NASA/TM–20250009042

ACERO TCL 1 Airspace Management Data Sharing Performance Analysis

M. Gilbert Wu
Alexey A. Munishkin
Shawn R. Wolfe

NASA Ames Research Center, Moffett Field, CA, 94035

Xavier Elder

HX5, LLC, NASA Glenn Research Center, Cleveland, OH, 44135

December 2025

NASA STI Program . . . in Profile

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

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

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

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

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

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

NASA/TM–20250009042

ACERO TCL 1 Airspace Management Data Sharing Performance Analysis

M. Gilbert Wu
Alexey A. Munishkin
Shawn R. Wolfe

NASA Ames Research Center, Moffett Field, CA, 94035

Xavier Elder

HX5, LLC, NASA Glenn Research Center, Cleveland, OH, 44135

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California, 94035-1000

December 2025

Abstract

1 Introduction

NASA's Advanced Capabilities for Emergency Response Operations (ACERO) project aims to develop, integrate, demonstrate, and transition airspace management technologies, involving NASA and industry partners, to wildland firefighting operational use. ACERO held its first Technical Capabilities Level (TCL) field evaluation in California in late 2024 and early 2025, the first of a series and thus named TCL 1. These field evaluations are instrumental in validating key aspects of ACERO's prototype airspace management system for wildland aerial operations.

ACERO's concept of an airspace management system focuses on enabling a digital"common operating picture" that provides participating users access to aerial operations' intent, traffic data, as well as supplementary information such as weather, airspace restrictions, terrains, and ground operations in the airspace of interest. This concept allows geographically dispersed teams to share data in real time, and should enable the Air Traffic Group Supervisor (ATGS) to manage and deconflict aerial operations without relying on visual observation from an aircraft, ultimately doing so on the ground. ACERO's prototype system is called Portable Airspace Management System (PAMS). Initial application of this prototype system focuses on drone operations, although the concept can be readily applied to manned flights as well. Detailed design and considerations of PAMS is described in [1] for interested readers.

Wildland fires progress rapidly and are not entirely predictable. Real-time data sharing adds tremendous value to the effectiveness and safety of aerial operations and plays a crucial role in the enabling of a real-time common operating picture. Communication technologies for sharing digital data among geographically dispersed teams exist but are not all affordable or feasible. Terrain poses a challenge to connectivity, too. TCL 1 field evaluation leverages the use of an ad hoc radio mesh network with an aerial communication relay node, which are expected to overcome most of the restrictions terrains impose on ground-based platforms. This ad hoc radio mesh network can work alone or connect with cellular or satellite networks.

This technical report examines the effectiveness of the radio mesh network in TCL 1 field evaluation in terms of it transmission success rate and latency. Section 2 describes the TCL 1 field evaluation. Section 3 presents the data analysis approach. Section 4 presents calculation results and findings. Section 5 concludes the findings and discusses implications.

2 TCL 1 Field Evaluation

ACERO's TCL 1 field evaluation aims to accomplish the following technical objectives:

- Introduce airspace management technologies into wildland fire operations
- Validate information sharing and coordination capabilities
- Validate use of aircraft as a communication relay

Table 1 summarizes the two phases of TCL 1. This report examines primarily data from Phase 2, but also does a limited scope analysis of selected Phase 1 data.

Table 1. TCL 1 field evaluation summary

Description \ Phase	**1**	**2**
Location	La Selva beach, California	Salinas Valley, California
Dates	11/13/2024 to 11/24/2024	3/19/2025 to 3/27/2025
Description	system checkout and proof-of-concept	data collection in an operationally relevant environment
# of flight sessions	10	13
# of PAMS cases	3	4 or 5
# of radio nodes	4	5

The baseline scenario for TCL 1, Incident Monitoring, demonstrates how information sharing aids to establish a common operating picture. This scenario consists of three drone operations described as follows:

1. A simulated Temporary Flight Restriction (TFR) is entered into the system as the bounding airspace for all aerial operations.
2. The first drone is launched to act as an aerial communication relay station.
3. Once the aerial communication relay is established, two more drones are launched from the second and third sites, to collect image data and perform simulated aerial ignition, respectively.
4. The second and third drones complete their operations and return to base.
5. The first drone returns to base.

The first drone sometimes stayed in the air in between flight sessions so as to save the time needed for landing and take-off. Fig. 1 illustrates this baseline scenario. Several variations of the baseline scenarios were also conducted to exercise the TFR-violation alert, strategic deconfliction, and conformance monitoring features of WFSS. In several sessions of Phase 2, an additional optionally-piloted aircraft joined the PAMS and shared its operational data.

Each drone operation team used a collocated PAMS case (also shown in Fig. 1) to submit, update, activate, and close an operationl intent (OI). Telemetry data were sent from each drone operation team's ground control station (GCS) to the PAMS case. TFR, OI, and telemetry data were all broadcast through the radio mesh network by each PAMS case. The PAMS case is expected to be used by the drone operators once it becomes mature enough. For TCL 1, PAMS researchers operated the PAMS case right next to drone operators, who would review and comment on PAMS display when they were free.

Silvus Streamcaster 4200 E+ radios (see Figure 2 below), operating at S-band, served as the primary radio mesh network to be evaluated. NASA's partner, Overwatch Aero (OWA), uses the radios to stream live video and data from their drone down to users connected through this radio network. DoodleLabs radios, a back-up

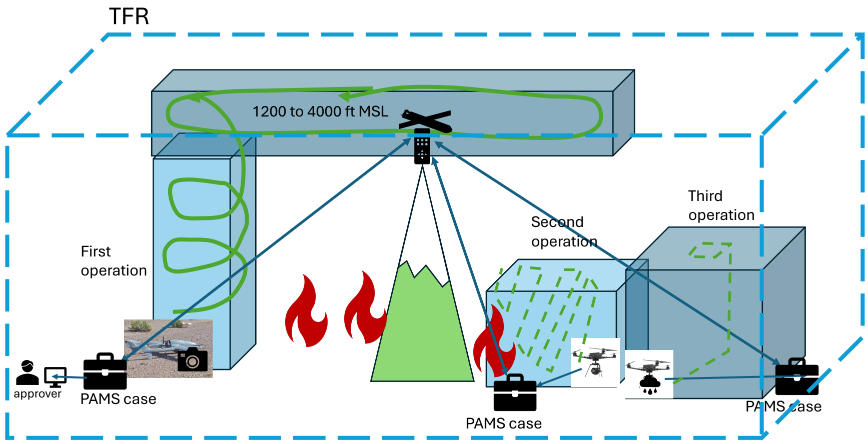

Figure 1. TCL 1 baseline scenario: Incident Monitoring

radio option operating at S-band, were also used for integration testing and several of the flight sessions of TCL 1. Silvus radios are powerful radios that have been in operational use by OWA and are expected to cover broader range and provide higher bandwidth than DoodleLabs radios. Given the testing environment in TCL 1, the Silvus radios had a data bandwidth of about 20 Megabytes per second (MBps) and the DoodleLabs radios had about 10 MBps bandwidth.

Figure 3 shows how the radio mesh network connected to the PAMS cases on the ground. The top diagram is for flight sessions 8, 9, 10, and 11, that used Silvus radios. PAMS cases 1, 2, 3, and 7 were each directly connected to a radio node. The approver PAMS case, 8, was connected via an ethernet cable to Case 1 (between their network switches). The bottom diagram is for flight sessions 5, 12, and 14, that used DoodleLabs radios. PAMS cases 1, 2, and 3 were each directly connected to a radio node. The approver PAMS case, 7 for these sessions, was connected via an ethernet cable to Case 1 (between their network switches). For these sessions, Case 7 was also directly connected to a radio node that connected with the radio mesh network. Note the approver PAMS case was not there in Phase 1.

Figure 2. Silvus radio

Figure 4 shows the components inside a PAMS case. At the top-center is a tablet computer hosting WFSS's core airspace management capabilities. Additional software on the tablet computer processes ADS-B and fire line data. The ADS-B receiver injects air traffic information to the system. A portable battery shown to the left of the case is also available. Each PAMS case is directly connected to a radio (not shown in this figure) via an ethernet cable from the network switch.

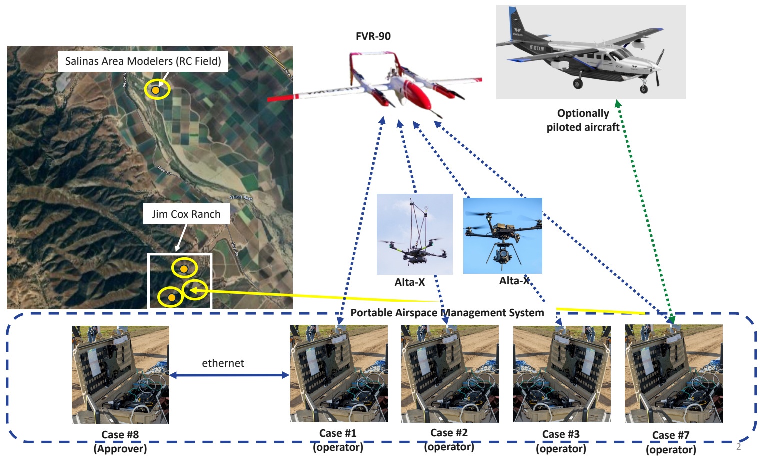

(a) Silvus Radio

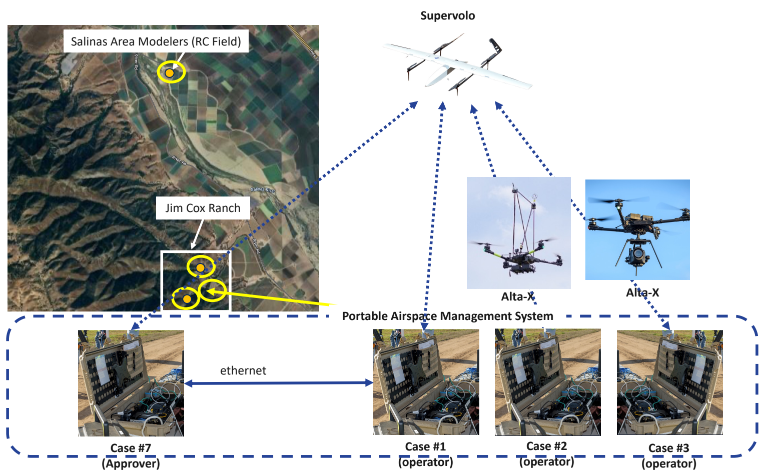

(b) DoodleLabs Radio

Figure 3. Radio connectivity, (top) Silvus radio sessions, (bottom) DoodleLabs radio sessions

3 Network Performance Analysis Approach

3.1 Data Shared across the Network

Table 2 shows the two categories of data shared across the network, Wildland Fire Service Supplier (WFSS) and Data Processing Tool (DPT). These messages were encoded as JavaScript Object Notation (JSON) strings.

The following data were used for the analysis:

1. WireShark recordings
2. WFSS log files
3. UTM Gateway (connecting WFSS to the network) log files
4. DPT log files
5. Screen capture video files (used to ascertain events)

The following user-level performance metrics were calculated:

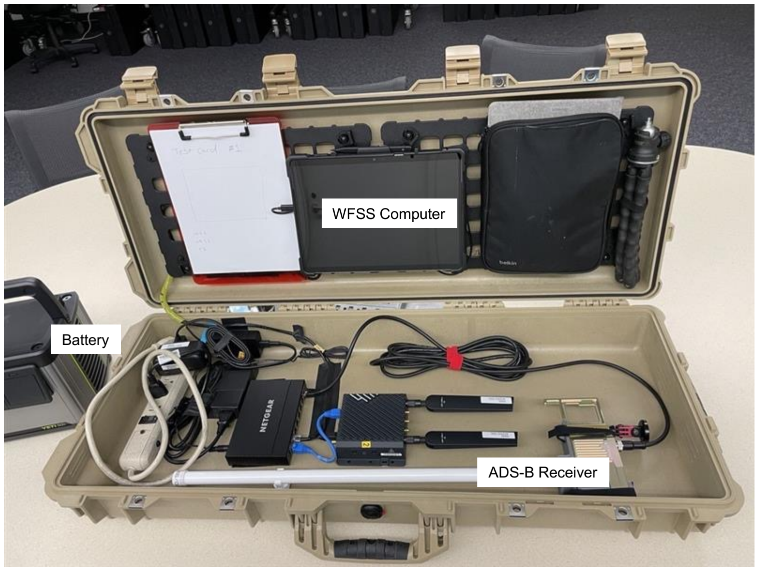

Figure 4. PAMS Case Components

Table 2. Types of data shared across the network

Category	WFSS	DPT
Type	TFR (1/60 Hz + event) OI (1/60 Hz + event) Telemetry (1 or 5 Hz) Heartbeat (0.1 Hz) Message Protocol: UDP	Fire line data (1 Hz*) ADS-B data (1 Hz*) Message Protocol: REST API

1. Update success rate for TFR, OI, and fire line, respectively – if a TFR update is successfully transmitted from, for example, PAMS Case 1 to PAMS Case 2, updating PAMS Case 2's TFR status, it is counted as one successful update event. There could be several identical TFR update messages that failed to be transmitted across the network before a successful one. As long as one of such messages made it to its target, it is counted as one successful update event.

2. Latency for TFR, OI, and fire line, respectively – a latency is calculated for each and every successful update event between the sent time and received time (for a fire line update the latency is the time between the sent request and the received response). These latencies are put into a list. The median of these values per message type is reported.

The following network-level metrics were calculated:

1. WFSS Data transmission success rate per session, TR, is defined as

$$\text{TR} = \frac{\text{total data received per session}}{\text{total data sent per session}} \tag{1}$$

Depending on the level at which the measurement was done, there are

(a) TR_L from the WFSS log files, and

(b) TR_N at WireShark (Network/Link) layer.

Note only data transmitted via the radio mesh network were included. For example, data transmission between Case 1 and Case 8 for Silvus radio sessions were excluded because Case 1 and Case 8 were connected via an ethernet cable. Silimarly, data transmissions between Case 1 and Case 7 for DoodleLabs Radio sessions were excluded. Additional analysis breaks down TR_L into values for each case pair for certain flight sessions.

2. Latency (LT_{WFSS}) – LT measures the time between when a WFSS telemetry message is sent and when a message is received.

3. DPT Response Success rate (RR) – the percentage of DPT requests that resulted in successful responses.

4. Total Response Time (ΔT_R) – For DPT messages, ΔT_R measures the time between when a message is requested and when a response is received, both by the requestor of the message. Therefore LT accounts for the total time for the request to travel to the receiver, the time taken by the receiver to process the request and reply, and the time for the response to travel to the requestor.

3.2 Flight Sessions Analyzed

Table 3 shows seven Phase 2's flight sessions whose data were analyzed. The other flight sessions were omitted due to either flight issues or data collection issues. Two flight sessions from Phase 1, S-9 with Silvus radios and S-10 with DoodleLabs radios, were also analyzed to some extent in comparison.

Table 3. Flight session analyzed

Session #	**Cases #**	**Approver Case**	**Radio**	**Missing Data**
5	1, 2, 3, 7	7	DoodleLabs	Case 1 video
8	1, 2, 3, 7, 8	8	Silvus	
9	1, 2, 3, 7, 8	8	Silvus	Case 7
10	1, 2, 3, 7, 8	8	Silvus	Case 7; Case 1 video
11	1, 2, 3, 7, 8	8	Silvus	Case 7
12	1, 2, 3, 7	7	DoodleLabs	Case 2 WireShark; Case 7
14	1, 2, 3, 7	7	DoodleLabs	Case 7
S-9	4, 5, 6	–	Silvus	
S-10	1, 2, 3	–	DoodleLabs	

3.3 Data Issues

During this analysis, some data were found to be missing (see the rightmost column of Table 3). These missing data reduced the scope of the analysis.

In addition to missing data, two additional issues were identified.

- Clock discrepancy – in several flight sessions, a computer's clock was found to lag behind, leading to incorrect timestamps for the recorded messages. These offsets were estimated and corrected in the analysis.
- Duplicate WFSS messages – due to the side effect of an issue fix, UTM Gateway sends two identical WFSS messages to each recipient, one via IPv4 and one via IPv6.
- Misconfiguration – Session 8's Case 8 had DPT misconfigured. As a result, Case 8 did not communicate DPT messages with the rest of the system.

It was also found that messages sent by UTM Gateway to the network could be lost before reaching the WireShark level. Figure 5 illustrates the message duplication and loss. Suppose two WFSS instances are connected via a radio mesh network (there were more instances in the field evaluation). It appeared that, if connectivity was down between a case pair, the messages sent by UTM Gateway would fail to reach the network/link layer recorded by WireShark. TR was computed at both the WireShark level or the WFSS level, and their values were different.

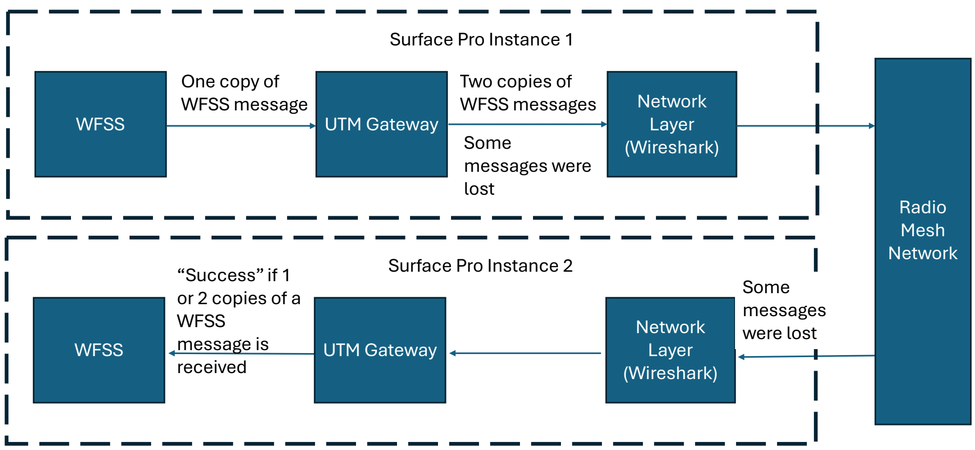

Figure 5. WFSS data duplication and losses

4 Analysis Results

Table 4 shows the TFR update success rate and latencies, aggregated by radio type. TFR update had a perfect success rate, although several update took a considerable amount of time that impacted the timeliness of TFR sharing.

Table 4. TFR update success rate and latency (median followed by maximum in parentheses)

\ Phase	1		2	
Radio \ Metric	**Success %**	**Latency (sec)**	**Success %**	**Latency (sec)**
Silvus	100	0.2 (0.24)	100	1.2 (720.9)
DoodleLabs	100	120.1 (239.5)	100	0.5 (540.6)
All	100	0.5 (239.5)	100	0.9 (720.9)

Table 5 shows the OI update success rate and latencies, aggregated by radio type. The OI update success rate is regarded high. Interestingly, Silvus radios did worse in Phase 2 sessions compared to DoodleLabs radios. Several short-lived OI states (e.g., drone transitioning back-and-forth between Activated and Non-Conforming) appeared to account for most of the failures.

Table 5. OI update success rate and latency (median followed by maximum in parenthese)

\ Phase	**1**		**2**	
Radio \ Metric	**Success %**	**Latency (sec)**	**Success %**	**Latency (sec)**
Silvus	100	0.2 (0.5)	84	10.5 (1628.9)
DoodleLabs	88	0.3 (180.5)	91	0.8 (1192.4)
All	92	0.2 (180.5)	87	1.7 (1628.9)

Table 6 shows the fire line update success rate and latencies, aggregated by radio type. What's consistent between TFR, OI, and fire line results is the Silvus radio sessions performed worse than the DoodleLabs sessions in Phase 2. The maximum fire line update time was less than those for the TFR and OI update. This was likely due to more frequent attempts (every 10 seconds for the fire line vs every 60 seconds for TFR and OI).

Table 6. Fire line update success rate and latency (median followed by maximum in parenthese)

\ Phase	**1**		**2**	
Radio \ Metric	**Success %**	**Latency (sec)**	**Success %**	**Latency (sec)**
Silvus	Not calculated		70	4.4 (290)
DoodleLabs			100	2.1 (200)
All			77	3.6 (290)

Table 7 shows the network-level metrics per flight session. The most obvious observation of these metrics is the four Silvus radio sessions resulted in lower TRs than the three DoodleLabs radio sessions. This is contrary to the team's expectation, because Silvus radios are known to provide longer range and larger bandwidth. Due to the design of these flight sessions, no single environmental variable could be attributed to this performance degradation. During the flight sessions in which Silvus radios were utilized, the drone carrying the aerial relay circled at higher altitudes and farther away from the ground sites. In all flight sessions, four ground radios were in the network, although there were five PAMS cases in the Silvus radio sessions and four PAMS cases in the DoodleLabs radio sessions.

The following analysis were done on Phase 2 flight sessions only.

Data rate across the network is likely to impact the performance of the network if the data rate gets close to the network's bandwidth. Only Session 8 had all data collected and thus was further analyzed. Figure 6 shows TR_N as function of the session time. The data rates increased significantly after minute 780 of the day after a fourth drone joined the system and started broadcasting its telemetry data at 5 Hz

Table 7. Performance metrics per session

Session #	Radio	WFSS			DPT		All Data
		TR_L	LT	TR_N	RR	ΔT_R (sec)	TR
5	DoodleLabs	54%		68%	94%	1.0	68%
8	Silvus	28%		37%	32%	7.1	59%
9	Silvus	27%		41%	26%	4.2	55%
10	Silvus	35%	$\leq$ 1 sec	44%	58%	7.7	51%
11	Silvus	13%		41%	52%	10.5	52%
12	DoodleLabs	30%		72%	93%	3.0	73%
14	DoodleLabs	36%		69%	92%	0.6	76%
S9	Silvus			99%	~100%	0.4	
S10	DoodleLabs			89%	67%	0.7	

(the other drones broadcast their telemetry at 1 Hz). TR_N appeared to have reached a minimum of 25% but bounced back to about 45% around minute 800. There does not appear to be a strong correlation between TR_N and the data rate.

Figure 6. TR_N as a function of time for Session 8

Figure 7 shows the same type of plot for Session 10. Even though recording of Case 7's sent data were not available, other cases received Case 7's data throughout the duration of the time window shown in this figure. It can be assumed Case 7 added a constant amount of data rate to the system. The results seems to indicate a reverse correlation between TR_N and the average data rate by sent requests.

Another interesting finding is the asymmetric TR_N for case pairs as shown in Fig. 8. The values of TR_N from Case 1 to Case 2 is higher than that from Case 2 to Case 1. Similarly, the values of TR_N from Case 1 to Case 3 is higher than that from Case 3 to Case 1. The reason of this asymmetric performance is unknown.

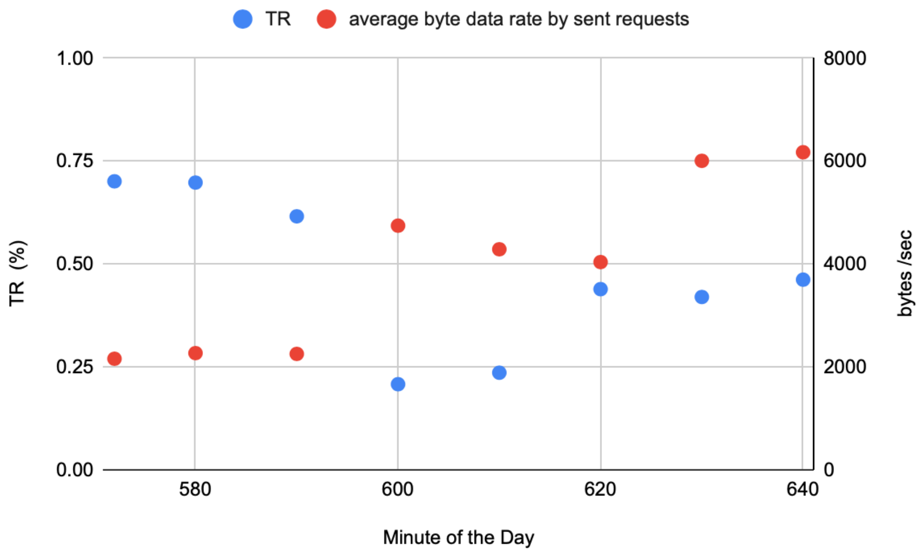

Figure 7. TR_N as a function of time for Session 10. Note that Case 7's sent data were not recorded and therefore not included.

To better understand how frequently the network connectivity went on and off, Figure 9 shows the connectivity per Case pair for Session 8. The heartbeat messages received by each case serve as an indicator of connectivity. The heartbeat message was sent once every ten seconds from each Case to each of the other Cases. Note the dots representing the heartbeat messages appear to be a line for the 8 to 1 and 1 to 8 Case pairs just because the size of the dots being too big to separate. Several observations can be drawn from this plot:

1. The connectivity between Case 1 and Case 8 was robust as expected, since they were connected via an ethernet cable.
2. Connectivity was not asymetric between a Case pair, as previos plots already illustrated.
3. Connectivity could be "off" for several minutes (for example, Case 1 and Case 2 around minute 783)
4. The connectivity "off" time appeared to be correlated across Case pairs. However, some connectivity remained when most Case pairs were off. For example, around minute 791, when most Case pairs were missing their connectivity, Case 7 was able to send heartbeat messages to Case 2 and Case 3.

These findings highlight the challenge of building a network for effective sharing of digital airspace management information in a wildland fire operational environment. Nonetheless, instability of network in an austere wildland firefighting environment is a very likely, sometimes unavoidable, situation. PAMS design should take this into account and be as robust as it can be against such challenges.

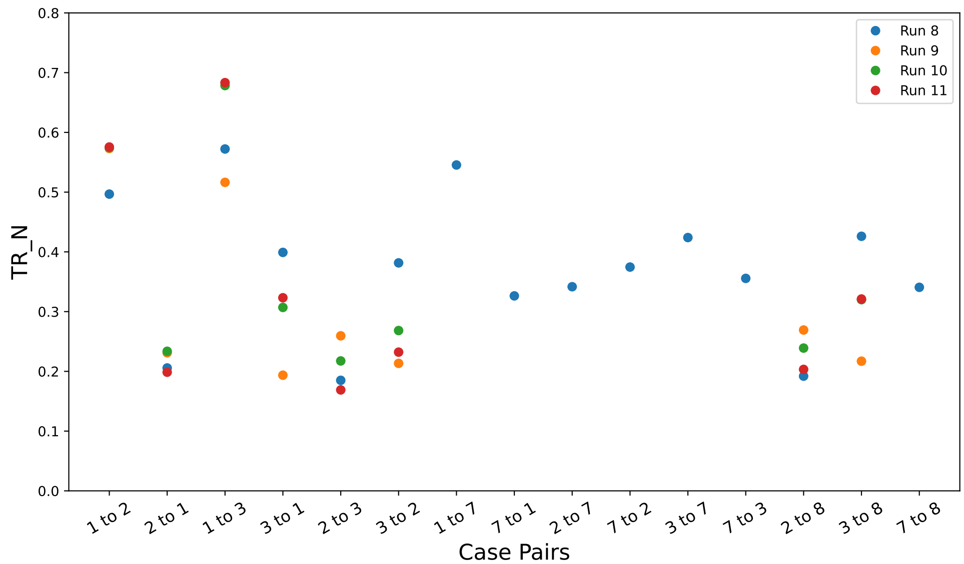

(a) Silvus Radio Sessions

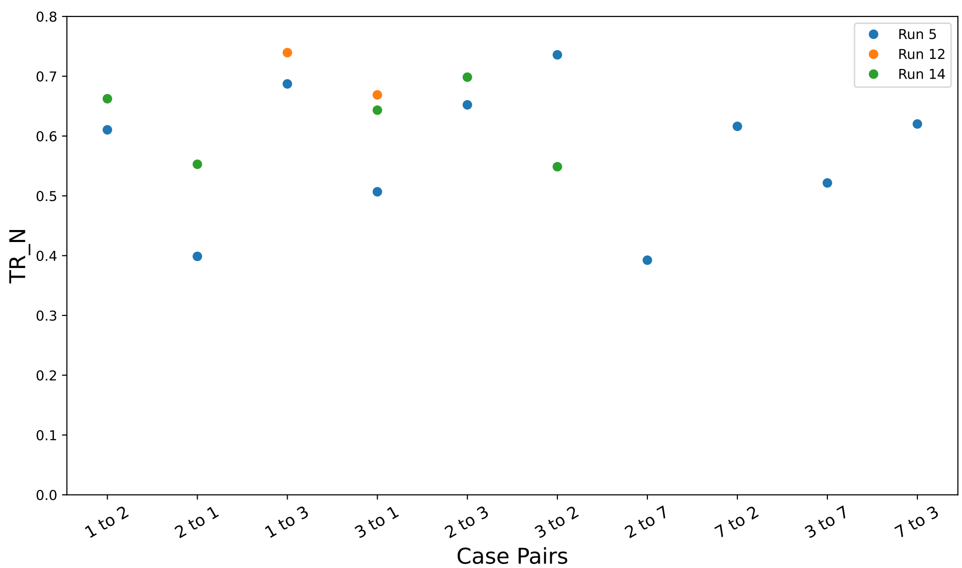

(b) DoodleLabs Radio Sessions

Figure 8. TR_N per case pair, (top) Silvus radio sessions, (bottom) DoodleLabs radio sessions

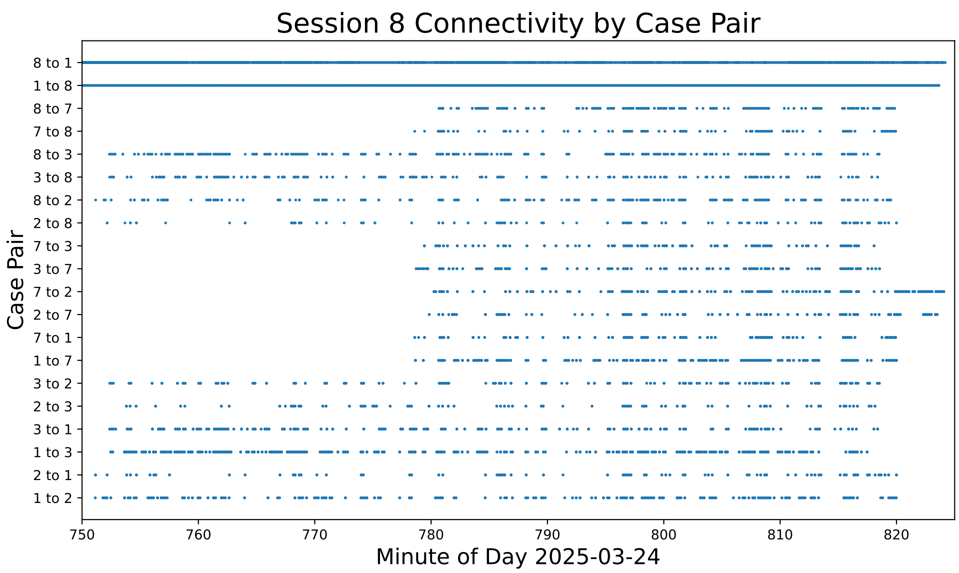

Figure 9. Session 8 connectivity by Case pair

5 Conclusion

Analysis of the TCL 1 Phase 1 data showed the radio mesh network performed well in sharing airspace management and mission planning information. TCL 1's Data Phase 2, however, saw noticeable delay or even loss of data during transmission across the network. The results suggest the performance of radio mesh networks and their associated concept should be further examined and assessed.

It is likely the single aerial relay radio node does not scale well with the number of ground radio nodes. The Phase 1 session had three ground nodes, and the Phase 2 session selected for comparison had four ground nodes. The increased distances between the aerial communication relay aircraft and the ground nodes were likely to contribute to the degradation. Additional analyses will have to be performed to better assess the cause. Alternative concepts or technical options should be considered as well.

Regardless of the network performance, PAMS should be designed to be as robust as possible against various levels of network stability. For example, in case of intermittent network connectivity, PAMS can be more conservative in its operation planning, reserving more airspace buffers between different operations' volumes. If PAMS misses a certain drone's tracks for an extended amount of time, a growing buffer around the missing drone's last location should be calculated so as to maintain safety. These lessons learned will be applied to future versions of PAMS.

References

1. M. G. Wu, A. A. Munishkin, J. Jung, and M. Xue, "Development of a Portable Airspace Management System for Wildland Fire Operations," AIAA Aviation Forum, AIAA-2025-3305, 2025.

Appendix

Figures 10, 11, 12, 13, 14, and 15 show the connectivity by case pair for other sessions than Session 8 from the TCL 1 Data Phase 2. Interestingly, for Session 9, Case 1 and Case 8 struggled to connect with each other for the first half of the session, even though they were connected via an ethernet cable. Similar unstable connectivities were observed also for the first half of Session 10. The reason for this issue is unknown.

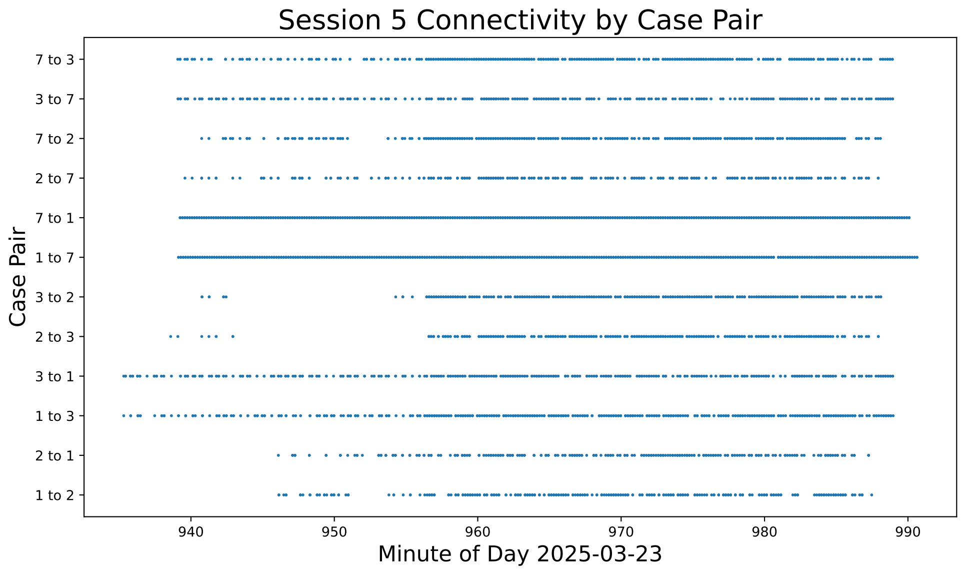

Figure 10. Session 5 connectivity by Case pair

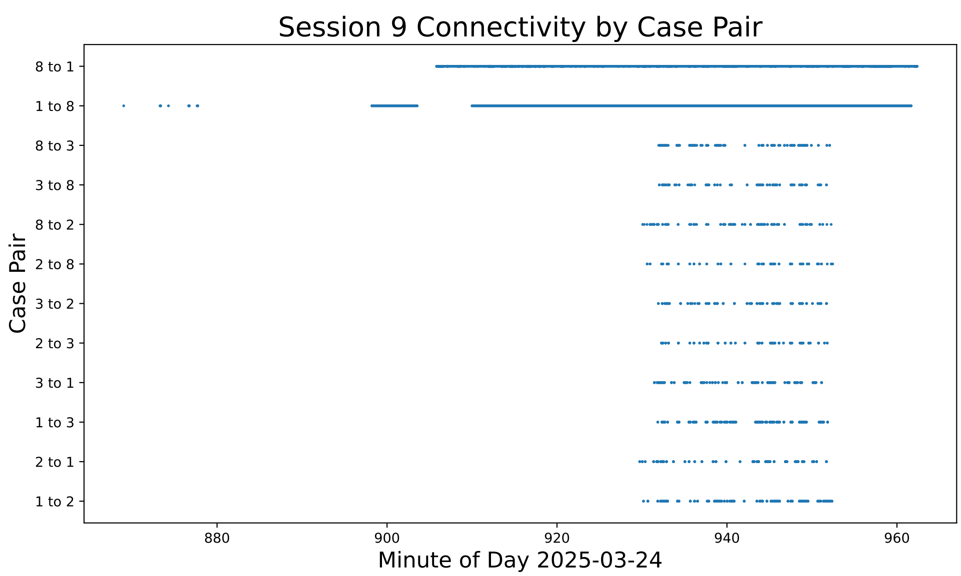

Figure 11. Session 9 connectivity by Case pair

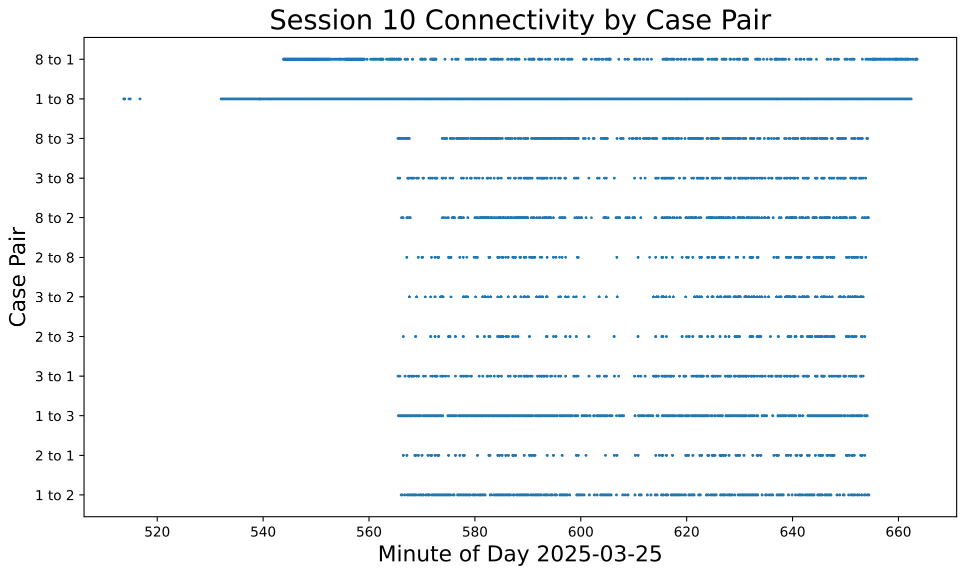

Figure 12. Session 10 connectivity by Case pair

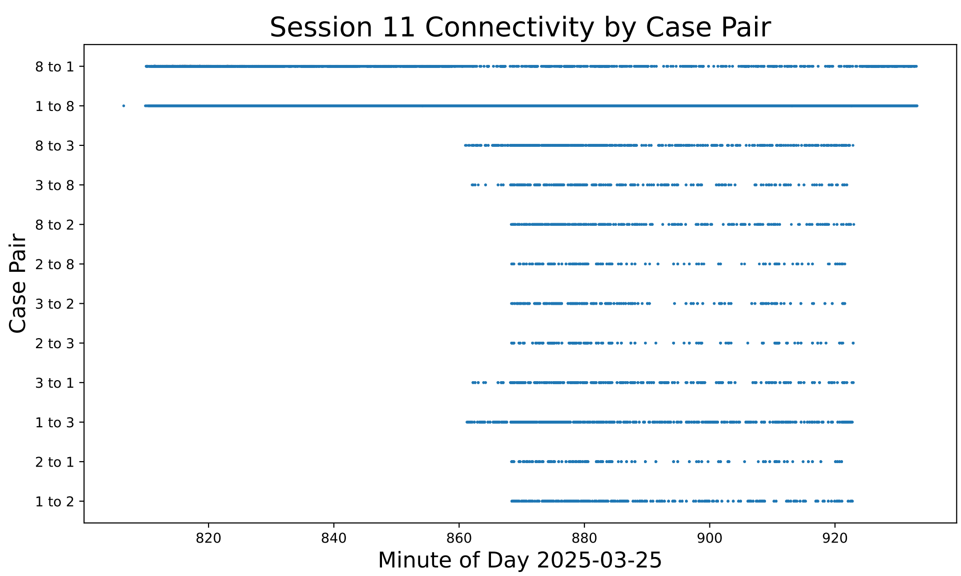

Figure 13. Session 11 connectivity by Case pair

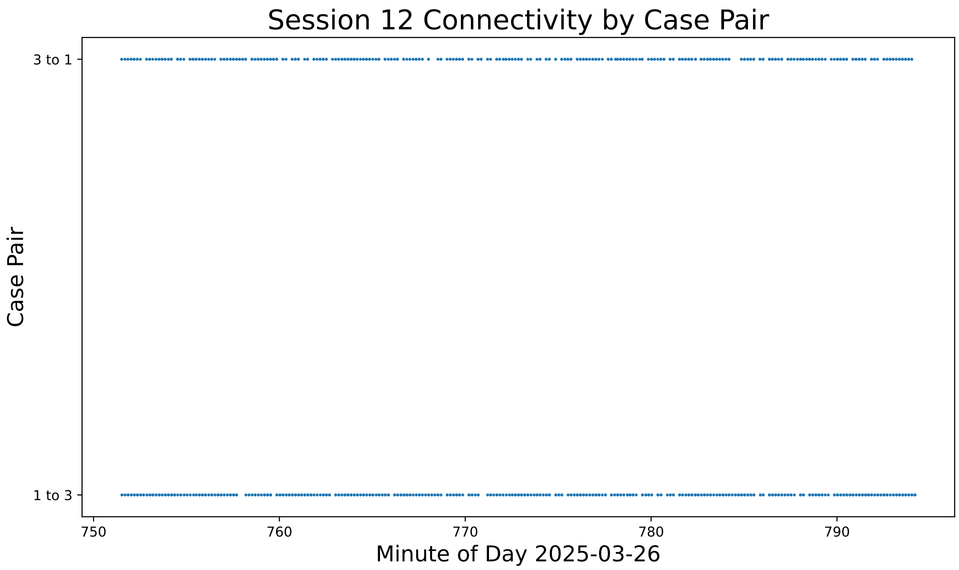

Figure 14. Session 12 connectivity by Case pair

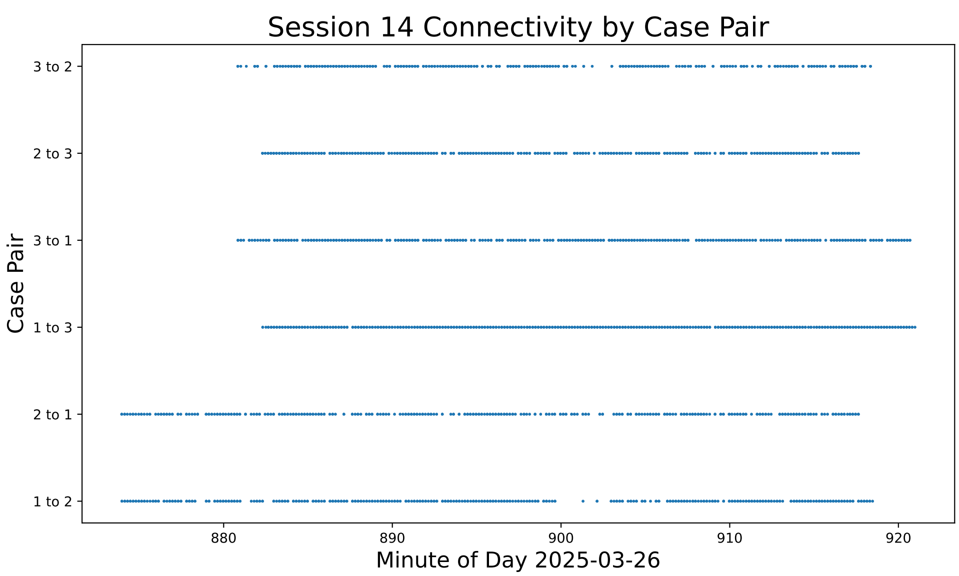

Figure 15. Session 14 connectivity by Case pair