Modeling and Simulation of an Aeronautical CSMA Subnetwork

Brian T. Hung
The MITRE Corporation
Center for Advanced Aviation System Development (CAASD)
McLean, VA 22102
bhung@mitre.org

Abstract
This paper presents results of recent modeling and simulation of an aeronautical subnetwork using OPNET. This subnetwork uses multiple 25-kilohertz (kHz) channels in the 118-137 megahertz (MHz) Very High Frequency (VHF) band to provide data communications services between aircraft and ground facilities. It uses the Carrier Sense Multiple Access (CSMA) protocol to control access to the 25-kHz channel between the aircraft and a ground station. The main purpose of these simulations is to determine the performance of this subnetwork under various Aeronautical Operational Communications (AOC) and Air Traffic Services (ATS) application message traffic and aircraft loadings.

1. Introduction
This paper presents the results of recent modeling and simulation of the International Civil Aviation Organization (ICAO) VHF Digital Link (VDL) Mode 2 subnetwork using OPNET. ICAO developed VDL Mode 2 standards to provide data communications services for AOC. The simulations were performed to determine the suitability of using this subnetwork to provide air/ground data communications services for ATS. The VDL Mode 2 subnetwork is one of the air/ground subnetworks in the Aeronautical Telecommunication Network (ATN) [1], which is a global network to provide data communications services for civil aviation. A range of application message traffic and numbers of aircraft were used in the simulations to determine the performance of the VDL Mode 2 subnetwork under different loads.

The U.S. air traffic control is divided into four domains: airport, terminal, en route, and oceanic. In the ATS network environment, which includes the VDL Mode 2 subnetwork and the ATN, certain message delay requirements are imposed on each portion of the network. For the VDL Mode 2 subnetwork portion of the network, the 95th percentile end-to-end delay requirement is 3 seconds in the terminal domain. An important goal of the simulation is to determine the number of aircraft a VDL Mode 2 ground station can support while meeting this requirement.

2. VDL Mode 2 Subnetwork
The VDL Mode 2 subnetwork provides data communications services between the aircraft and ground entities. Its standards are defined in the ICAO VDL Mode 2 Standards and Recommended Practices (SARPs) [2]. It uses a network of ground stations as shown in Figure 1 to provide the desired airspace coverage utilizing part of the VHF band allocated for aeronautical communications. This band is divided into multiple independent 25-kHz channels. It is expected that VDL Mode 2 will have approximately 20 channels allocated in the 136-137 MHz band. A CSMA protocol controls the media access between the aircraft and the ground station for each 25-kHz channel allocated for VDL Mode 2. As specified in the VDL Mode 2 SARPs, this subnetwork consists of a physical layer, a data link layer, and a subnetwork layer of the International Standards Organization (ISO) Open Systems Interconnection (OSI) 7-layer model.

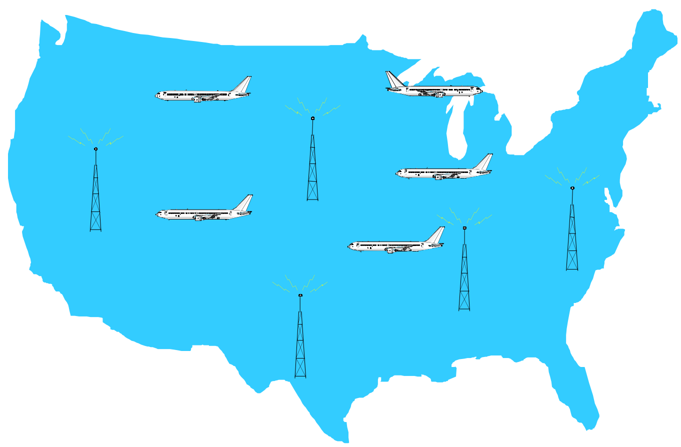

Figure 1. VDL Mode 2 Subnetwork

2.1 Physical Layer
Using a differentially encoded 8-phase shift keying (D8PSK) modulation, VDL Mode 2 operates at a burst rate of 31.5 kilobits per second (kbps) for each 25-kHz channel. The physical layer generates a 108-bit training sequence for each frame transmitted.

2.2 Data Link Layer
The data link layer consists of two sublayers and a management entity.

The Media Access Control (MAC) sublayer uses a *p*-persistence CSMA protocol to control access to a shared 25-kHz channel between the aircraft and the ground stations.

This work was performed under the Federal Aviation Administration contract number DTFA01-01-C-0001. The contents of this material reflect the views of the author. Neither the FAA nor the Department of Transportation, makes any warranty or guarantee, or promise, expressed or implied, concerning the content or accuracy of the views expressed herein.

The Data Link Service (DLS) sublayer provides error detection, error recovery, and address identification of frames. It supports unicast and broadcast addresses, and uses a variant of the High-level Data Link Control (HDLC) protocol to provide sequencing and acknowledgement of frames.

The Link Management Entity (LME) provides link management and release services between the local DLS and the remote DLS.

2.3 Subnetwork Layer

The subnetwork layer uses the ISO 8208 protocol to provide Switched Virtual Circuit (SVC) service between the aircraft and the ground station.

3. Technical Foundation for CSMA and Previous Work

A review of previous analysis of the *p*-persistence CSMA protocol and VDL Mode 2 subnetwork is presented in Appendix A.

An ARINC report [3] presents the results of a service provider's modeling and simulation of VDL Mode 2 subnetwork. This model includes all the subnetwork protocols as well as the transport protocol to account for part of the ATN overhead.

4. Application Message Traffic Model

The simulations presented in this report use a terminal domain application traffic model summarized in Table 1 [4]. This terminal domain model includes projected ATS and AOC message traffic for 2015. The ATS traffic includes pilot and controller communications, Traffic Flow Management (TFM), Flight Information Services (FIS) planning services, aircraft originated meteorological observations, advanced Air Traffic Management (ATM), and delivery of route deviation warning services. These services and the AOC traffic are combined into four applications in each ground and aircraft End System (ES). The message size, arrival rate, arrival distribution, and priority describe these applications. These traffic characteristics are assumed to be equal for all aircraft. The combined uplink and downlink air/ground data link traffic for each aircraft in Table 1 is defined as one Load Factor (LF).

Table 1. Terminal Domain Application Message Traffic Model

Application Message Distribution	Priority	Uplink (From Ground Station)		Downlink (From Aircraft)	
		Average Message rate	Average Size in bits	Average Message rate	Average Size in bits
Exponential Inter-arrival With Poisson Message size (Point-to-point)	High	0.0179	127.9	0.0238	107.6
	Medium	0.00083	800	0.00083	100
	Low	0.001	2400	0.0017	2400
Constant	Low	0.0166	3325	0.0033	1760

Notes:
1. Message rates are in number of messages per second per aircraft
2. Each message is acknowledged at Data Link Sublayer, except broadcast messages
3. Uplink broadcast messages are represented by constant uplink messages
4. Downlink periodic meteorological observations are represented by constant downlink messages

5. VDL Mode 2 Subnetwork OPNET Models

The VDL Mode 2 OPNET models used in the simulation were based on the models developed previously by CAASD. Some changes were made to these models to make them more closely reflect the ICAO VDL Mode 2 SARPs.

5.1 Network Models

Five network models were created. Each of these models contains one ground station with 40, 60, 80, 100, or 120 aircraft. The number of aircraft in each model remains constant during each simulation. To simplify the models and as a rough approximation of a real system, aircraft are assigned a predetermined fixed position at varying distances from the ground station during each simulation.

5.2 Ground Station and Aircraft Node Models

The modules implemented in the ground station node model are shown in Figure 2. The aircraft node model is similar to the ground station node model. Each ground station or aircraft contains four application message generator modules, a sink process module, a subnetwork layer process module, a DLS sublayer process module, a MAC sublayer process module, a transmitter module, a receiver module, and an antenna module.

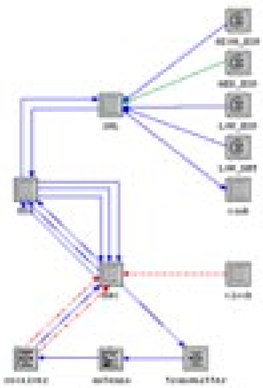

Figure 2. VDL Mode 2 Ground Station Node Model

5.3 Process Models

The four applications generate high-priority exponential, medium-priority exponential, low-priority exponential, and low-priority deterministic application messages from the ground station or aircraft; each generates application messages with the message inter-arrival distributions and message length distributions specified in the traffic model in Table 1.

The sink accepts the frames received from the subnetwork layer.

The aircraft and ground subnetwork layers do not currently implement the ISO 8208 protocol as required in the SARPs. They receive messages from the applications and send packets to the DLS sublayer. They receive frames from the DLS sublayer and record either uplink or downlink end-to-end delays and forward them to the sink. The ground subnetwork layer calculates and prints statistics including offered load (*G*) and throughput (*S*), which are defined in Appendix A.

The DLS sublayer does not currently implement the HDLC protocol as required in the SARPs. The DLS sublayer receives packets from the subnetwork layer, and sends a new frame for each packet received to the MAC sublayer. For each frame requiring an acknowledgement, the DLS sublayer starts a retransmission timer, T1 (see Section 6.3). If the timer expires, the DLS sublayer retransmits the same frame for up to N2 (see Section 6.3) times. The DLS sublayer receives frames from the MAC sublayer and forwards them to the Subnetwork Layer (SNL) sublayer. For each frame received that requires an acknowledgement, the DLS sublayer sends an acknowledgment frame.

The MAC sublayer receives frames from the DLS sublayer. A transmitting entity, aircraft or ground station, checks for channel status before proceeding. If the channel is busy, the transmitting entity backs off for a fixed interval of time, TM1 (see Section 6.4), before rechecking the channel status. If the channel is not busy, the transmitting entity attempts to transmit with a probability of p and backs off with a probability of 1-p. Figure 3 shows the MAC sublayer process model. It forwards frames received from the receiver and addressed to this MAC sublayer to the DLS sublayer.

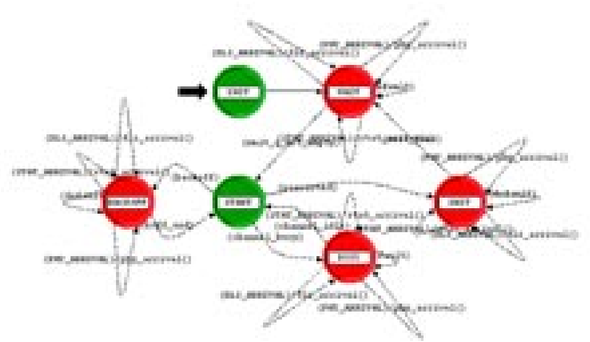

Figure 3. VDL Mode 2 MAC Sublayer Process Model

The transmitter receives frames from the MAC sublayer and transmits them over a 25-kHz channel. The receiver receives frames over a 25-kHz channel and sends them to the MAC sublayer. The transmitter and receiver operate at 118 MHz with a 31.5 kbps channel bit rate.

6. VDL Mode 2 Simulation

The simulations are performed with different numbers of aircraft and values of Adjusted Load Factor (ALF), which is defined in Section 6.2, to study the effect of changing the values of these parameters. Default values of parameters are used in the DLS and MAC sublayers.

The performance metrics include end-to-end application message delay and system throughput.

6.1 Number of Aircraft

Models with 40, 60, 80, 100, and 120aircraft were used.

6.2 Application Message Traffic

In the simulations performed for this paper, the traffic is made up mostly of application message traffic. The only protocol overhead modeled in the simulation, however, is the DLS acknowledgement, which is transmitted for every frame received. This overhead is less than 20 percent of the total traffic in the simulations. In a previous simulation study of the VDL Mode 3 subnetwork in the ATN environment [5], it was found that the protocol overhead was about 50 percent of the total traffic. In order to account for the overhead of the ATN protocols, we define an ALF as equal to twice the size of LF. Therefore, 1 ALF is equivalent to 2 LF. In the simulations, ranges of values of ALF between 0.5 and 4 were used.

6.3 Data Link Service Sublayer Parameters

The DLS sublayer used the maximum number of transmissions, N2, counter value of 6. It used an adaptive delay before retransmission timer, T1, defined by an equation in the SARPs. The value of T1 is between 2.2 seconds (s) and a very large number.

6.4 MAC Sublayer Parameters

The MAC sublayer used the default inter-access delay, TM1, value of 4.5 milliseconds (ms). The maximum number of access attempts, M1, has a default value of 135, but it is not implemented in the models.

It was discovered during the simulations with ALF 1 that if both the ground station and the aircraft used the same persistence value, *p*, the uplink delays will eventually increase without bound (see Figure 4). This is because the total amounts of uplink and downlink message traffic are about the same. However, because there are more aircraft than ground stations, the ground station has more messages to send while competing for access to the radio channel with the aircraft on an equal basis. This caused the uplink queue in the ground station to grow with time, and thus caused the long delays. By increasing the persistence of the ground station by an order of magnitude, the ground station was allowed to send messages more often than the aircraft and the issue was resolved. In the simulations, the aircraft used the default *p*-persistence value of 0.05078 while the ground station used a value of 0.609.

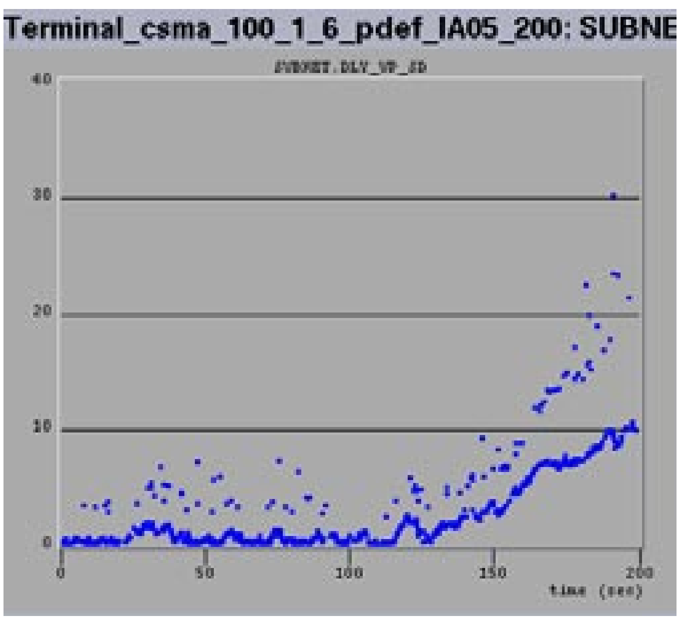

Figure 4. Uplink Delay with Ground Station and Aircraft Using the Same p Persistence

6.5 Physical Layer Parameters

The ground station transmitter power is 25 Watts (W) and the aircraft transmitter power is 20 W.

7. Simulation Results

All simulations are run with 3600 seconds of simulation time. Data is collected starting at 500 seconds.

7.1 Throughput

7.1.1 A Comparison of Theoretical and Simulated Throughputs without the Training Sequence

This section compares the theoretical throughput with the simulated throughput without the training sequence because the throughput analysis in Appendix A does not include the training sequence.

Equation A1 in Appendix A calculates the theoretical throughput as a function of various parameters including the offered load and *a*, which is the ratio of propagation time and message duration *T*. Figure 5 plots two theoretical curves from this equation. Both curves assume the default p value of 0.05, and *T*, of 5.5 ms[1]. The first curve assumes that the distance between the stations are 100 nautical miles (nmi) which gives *a* = 0.112. The second curve assumes a distance of 200 nmi which gives *a* = 0.224. The curves show the classic behavior of a CSMA system. For small values of offered load, the throughput increases in proportion with the offered load. Eventually, the throughput reaches a maximum and decreases with increased offered load. The capacity of a system is defined as the maximum throughput. Figure 5 shows that the value of *a* can be a significant factor in the system performance. It shows that the difference in the theoretical performance between *a* = 0.112 and *a* = 0.224 is large. For *a* = 0.112, the capacity is about 0.51, and it occurs around offered load of 5. For *a* = 0.224, the capacity is about 0.39, and it occurs around offered load of 5.6.

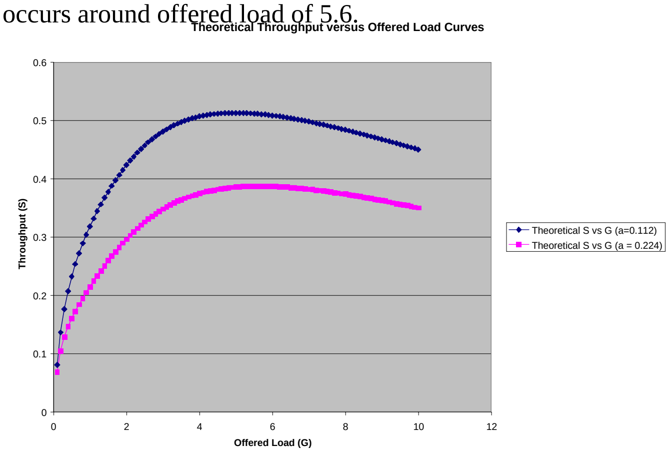

Figure 5. Theoretical Throughput versus Offered Load Curves

Simulations were performed without the training sequence for comparison with the theoretical results. Figure 6 shows the combined simulation and theoretical results. The simulated curves for 100, 120, 130, and 150 aircraft are on the left side of the figure. The shapes of these curves are fairly different from that of the theoretical curves. This difference in the shapes is probably due to the more detailed modeling of the VDL Mode 2 subnetwork in the simulation compared with the theoretical analysis. The simulation curves reach the capacity

[1] 5.5 ms was average frame transmission duration collected from the simulations without the training sequence.

at much lower offered loads compared with the theoretical curves. The capacity of the simulation curves are very close from that predicted by the theoretical curve for a = 0.224.

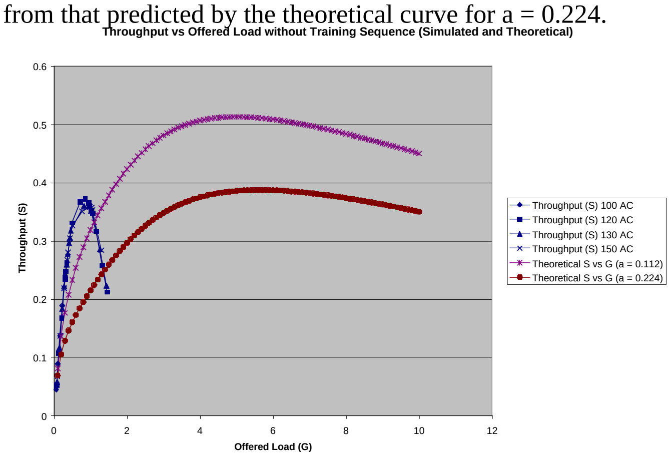

Figure 6. Simulated and Theoretical Throughput versus Offered Load Curves without the Training Sequence

7.1.2 Simulated Throughput with the Training Sequence

Figure 7 presents the throughput versus offered load curves for 40, 60, 80, 100, and 120 aircraft with the training sequence. In these simulations, the training sequence is considered an overhead and is not counted as payload when the throughput is calculated. Again these curves are very similar. As seen from Figure 7, the capacity of the VDL Mode 2 subnetwork when the training sequence is included is less than 0.3.

Under the anticipated aircraft and message traffic loads, the system will be operating in the low offered load and throughput region. For 80 aircraft with ALF 1, the offered load is about 0.15 and the throughput is about 0.07. This corresponds to a stable operating region.

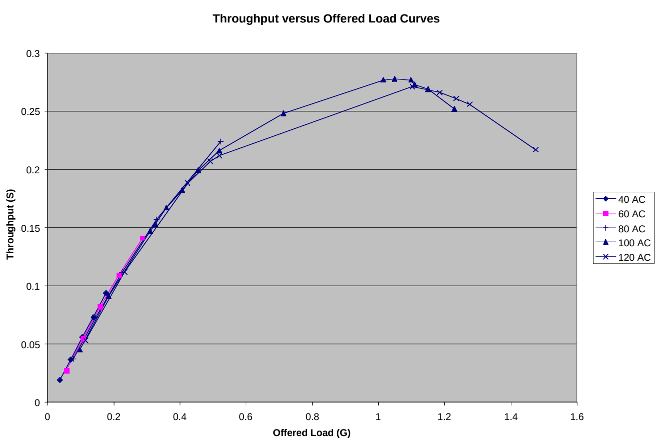

Figure 7. Simulated Throughput versus Offered Load Curves with Training Sequence

7.2 95th Percentile Delay

7.2.1 A Comparison of Analytical and Simulated 95th Percentile Delay without the Training Sequence

This section compares the analytical 95th percentile delay with the simulated delays without the training sequence because the delay analysis in Appendix A does not include the training sequence.

Equation A3 in Appendix A calculates the analytical 95th percentile delay as a function of several parameters including a, T, and d. We assume that a = 0.224, T = 5.5 ms, and d = 10, i.e., the average retransmission delay is ten times the packet transmission length T. Figure 8 displays the simulated high-priority uplink and analytical 95th percentile delays without training sequence versus throughput curves. It shows that the simulated and analytical curves are in fairly good agreement until they reach the maximum throughput. Thereafter, the simulated throughputs fall sharply with increased load.

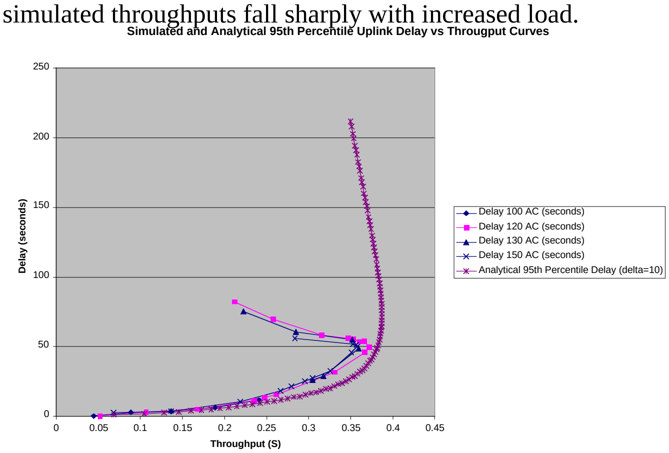

Figure 8. Simulated High-Priority Uplink and Analytical 95th Percentile Delays without Training Sequence versus Throughput Curves

Figure 9 displays the simulated high-priority downlink and analytical 95th percentile delays without training sequence versus throughput curves. Like Figure 8, it shows that the simulated and analytical curves are in fairly good agreement until they reach the maximum throughput. It is believed that the reason the simulated delays do not increase as much as the theoretical curves after reaching the maximum throughput is because the packets are discarded once they have reached the maximum retry count of N2 times.

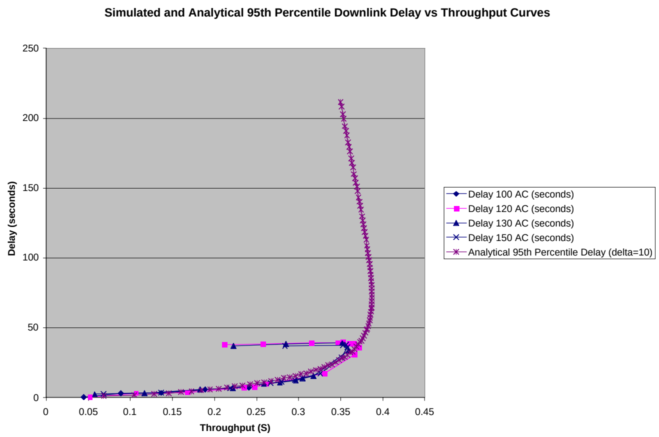

Figure 9. Simulated High-Priority Downlink and Analytical 95th Percentile Delays without Training Sequence versus Throughput Curves

7.2.2 Simulated 95th Percentile Delays with the Training Sequence

Figure 10 shows the simulated high-priority 95th percentile uplink delay versus throughput curves for the VDL Mode 2 subnetwork with the training sequence.

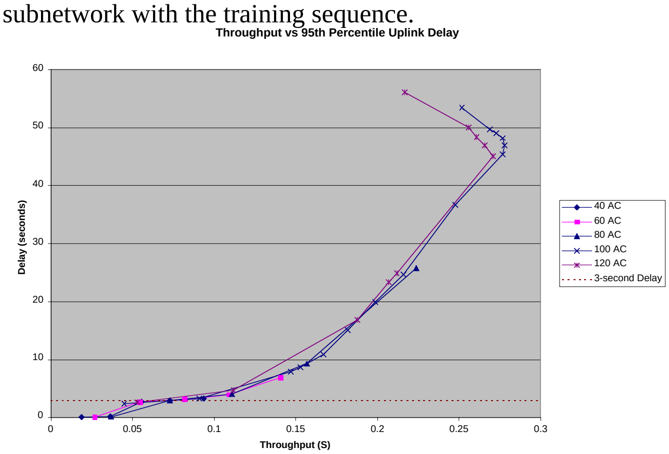

Figure 10. Simulated High-Priority 95th Percentile Uplink Delay versus Throughput Curves

Figure 11 shows the simulated high-priority 95th percentile downlink delay for the VDL Mode 2 subnetwork with the training sequence.

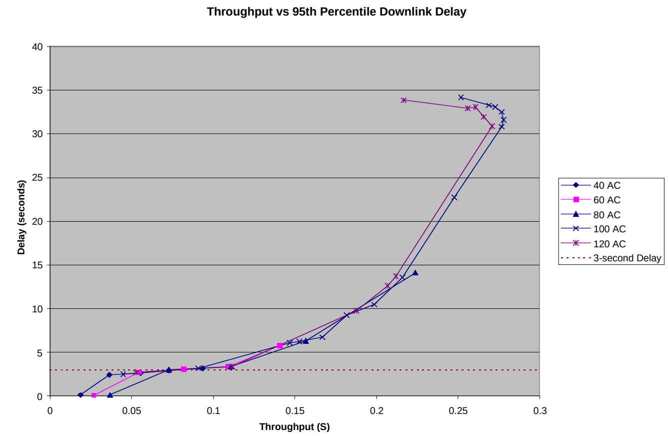

Figure 11. Simulated High-Priority 95th Percentile Downlink Delay versus Throughput Curves

7.3 Mean Delay

7.3.1 A Comparison of Theoretical and Simulated Mean Delays without the Training Sequence

This section compares the theoretical mean delay with the simulated delays without the training sequence because the delay analysis in Appendix A does not include the training sequence.

Equation A2 in Appendix A calculates the theoretical mean delay as a function of several parameters including S, G, a, T, and d. Again, we assume that $a = 0.224$, $T = 5.5$ ms, and $d = 10$. Figure 12 displays the simulated high-priority uplink and theoretical mean delays without the training sequence versus throughput curves. It showed that the simulated and analytical curves are in fairly good agreement for small values of throughput.

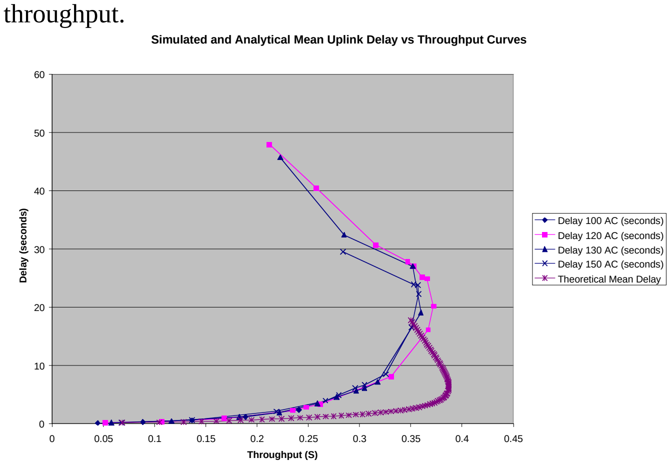

Figure 12. Simulated High-Priority Uplink and Theoretical Mean Delays without Training Sequence versus Throughput Curves

Figure 13 displays the simulated high-priority downlink and theoretical mean delays without training sequence versus throughput curves. Like Figure 10, it showed that the simulated and analytical curves are in fairly good agreement at low throughput.

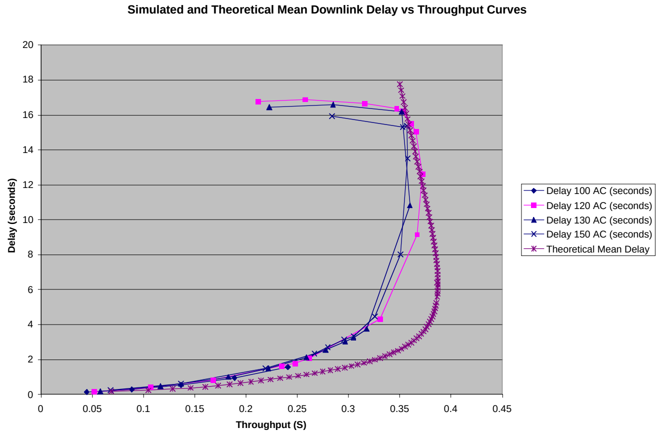

Figure 13. Simulated High-Priority Downlink and Theoretical Mean Delays without Training Sequence versus Throughput Curves

7.3.2 Simulated Mean Delays with the Training Sequence

Figure 14 shows the simulated high-priority uplink mean delay versus throughput curves for the VDL Mode 2 subnetwork with the training sequence.

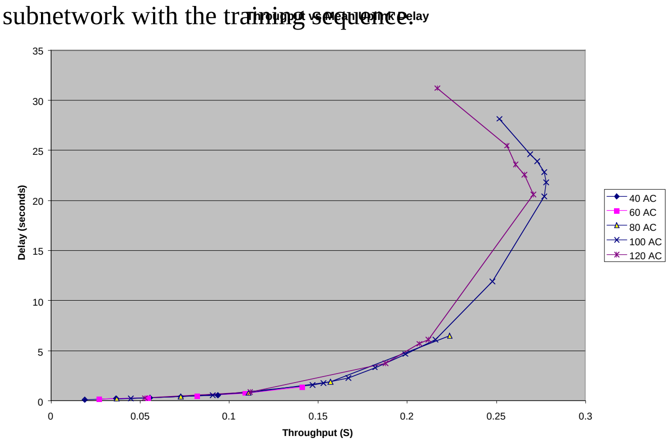

Figure 14. Simulated High-Priority Uplink Mean Delay versus Throughput Curves

Figure 15 shows the simulated high-priority downlink mean delay versus throughput curves for the VDL Mode 2 subnetwork with the training sequence.

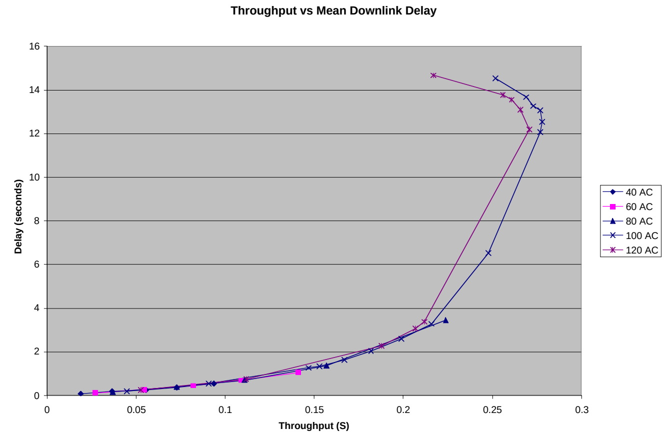

Figure 15. Simulated High-Priority Downlink Mean Delay versus Throughput Curves

7.4 95th Percentile Delay versus Number of Aircraft

Figure 16 with ALF 1 corresponds to the case of LF 1 with ATN overhead. This figure showed that the number of aircraft that can be supported while meeting the 95th percentile end-to-end delay of 3 seconds is approximately 80 aircraft.

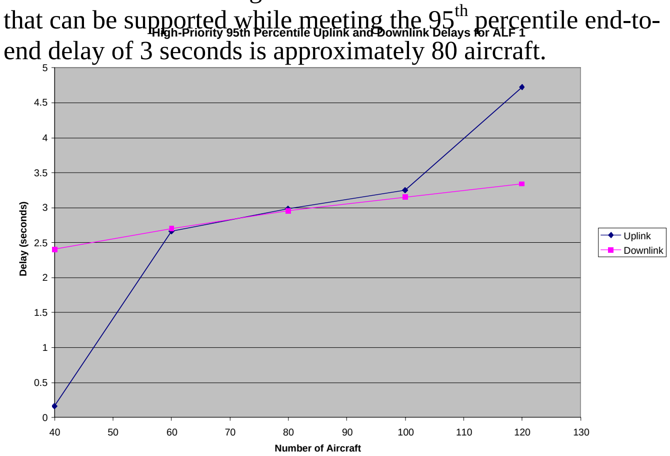

Figure 16. High-Priority 95th Percentile Uplink and Downlink Delays for ALF 1

8. Observations and Conclusions

This paper presents the results of recent simulations on the VDL Mode 2 subnetwork using OPNET. The following are the conclusions:

1. There is fairly good agreement between the analytical and simulation results in lightly loaded conditions. The simulation results tend to be worse compared to the analytical results at higher loads because the analytical approach makes many simplifying assumptions.
2. The end-to-end delays are dependent on the number of aircraft in the system and the traffic load.
3. The MAC sublayer may have difficulty operating as defined in the SARPs if the ground station and the aircraft use the same p persistence. Choice of the p persistence value of the ground station should be carefully evaluated.

References

[1] ICAO, "Manual of Technical Provisions for the Aeronautical Telecommunication Network (ATN)", First Edition, Doc 9705-AN/956, International Civil Aviation Organization, 1998

[2] ICAO Annex 10 Aeronautical Telecommunications, Volume III Part I Digital Data Communication Systems.

[3] ARINC Incorporated Aviation Systems Division and Federal Aviation Administration Technical Center, "Cooperative Research and Development Agreement No. 96-CRDA-0090 VDL Mode 2 Simulation Report", November 1997.

[4] The MITRE Corporation, "A Proposed VDLT Traffic Model for Capacity Simulation", AMCP/WG-D/5 WP-16, January 1996.

[5] Hung, B., and Hong, Y., "Modeling and Simulation of the VDL Mode 3 Subnetwork in the ATN Environment", 19th Digital Avionics Systems Conference, October 2000.

[6] Kleinrock L., and Tobagi, F.A., "Packet Switching in Radio Channels: Part I-Carrier Sense Multiple-Access Modes and Their Throughput-Delay Characteristics", IEEE Transactions on Communications, Vol. COM-23, No. 12, pp. 1400-1416, December 1975.

[7] D'Amours C., and Mazur, B., "Throughput and Delay Characteristics of AVPAC Radio", Second Meeting, Aeronautical Mobile Communications Panel, International Civil Aviation Organization, November 1992.

Appendix A
Analysis of CSMA and VDL Mode 2 Subnetwork

The CSMA protocol with its many variations including *p*-persistence has been analyzed in [6]. This protocol reduces the frequency of message collisions over a radio channel by having the ready stations sense the carrier from the other stations before attempting to transmit a message. The *p*-persistence protocol works as follows. A ready station senses the channel before transmission. If the channel is sensed idle, then it transmits with probability *p*, or with probability *1-p*, it postpones the transmission by τ seconds. At the end of each postponement, the station repeats the same process if the channel is still idle. Otherwise, the channel is busy, and the station schedules the retransmission of the packet according to a random retransmission delay with a mean of $\bar{X}$. The average delay $\bar{X}$ is large compared to *T*. We define δ (delta) as the ratio of $\bar{X}$ over *T*, i.e., $\delta = \bar{X}/T$. If a ready terminal senses a channel busy, it continues to wait until the channel becomes idle and then operates as above.

Throughput, *S*, and delay, *D*, are normally the two performance characteristics of interest in a communications system. Throughput is defined as the successfully received traffic bit rate normalized with respect to the channel bit rate. It measures the utilization of the channel bandwidth, which is always less than or equal to 1. Throughput is a function of many parameters including the offered load, *G*, which is defined as the total amount of traffic transmitted by the MAC sublayer. This includes newly arrived traffic as well as retransmissions. The offered load is normalized with respect to the channel bit rate and can be larger than 1.

Several assumptions were made in [6] to facilitate the analysis of p-persistence CSMA. They are:

1. Zero delay in switching from transmit to receive modes and vice versa;
2. The time to detect the carrier is negligible;
3. All messages (packets) are of constant length *T* in duration;
4. Channels are noiseless;
5. Noncapture (i.e., the overlap of any fraction of two packets results in destructive interference and both packets must be retransmitted);
6. Propagation delay to be identical for all source-destination pairs;
7. A separate channel is used for acknowledgement; and
8. The system is considered to be slotted into mini time slots, τ. The duration of τ is equal to the propagation delay.

From [3], for given offered load *G* and a given value of persistence *p*, the throughput *S* is given by:

$$S(G,p,a) = \frac{(1-e^{-aG})[P_s'\pi_0 + P_s(1-\pi_0)]}{(1-e^{-aG})[a\bar{\tau}'\pi_0 + a\bar{\tau}(1-\pi_0)+1+a]+a\pi_0} \quad \text{(A1)}$$

where:

a is the ratio of propagation time over message duration, i.e., $a = \tau/T$;

P_s' is the probability of a successful transmission over the first Transmission Period (TP) during a busy period;

P_s is the probability of a successful transmission over a TP other than the first TP during a busy period;

$\bar{\tau}'$ is the average period of wasted mini slots between the first and second TPs;

$\bar{\tau}$ is the average period of wasted mini slots between the TPs besides the first period; and

π_0 is the probability that no stations are waiting to transmit.

Another paper [7] analyzed the throughput and delay characteristics of the p-persistence CSMA of the VDL Mode 2 subnetwork. The analysis is based on [6] with a few modifications. It developed a mean delay D_m given by:

$$D_m = \left(\frac{G}{S}-1\right)\left[1+2a+\delta+\frac{1-p}{p}(a+P_b\delta)\right]T + \left[1+a+\frac{1-p}{p}(a+P_b\delta)\right]T \quad \text{(A2)}$$

where P_b is the probability of a deferring station resensing the channel to be busy. It is equal to:

$$P_b = \frac{1}{1 + \bar{t}(1 - \pi_0) + \bar{t}' \pi_0}$$

Paper [7] also developed an equation for 95th percentile delay D_{95} given by:

$$D_{95} = (M - 1)[1 + 2a + \delta + W(a + P_b\delta)]T + [1 + a + W(a + P_b\delta)]T \quad \text{(A3)}$$

M is found by solving the equation

$$\sum_{N=1}^{M} p(N) = \sqrt{0.95}$$, and

W is found by solving the equation

$$\sum_{Z=0}^{W} p(Z) = \sqrt{0.95}$$

p(N) is given by

$$p(N) = (1 - \frac{S}{G})^{N-1} \frac{S}{G}$$, and

p(Z) is given by

$$p(Z) = (1 - p)^{Z} p$$

The assumptions in [6] were made to facilitate analysis. In a real system, such as the VDL Mode 2 subnetwork, most of these assumptions do not hold, and resulting network performance will differ from that predicted by the analysis. In addition, the *p*-persistence CSMA protocol in the VDL Mode 2 subnetwork does not perform retransmission. The DLS sublayer performs such a function. Also, the acknowledgements are transmitted in the same 25-kHz channel and not in a separate channel as assumed in assumption 8. Therefore, the power of simulations is the ability to model the details of the subnetwork, which are very difficult to analyze analytically.