plan generated by Spectrum Prospector and the PIAC from the aircraft traffic analysis as inputs. For each SV, the aircraft are initially equally-distributed amongst all frequencies assigned to that SV. The next step is to create OPNET simulation models, using a separate mode for each frequency. These models are developed with the additional inputs of the GS locations, SV polygon shapes (to limit aircraft

locations), altitude profiles for aircraft in each SV polygons, the data traffic profile, and the equipage division between FANS and ATN.
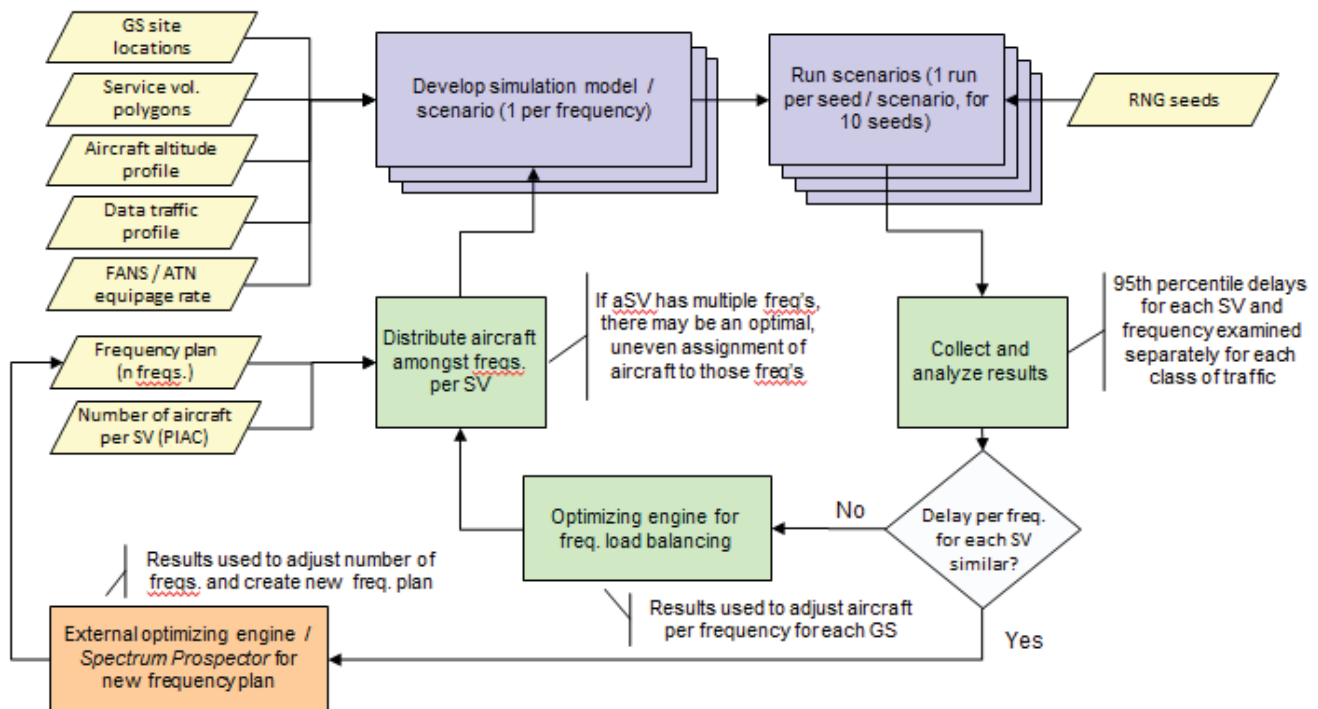
Figure 10. NASA VDL-2 Simulation Process

Simulations are then run for each simulation model, with 10 independent runs per frequency using different Random Number Generator (RNG) seeds. Once all simulations have completed, the results (per SV and frequency) are collected and analyzed.

For scenarios with more than one frequency per SV, the results for each SV are compared amongst its allocated frequencies to determine if, e.g., one frequency performed poorly while another performed well. This is an indication that the load distribution could be optimized (assign more aircraft to the better frequency and fewer to the poorer). After re-assignment, the simulation process is re-performed.

When the load is optimized (or when there is only a single frequency assigned per SV), the results from the simulation are used to drive the next iteration of the frequency plan generation process.

The VDL-2 simulation assumes that ATS traffic is given higher priority than AOC traffic in the ACARS protocol model. However, VDL-2 does not support priority, i.e., messages from all categories (ATS, AOC, network) are modeled to be treated with equal priority at that layer.

A performance delay requirement threshold is used as pass/fail at the simulator for ATS network traffic (95[th] and/or 99[th] percentiles). Pass/fail criteria are not applied to AOC traffic.

The latency threshold requirement refers to the air-ground delay portion. Ground network delays are excluded. Significant delays at the air-ground portion in CSMA result primarily from re-transmission due to collision, which is the event we aim at reducing by optimizing the network and therefore reducing the likelihood of co-channel interference at neighbor stations. Retransmissions also impact the operation of the ACARS protocol. Its message segmentation and stop-and-wait behavior, where one transmission must be acknowledged before the next can occur, causes additional delays when collisions occur.

In order to take into account the effects of aircraft distribution, the tool uses as input the expected number of equipped aircraft (PIAC values), as well as the relative distributions of aircraft with altitude for each SV.

The methodology to estimate spectrum required for each year consists of simulating selected frequency plans as starting points and moving to one

with a larger pool size if performance (latency) is not met in at least one SV in the system; or, alternatively, simulating a plan with a smaller pool size if performance is met. The final selected frequency plan for each year consists of the plan with the smallest pool size that meets performance for all SVs, for the demand expected that year. That smallest pool corresponds to the spectrum requirement for that year.

The simulation allows verification for each SV of the percentile of ATS messages meeting the latency requirement threshold, taking into account the combined effect of the ATS+AOC traffic loading of the channel at each location. Using simulation results and comparing them with the C/I maps generated for the same frequency plans, we observed the following trends:

- SVs showing co-channel C/I values above 17 dB (at boundaries with other SVs) consistently met the 95[th] latency percentile well above required threshold.
- SVs with areas of C/I values between 5dB and 17dB presented results that were marginally meeting the 95% latency requirement.
- SVs with areas of C/I values below 5 dB (typically occurring in co-channel neighbor situations) have mostly failed to meet the 95% latency requirement, unless the said SV had very low levels of traffic.

Summary and Observations

In this paper, we focused on presenting the methodology developed to assess multiple questions of interest that arise during the planning, design and optimization phases of implementing a VDL-2 network. The framework addresses questions such as how to handle traffic as it grows, how much capacity is achieved per channel for different architectures, and how to utilize spectrum efficiently while meeting performance requirements of coverage and latency.

Across the multiple studies performed in support of the Data Comm program, for multiple scenarios of traffic, phases, years, and architecture configurations, we gained further insight into key trade-offs related to designing VDL-2 networks. While numerical results obtained from those studies are not part of the

scope of this paper, some of those insights are shared in the following paragraphs.

The effect of hidden terminals in the en route environment directly impacts the ability to use spectrum efficiently. Although data link systems may begin operations with a single frequency, as traffic grows in the en route environment, it is expected that more frequencies will need to be added to the system (referred as multi-frequency operation) in order to continue to meet latency performance requirements.

Based on the scenarios simulated, we found that as the traffic continues to grow to the point where additional frequencies are needed, the cellular architecture performs with higher spectral efficiency than the multi-layer reuse-1 architecture. We expect that a larger pool of frequencies is needed if multiple layers of reuse-1 are implemented than if a cellular frequency plan is adopted. Additionally, for the same pool of frequencies, we expect that a higher capacity can be achieved per SV with a cellular architecture than with reuse-1. We have not attempted to assess an exact threshold (pool size) for when that transition (to cellular being more efficient) occurs.

In the multi-frequency scenarios, an operational advantage of multiple reuse-1 layers is the ability to keep aircraft on the same frequency while they cross multiple service volumes (i.e. less frequency switching at the aircraft). A disadvantage of this approach is that multiple frequencies are needed earlier at each ground station than what would be required in a cellular approach. Besides the additional equipment needed, this calls for more sophisticated radio resources management mechanisms such as load balance algorithms [5] in order to maximize capacity per frequency by appropriately assigning the new connections to frequencies that have better local channel conditions (less load on the RF channel). This may also require better mechanisms of estimating uplink RF channel utilization (from the air perspective) in real-time and/or utilizing pre-processed network topology data if those mechanisms are not available.

On the other hand, the cellular scenario has the advantage of allowing the system to continue to operate with a single frequency per ground station for a prolonged time. For all scenarios we simulated, we did not come across any scenario where a cellular architecture with one channel per ground station was not able to handle all traffic demand for each ground

station. This is because the channel capacity is greatly increased through the protection added by a large reuse distance.

We observed that the value of the latency requirement has a strong impact on the spectral efficiency of VDL-2 and, therefore, on spectrum requirements for a given traffic demand. For instance, exploring sensitivity to latency, we noticed that a relaxation of the latency requirement by approximately 150% (increase in the latency time threshold) improved the channel capacity by about 40% (i.e. 40% more traffic could be handled per VDL-2 channel). Although the impact is not expected to be constant for different systems or to vary linearly over different ranges, this variation gives an idea of the significance of the latency requirement in the dimensioning process.

In the en route domain the uplink is clearly the weakest link with respect to latency performance, i.e., for any proposed solutions we simulated, as traffic grows, the uplink fails first (triggering additional spectrum requirements). This is consistent with the fact that hidden node problems affect more heavily the uplink receptions, combined with the fact that the traffic models used were also heavier in that direction.

Since the channel is shared between AOC and ATS, the usage of AOC communications over VDL-2 directly impacts performance of ATS communications. While prioritization mechanisms can possibly be implemented by the CSPs at the ground stations for uplink transmissions, given the random access nature of p-CSMA systems, more sophisticated mechanisms would need to be in place to implement priority in the downlink direction. This need may be further aggravated by the use of the new Media Independent Aircraft Messaging (MIAM) protocol that may potentially have large impacts on the AOC traffic characteristics and therefore on ATS performance [9]. The accurate forecasting of AOC traffic growth is one of the challenges in dimensioning spectrum requirements for shared ATS and AOC networks.

As spectral efficiency continues to be an important priority for the FAA, and for the aviation community in general, we expect that the framework described here can be helpful in the assessment and optimization of achievable scenarios and configuration architectures, with the ultimate intent of utilizing spectrum in the most efficient manner.

References

[1] L.Z. Ribeiro, "Spectrum Dimensioning and Optimization for Air-Ground Data Communications", The MITRE Corporation, McLean, VA, Tech. Rep. MTR110116, March 2011.

[2] L.Z. Ribeiro et al., "Application of Soft Frequency-Assignment Rules to FAA's Data Communications Service", The MITRE Corporation, McLean, VA, Tech. Rep. MTR120474, Sep. 2012.

[3] Box, Frank et al, "Minimizing Self-Interference in an Aeronautical Data Communications Network" in Proceedings of the 2012 Integrated Communications Navigation and Surveillance (ICNS) Conference, Herndon, VA, 2012.

[4] L.Z. Ribeiro et al, "En Route Data Communications Demand and Spectrum Analysis - Study Update", The MITRE Corporation, McLean, VA, Tech. Rep. MTR130479, Sep. 2013.

[5] L.C. Monticone, R. Snow and L. Ribeiro, "Load-Balancing Algorithms for Multi-frequency Data Communications", The MITRE Corporation, McLean, VA, Tech. Rep. MTR130451, Sep. 2013.

[6] FAA Document, "Data Communications Network Service Description (DCNSD) – Requirements Applicable to Segment 1, and 2", June 2013.

[7] RTCA DO-224B "Signal-in-Space Minimum Aviation System Performance Standards (MASPS) for Advanced VHF Digital Communications Including Compatibility with August 3, 2005.

[8] Communications Operating Concept and Requirements (COCR) for the Future Radio System, Eurocontrol/FAA, Version 2.0.

[9] S. Deric, B. Nguyen, and S. Bretmersky, "A Study on Media Independent Aircraft Messaging (MIAM)" in Proceedings of the 2011 Integrated Communications Navigation and Surveillance (ICNS) Conference, Herndon, VA, 2011.

Disclaimer

Acquisition Management System Clause 3.5-13, Rights In Data-General, Alt. III and Alt. IV (Oct. 1996).

The contents of this document reflect the views of the author and The MITRE Corporation and do not necessarily reflect the views of the Federal Aviation Administration (FAA) or the Department of Transportation (DOT). Neither the FAA nor the DOT makes any warranty or guarantee, expressed or implied, concerning the content of accuracy of these views.

All network illustration examples in this document refer to hypothetical, postulated networks, not specific for a given Communication Service Provider (CSP).

Email Addresses

lribeiro@mitre.org
lmontico@mitre.org
rsnow@mitre.org
fbox@mitre.org
rafael.d.apaza@nasa.gov
steven.c.bretmersky@nasa.gov