Investigation results of the latency increase after the Ridgecrest M7.1

August 15, 2019

Igor Stubailo, Caltech (stubailo@gps.caltech.edu), Marcos Alvarez, USGS (malvarez@usgs.gov),
Glenn Biasi, USGS (gbiasi@usgs.gov), Mike Watkins, Caltech (watkins@gps.caltech.edu),
Egill Hauksson, Caltech (hauksson@gps.caltech.edu), Bhadha, Rayomand, Caltech (rayo@caltech.edu)

Executive summary

For the 2019 M6.4 and M7.1 Ridgecrest earthquakes, 40 to 50 seconds after their origin time for cellular links and much sooner for the radio connected sites, a number of SCSN telemetry links suffered from moderately increased data latency due to decreased efficiency of data compression and corresponding increased data and packet volumes. The problem could be solved by substantially increasing the available bandwidth.

Summary

Overall, the Southern California Seismic Network (SCSN) functioned well during the 2019 M6.4 and M7.1 Ridgecrest sequence. The initial M6.4 and M7.1 mainshock magnitudes and epicentral locations were issued after 80 seconds and verified within 5 minutes. Subsequent notifications of aftershock activity were similarly successful. Archiving of data from ~600 seismic stations was very successful with no data loss or delays. ShakeAlert messages using SCSN data were issued within 7 sec of the occurrence of the events.

Because the first seconds of the M7.1 (and M6.4) were not affected by the increased latency of the nearby stations, the corresponding data were delivered to the processing center promptly and the EEW algorithms detected the triggers and identified the major event picks successfully. EPIC, the main ShakeAlert algorithm, relies on the first 4 seconds of P wave data received from 3 stations to create a solution and 4 stations to issue an alert.

Later, a number of SCSN telemetry links suffered from cellular network disruptions and from moderately increased data latency due to decreased efficiency of data compression and corresponding increased data volume. The data links recovered as the compression became more efficient and data volume decreased with time.

Because we have implemented the necessary infrastructure to track and archive data latencies as L1Z channels, we are to identify several key factors that affected the data transport time: increased data and packet rates caused by less efficient strong event data compressibility combined with poor wireless signal due to significant shaking, mostly present in the radio connected sites; artificially decreased (not telemetry related) bandwidth; and short-term connection losses starting about 40 sec after the origin time (not directly shaking related) on many of the cellular connected sites possibly due to increased cell network stress. The latency increase is also visible during some of the smaller earthquakes but on a much lesser scale.

The changes in latencies were detectable on many sites (due to fully implemented and archived L1Z channel data timestamps) that use antenna devices (cell modems, radio, VSAT). The changes were not present or barely noticeable on hardwired stations (fiber, Ethernet) and high bandwidth-to-demand antenna sites (USGS microwave, 4.9GHz links, some 900MHz connections).

Completeness of the archived data was not affected by the latency increase but the data delivery time was affected; timeliness is crucial for the ShakeAlert EEW system. If large events occur in quick succession, the data delay may affect the EEW alerts, which could be delayed and some events could be missed. For example, if the M5.0 foreshock to the M7.1 event had happened closer in time than the actual 200 sec, the warning time for the M7.1 could have been less or the magnitude could have been underestimated.

We plan to use these findings to establish new data communications requirements for ShakeAlert. Such standards will provide guidance for improvements of existing links and engineering of new ones. We also strongly recommend that all seismic networks that participate in ShakeAlert implement and archive L1Z channels. The L1Z data are necessary to perform latency analysis provided in this report. The L1Z channels also are an essential part of both short-term and long-term station health monitoring.

Latency definition

Recognizing the latency data as one of the crucial parameters in EEW, we archive the per-packet latencies in mseed format as it is traditionally done for the seismic data. We implemented the procedure as an earthworm module with a possibility to archive 3 types of latencies: from the end of the packet (L1Z named channel), middle of the packet (L2Z), and start of the packet (L3Z).

The transport latency L1Z in SCSN is archived as a 1 sample/second time series and is calculated as the difference between the "endtime" parameter in the TraceBuf2 header and the current system clock value at the time the input record is processed at Caltech for one Z channel per site (either HHZ or HNZ). Or in other words, from the time a sensor generates a signal to the time a data packet is being accepted by a server at Caltech (the datalogger compression and filtering time is included but the data center archival time is not).

Benefits of the latency archival as it will be apparent from the work presented here:

- Archive the latencies to better understand and observe the telemetry long term changes caused by the weather, telemetry hardware upgrades, malfunction, and route changes;
- Ability to go back and investigate how latent waveform data was during a specific event or during a specific time period. An event may be an earthquake or a GPS download or a download from a Baler or something else;
- Employ seismic software tools to process the latency archived data, like filtering, cross correlation, etc. This would help to identify the behavioral similarities in a large pool of telemetry links;
- Verify that the sites and telemetry behave as expected and pass the EEW transport latency requirements.

Examples of increased latency

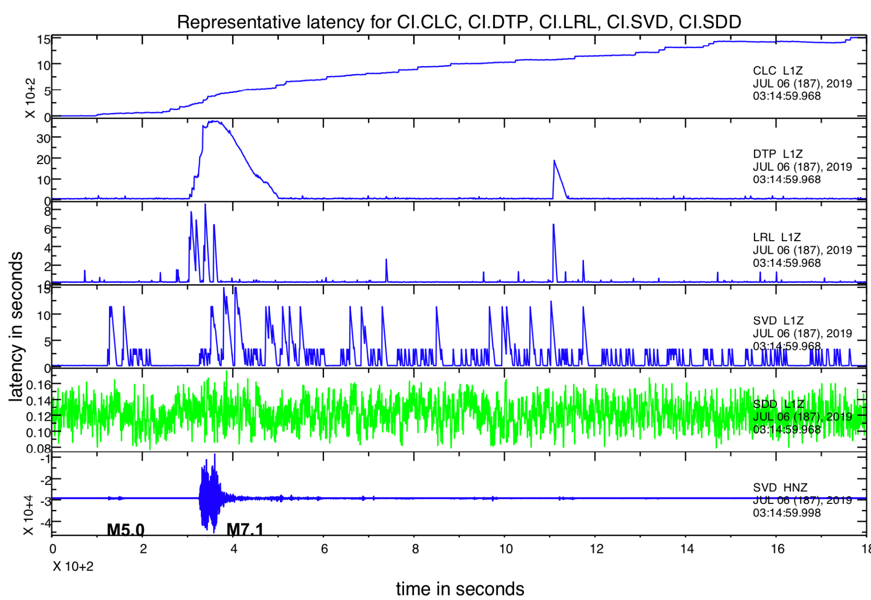

Fig. 1. Bottom trace shows the CI.SVD.HNZ seismic waveform with an M5.0 foreshock and the M7.1 main shock. Blue traces are example stations with increased latency, green shows a station with no increase in latency. Latency for station CLC started with the foreshock and did not recover.

What can be done to eliminate or minimize the latency increases during strong events:

1) Increase the bandwidth of the telemetry links to accommodate the data rate increase. "Peace time" measured bandwidth should be <u>at least</u> 10 times per site that nominally required. As suggested in Steim, 2014 paper, a sustained data rate can be as high as 50 kbps or 6250 Bps (Bytes per second) for a *single* 6 channel Q330. Our observed numbers at Caltech were as high as 3000 Bps averaged over 100 seconds. Increasing radio bandwidth may require shorter connections. During the new station installations make an attempt not to go over 20 miles (32 km). Make a proper estimate of the distance vs. the bandwidth for all types of radio in our use.
2) Measure the bandwidth on a regular basis to get the baseline and to detect the change when it happens. Use the Q330 flooding capability and compare with 4 MB file transfer results obtained previously. Basalts do not have that built-in capability, but Q330 data can be used as a proxy. Anticipate some connection quality decrease as the link ages, and plan on some periodic maintenance.
3) Ensure that high-capacity link stations are spatially distributed so as to record a large event anywhere in the network area.
4) Monitor the bandwidth on a regular basis (bi-weekly). Install/replace passive Adam and other switches with Microtik routers or similar devices that would allow traffic monitoring. Implement traffic monitoring on IP radios. The periodic monitoring will help us to identify the telemetry changes and plan for the anticipated maintenance or upgrades.
5) Look into changing all the Q330 configurations to low latency as suggested in the Stem, 2014 paper to Group_Timeout=0.2, Group_Size=0; ACK_Timeout=0.2, ACK_Grouping=1. Our current settings: Group_Timeout=0.2, Group_Size=0; *ACK_Timeout=0.3*, *ACK_Grouping=4*. Verify they are applicable to us.
6) Develop a procedure to better analyze and understand the 900MHz connections. The Ridgecrest earthquakes exposed the importance of radio packetization, transmission and retransmission protocols, handshaking with the datalogger. We need to be able to tell what is happening and how to improve a connection quality with a greater certainty. Find proper software and hardware tools for this.
7) Possibly decrease dependency on the cell telemetry. As visible on the map (Fig.2) and latency plots (Figures 20 and 21), the cell connections didn't have sufficient quality and stability in data transmission *probably* due to tower congestions and/or people calling after a strong event or some network disruption. Also, we have no control over the end-to-end sensor-to-Pasadena connections when a mid-point is commercial cellular and no ability to clearly detect the weakest points.
a) Have a discussion with Verizon to ask whether they implemented a throttling procedure and when/how it comes into effect on their cellular network.
b) Look more into the performance of devices at Caltech to verify they were not involved in the data delay (instead of VZ). It is possible that something else is responsible. Note: so far we weren't able to find a bandwidth or system bottleneck at Caltech.
c) Compare the cell network performance with what is happening during large sport competitions or performances on stadiums. The stadium-goers are always competing with a limited pool of bandwidth.
d) Consider moving a subset of cellular connections to a private network which allows data prioritization.

8) Find high-capacity alternatives for more SCSN sites. One candidate is CalOES towers (this will allow us to have more bandwidth and to minimize cellular dependency as well). The test site CI.LMS on the CalOES system had excellent performance.
9) Separate the seismic and GPS collocated sites telemetry. When the bandwidth and latency are in high demand during the events, the shared links may not accommodate 2 or more streams of data. Note that the GPS sites also experienced the higher data transport times.

Alternatives Requiring Modification to EEW Algorithms
10) Make an attempt to calculate the algorithm solutions onsite and deliver the results only (if the telemetry is not strong enough for the complete waveform, this might work). This could be tried on the dataloggers or Raspberry PIs (RPIs), for example. With the distributed intelligence, the EEW network becomes more scalable and reliable. The RPI approach is a preferred one. This will allow having identical devices at all sites and developing/building the algorithms for only one platform.
11) If the 10) above is implemented, use different independent real time waveform and EEW channels pathways.
12) Develop/improve algorithms that would be less sensitive to the data delivery irregularities.
13) Include other parameters for the event detection, for example data rate increase that can be reported by each site.
14) To preserve bandwidth during a strong shaking, maybe we can rely mostly on the strong motion HN channels and hold the HH broad band channels. Come up with a priority for which channels should be delivered first to satisfy the EEW requirements and minimize the alert times.
15) Perform an EEW system stress test to verify the increased data rate (especially after the latency increase issues are corrected) can be handled by the servers and the algorithms. The test can be done for an event happening in a high density station area (an LA urban area) that would trigger many sites almost at the same time making them to send increased data volumes to the processing centers.
16) Replay the event data with latency incorporated to test the algorithms.

Latency observations

The data transport latency was different for different intervals after the M7.1 origin time. We present the station latency performance for 10 sec, 30 sec, and 30 min intervals after the event. As it was already mentioned, the first seconds of the M7.1 were not affected by the increased latency of the most of the nearby stations (Figure 3) although the data delays started increasing quite rapidly as can be seen in the Figure 4.

For this sturdy, we will concentrate on the long term 30 min interval since this data is showing us what the network telemetry is capable of.
Six different types of latency behavior were observed. The same six types can be identified for both M6.4 and M7.1, but the number of sites involved is different (probably because of different energy release directions). The two features we use to place the sites in different categories are the *duration* of the latency increase and the *amplitude* of the increase. Latency amplitude refers to the actual latency increase; it can be negligible,

slight (<5 sec), large (>> 5 sec), or huge (delays of 100s of seconds). Latency duration refers to how long the latency increase lasts before returning to pre-event conditions. Duration of latency increases are categorized as short (<<600 sec), or long (>600 sec) for the largest events. For smaller events, they form more of a continuum.

Below we categorize 320 SCSN sites for the M7.1 event. For convenience, we mention the site colors from the maps in Fig. 2 where we show amplitude and duration of latency observations in map view for the first 30 minutes after the origin time. Symbols are used to distinguish telemetry type. See the caption for details. Two points are evident. First, latency increases were observed in an irregular pattern across the network. Second, a clear concentration is observed in the epicentral area. For the wider area view, we study impact by telemetry type. For the epicentral area, the discussion focus more closely on individual links and network topological effects.
Examples of most categories are given in Figures below (24-28). Note that both Basalts and Q330s are represented in the station lists and there is no clear behavioral distinction between them.

How the mapped telemetry types were selected:
- for non-mixed telemetries involving an antenna (radio only or cell only or radio connected to Ethernet), that type is used to label a connection;
- for the radio->cell mixed telemetries, up to 30 sec of P wave travel time (~180 km distance) the *radio* was used as a type name, beyond 30 sec, the *cell* type was used. This decision was based on the observation that the latency increase for the cell connections happened almost at the same time, in the range of 30-50 sec after the M7.1 origin time (discussed in more details below). Please note that the number of radio->cell mixed telemetries is fairly small;
- for the mixed telemetries involving a landline (fiber, Ethernet, etc), the antenna connection type was selected as a label (i.e., *radio* or *cell*);
- all landlines are labelled as *fiber*.

Color range corresponding to latency increase	Telemetry type (if mixed, largest latency used)
0-1 : light blue, no L1Z increase 1-2 : light green, slight short 2-3 : green, slight long 3-4: yellow, large short 4-5 :orange, large long 5-6 : red, huge increase in time (100s sec)	F: fiber, edison/dwp, frad, Ethernet R: radio 900Mhz or 4.9GHz C: cellular V: VSAT

1) No latency issues, color: light blue (all colors are taken from the map in Fig.2).
lblue**F**="PSR SDD BAK CHR PLS GR2 IPT WGR THM SBPX CHF DAN LUG WLT EDW2 RRX CBC APL MLS CHN DLA ALP LAF NWH DEV KIK CAC PDU MMC VCS MOP RUS USB DZA GSC WTT2 LFP ADO HLL LCG VES SRN STS SCI2 BFS BCW SUN CWC DEC RSS CFD USC"
lblue**R**="RHC2 SMI SSS LVO YUH2 GVR MWC SRI WWF HYS WMD LDR FOX2 GRA SWS THC TOR CRR JNH2 DRE CFT SAL ERR NCH LMS IDQ BOM CLI2"
lblue**C**="TJR GLA SS2 TPC GFF SAN CPO RCU LAT MIS VCP TEJ VLY CRF CSL BBS HLN IVY NPN FRM CVW SQC YUC CDM GOU DTC QAD PGA SNR CWP RHR CKP HMT2 SOC"

2) Slight latency increase (up to 5 sec), over a short (<600sec) time interval: color light green.
lgreen**F**="OLI MPP PDE LGB FUR STG LLS CLT MGE BRE RIO RSB WWC WNS MIK MIKB LTP SPG2 MSJ BLY"
lgreen**R**="TA2 OCP BAC HOL DSC CJV2 LEO CGO PDW"
lgreen**C**="RSI GFS OSI YEG2 BLA2 HAY CFS HAR LKH LMH CRG CAR CZN TPO EMS RPV CLO SHU FON EOC SMT TIN RUN FMP BBR SLV VDJ IMP PER LMY IRM MTG SYN SLR SMW QLC RKMO ELS2 SDR SMF2 VTV FHO MLAC SBB2 BAR QUG PHL COA CSH WSS ARV WRV2 NBS WMF JVA SMM"
lgreen**V**="TEH RFR HIW"

3) Slight latency increase (up to 5 sec), over a long (600sec) time interval: color green.
green**F**="RVR PMD GMR LYP"
green**R**="HDH RXH SIL"
green**C**="MAG CRN FUL MTA GCC MES BUE SDG LBW2 SLM MNO DJJB ASP AGO WLS2 FMO OGC PUT MOR RMM BAI LVY DPP RAG HEC FDR LOC MCT PDM TFT CPT2 CHI OAT LRR2 SRA SHO AVM SWP JTH POR BLC CTW GOR PALA DGR BOR SBI SLH"

4) Large latency increase (> 5 sec), over a short (<600sec) time interval: color yellow
yellow**F**="PLM CTC"
yellow**R**="LRL IDO SNO WBS MSC CCC CCA GATR DTP"
yellow**C**="NEN ISA LUC2 WES TUQ CJM NEE2 VOG WHF WOR LMR2 LBW1 FIG POB2"
yellow**V**="KYV NJQ LDF CYP"

5) Large latency increase (> 5 sec), over a long (600sec) time interval: color orange
orange**C**="CIA SVD MUR MTP OLP NSS2 GMA IKP BC3 DJJ LJR JEM SMR"
orange**F**="PSD"

6) Huge latency increase (100s sec): color red
red**F**="SLA LPC SLB MPI"
red**C**="SPF PDR WBM WAS2 SRT IDY BHP"
red**R**="MPM TOW2 WRC2 WCS2 WNM CLC WVP2 BTP"
red**V**="EML"

If we consider the sites in categories 1) and 2) above as having an acceptable latency, then the success rate of sites per telemetry type during the 30 min interval would be the following (from best to worst):

fiber, edison/dwp, frad, Ethernet (F, 72 out of 83):	86.7%
radio 900Mhz and 4.9GHz (R, 37 out of 57):	64.9%
cell (C, 90 out of 172):	52.3%
VSAT (V, 3 out of 8):	37.5%

By far the best performing links were wired and high capacity backbones including the 4.9GHz radios. 900 Mhz radio links showed mixed performance, but still only 2/3rds sustained adequate throughput. Only half of cell systems sustained low latencies and capacity throughout the largest events, and few of the VSATs performed well.

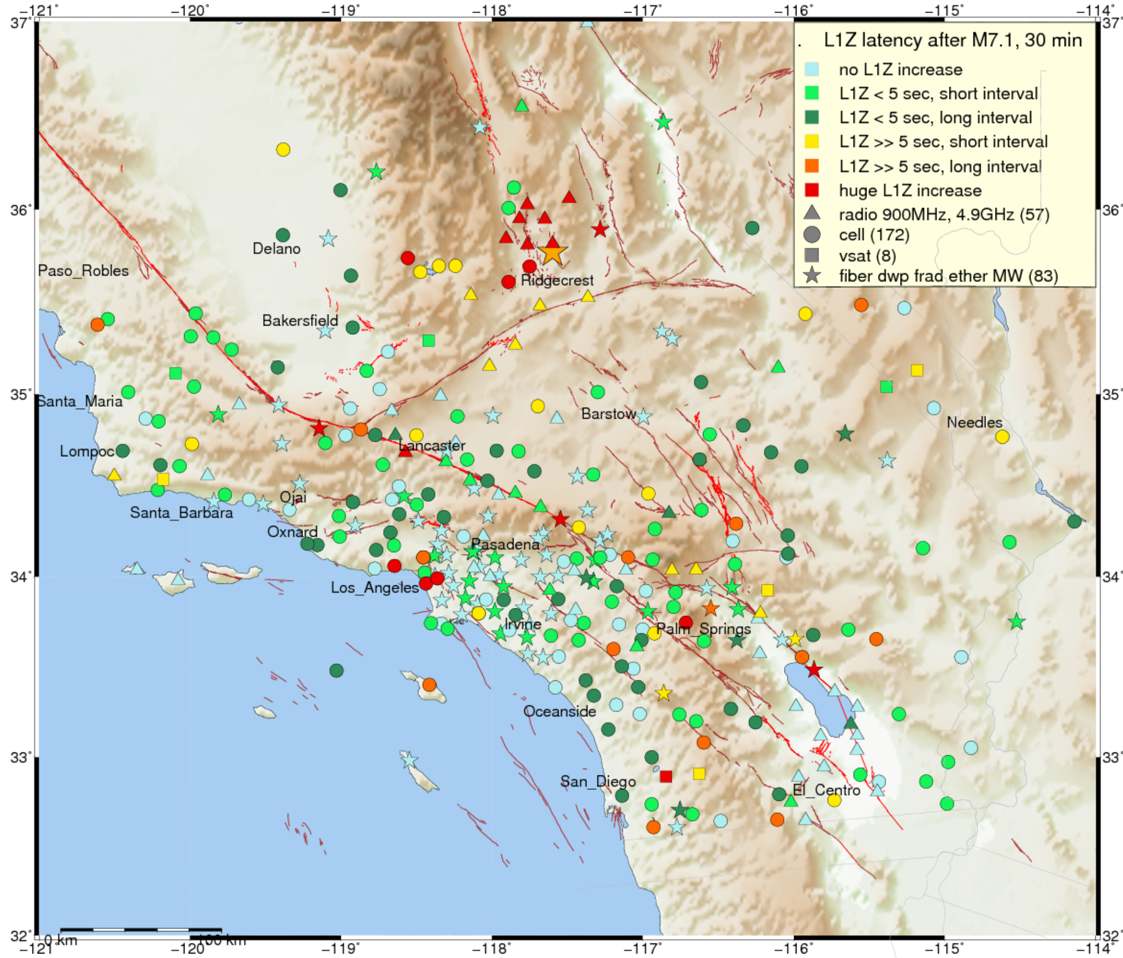

Fig. 2. Map displaying maximum latency (colors) and telemetry types (shapes) for SCSN sites observed *during the first* ***30 min*** *after the M7.1* origin time. The number of sites per telemetry type is given in brackets. It is surprising to see so many cell-connected sites affected by the shaking.

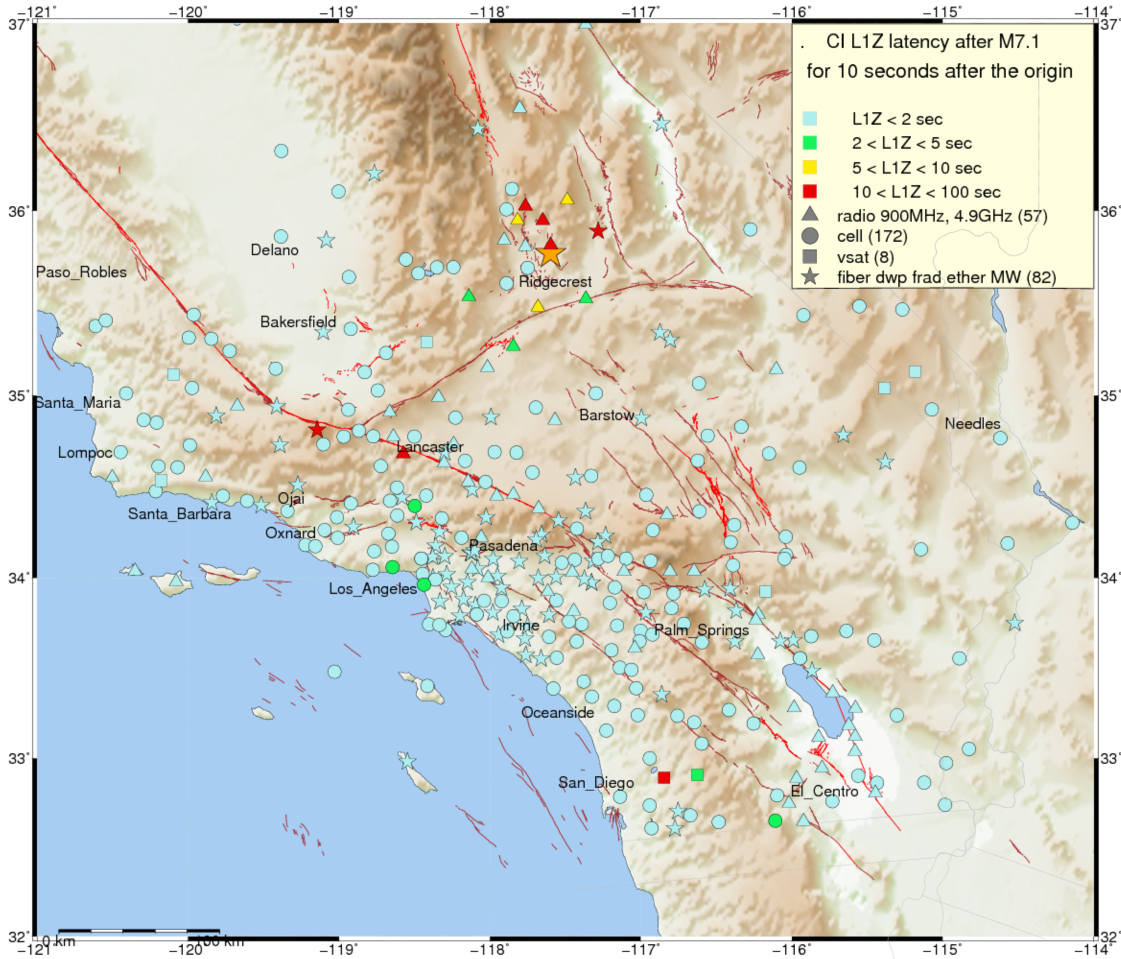

Fig. 3. Map displaying maximum latency (colors) and telemetry types (shapes) for SCSN sites observed *during the first **10 sec** after the M7.1* origin time. The first seconds of the M7.1 were not affected by the increased latency of the most of the nearby stations which is important for the EEW.

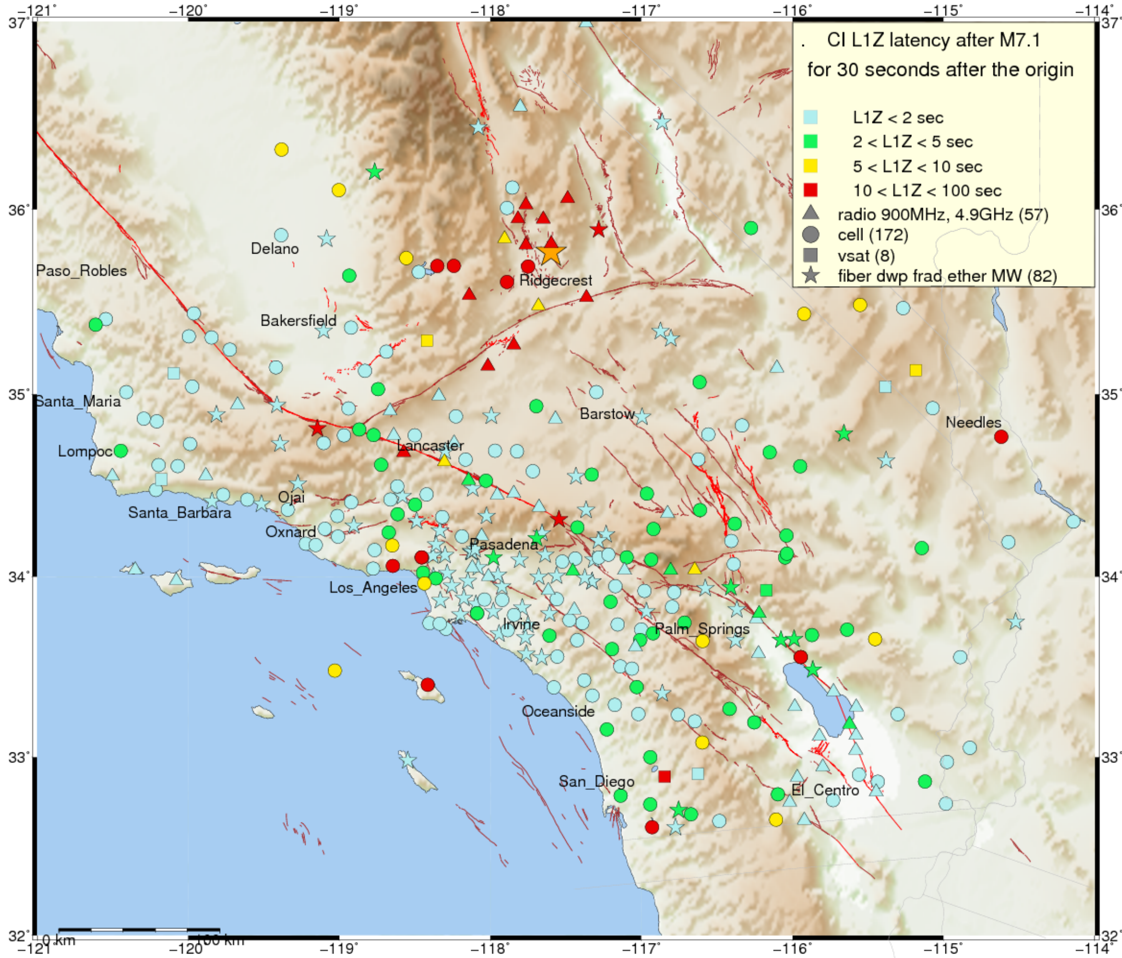

Fig. 4. Map displaying maximum latency (colors) and telemetry types (shapes) for SCSN sites observed *during the first **30 sec** after the M7.1* origin time. The latency of the nearby stations was affected because of the strong shaking.

A lab test that shows the data delays

In order to simulate the decreased bandwidth due to antenna shaking (either the shaking at the seismic sites, or seismic site antennas and microwave antennas on the cell towers or something else, or people calling and overloading the cell towers), we did a large file transfer from a desktop computer to a Basalt in the lab while collecting seismic data from it during a quiet time and simulated events.

What was done:

a) started scp a 200Mb file from the server *crust* to the lab Basalt CI.DSN8

b) started scp a 200Mb file (at the same time) from the server *pine* to the lab Basalt CI.DSN8

c) generated the M7.1 (Fig.29-30) by injecting the voltage changes generated by the computer sound card (functioning as a wave generator) during a M7.1 HNZ channel playback and monitored the data latency on plume. For the sound voltage input we built a special sensor cable.

d) to prove the concept, we also manually shook an Episensor (instead of a sound input) or jumped next to it and observed the increased data latency as well.

We noticed several different behaviors:

1) the latency *was not* increasing when the table with the Episensor was moving periodically and the data high frequency content was low
2) the latency *was* increasing when the signal was generated by a jump or simply walking (complex irregular waveform with different frequencies and amplitudes) or a simulated M7.1 playback.
3) while trying to move the table (manually) with sensor in this way: periodically->randomly->periodically, the latency was mostly increasing during the periodically->randomly and randomly->periodically episodes but was not high during the periodic motion. The observed latency times in this case were on average: ~0.3 sec vs. ~1.2 sec, with max latency jumps at ~6 sec.

The above observations could be explained by different data compression efficiencies during an event and before/after it (see the Appendix for email correspondence with J.M. Steim). It seems it works similar to the motion picture frames: i.e., when a video is playing, only the updates are delivered, not the complete frames. The updates during the periodic motions were small and hence the latency was not increasing; in the case of an earthquake, the data is more complex in the frequency domain and the compression is less efficient delivering more data with every frame update (using the video analogy). The theoretical explanations for this effect were provided by Joseph Steim and can be found on his poster (see the References section).

What was tested but was not confirmed:

We noticed that during the events, the data rate (Bytes/sec) and number of packets increased dramatically (as reported by *qmaserv* running at Caltech and logging the values every 100 seconds for each Q330 delivering data). During that data rate increase, the latency also increased on many sites, not necessarily close to the epicenter (Fig.19-20). To find out what was happening, we tried to verify several hypothesis.

1) STEIM2 compression only.

Initially we thought that the STEIM2 compression was requiring more time to package the high magnitude event data. This hypothesis was eliminated after comparing the latency between the radio/cell and fiber/Ethernet sites. The hardwired sites didn't display any latency increase, i.e. the compression side effects were not present (Fig. 21).

2) Radio and cell modem buffering.

Another assumption was associated with the radio and cell modem hardware. We thought that the devices could not keep up with the increased data rate.
To test this hypothesis, we tried to replicate the problem in the lab but failed (the devices worked every time).
We executed numerous trials for 3 setups and compared the outcomes:
Basalt -> Ethernet -> desktop switch -> data is delivered to plume/CWB

Basalt -> Ethernet -> *cell modem* -> data is delivered to plume/CWB

Basalt -> Ethernet -> *radio1 -> radio2* -> Ethernet -> desktop switch -> data is delivered to plume/CWB

How we generated the data:
a) We converted an HNZ seismic data channel where the M5.0 before the M7.1 and 3x M5s after where clearly visible into the sound file (wav and mp3), Fig. 27-28. The audio file is posted at *http://crust.caltech.edu/~igor/events/Ridgcrest7.1/sound*
A sound card, during the sound play, generates voltage fluctuations that are similar to the voltage changes generated by a sensor during an event. We connected the PC sound card output to the custom made Basalt sensor cable (providing 4 channels of HNZ, HNE, HNN, EHZ data). Please note that by playing the quake sound waveform with different volume, one can control the event amplitude.
The seismic data collected by a datalogger looks very similar to the played sound wavefom.

b) We simulated a strong shaking by placing an Episensor on a lab made shake table. The waveforms differ for different samples collected but the outcome is not affected by this.

The latency of data collected in cases involving antenna devices looked identical to the one when the Basalt was connected via an Ethernet cable (without any wireless devices involved). The latency was monitored in a sniffwave on plume as well as inspected in the L1Z channel (retrieved from CWB).

The benefit of having a converted to sound M7.1 quake was in an exact repeatability of the tests that allow us to compare all results properly.

3) 12V Adam switches only.
We noticed that the data rate was reported to be increasing by the dataloggers but the rate reported by the cell modems (in AVMS) was not increased. The hypothesis was that the Adam 12V switches were limiting the data throughput hence causing the delay. During the lab tests, several new switches performed well and accommodated any data rate that would be seen in the field. We have not tested the switches that are installed at the sites with the highest latency though, the results for those could be different. We decided not to blame the switches at this time.

Case studies

Below we are presenting our research focused on details of particular telemetry types and different latency types trying to eliminate particular technologies and find common features that would allow us to isolate the issues.

Radio connected links

a) Site locations near Laurel Mountain.
Comparing similar radio links:
CI.HYS: displays no latency increase, the site connects to Shadow Mtn (part of the USGS MW network), 12 km away.
CI.CLC: the site had major problems delivering the data. The data relays to T1 in the program office. The increased latency can be explained by the insufficient link capacity as a consequence of the poor telemetry quality that could not accommodate the event data rate. Both sites had the same FGR+ radios.

Conclusion:
Comparing the L1Z before and during the M7.1, we show the difference in link quality and available bandwidth during the data delivery. The shaking is similar during the M7.1 based on the counts. The dataloggers can compress as fast as it is needed because there was no issue with HYS (i.e. the datalogger is exonerated). The only thing left is the communication link, it appears to be insufficient to accommodate the increased data flow. Note that before the M7.1, the CI.CLC's latency was good.

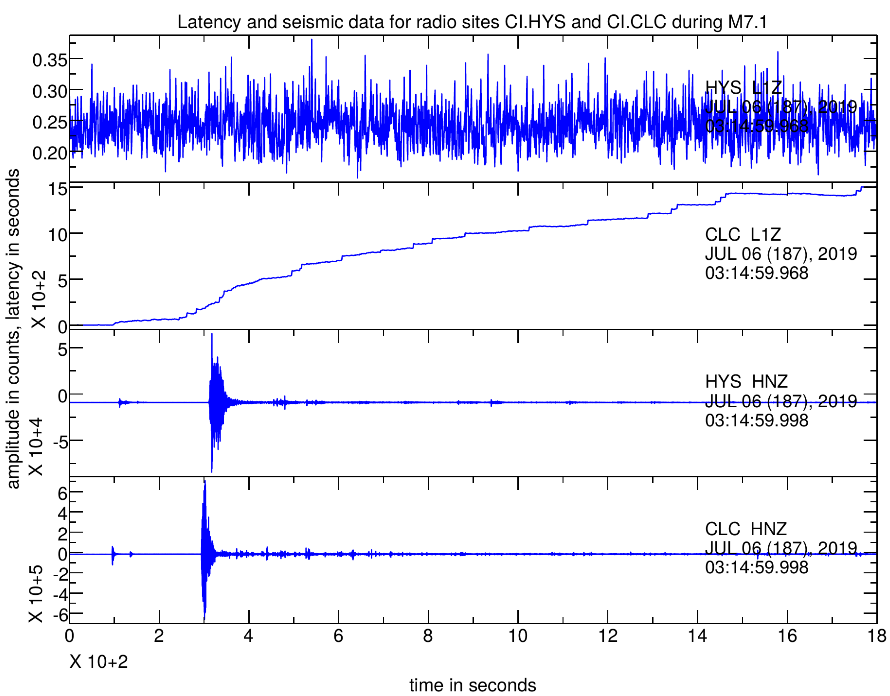

Fig. 5. Comparing latency and seismic data for similar radio links near Laurel Mtn

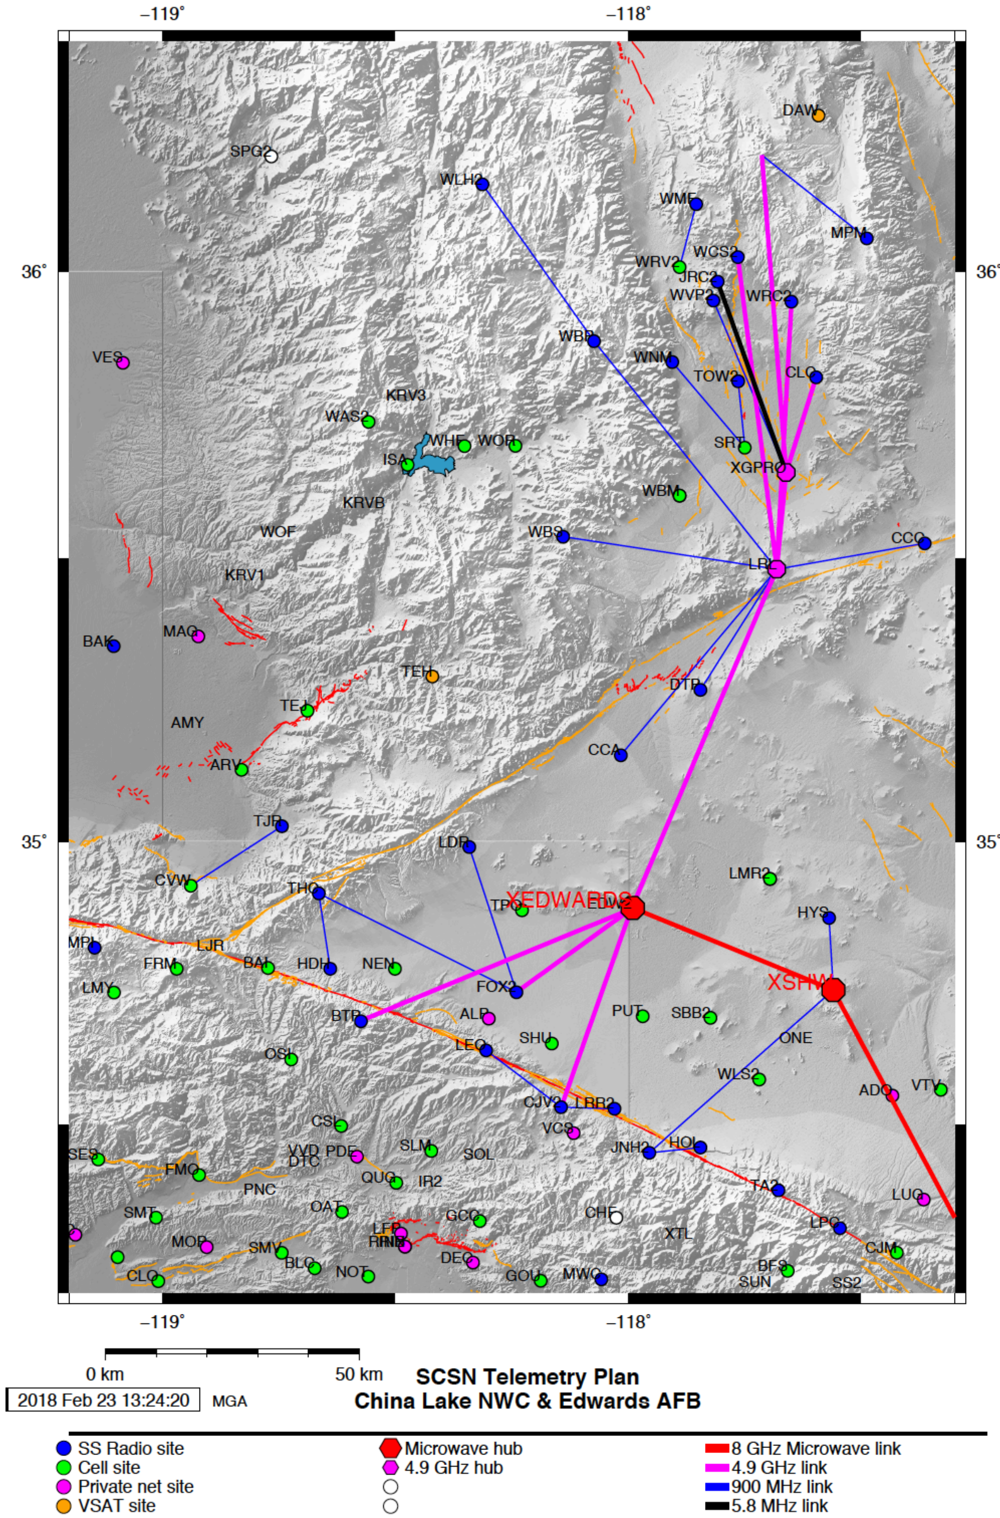

Fig. 6. Map of the site locations and telemetry connections near Laurel Mountain

b) Looking into the data transferred over the T1 connection.
The sites CLC, WRC2, DTP, CCA, WBS, WCS2 all go through T1 in program office.
But TDP+CCA+WBS are different from WCS2+WRC2+CLC.
Difference is the distance and link quality.

Conclusion: T1 is up and is working properly.

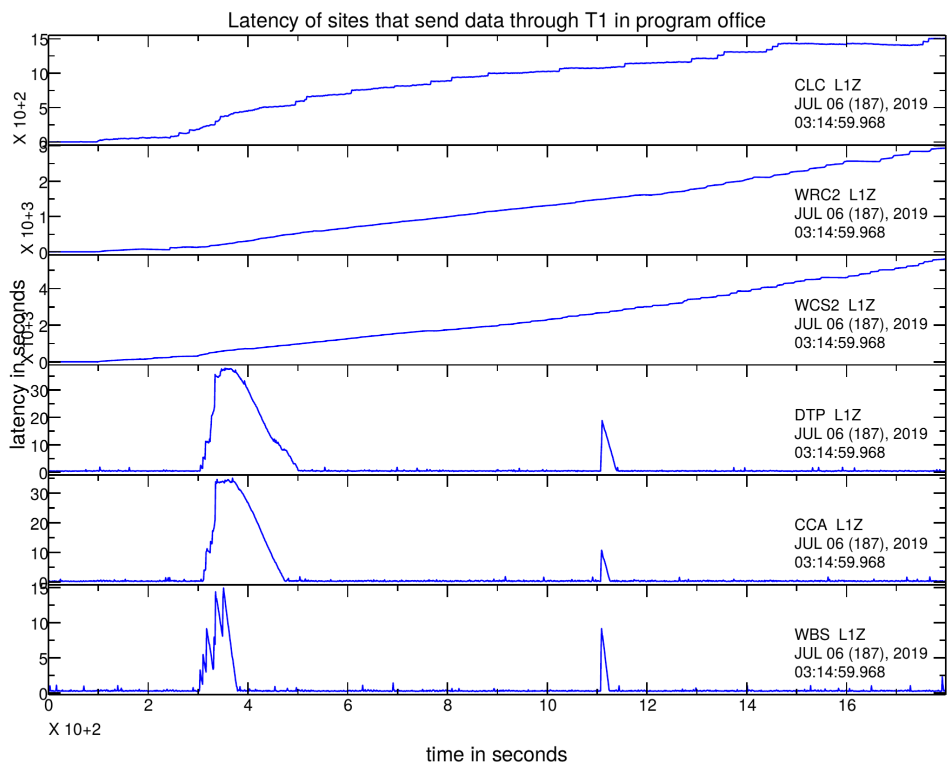

Fig. 7. Example of station latencies that move data through T1 in program office

c) CI.XLRL relays data from from 9 sites to XGPRO over an HTplus high capacity radio link. We can see the latency spike during M7.1. The L1Z before the event was good. The radio shots to LRL are of different distance and telemetry quality but the XLRL-XGPRO line effect is very similar for all of LRL, CCC, WRC2, DTP, WBS, and WCS2.
The compression is not a problem for smaller events with low ground motion (2 bytes/sample, (Steim,2014)) but a larger event becomes an issue with 4 bytes/sample with less efficient compression.

Conclusion: The latency from XLRL was affected by the link to XGPRO capacity (400 kbps, 6 considered sites).

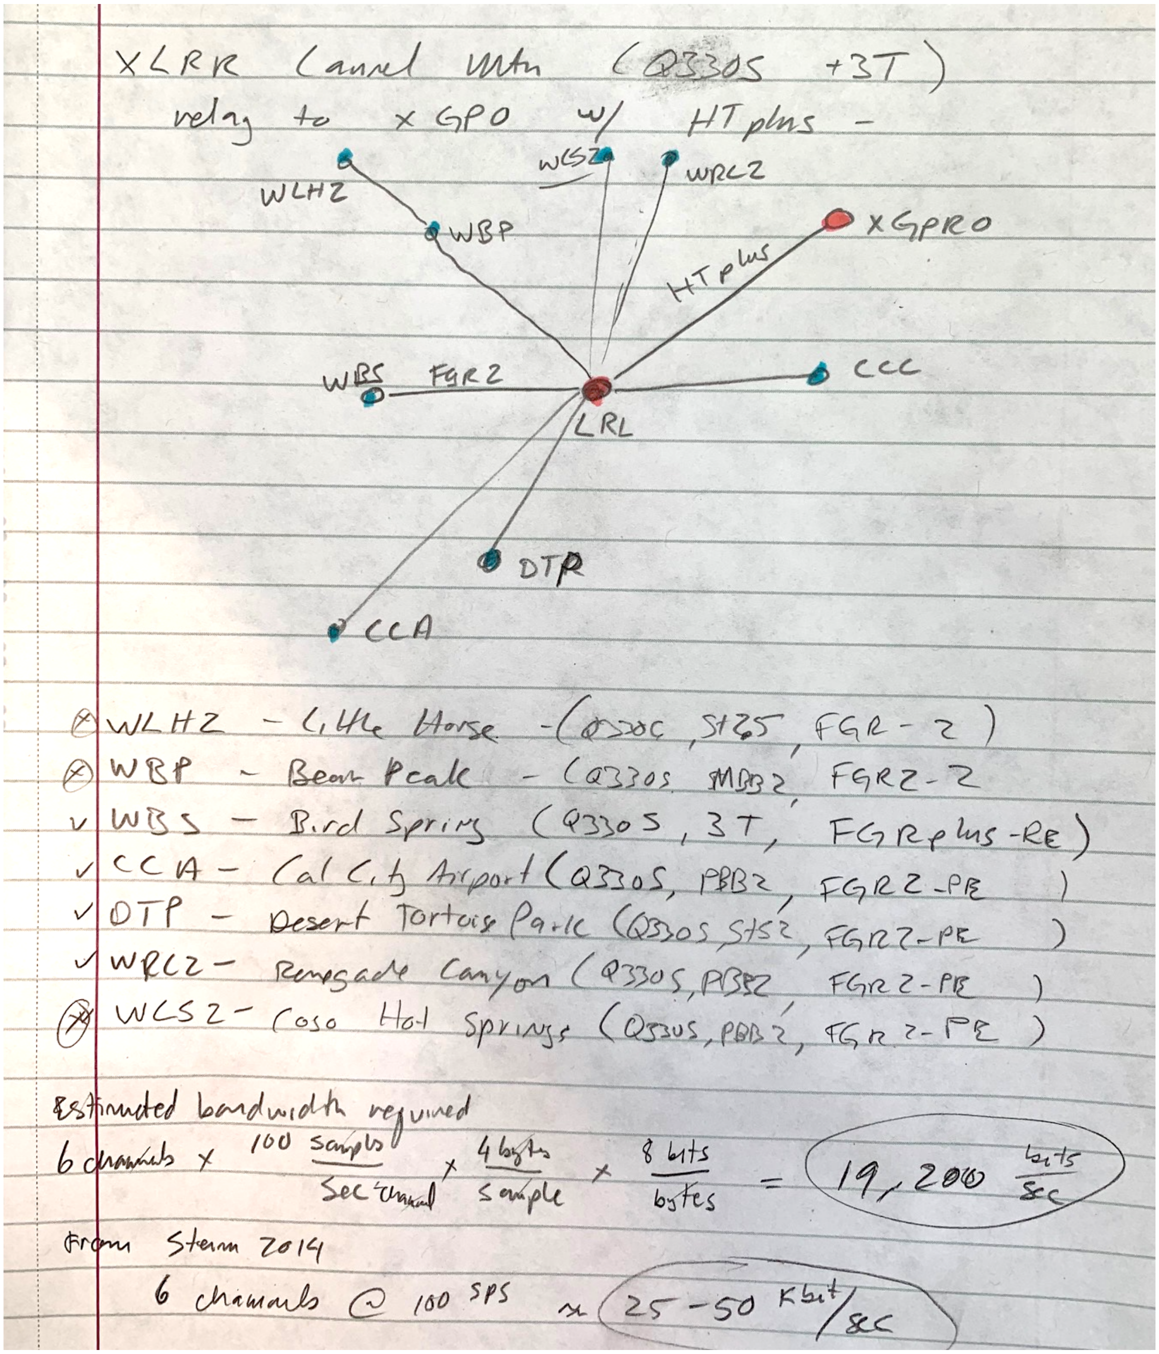

Fig. 8. Schematics of CI.XLRL to XGPRO data relays. The HTplus corresponds to a High Throughput radio installed to route data from CI.XLRL.

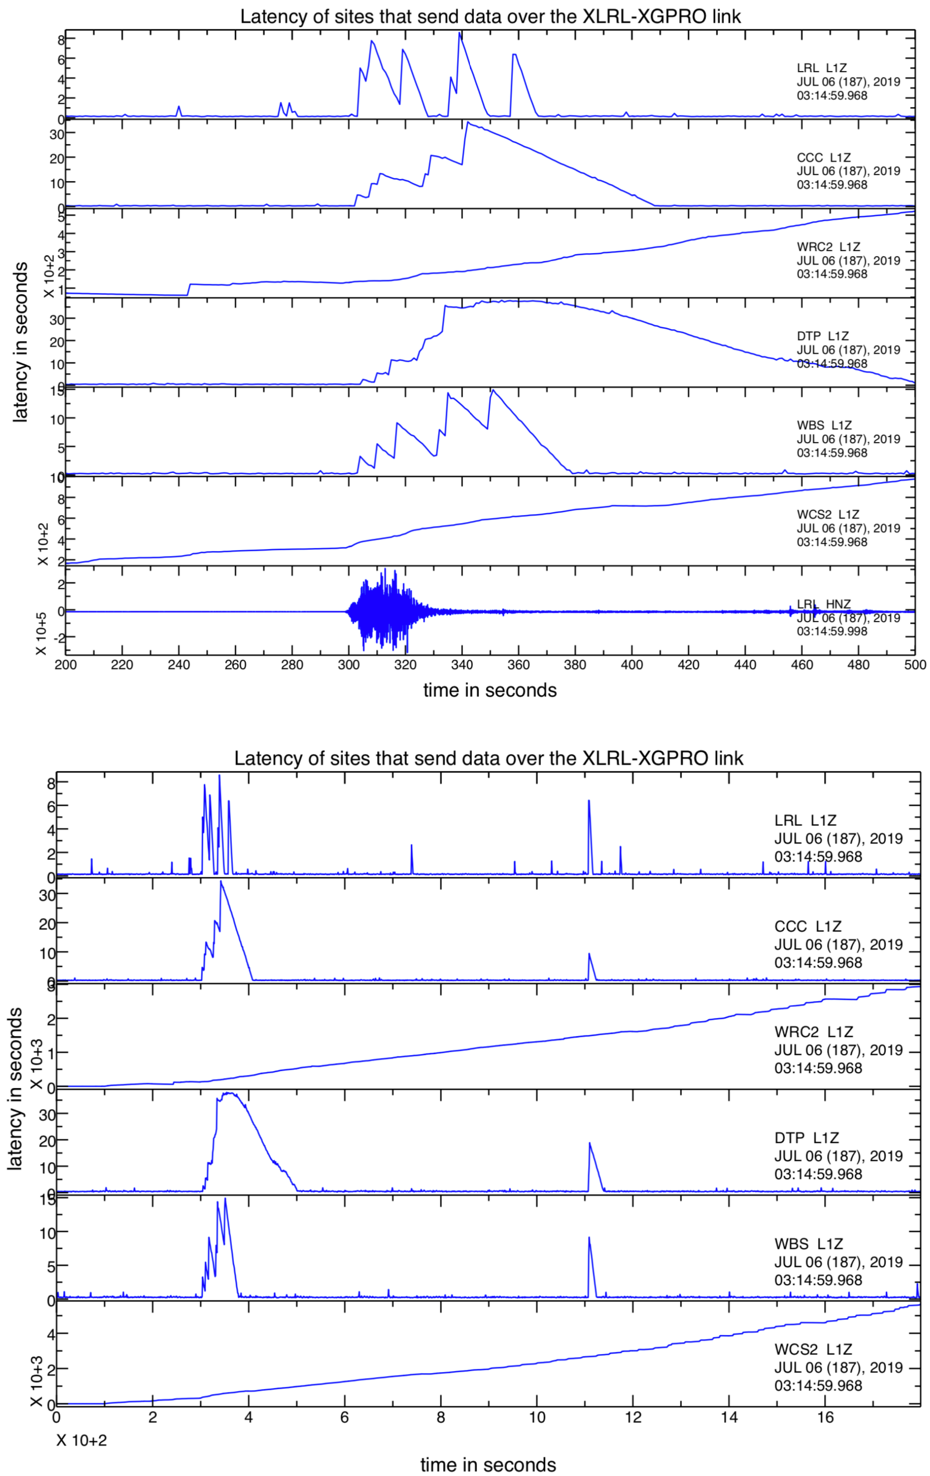

Fig. 9. Data latencies in two zoom levels for sites sharing CI.XLRL to XGPRO data link

Sites connected with cellular modem links

a) WMF+WRV2
WMF radios to WRV2 where they share a cell modem uplink. The stations are at the north end of the M7.1 rupture. Data are timely throughout the main shock. Latencies increase briefly to 5-6 seconds 100 s after shaking started. The L1Z shape suggests a cell outage of a few seconds. Latencies return to under 1 second after 10-15 s.

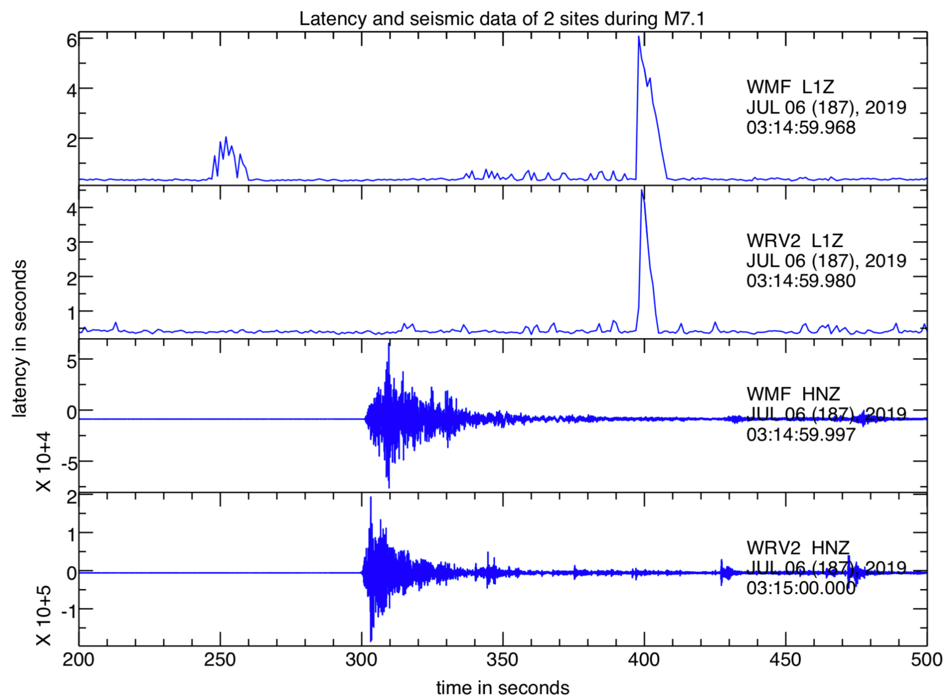

Fig. 10. Data latencies compared to seismic waveforms. WMF radios to WRV2 where the two share a cell modem uplink. Cell data are timely throughout the M7.1. A short duration latent period occurs after about 100 s.

b) CI.SRT. This station receives data from two other seismic stations and the three share a single cell modem uplink. Data are timely for 30 seconds (Fig. 11) after arrival of the main shock, then are steadily latent by 10-15 s. At 116 s after the main shock, some further change occurs, and the link begins losing ground. L1Z increases to over 25 minutes. The slow increase in latency indicates insufficient capacity and not a simple outage. Because the onset of serious latency began well after strong shaking had subsided, a cultural source (heavy cell usage or a cell-system reaction) is suspected.

c) ISA+WHF: Each station has a cell modem. Both are in the Sierras in sparsely populated areas. Latencies during the M7.1 are low and not affected by shaking. The connection with WHF was interrupted once for ~15 s and a second time for ~12 s, and recovers completely after 40 s. ISA has one interruption of 6 seconds, but recovers and delivers backlogged data in about 12 s, indicating that the link has substantial capacity beyond the steady-state demand. Confirmation of this assessment is described below.

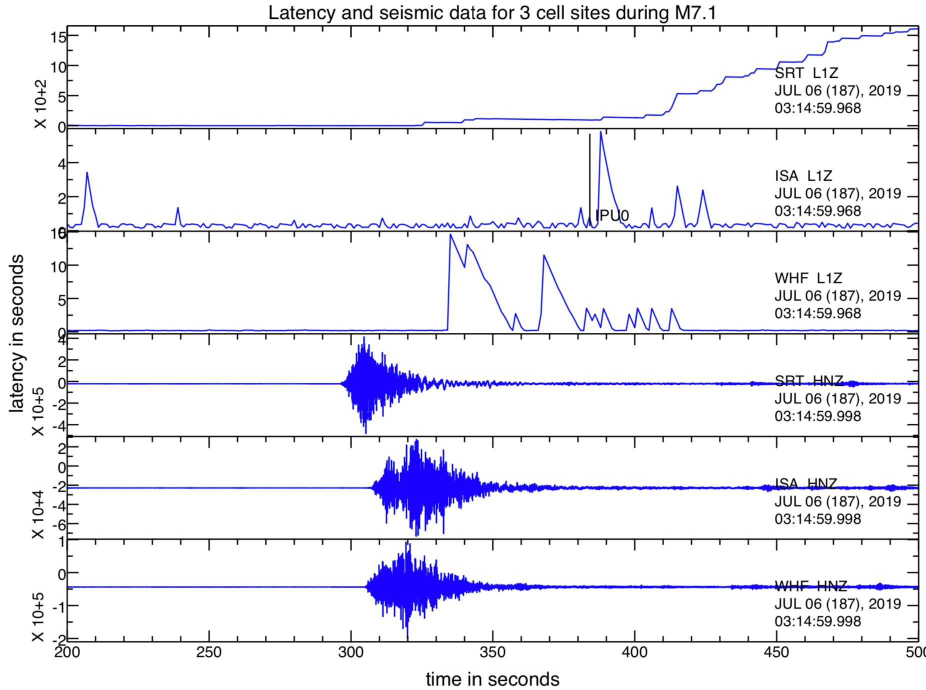

Fig. 11. Data latencies of three cell modem sites compared to seismic waveforms. SRT is located just NW of Ridgecrest, and receives data by radio from two other stations. Data remain timely for ~30 seconds before falling minutes behind real time. ISA and WHF have their own modems and do not share. ISA has a brief period of 5 s latency, but recovers quickly. The WHF connection was open for 12 s, and <40 seconds total with latencies >3 s.

d) We moved several large files back to back (over scp) from ISA and WHF to a computer at Caltech while they continued to record seismic data as a way to actively test their links. The procedure was performed on August 2, almost a month after the M7.1. No problems were found, and the bandwidth was consistently good at 1.7 Mbps. Unexpectedly we found that the seismic data latency for CI.ISA actually improved (Figure 12) while we were transferring our test data, in addition to the real time regular seismic 1-sec packets. The impact of simultaneous data transfers had little or no effect on CI.WHF latencies.

One possible explanation of the ISA improved latency could be the cell modem packetization implementation. It appears the packets need to be of a particular size before being sent. When an additional data was generated by our scp test, there was no need to wait as opposite to the seismic records of a smaller data volume.

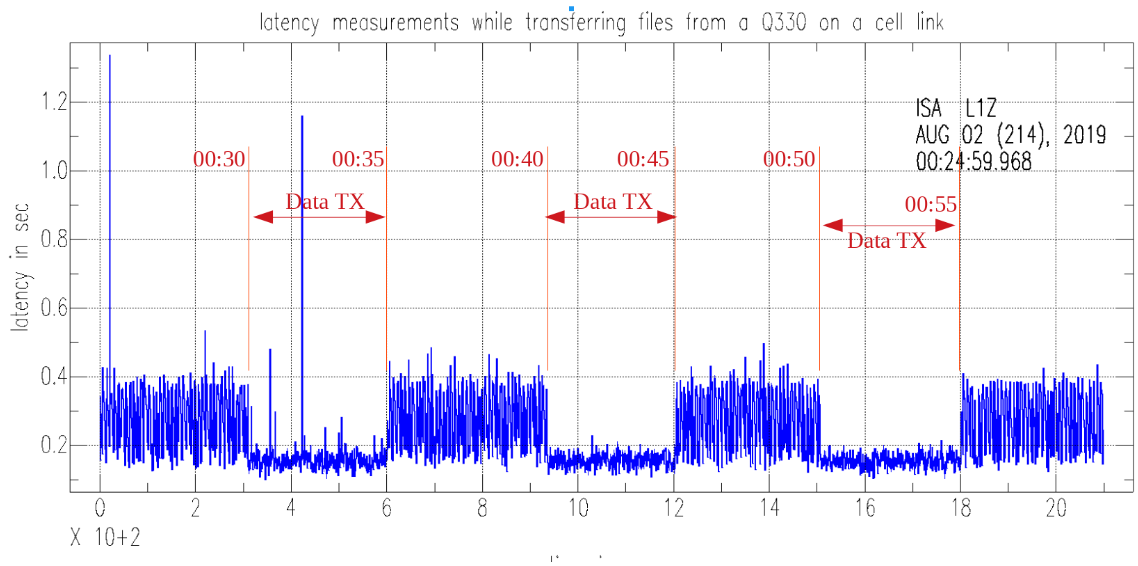

Fig. 12. Transport latency during the data transfer via scp from CI.ISA. Notice that the seismic data latency is actually *improving* during the additional data traffic.

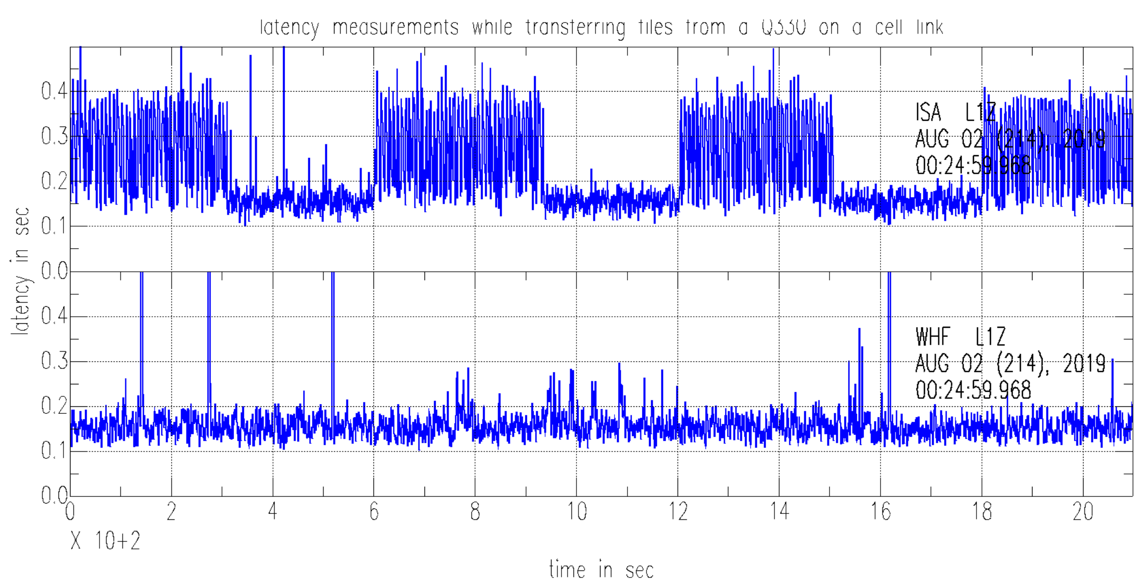

Fig. 13. Transport latency comparison during the data transfer via scp from CI.ISA and CI.WHF. The latency is *decreasing* for CI.ISA and had little or no effect on CI.WHF

e) Data throttling or competing for bandwidth or network outages

Figures 14 and 15 illustrate the timing of cell station data latencies relative to shaking across the seismic network. In Figure 14 seismic waveforms are shown for the M7.1 earthquake for the nearest and farthest 5 stations (for the plots of all 62 traces sorted by the distance and amplitude see the Figures 22 and 23). Seismic P-waves take ~50 seconds to reach the edge of the array. Figure 15 shows that the response of cell sites was variable. For station IMP there were no delays of even one second. Stations SWP and SBB2 may have had very brief interruptions on the scale of seconds. Stations SHO and AVM interruptions of a few seconds, but recovered quickly. A few stations (LMR2, WES, IKP, OLP) sustained longer L1Z increases. In L1Z plots (Figure 15) simple triangular shapes with an abrupt onset and linear sloped recovery are diagnostic of a sudden loss of connection, with the duration similar to the height of the triangle, followed by recovery given by the width of the base. Struggling connections recover but in a more saw-toothed pattern. Figure 15 shows that the onset of latency, where it is observed, is not keyed to shaking, but affects stations in a narrower window of time than the phase velocity of shaking (the P wave arrival moveout characteristic for the seismic data is not observed for the latency increase starting times). This indicates that the cause is not shaking and not telemetry related, but has a source that can affect the cell system more quickly and approximately at the same time.
There could be several potential explanations:
a) since the shape indicates a dropped or disrupted connection, one possible reason is *competition with users for cell tower access*;
b) the plots below may show that *Verizon could have throttled* our data flow, about 40-70 sec after the M7.1, when their cell towers detected in-network congestion, i.e., it seems the entire network must be under stress for this to take effect. They either reduced the general population bandwidth to give the rest to the first responders, or to protect the hardware or else... It would be interesting to meet with a Verizon rep and convince that our stations need a preferential treatment, as we learn more about what we experienced;
c) a *device at Caltech* (for example a switch, or an overloaded system) that was processing arriving packets and experiencing some issues. This item is included to keep the list complete but it was not probably a reason for the increased latency since the spikes happened over a long time interval, not instantaneously for multiple sites;
d) brief cell *network outage* (or outages) that caused the connections dropped completely for a few seconds. The dropped connection on a latency plot looks as a vertical straight line that gradually decreases to the pre-spike time: 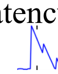 . For a link that can support the requested data rate, the line amplitude (duration of the outage) is approximately the same as the recovery time.

At this time, we don't have a strong evidence that would explain what was observed and more research is needed to complete the study.

For comparison, we show data and L1Z time series for an M4.7 aftershock on July 26th. Latencies are similar or somewhat smaller at most stations, and no shaking-related changes are seen.

Notice that the majority of the cell sites near the epicenter transferred all the quake data just fine and the L1Z jumps started about 50 sec after the P wave arrival, well after the major shaking happened.

Here is an example of the P arrival at 5 nearest and 5 far away cell sites compared to the latency increase on those sites. While the P wave arrives at different times, the L1Z increases approximately at the same time. This discrepancy cannot be explained solely by the telemetry effect.

The list of 62 cell sites (no VPN, GPS, radio, on Verizon, singular sites) where the L1Z increase is present (see also Figures 22 and 23):

AGM, ARV, AVC, AVM, BAR, BBR, BHP, CHI, CIA, CPT2, CRN, CSH, DGR, DNR, ELS2, ESI, FIG, FMO, FRK, GOR, IKP, IMP, IRM, ISA, JEM, JTH, LCP, LGU, LMR2, LMY, MLAC, MTA, MTP, MUR, NEE2, NOT, NSS2, OLP, PDM, PER, PTD, SBB2, SDG, SDR, SES, SHO, SLR, SLV, SMR, SMT, SMV, SMW, SOF, STC, SVD, SWP, TFT, TIN, VLY, VOG, VTV, WES

No L1Z increase, 29 sites:

TEJ, PUT, WLS2, SLM, FRM, HLN, RCU, MNO, AGO, WSS, FUL, CZN, IVY, LVY, QAD, GFF, PGA, MOR, SQC, GFS, CPO, SHU, NEN, FON, CLO, SAN, MIS, IDY, EMS

Top 5 traces (nearest cell sites, Figures 14 and 15): ISA, LMR2, AVM, SHO, SBB2
Bottom 5 traces (farthest cell sites): WES, IKP, IMP, SWP, OLP

To compare the latency for the same 10 sites much later after the M7.1, we plotted the same waveforms for a M4.7 from July 26 (Figure 16). There are no obvious L1Z issues present.

The M4.7 event used:
Time : 2019-07-26 00:42:47 (UTC)
Location: 35.926°N 117.707°W

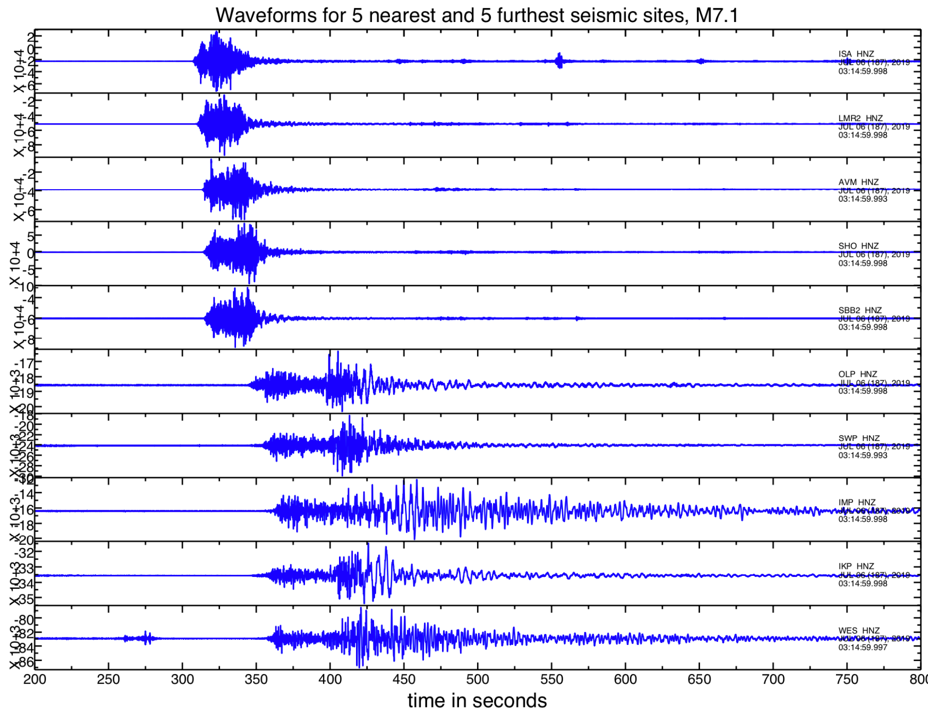

Fig. 14. Waveforms for seismic sites located approximately 300 km apart. The top 5 traces correspond to the stations located not far from the epicenter.

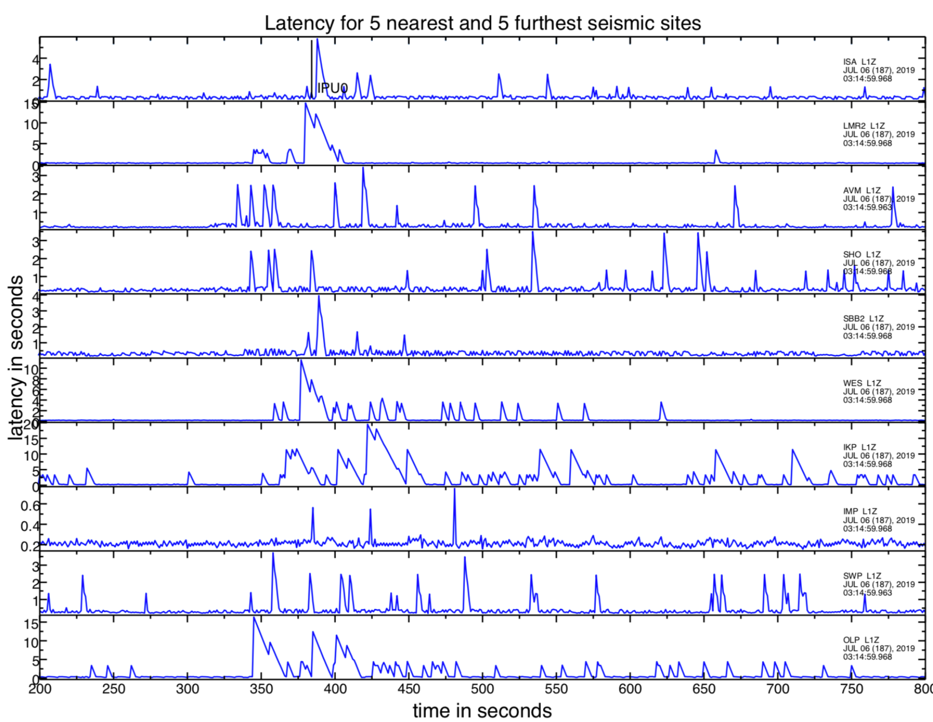

Fig. 15. Latencies for the sites from the Figure 14. There is no clear distinction between the spike times for the top 5 and bottom 5 locations unlike those present in seismic data.

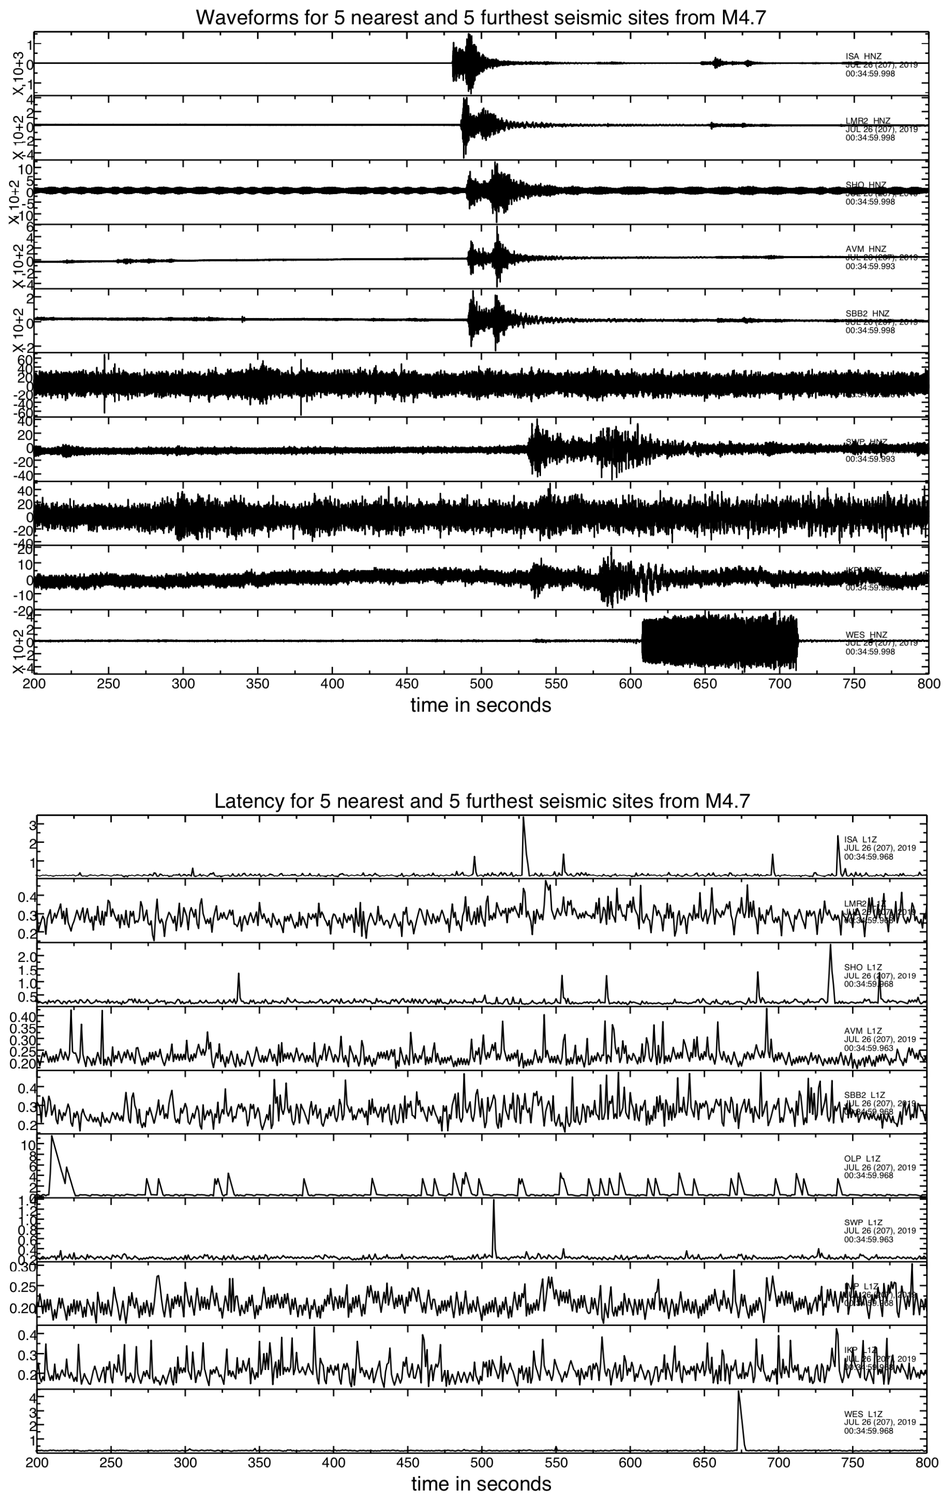

Fig. 16. Waveforms and latencies for the same cell sites located at different distances from an M4.7 that happened in the same location as the M7.1

f) The case where we show the benefit connecting to a CalOES tower (CI.LMS). The data experienced no latency issues. The other two sites, CI.CRN (cell) and CI.LMH had the latency issues over a long interval (110 and 50 sec respectively).

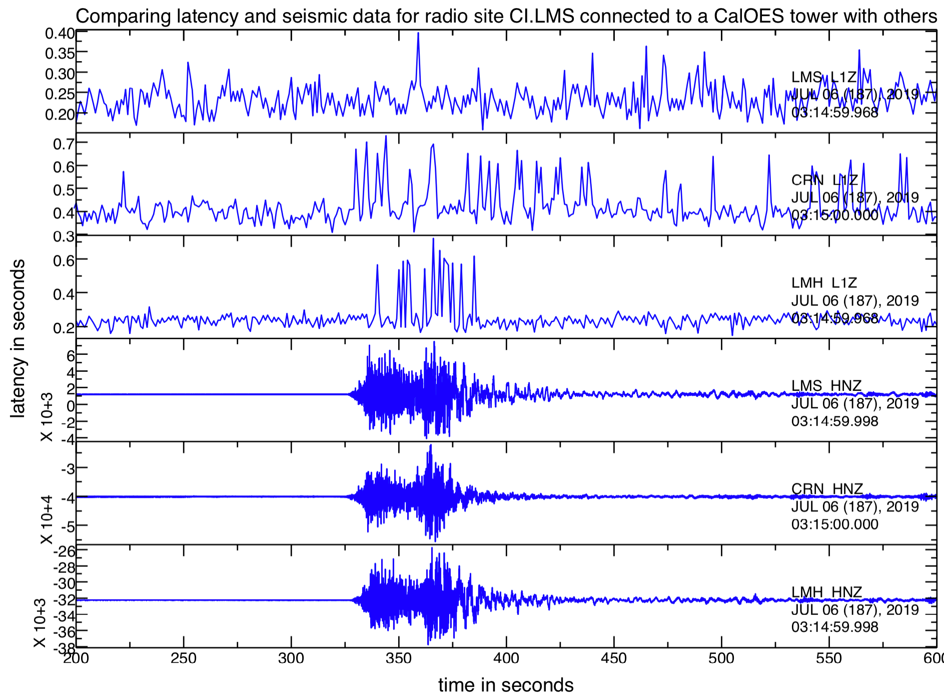

Fig. 17. CI.LMS is connected to a CalOES tower, it has very good transport latency.

Possible complications of seismic data delivery for the collocated with GPS sites

After the latency analysis, we noticed that many seismic sites that share telemetry with GPS stations, displayed the latency issues. The range of delays is from barely noticeable to a substantial. The scale depends on several factors but the distance to the epicenter and the telemetry quality would be the major reasons. Out of 66 seismic collocated sites, 31 have latency delay markers (or nearly 50%). To the opposite of the seismic data latency, the GPS times were not that heavily affected if any at all. One of the explanations to this phenomenon might a difference in data delivery methods. It seems that the tools/protocol of the GPS data delivery is much more greedy and it takes all the bandwidth it needs while the seismic data has to get by with whatever was left available (and in the case of a strong shaking, not enough to deliver seismic data seamlessly).

GPS and seismic collocated sites that share telemetry links with M7.1 latency issues (31 sites):

SNO, TOW2, WHF, WOR, RKMO, RAG, PSD, PHL, PALA, MSC, LRR2, LJR, KYV, IDO, HAR, FHO, DTP, CRG, CJV2, CJM, CHI, WAS2, SYN, RUN, POB2, MTG, LMH, LKH, GMA, ELS2, COA

The entire list of collocated (66 sites):

CRG, COS, COK, CJV, CJM, CHI, CAC, BUE, BOM, VTI, LVM, CUH, WAS, TPO, SYN, RUN, POB, MTG, LOM, LMH, LKH, GMA, EOC, ELS, COA, LRR, LPC, LMS, LJR, KYV, JNH, IDQ, IDO, HOL, HMT, HAR, FSH, FOX, FHO, ERR, ELT, DTP, DHL, TAB, SPK, SNI, SLH, SGL, RKMO, RHC, RAG, PVE, PSD, PHL, PAL, MSC, YUH, WWF, WOR, WMD, WLH, WHF, TOW2, THM, TAB, SNO

Data and packet rate increases for all types of telemetry

a) Fiber and Ethernet displays the rate increase but not the latency increase.
Please see below 2 figures for station CI.PDU that is delivering data to Caltech over a fiber line (no antennas involved, we wanted to eliminate a possible packetization overhead when many devices are involved on a telemetry link and telemetry bandwidth causes packets resends). The M7.1 quake happened at 20:19 local time (as on the upper plot), or 3:19 UTC (lower seismic plot).

The upper plot is displaying the data rate in Bytes/sec (with the Y axis on the left) and packets/sec (with the Y axis on the right). The rates are measured every 100 sec with qmaserv and are showing the total rate of all the data arriving at Caltech from the Q330. As you can see, the data rate went up near the P wave from ~1200 Bps to ~2600 Bps then lowered to ~1500 Bps and remained at that elevated level for long time after the major event was over.

The lower figure is displaying the latency L1Z in seconds and HHZ/HNZ channels corresponding to the orange time interval box from the data rate plot. For the 400 seconds shown, the earthquake causes no material change in latency. We think this means that telemetry was consistently good. We expected the packet and data rates to be constant on a solid communication link.

For CI.PDU, the data rate increase is more than 2x compared to the pre event rate. Please note that the rate values plotted are the means per 1 min, the actual instantaneous values would be much higher, above 2x (this is how the data rate was setup to collect, we changed it to every 5-20 sec already).

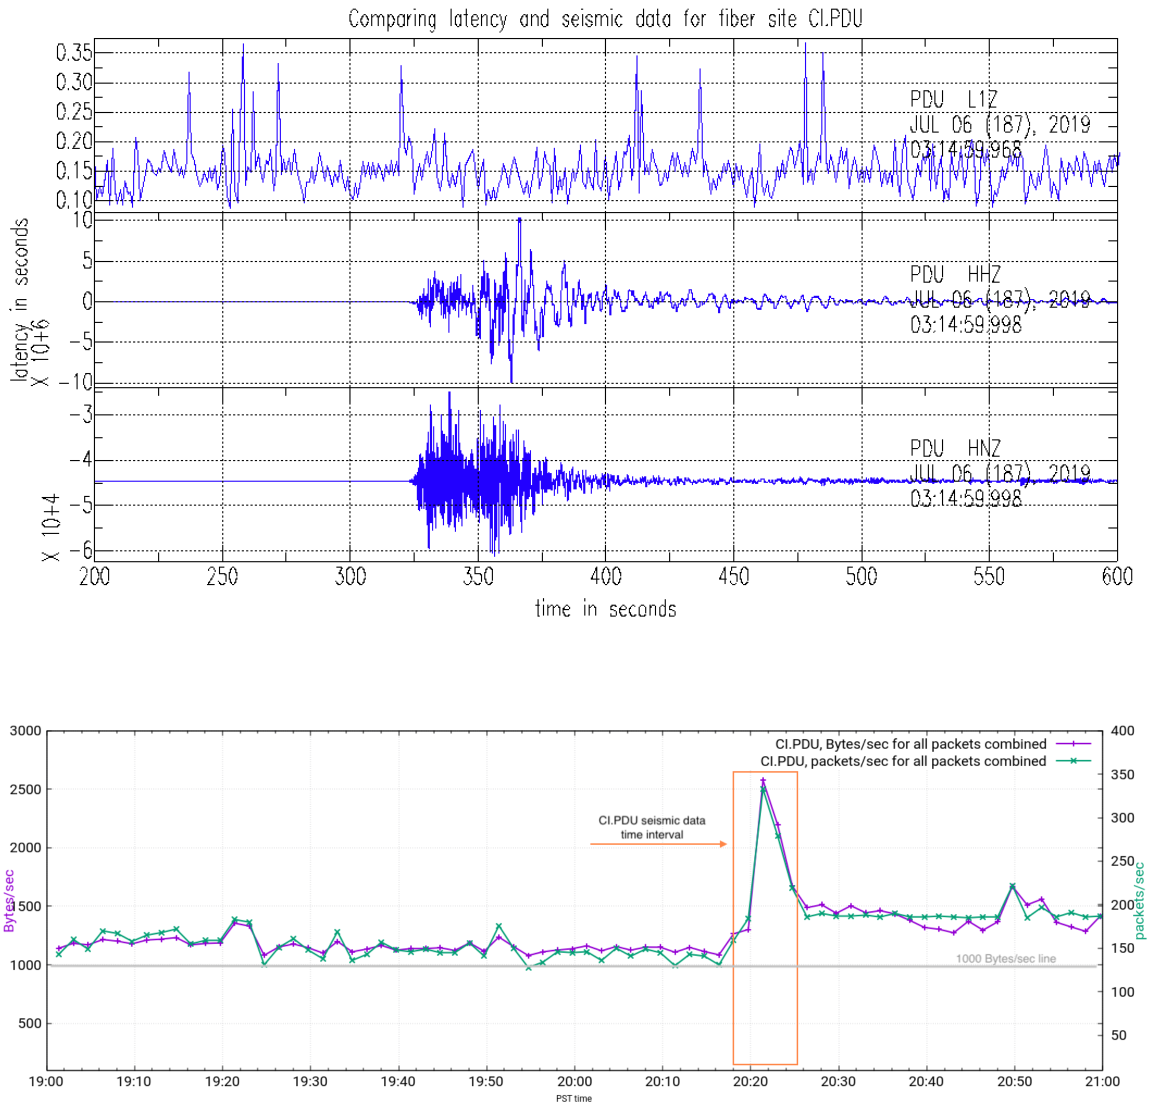

Fig. 18. Latency, seismic data, data rate, and packet rate for CI.PDU. The x axis time interval on the blue waveforms in the upper plot corresponds to the orange box in the lower figure.

Additional examples

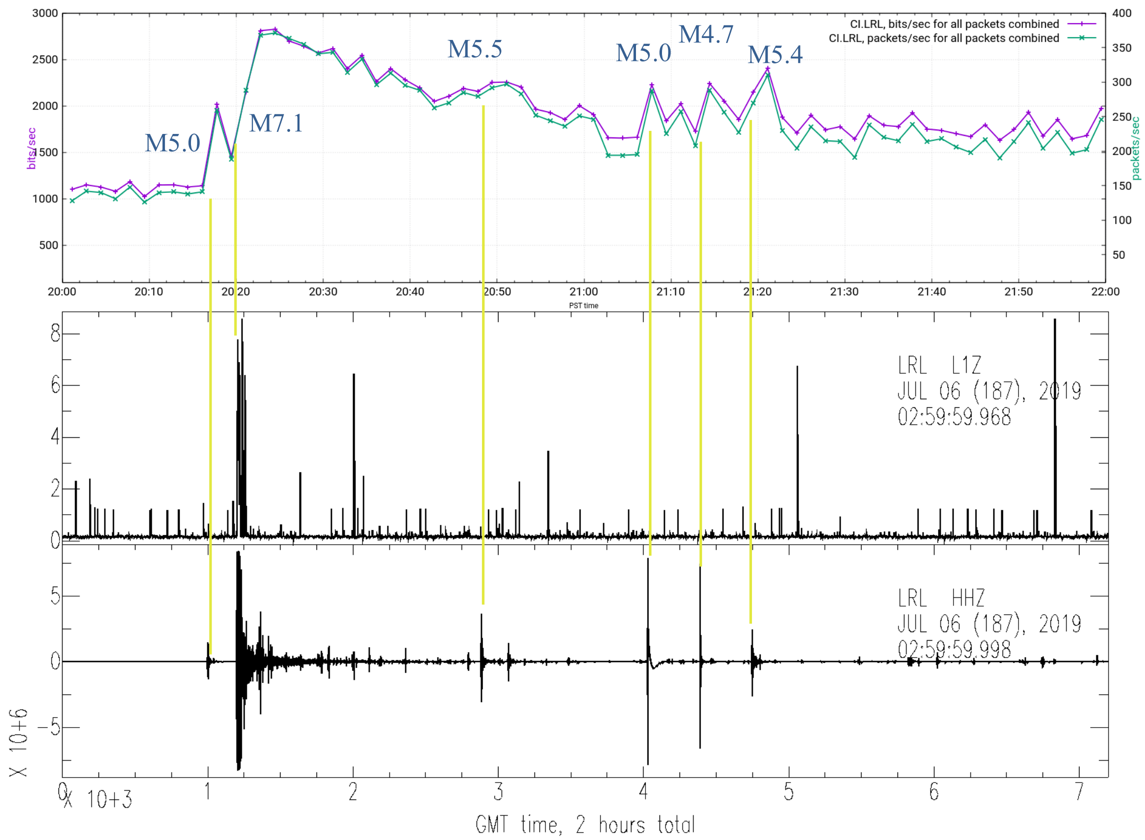

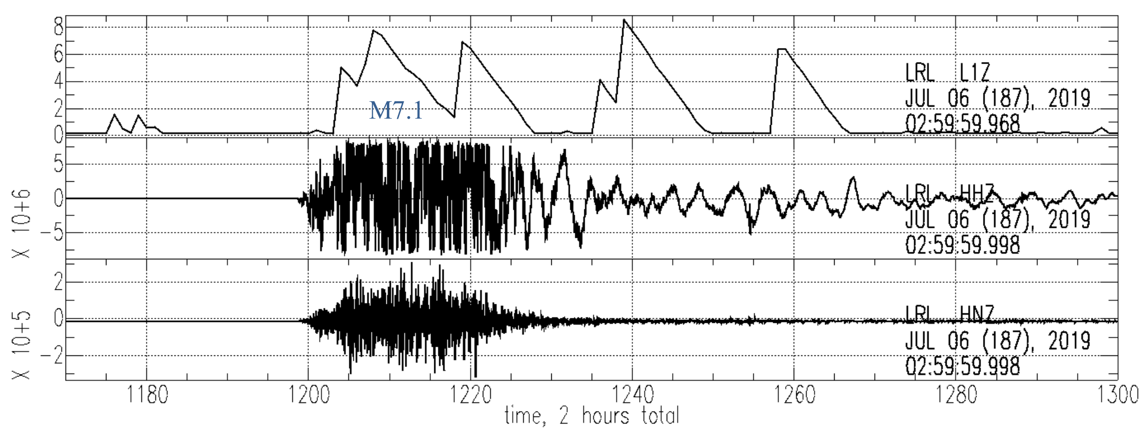

Fig. 19. Comparison of data rate and number of packets increases, seismic data, and L1Z latency during the event sequence M5.0, M7.1, and other significant events that follow (top) and only during the M7.1 (bottom). The rate results are based on the qmaserv polling of the Q330 dataloggers every 100 sec (obtained from cs-import1)

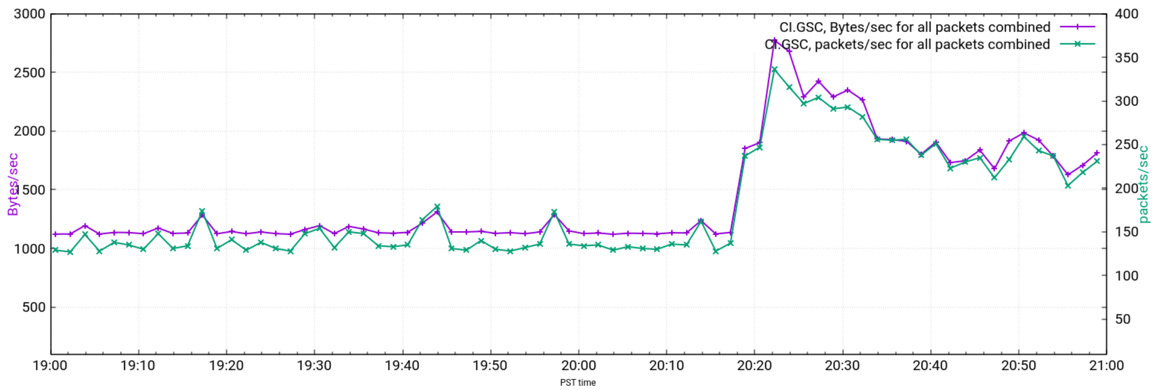

Fig. 20. Data and packet rates for a different site, CI.GSC, delivering data over Ethernet. Note two different Y scales for data rate (left) and number of packets (right). The rates are per second, calculated as averages per minute, increase up to 3 times.

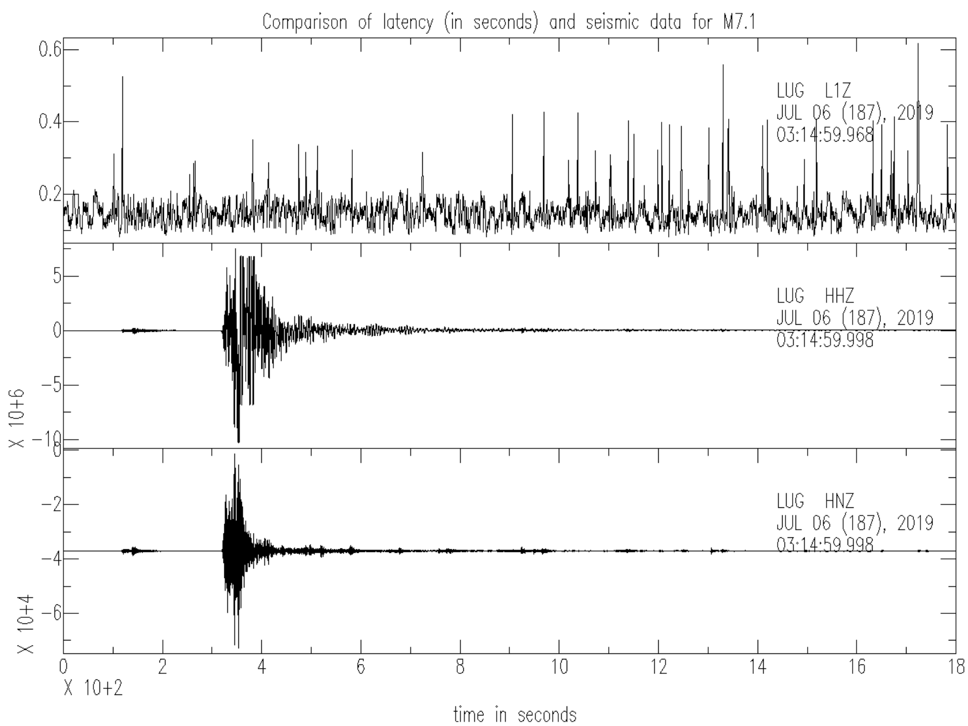

Fig. 21. Comparison of latency and seismic data for a fiber connected site. The sites similar to CI.LUG (no antennas involved, CI.ADO, CI.GSC, CI.APL, CI.CTL, CI.HLL) didn't show any latency issues.

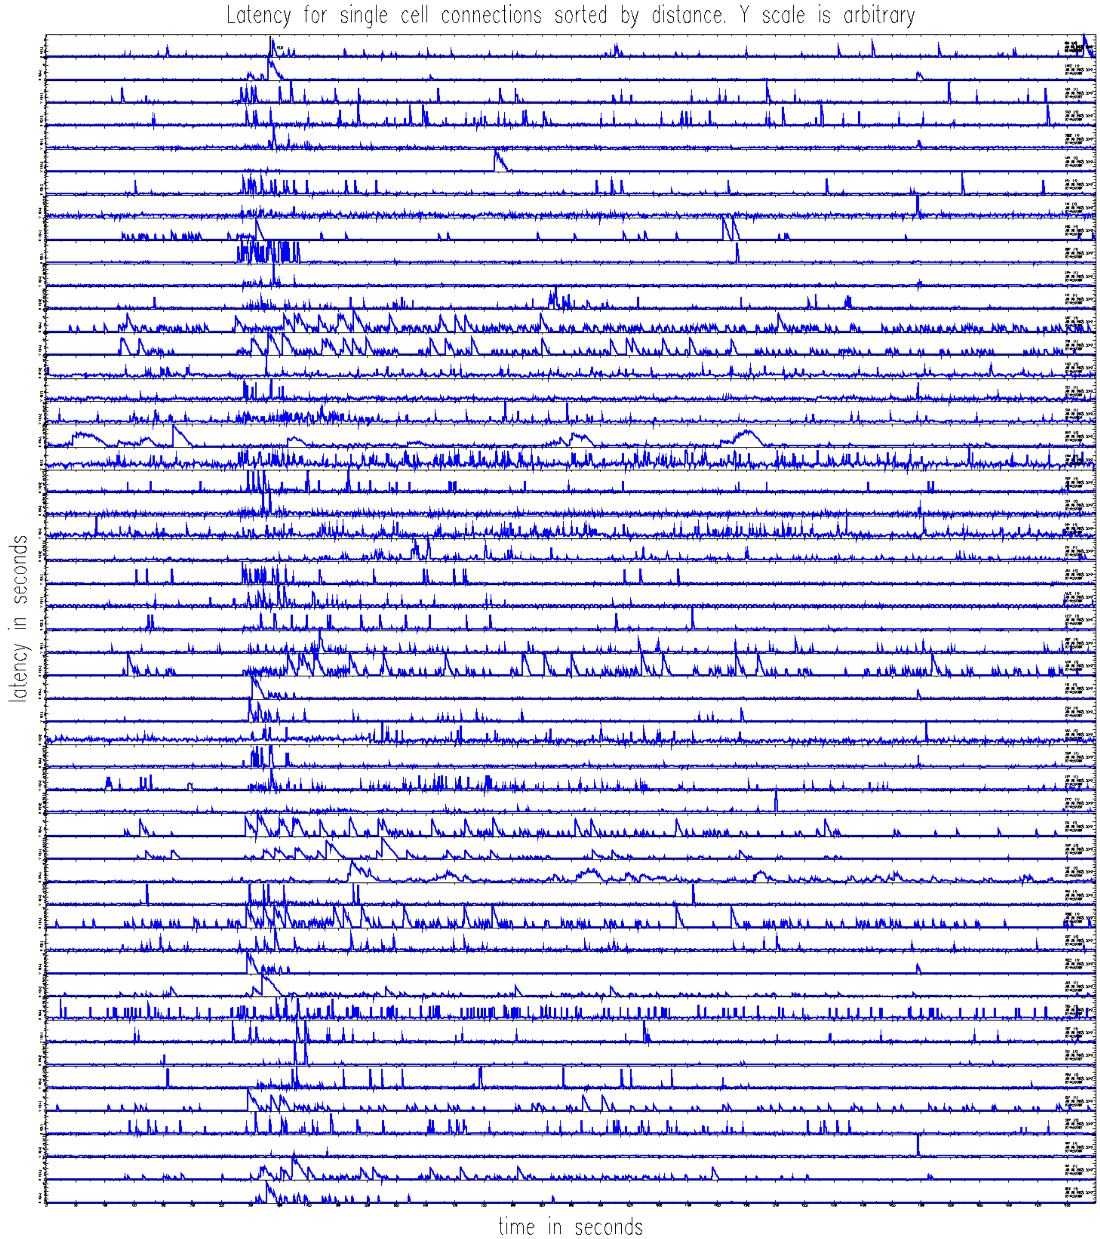

Fig. 22. Latency records for 62 cell sites that have no VPN, collocated GPS or radio stations and use Verizon service. The data is sorted by the distance from the M7.1 epicenter. The nearest site is at the top, the distance to the farthest site is approximately 300 km. Please notice that the service disruption (visible as the latency increases) happened almost at the same time in the entire Southern California.

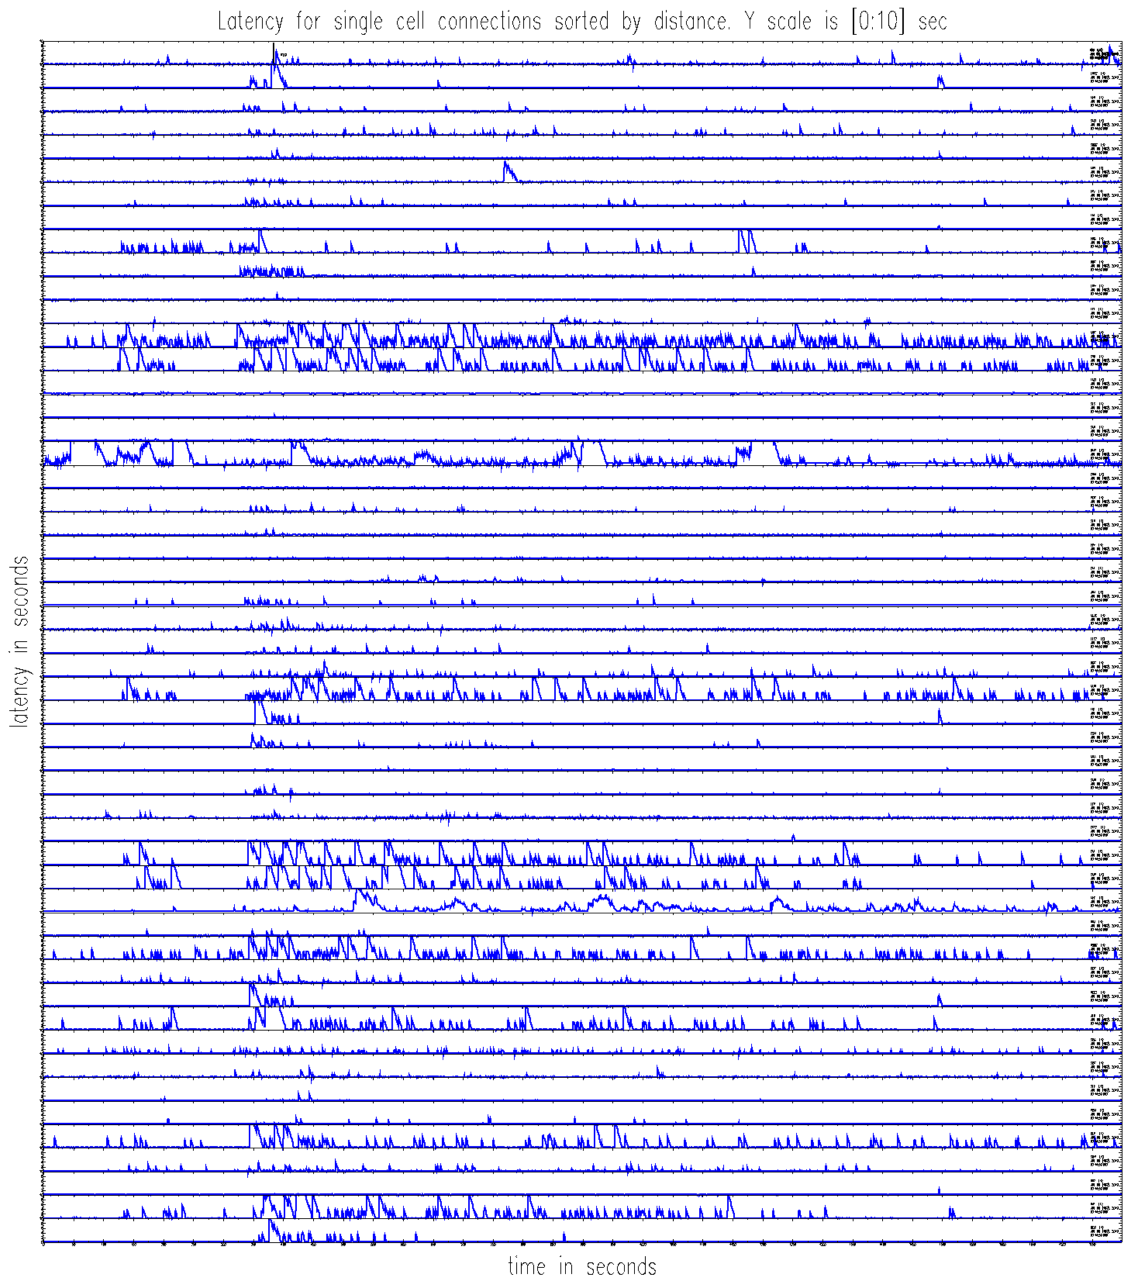

Fig. 23. The same latency data sorted by the distance from the epicenter as in the Figure 20 but the Y scale is now fixed at [0:10] seconds. We show that the latency increase has no correlation with the distance.

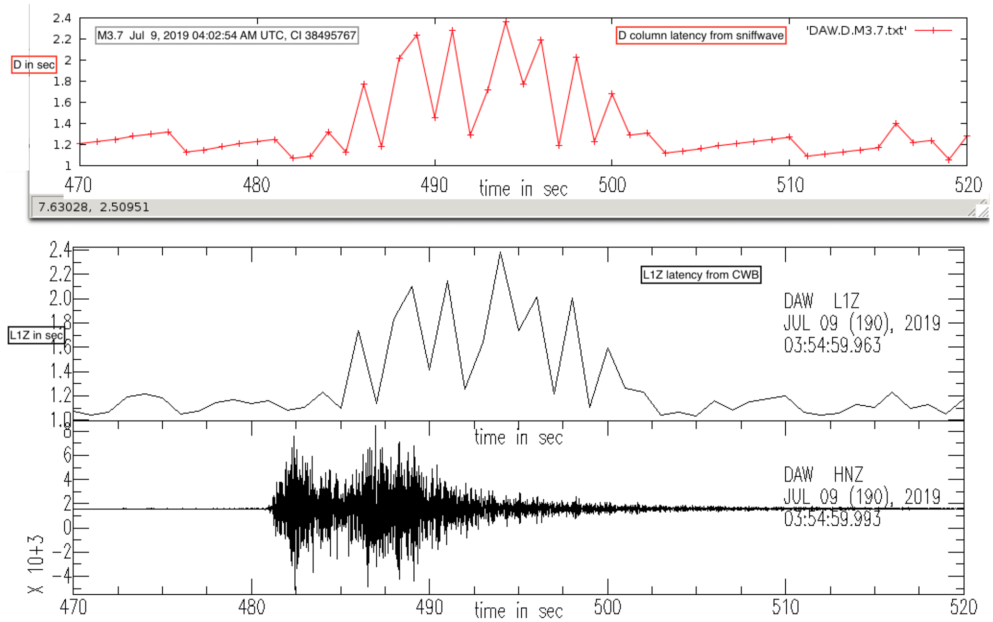

Fig. 24. Latency increase for a VSAT site during a M3.7

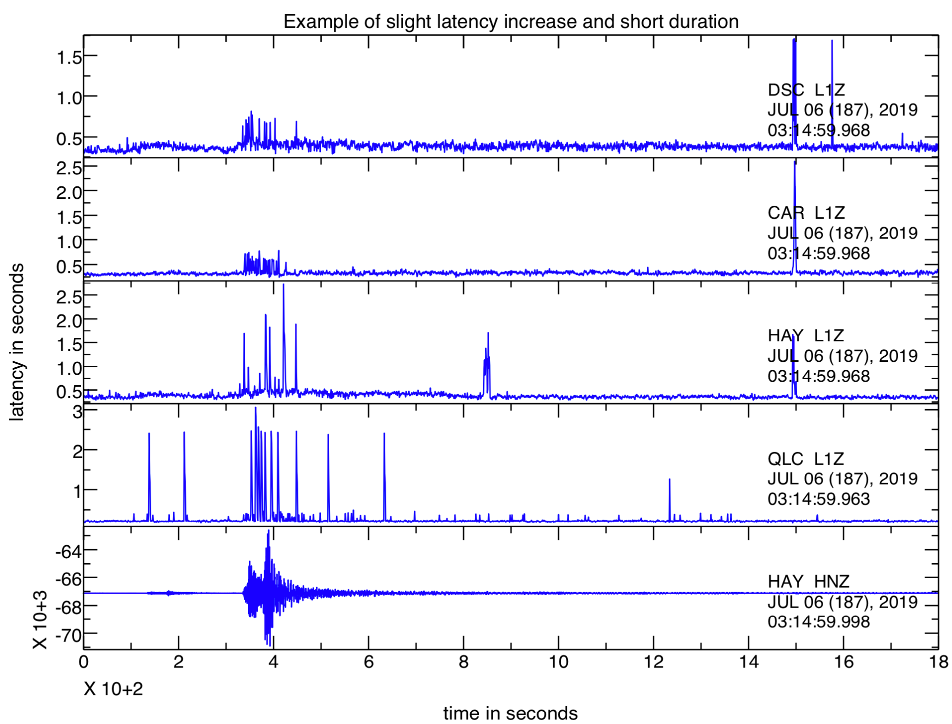

Fig. 25. Slight latency increase and short duration

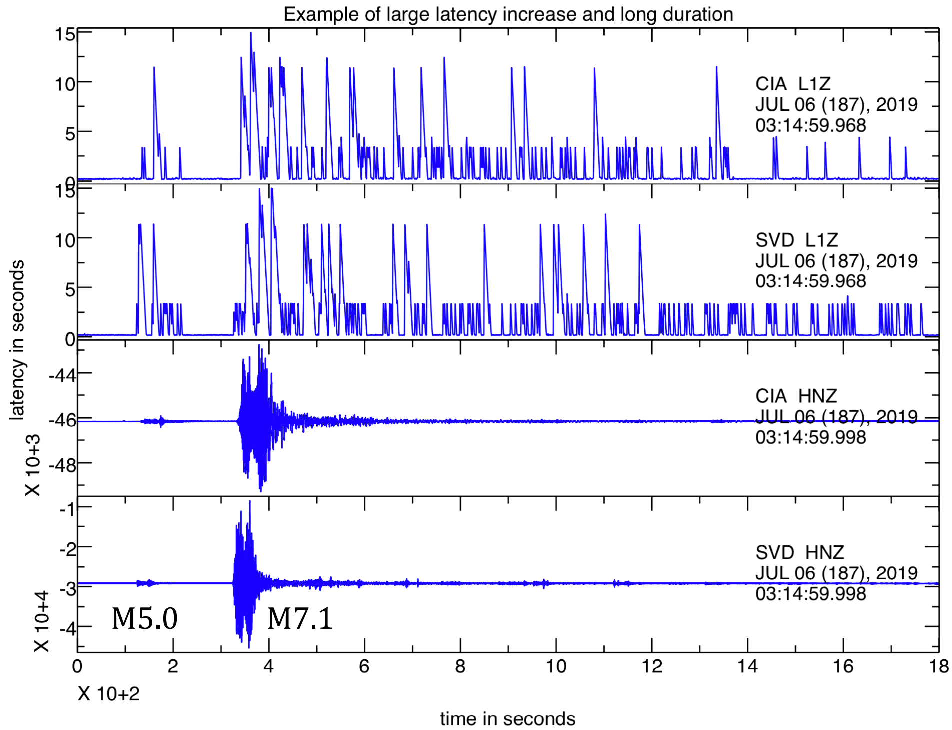

Fig. 26. Large latency increase and long duration, an M5.0 had the impact as well

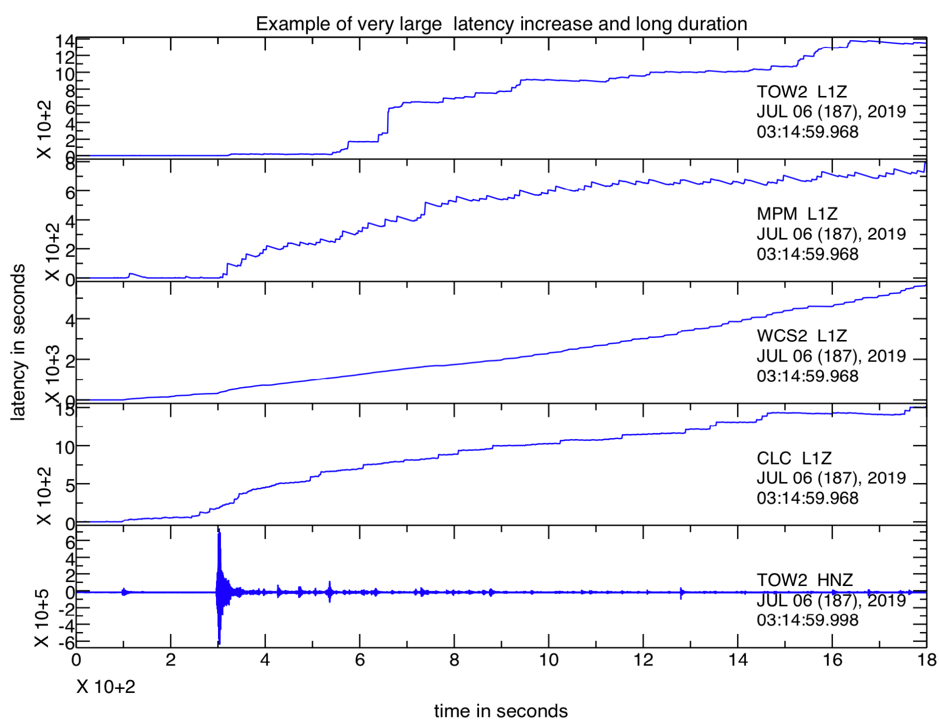

Fig. 27. Very large L1Z increase and long duration

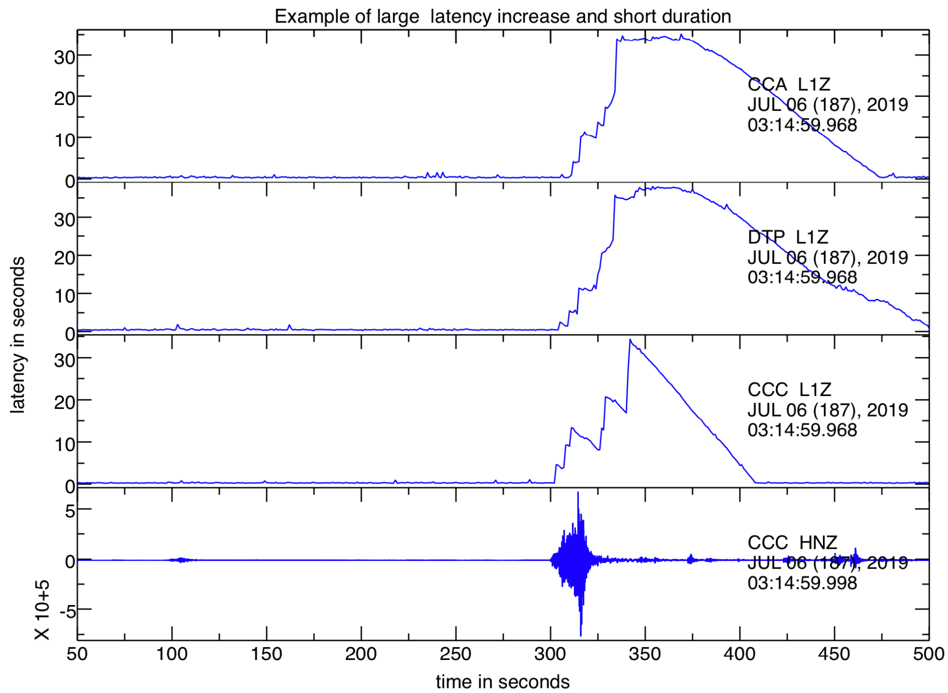

Fig. 28. Large L1Z increase and shorter duration

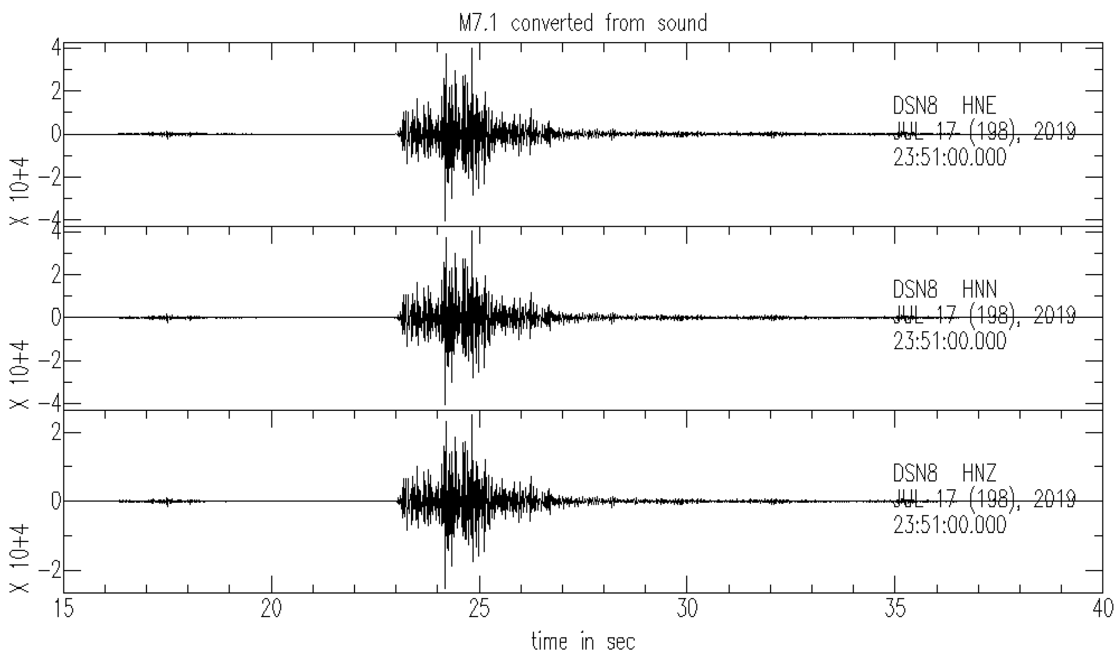

Fig.29. Converted to sound M7.1 collected on the CI.DSN8 Basalt. The same waveform used for all 3 channels, the amplitude is comparable to the actual seismic waveform (the amplitude was regulated by the volume knob and was set to almost the maximum).

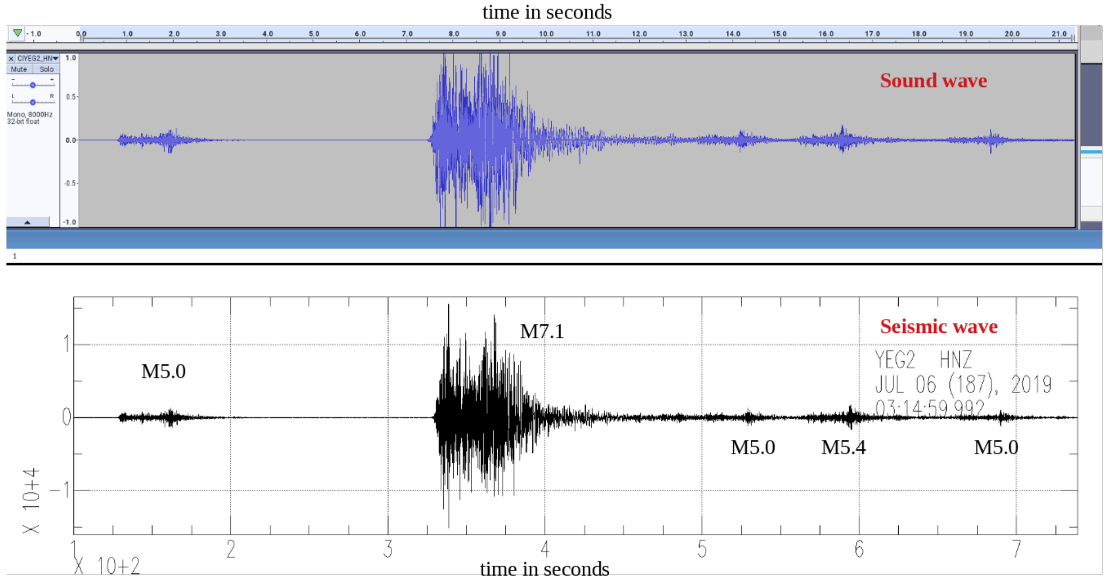

Magnitude: 7.1
Time: 2019-07-06 03:19:53 (UTC)
Location: 35.770°N 117.599°W
Depth: 8.0 km

Fig.30. Sound wave and the actual seismic HNZ channel that was used for conversion.
To listen: *http://crust.caltech.edu/~igor/events/Ridgcrest7.1/sound*

References

1) Timeliness of Data Delivery from Q330 Systems, *J. M. Steim; R. D. Reimiller, Seismological Research Letters* (1 July 2014) 85 (4): 844–851

2) Data Latency and Compression, *J.M. Steim, E.N. Spassov* , IRIS poster , obtained via personal communication,
http://crust.caltech.edu/~igor/events/Ridgcrest7.1/L1Z/LatencyAndCompression_2016.pdf

Appendix

Email correspondence with Joseph Steim

We asked Joseph Steim several questions about why the data rates were increased during the events. Please see below his response.

Questions:

…

What explanations should we consider for the data and packet rate increases?

What is the nature of this symptom?

What are the expected maximum limits of the increase? (In your 2014 SRL paper you mentioned a sustained data rate of 50 kbps = 6250 Bps, is this still valid?)

We are wondering if either origin queuing (OQ) or transport layer packet aggregation (TLPA) may be responsible for this increased latency as described in your paper?
I would like to mention our latency relevant Q330 settings during the event, just in case they matter:
Group_Timeout=0.2, Group_Size=0; ACK_Timeout=0.3, ACK_Grouping=4.
We are in the process of changing ACK_Timeout=0.2, ACK_Grouping=1 as you suggest in your paper.
…

Response:

Dear Igor

These results are expected and do indicate as you suggested that your telemetry bandwidth is sufficient for the traffic.

The average bit rate, including the number of packets/s, is expected to increase during an event because the data are less compressible. The Q330 sends data in 1s aggregates. An aggregate comprises one or more packets having a maximum fixed payload of around 500 bytes. Packets may be less than 500 bytes if that volume is not required. Recall that the Level 2 compression used in the Q330 comprises several object sizes to store 1st differenced data (see the attachment). First-differenced quiet data typically requires no more than 8-bit objects to represent each sample. Therefore at 100sps, with 3 axes, one second of data would fit into a single packet of maybe 350 bytes. However, event data is both larger, and less compressible as it contains higher frequency content. The object size necessary to store each sample of first-differenced data may increase to 15 bits or briefly event possibly 30 bits. This may effectively double the number of bits required during the largest amplitude part of the event, and may double the number of packets. You can see also in the attachment in the figures in the center showing "64-byte frame latency". These figures show the corollary of data rate increase during an event. These figures show that the time duration of a fixed *volume* of data decreases during an event by a large amount – similarly corresponding to the interval of reduced compressibility. As soon as the largest amplitudes subside, compressibility improves and data rate decreases. The smaller the event, and the farther from the source, the smaller is the increase in data rate accompanying the event.

So what you're seeing is expected from a link with sufficient telemetry bandwidth. If latency *were* increasing during the event, and the data rate *were not* increasing, that would indicate that your telemetry is limiting the transmission bandwidth. In that case, OQ, would add latency because of the inability to transmit data at least at rate at which the data are generated.

This example is a good observation indicative of a healthy telemetry system.

FYI – the "flooding" capability allows you to assess bandwidth "headroom" at any time. When "flooding" is briefly enabled, the Q330 attempts to send as much data as possible through the link without inducing OQ delays. "flooding" essentially fills any void time with empty data packets. "flood" packets are also totaled by qmaserv. An estimate of link rate capacity can be obtained when the rate of flood packet reception is added to the rate of actual data packet reception. If a link is not capable of sustaining a rate of flood packet reception at least 2 times or so the rate of actual data packet reception, the link has insufficient headroom to respond to the increased data rate expected during a large event. I suggest that you use flooding according to some plan to assess telemetry capability.

Cheers
Joe