Analysis of ALERT2 Performance
in the Urban Drainage and Flood Control District

Submitted by: Don Van Wie, Telos Services Inc. and R. Chris Roark, Blue Water Design LLC
Oct 30, 2010

The Urban Drainage and Flood Control District (UDFCD) is contemplating the use of ALERT2 Message Concentration as a means of reducing ALERT contention in the District's ALERT Flood Detection Network. Since this is to be an early adoption of ALERT2 technology, extensive testing has been performed over the last 2 years by operating ALERT and ALERT2 on parallel output channels of the Network. Since February of 2010, data from the Diamond Hill receive site has been collected from the NovaStar 5 platform and analyzed on a regular basis. This report summarizes our findings.

The purpose of the testing is threefold:

1. Demonstrate the reliability and availability of the ALERT2 hardware/firmware while operating under actual environmental conditions
2. Quantify the relative performance of ALERT2 and the degree to which contention data losses are reduced in an operational environment
3. Verify and adjust operating parameters of the ALERT2 process based on analysis of the operational data on a production scale.

Data and Methodology

More than 2 million records are in each of the UDFCD ALERT and ALERT2 datasets covering the period from February 23 through October 18, 2010. Analysis has typically taken place using two-week blocks of data as they become available. The primary method of analysis has been to separate the data into ALERT and ALERT2 streams, then import them into Microsoft Access. The data are next sorted by repeater using queries with identical repeater pass lists. A few IDs were excluded from the study because of irregularities in repeater configuration. Examples are the repeater status reports, some of which are programmed only for ALERT, and the Magnolia weather station, because it reports only on ALERT2. The data sets from each repeater were then summed into hourly counts of reports. There are 5,654 hourly totals in the analysis period presented here, with counts from the two data streams split into 5 repeater sources each.

During the spring and summer of 2010, we observed an anomaly at Blue Mountain in which the ALERT2 data from a repeated weather station was lost if its transmission was coincident with that of the ALERT2 transmitter. This will be discussed in greater detail below. However, in order to preserve data integrity and an accurate comparison between ALERT and ALERT2 data sets, the data from Blue Mountain before August 1 was excluded from the study. Another anomaly occurred at West Creek when power to the repeater was cycled during maintenance. On August 24 at 11:00, the ALERT2 pass list was partially deleted, and the problem was corrected on August 26 at17:00 when the repeater executed a re-boot. West Creek data from this period were excluded from the summary statistics.

Hourly data were transferred to Microsoft Excel and used to co-plot the ALERT and ALERT2 message counts, giving a continuous graph of relative performance for each repeater and for the overall system.

Below is an hourly graph of all included sites for the duration of the study period. The period of no reports on June 1 resulted from a brief base station outage.

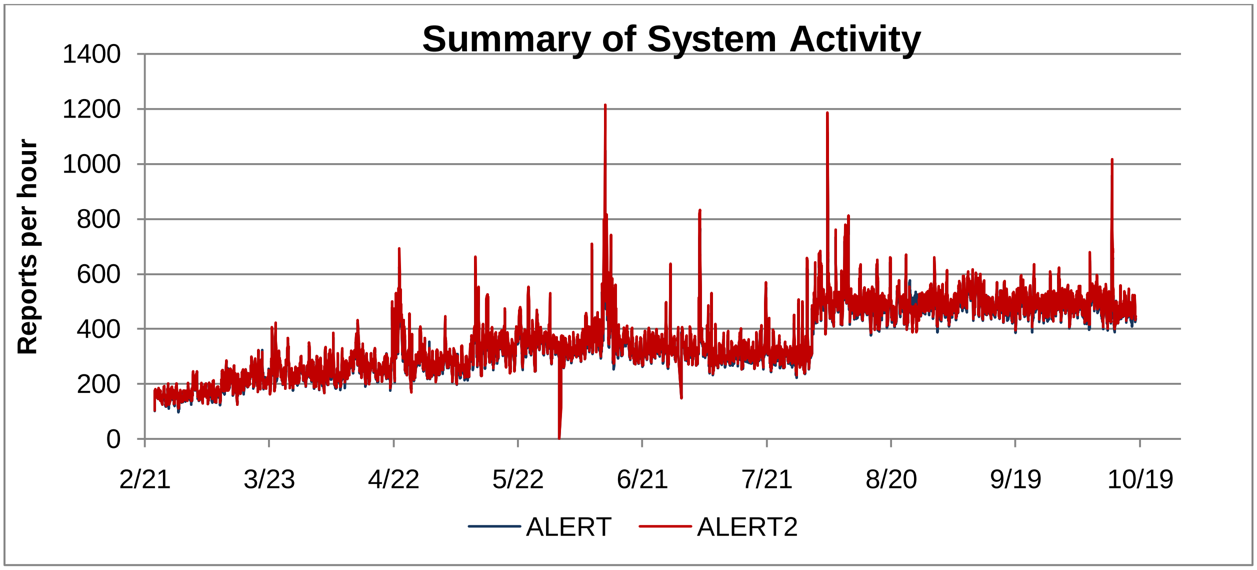

Additional analyses were performed on subsets of data to assess performance of individual gages or individual events. Another analysis method consisted of matching individual reports from each of the data sets. The result was two counts: reports found in the ALERT2 list but not ALERT (ALERT failures) and reports found in ALERT but not in ALERT2 (ALERT2 failures). There is also the possibility that a message was missed by both ALERT and ALERT2. However, if the losses on ALERT and ALERT2 are independent of each other, then the probability of a message being missed by both is the product of their individual loss percentages and is extremely small.

ANTENNA AND PATH CONSIDERATIONS

In adding ALERT2 in parallel to ALERT, it was necessary to install a new antenna, transmission line and transmitter at each repeater site. We attempted to configure the ALERT2 installation so that the effective radiated power (ERP) was close to the same for both systems. Space considerations prevented us from using exactly the same antenna type for ALERT2 as was used for ALERT and some minor differences undoubtedly remain.

- At Smoky Hill, the ALERT2 antenna was installed on the new tower at the same time as the ALERT antenna. Both are Yagi antennas oriented north-northwest. Both installations use a 25 W power amplifier and the match between the two systems is thought to be very good.
- Blue Mountain ALERT has a DB222 antenna with the radiators oriented northeast and southeast. Given limits on available mounting space, we approximated the same pattern with a pair of Yagis in a phased array, one pointing northeast and the other pointing southeast. There is clear sight to Diamond Hill from Blue Mountain, so modest differences in ERP should have little or no effect.
- The West Creek ALERT2 antenna is a tower-mounted Yagi oriented north. The ALERT antenna is mounted about 60 feet higher than the ALERT2 antenna, giving ALERT some advantage here.
- At Lee Hill, the ALERT antenna is mounted slightly higher on the tower, but has a slightly lower gain than the ALERT2 antenna. Both systems use a 25 W power amplifier. The match here should be close.

- The Gold Hill repeater uses a DB222 antenna for the ALERT transmitter. It is configured with one radiator pointing ENE, and the other SE (toward Diamond Hill). The ALERT2 antenna is a 7 dB gain Yagi oriented to Denver. The path is obstructed by Green Mountain. The ALERT2 antenna may have an advantage here of 2 to 3 dB.

Hardware reliability

Blue Water Design Encoder-Modulators and Demodulator-Decoders are installed in two ALERT networks: UDFCD and Overland Park, Kansas-Kansas City, Missouri. For the latest version of hardware and firmware, combining the systems, there are 10 modulators operating with a collective service time of approximately 6.5 years. Six Blue Water Design Decoder-Demodulators have collectively accumulated 7.3 years of operational service. There have been no hardware failures, and no firmware problems have surfaced during operational deployment. Two receivers and one modulator have operated through both winter and summer months in an outdoor environment without problems.

ALERT2 Performance

The hourly data consistently show that more data reports were received over the ALERT2 pathway than the ALERT. ALERT2 outperformed ALERT at all 5 repeater sites. Performance improvement ((ALERT2-ALERT)/ALERT) ranged from 1.5% at Smoky Hill to 4.5% at Gold Hill. Over the period from February 23 to October 18, ALERT2 delivered 42,042 more reports than the ALERT channel. These gains were made primarily at baseline traffic levels, at an average of 479 messages per hour. During the study period, the traffic rate was less than 600 reports per hour 97.5% of the time.

Data Composition and Summary Statistics					
Repeater	Start Date	ALERT Records	ALERT2 Records	Difference (A2-A)	% Difference
Blue Mountain	8/1/2010	328,319	334,064	5,745	1.75%
Smoky Hill	2/23/2010	771,716	781,508	9,792	1.27%
West Creek	2/23/2010	478,256	489,544	11,288	2.36%
Lee Hill	4/9/2010	178,764	183,148	4,384	2.45%
Gold Hill	5/10/2010	250,250	261,083	10,833	4.33%
Total		2,007,305	2,049,347	42,042	2.09%

These summary numbers were obtained by looking at the difference between the hourly report totals from the two pathways. While this tells us something about the relative performance of ALERT and ALERT2, it does not speak to what the absolute losses are over the ALERT2 pathway. To assess this, we used a sample of 25,000 reports from each of Smoky, Blue Mountain and West Creek repeaters. The ALERT and ALERT2 records were matched up, and the number of messages found in the ALERT but not the ALERT2 list was interpreted to be the ALERT2 loss. The table below summarizes the results.

Estimation of Data Loss Rates on ALERT and ALERT2 Pathways						
		ALERT			ALERT2	
Repeater	Sample	Missed by Alert	Alert Errors	% Lost	Missed by Alert2	% Lost
Blue Mtn	25000	445	1	1.78%	1	0.004%
Smoky Hill	25000	149	0	0.60%	4	0.016%
Gold Hill	25000	966	5	3.88%	5	0.020%
Total	75000	1560	6	2.09%	10	0.013%

The study showed that10 ALERT2 reports were lost out of 75,000 sent, a rate of 1.3 reports per 10,000. The comparison is shown graphically below.

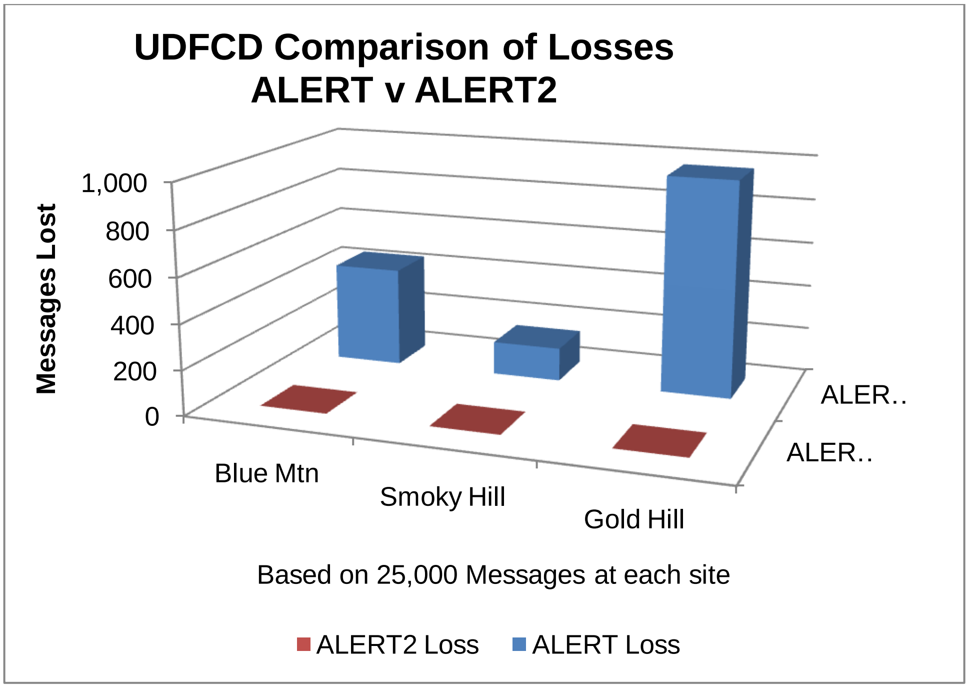

.

Across 25000 records, at Gold Hill there is a 19% chance that a single transmission was lost by both channels. For the other two sites, the probability is 2 %, so it is unlikely that the total count is off by more than one transmission.

Another test of the ALERT2 loss rate comes from the weather station that has been reporting in the ALERT2 Concentration protocol from Telos' office since June 10. It transmits every 15 minutes, plus each .01 inch rain tip. As of October 18, it has sent about 13,000 ALERT2 packets; 3 have been lost in 4 months over a 25 mile obstructed path between Magnolia and Diamond Hill.

CONTENTION LOSSES ON ALERT

In the ALOHA model (designed to predict collision losses in copper circuits), temporal contention between two messages always results in the loss of both messages in the collision. We are seeing that in the ALERT RF domain, it is possible for one message to survive a collision. Since the receiver-decoder can only process one message at time, at least one is always lost. In contention situations, the receiver will more often lock on to the stronger signal; we believe this is the primary reason for the difference in performance between Gold Hill and Smoky Hill. The Smoky path to Diamond Hill is 18.6 miles long and nearly line-of sight, with a 25 Watt power amplifier. Gold Hill has a 30 mile obstructed path and a 5 watt transmitter. The ALERT losses from Smoky average about 1.5% while 4.5% of messages are lost on the Gold Hill path.

There were no large rain events during the 2010 season and the ALERT losses reflect essentially baseline conditions. We know from the very low ALERT2 losses that the "improvement" seen with

ALERT2 is essentially the same as the ALERT contention losses. During one event, the system traffic reached approximately 2000 reports per hour and contention losses jumped from an average of 2% to about 15%. The plot below shows the range of expected losses (between the two solid lines), with the lower curve representing 1 loss per collision, and the upper representing the other limit of 2 losses per collision.

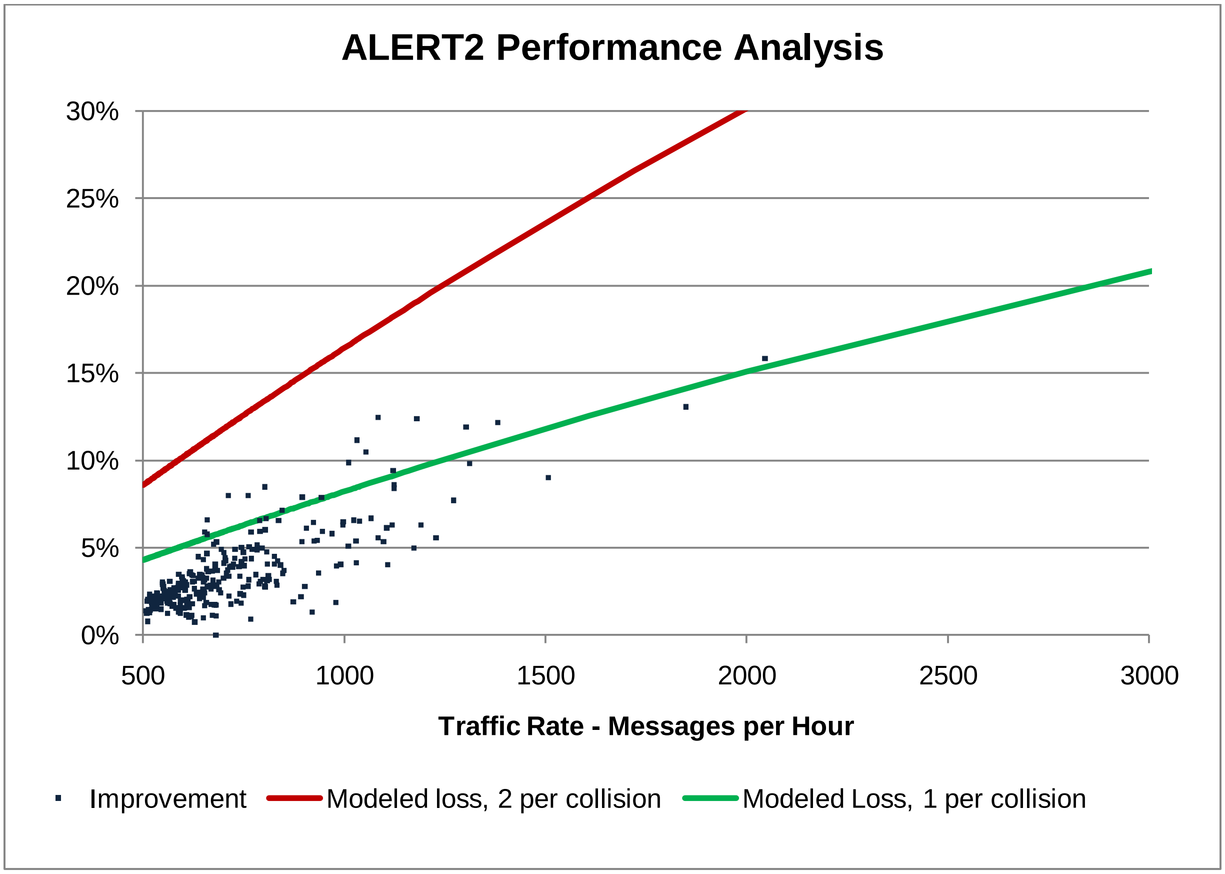

For this graph, identical hourly traffic values were grouped and an average improvement rate was determined for each category. These values were then plotted as a scattergram on the graph of expected losses. At low traffic rates, the actual ALERT losses are lower than the lower limit of the modeled losses. We believe this is because, at baseline rates, much of the traffic is in the form of timed weather station packets carrying multiple concatenated messages. Collision losses from the smaller number of larger packets will be lower than if they were all single messages. Further, the message pattern is non random because an effort is made to set RTU clocks to avoid contention. As the traffic level increases, single rain gage messages begin to predominate, and losses approach and exceed the one-loss-per-collision level. Extending this plot with information from higher traffic levels will yield useful information about the character of contention losses.

STORM DATA

There was a small storm during the night of June11 – June12 that raised traffic levels to nearly 2000 reports per hour for the peak period. It was the largest event of this season, so it was analyzed in detail. The graph below shows the ALERT and ALERT 2 traffic co-plotted.

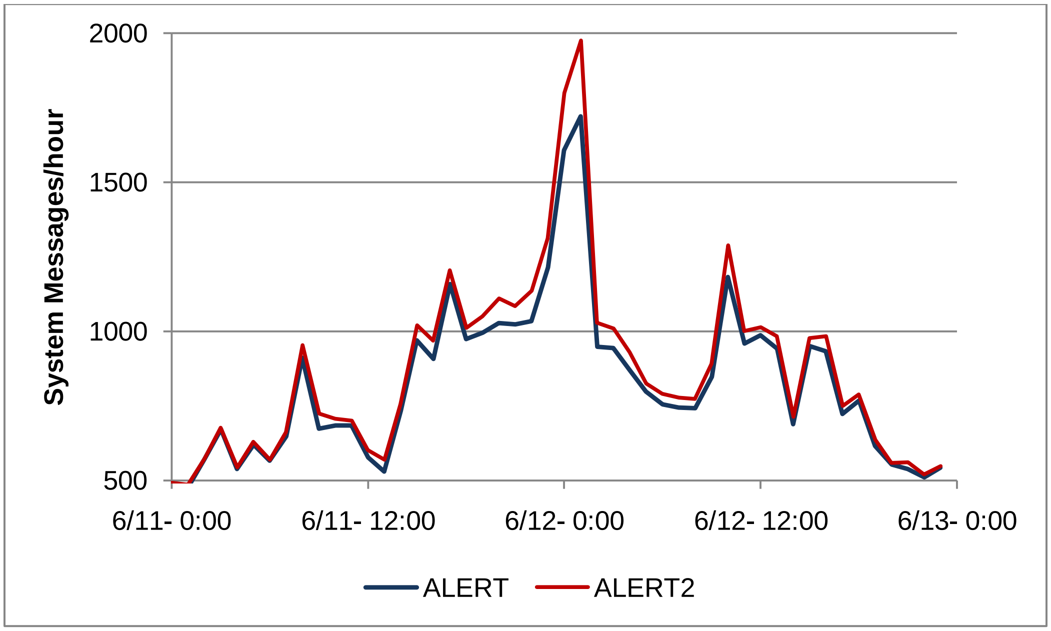

The ALERT and ALERT2 lines diverge as ALERT losses increase at higher traffic rates. The difference in message numbers for each hour illustrates this more directly, below.

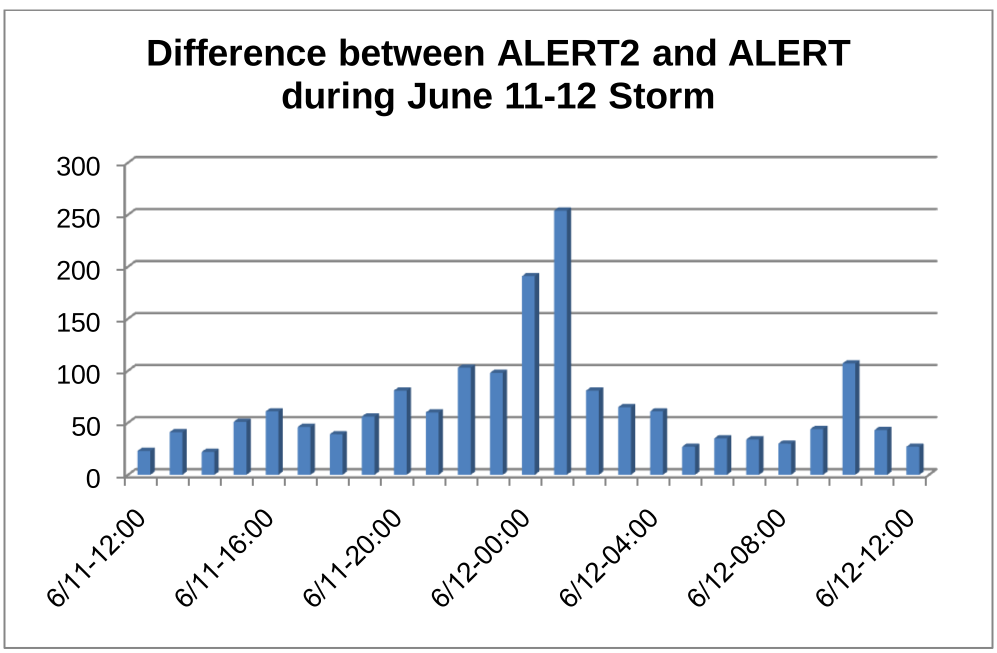

We know that relative performance of ALERT2 increases substantially as traffic increases, and we have seen that ALERT2 losses are in the range of 1 to 2 per 10,000 under baseline conditions. We also want to know what happens to absolute ALERT2 losses as traffic increases. In principal, the rate of ALERT2 losses would be unaffected by traffic load, unless there is degradation of the RF path under storm conditions. We used report matching on the ALERT and ALERT2 data from this storm to determine the actual ALERT2 losses. The analysis covers 6 hours of the storm from 20:00 to 02:00 on June 11/12.

Report Losses during June11-June 12 Storm		
	ALERT	ALERT2
Total Reports	5212	
Reports Lost	568	11
Errors	9	0
% Lost	11.07%	0.21%
Loss Reduction	98.10%	
Loss Ratio A:A2	52:1	

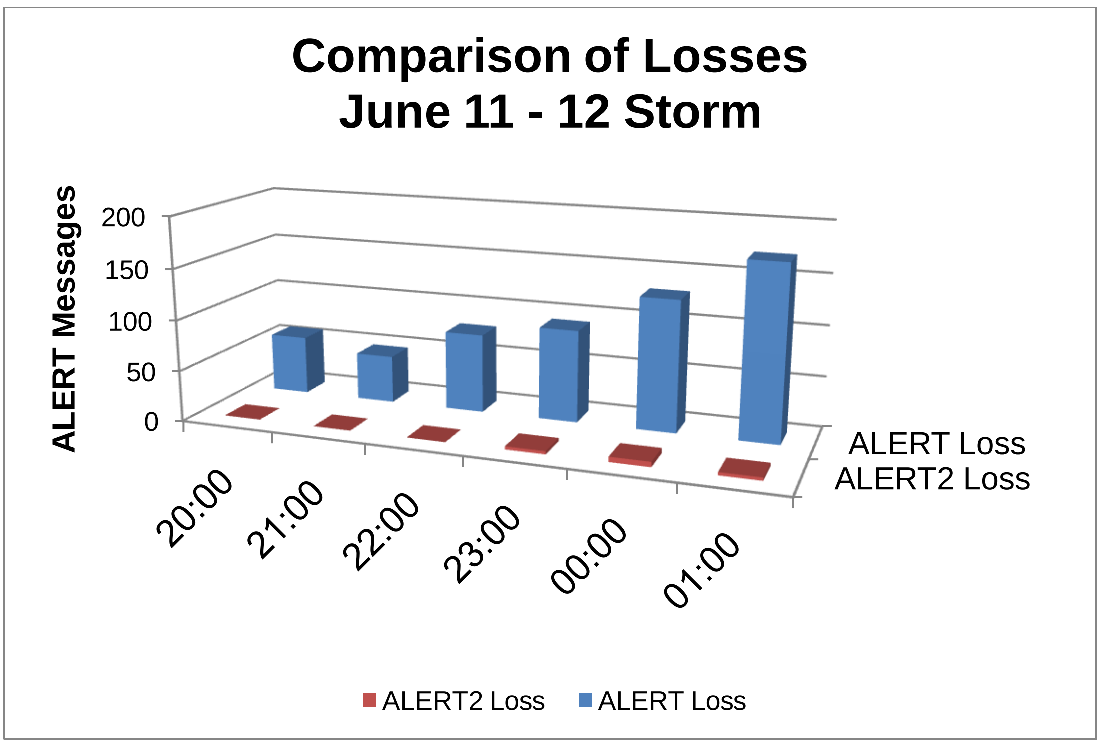

During this event, the ALERT2 losses were only 2% of the ALERT losses, but the rate of ALERT2 losses increased roughly tenfold over baseline to 21 per 10,000. This could suggest path degradation. We looked at the source of the errors by repeater: 2 were from Smoky Hill, 1 was from Lee Hill, and 8 were from West Creek. This is a very skewed distribution. The RF path from West Creek is the longest at 40 miles, and might be the most susceptible to degradation. However, the total number of losses is low enough that we don't believe we can draw any firm conclusions.

ECLIPSING

At Blue Mountain, we experienced one problem that affected ALERT2 operations. The symptom was that the data from a Campbell Scientific weather station would occasionally disappear from the ALERT2 data record for a period of a week or more, while the ALERT record was continuous. The problem occurred only when the weather station's time of transmission coincided with the ALERT2 "slot time" or assigned time to transmit. Although the slot time is tightly controlled by GPS synchronization, ALERT clocks are free to drift. Campbell clocks are more accurate than other manufacturers', and typically drift only a second or two in a month. The "eclipsing" problem occurred when the weather station's transmission drifted into slot time, and ended when it drifted out of it. The same thing was undoubtedly happening with any ALERT transmission that coincided with slot time, but its short length and rapid drift made the time period over which it was occurring too short to readily detect in the data. We learned that the problem occurred when ALERT and ALERT2 were transmitting simultaneously. We confirmed that there was no problem with the Modulator /Encoder, that the ALERT2 message was actually being transmitted, and that it could be heard on site even though it could not be decoded at a distant receive site.

We were able to reproduce the problem with a test gage producing long ALERT messages on the repeater input frequency transmitting at slot time. On an oscilloscope we could see that, in a blocked transmission, the modulation began before the RF signal had stabilized; at a distant receive site there was not enough preamble for the receiver to lock the signal before the data stream began. The cause of the problem is not certain, but it is related to the presence of the RF signal from the ALERT transmission, either at the ALERT2 transmitter or the ALERT2 receiver. The work-around was simply to extend the preamble by 55 milliseconds to a total of 90 from the earlier 35 milliseconds. The problem has not recurred, nor could we any longer reproduce it with test apparatus. The longer preamble has been incorporated into the ALERT2 protocol default settings.

LESSONS FROM OVERLAND PARK

Overland Park has an ALERT2 Concentrator network operating in production mode that is also being closely monitored. Analysis of data there is showing similar results; in particular, the ALERT2 losses "over the air" are typically in the range of a few in 10,000.

In Overland Park, there is no longer an ALERT stream from the repeater to the base. However, there are logging computers installed at each of the Modulator and Encoder sites which capture the incoming ALERT stream and the Modulator and Encoder console log. The latter includes each outbound report.

Two comparisons were made with this data. First, the serial data stream from the ALERT repeater was compared to the output list extracted from the console log. We were able to confirm, over tens of thousands of reports, that the Modulator and Encoder passed on every valid message it received on its input port. We found several instances where the repeater, a High Sierra 3300 series, output invalid 4-byte messages. The Modulator and Encoder does an independent check that the messages it receives meet the criteria of an ALERT protocol message, and these reports were not forwarded.

Second, the list of transmitted reports that we logged from the Encoder and Modulator console can be compared with the record of reports logged by the ALERT2 receivers. An analysis of data for the month of September 2010 showed that 127 messages were lost at the Fire Training Center (FTC) receiver out of approximately 295,000 reports sent. The loss rate is 4.3 reports per 10,000, even in the presence of some known intermittent RF interference at the FTC.

SLOT VARIANCE STUDY

Also at Overland Park, we studied the degree to which the Encoder and Modulator transmit-time varies from its assigned time slot. The Encoder and Modulator uses an internal clock to initiate transmissions, and this clock is checked against a GPS time signal every 30 minutes. If the clock has drifted more than 15 milliseconds, it is resynchronized. The resolution of the Modulator and Encoder clock is also 15 milliseconds, so there will also be a jitter of about +/- 7.5 milliseconds about the mean. The analysis relies on a receiver clock (doing the timestamping) that is itself closely synchronized to the correct time. We appear to have such a receiver at the OP Fire Training Center. The graph below shows a sample of the variance data that was analyzed for the entire month of September.

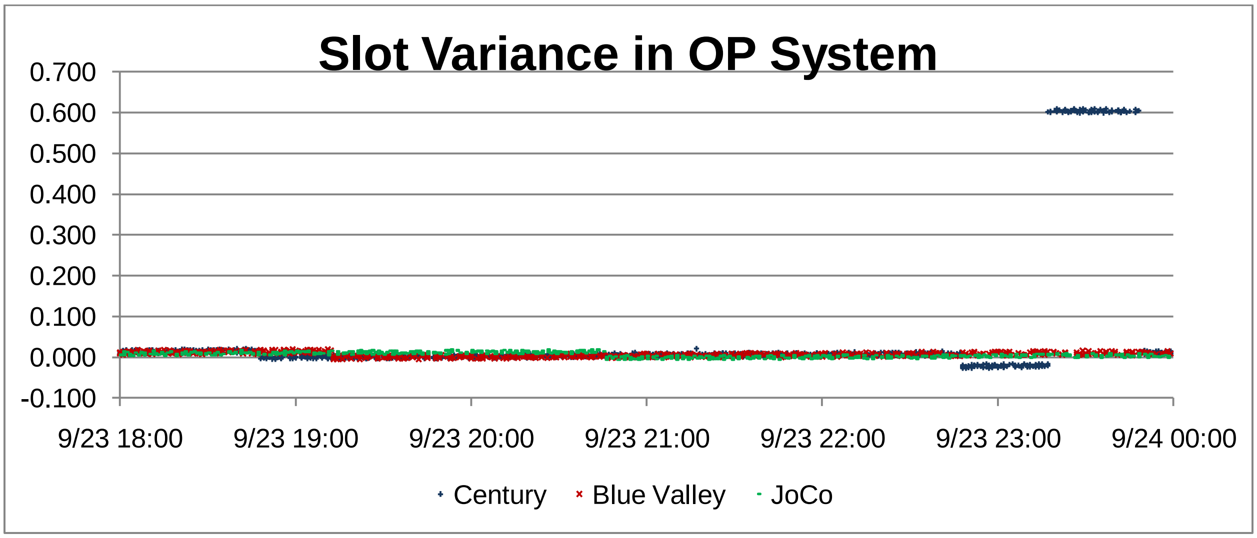

The graph above shows that all three Modulator and Encoders transmitted within a narrow slot for the great majority of the time. However, there were two instances during the month when the GPS time value was anomalous, and resulted in moving the Modulator and Encoder clock to the wrong value until the next GPS time check. We are confident that we can overcome this problem by modifying the time setting algorithm to limit the correction to a plausible value for the time elapsed since the last successful GPS check.

The graph below shows a detail of the time variance plot, expanded in both the X and Y dimensions. Each symbol represents a single ALERT2 transmission. Each of the three sites behave similarly, although independently. Note that each plots a line that is about 10 to 12 milliseconds wide; this is the jitter introduced by the resolution of the Modulator and Encoder clock. Each of the broad lines slopes upward to the right, which is a positive drift of the Modulator and Encoder clock. When the drift exceeds 15 milliseconds, it is corrected at the next GPS check, creating the saw tooth pattern. The oscillation appears to be contained in a band about 25 milliseconds wide.

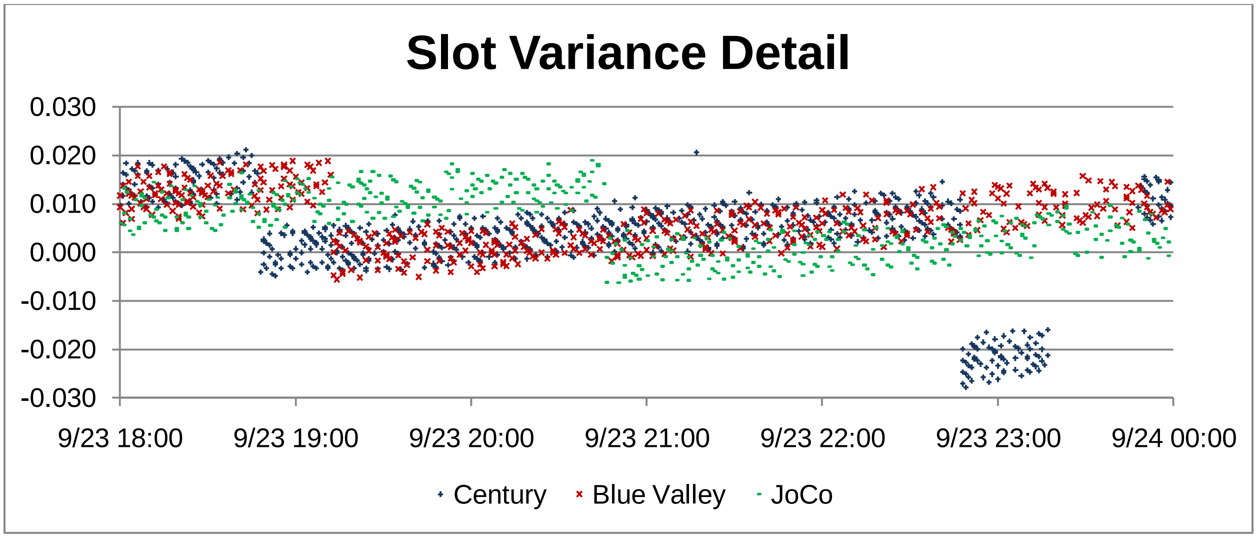

We have concluded that a dead band of 100 milliseconds between slots will be adequate. This means that a slot time of ½ second would have 400 milliseconds of usable transmit time, which is a capacity of 46 bytes. This would support a station with 5 sensors reporting 4-byte data values, along with a rain gage reporting up to 9 tips.

CONCLUSIONS

We have examined several million records in the UDFCD and Overland Park ALERT/ALERT2 systems. At the conclusion of 8 months of intensive testing, we are confident that ALERT2 hardware is reliable, and that Modulator/Encoder and Demodulator/Decoder firmware is robust and accurate. We have determined that ALERT2 performs well over any path that is suitable for ALERT, with "over-the-air" losses typically under 5 in 10,000. ALERT2 losses may increase under storm conditions due to path degradation, but we do not have compelling evidence for this as yet. It will be desirable to conduct an analysis on an event that exceeds 3000 messages per hour.

We are fully convinced that ALERT2, operating in ALERT Concentrator Mode, is ready for deployment in any ALERT system in production mode.