D2-01_08

Evaluations from accurately acquired and time stamped traffic to assess communication networks for IEC 61850 applications

by

Fred Steinhauser*, OMICRON electronics
Eugenio Carvalheira, OMICRON electronics

AT

SUMMARY

For the proper function of distributed applications in IEC 61850 systems, the timely delivery of the communicated data between the devices is crucial. The timing performance requirements vary from application to application, but there are always limits to be kept.

Once having Ethernet packets acquired from different acquisition points in such IEC 61850 communication networks, several evaluations regarding the time behaviour of the networks are possible. The primary precondition for meaningful evaluation and assessment of timing parameters is the accurate acquisition and time stamping of the packets. Typical PC based solutions without special capturing hardware cannot provide the required accuracy for most measurements. Time synchronized acquisition units make such measurements possible even in distributed scenarios. The usage of the Precision Time Protocol (PTP, IEEE 1588) provides an accurate and common time source for the devices involved.

In communication systems with redundancy, it is of interest how much the timings in the different redundant communication paths match each other or are different as well. In HSR systems, such timing differences are to be expected depending on the position of the senders and the receivers relative to each other in the ring. In PRP systems, the propagation properties of the redundant networks may as well be entirely different, depending on the topology and equipment applied.

KEYWORDS

IEC 61850, Ethernet, Local Area Network, HSR, PRP, PTP, Wide Area Network, Propagation Time, Delay Time

*   OMICRON electronics GmbH, 6833 Klaus, Austria
      Email: fred.steinhauser@omicron.at

1. MISSION CRITICAL COMMUNICATION NETWORKS

In the early days of the use of communication protocols in power automation systems, this was only used for SCADA purposes. First of all, the media and protocols used in the past were comparably slow, their performance would have been insufficient for applications as those implemented with today's fast network technologies. And secondly, the reporting of e.g. a switch position to be displayed for a human operator is not time critical. The person in front of the operator's screen will not notice if the message had been delayed by a couple of hundred milliseconds. Even if the status indication did not come through at all, the protection's function did not depend in this and was still be able to clear a fault and maintain the safety of the network operation.

With the communication network taking over all the tasks of the formerly hard wired connections, it becomes a mission critical component. The hard wired signal loops become replaced by GOOSE messages, that exchange the information between the protection relays and also towards the now intelligent switch controllers. Sampled Values deliver the data from the instrument transformers to the relays and meters. When all these elements of the so-called digital substation are utilized, a dedicated network for the real time data, often called a process bus, will be present.

To harden the process bus against failures, selected and reliable components are used and possibly redundancy is applied as well.

2. THE TECHNICAL REPORT IEC 61850-90-4

To support engineers in the design of communication networks for power utility automation, the technical report IEC 61850-90-4, titled "Network engineering guidelines" was published [5]. On more than 200 pages, it gives insight in all relevant aspects.

Of course, terms like performance and latency are frequently mentioned, e.g. in the context of the different network topologies.

3. MEASUREMENTS IN LOCAL NETWORKS

In a local network, such measurements can be relatively easy performed when it is possible to reach all locations from which traffic is to be captured from one single measurement device. The relative timing of the data packets captured with the device is always accurate and no additional time synchronization is required.
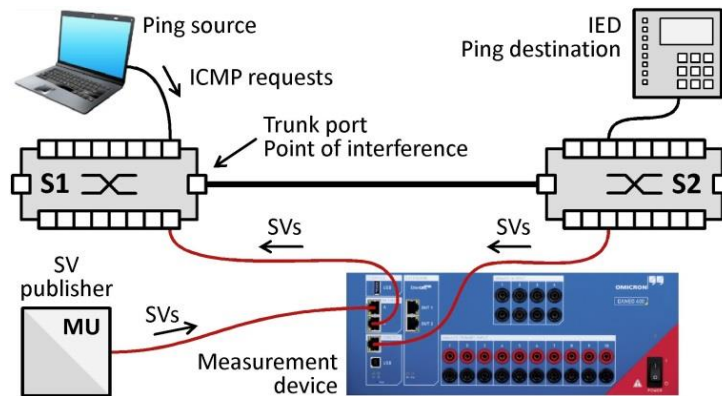
Figure 1: Possible measurement setup in a LAN

Figure 1 shows such a measurement setup for a minimal network with two switches S1 and S2 that are connected by a trunk link. A Sampled Value source (a merging unit) publishes one Sampled Values stream into the network. The measurement device captures the Sampled Values packets coming from the Sampled Values source before they enter the network at switch S1 and then again when they are broadcasted from the other switch S2, after traversing through the network. A PC connected to switch S1 generates load traffic by "pinging" an IED which is connected to switch S2. This forces ICMP messages to be exchanged over the trunk link and therefore interfering with the Sampled Values. The ping utility used for generating this load traffic allows specifying the size and the frequency of the ICMP packets.

In the used setup, all Ethernet links, also the trunk link, operated at 100 Mbit/s. The ICMP packets were be issued at a rate of about 1000 packets per second. At a packet size of 500 bytes, one ICMP packet occupies the network for 40 μs and the ping packets utilize 4 % of the total bandwidth.

The next figure shows the delay time distribution of the Sampled Values under the described conditions.
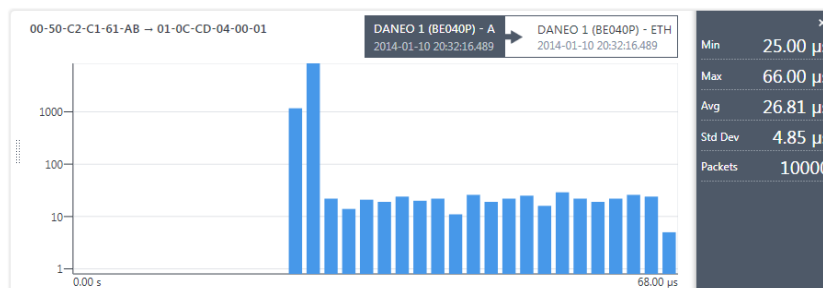

Figure 2: Delay time distribution with ICMP packets (size: 500 bytes) as network load

The left two bars, that contain most (96 %) of the measured packets represent the delay times between 25 μs and 28 μs. This is the time needed for a packet to traverse the little network undisturbed. The bars to the right, extending to delay times with a maximum of 66 μs, come from packets that got in conflict with one of the ping packets. The maximum is exactly 40 μs (the duration of the ICMP packets) higher than the average value for an undisturbed passage. This result matches perfectly the expectations that can derived from the theoretical examination of the situation.

This demonstrates how other traffic can for example influence the transfer times of mission critical information in such networks. A theoretical evaluation of the circumstances will not always be as easy as with this minimal setup with its few and well controlled parameters, but a measurement will always show the actual circumstances in a given installation.

4. MEASUREMENTS IN NETWORKS WITH REDUNDANCY

The two redundancy methods proposed for use with IEC 61850 are HSR (High-availability Seamless Redundancy) and PRP (Parallel Redundancy Protocol) [6]. The assessment of the propagation times in such architectures and the differences of the timings in the redundant paths are of particular interest. Also, it must be verified that not packets are lost when a switchover between the redundant paths takes place.

HSR packets can be identified by their distinct Ethertype and by stripping the HSR tags from the captured packets, the packets are converted to "normal" Ethernet packets that can be compared against those in regular Ethernet networks.
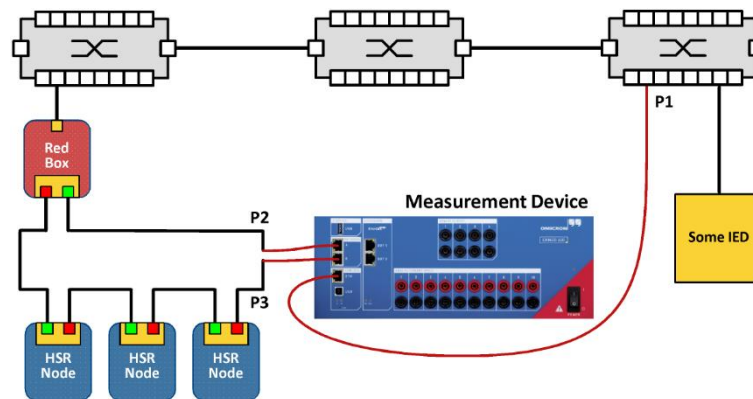

Figure 3: Measurements in a mixed HSR / Ethernet network

The setup shown in Figure 3 was used at the UCA IUG demonstration at the CIGRÉ Session 2014 in Paris. Traffic from the Ethernet network was captured at P1 and traffic from the HSR ring was captured at P2 and P3. The HSR ring remained intact, because the measurement device only tapped the traffic by still forwarding all packets. The timing from P1 to P2 showed how fast the packets were forwarded through the Ethernet network and the Red Box. The timing from P1 to P3 also included the propagation through the HSR nodes in the ring. The timing between P2 and P3 revealed explicitly the maximum time difference in the HSR ring.

PRP packets have no explicit Ethertype, they contain a redundancy control trailer (RCT) at the end of the packet, which makes their processing a bit trickier. But for timing measurements, the RCT has not necessarily to be processed, as long the packets can be unambiguously identified.

Figure 4 shows one possibility for a measurement in a PRP network, which will for example deliver the propagation time through the network B and that will work with any kind of Ethernet traffic.
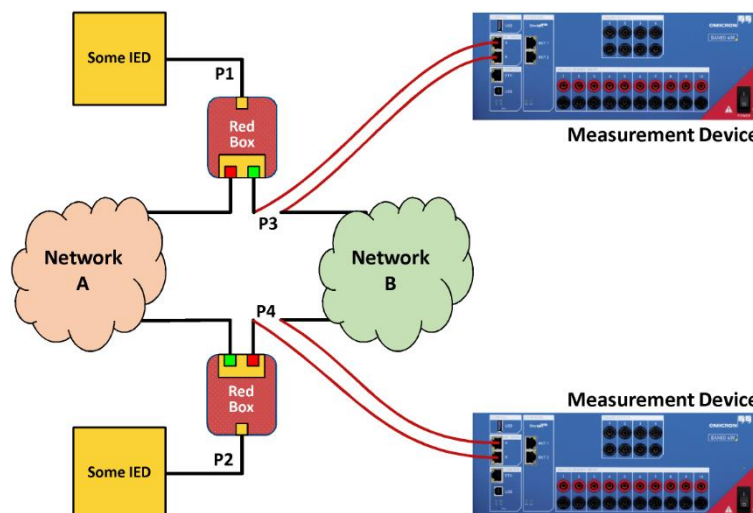

Figure 4: Measurements in a PRP network

Depending on type of traffic used for the measurements, also timings between other access points within the PRP architecture can be measured. For example, GOOSE or Sampled Values can be again identified by their Ethertype and the related processing may skip RCTs which might otherwise confuse the evaluation, making such packets ideally suited for measurement purposes.

Of course, in a setup such as in Figure 4 with multiple measurement devices, the measurement devices need to be accurately time synchronized to deliver properly time stamped data that can be used for an evaluation. The Precision Time Protocol (PTP) according to IEEE 1588 is well suited for such purposes, because it can be also distributed via the Ethernet network. The use of such PTP clocks is shown more explicit in the following example.

5. MEASUREMENTS IN WIDE AREA NETWORKS

When measuring delay times over WAN, the task becomes more complex. Multiple acquisition devices need to be used and they have to be precisely time synchronized to achieve the required timing accuracy. In the following picture, this is indicated by the GPS receivers connected to the acquisition devices.
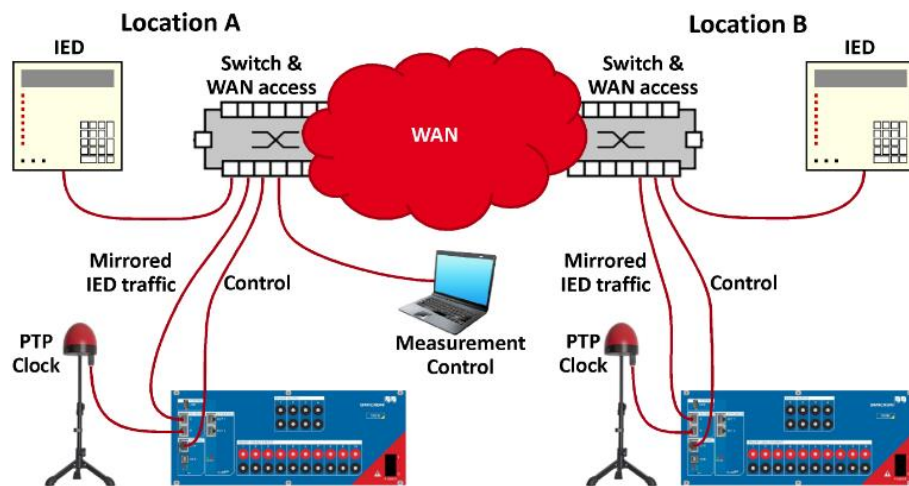

Figure 5: Propagation time measurement over a WAN

But still the whole measurement system is controlled from one single computer, making the operation as easy as possible. The expected range of the measured propagation delays is of course typically one order of magnitude larger than in LANs, but this depends very much on the bandwidth and WAN technology used.

Again, the measurement is the only way to find about the actual timing of the information exchange between the involved IEDs. A final assessment of the performance of the communication network cannot be made on theoretical evaluations alone.

Of course, such a distributed measurement setup as required with a WAN may be also used in a LAN (e.g. in a large substation) when the capture locations are too far from each other.

6. PROPAGATION DELAY OVER A MICROWAVE LINK

The following figure shows the results of a measurement performed between two locations which are connected by a microwave link. The measurement was made by capturing ICMP packets which were generated by "pinging" a device at the remote end.
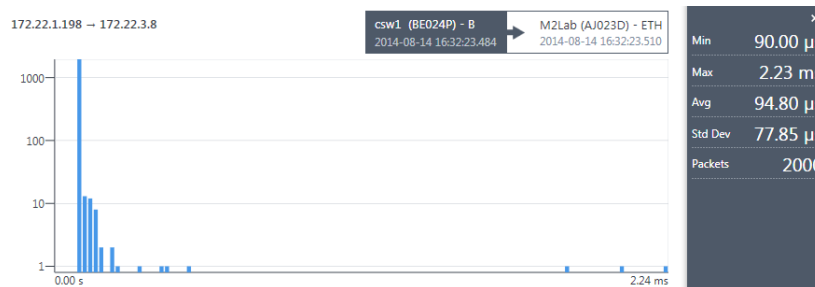

Figure 6: Delay time distribution over a microwave link

The base speed of the link is rather fast, the bulk of the delay times lies in the range of 90 µs to 180 µs. But there are outliers in the range of more than 2 ms. There was arbitrary load on the link, probably causing the outliers. If such a communication link was to be used for protection applications, the cause of the outliers needs to be investigated and removed to obtain consistent performance.

The two locations connected by the link were both in a common broadcast domain, so no routing was involved and the connectivity was as in a local network. This makes tasks like finding and configuring devices for remote control most easy.

7. TRANSATLANTIC LATENCY

To explore the extremes, a delay time measurement between two locations in two different continents was performed. It is unlikely that time critical data for protection, automation and control will be exchanged over such a distance, but it serves to show the orders of magnitudes that can be expected for the traveling time of information. The following figure illustrates the locations of the two involved sites in Austria and in Texas.
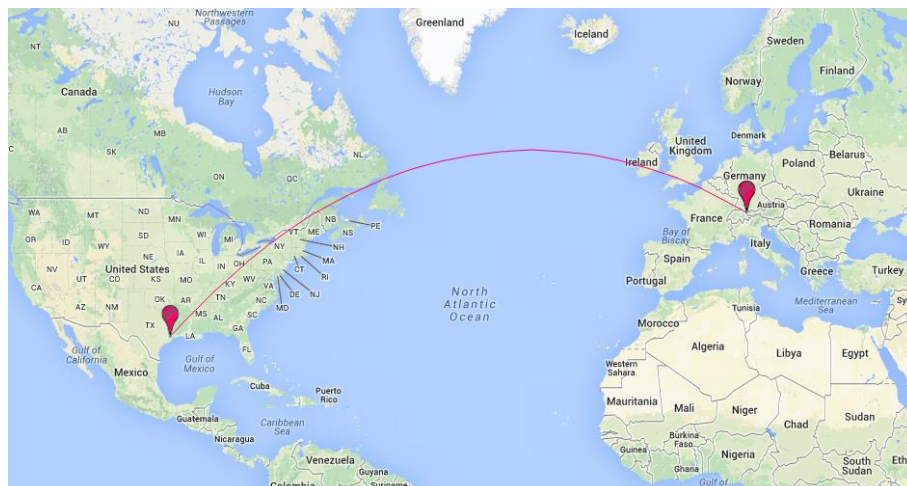

Figure 7: Locations and shortest path (beeline) in between
Map data © 2014 Google. Draft Logic distance calculator © 2014 draftlogic.com

Of course, such WAN connections involve routing, which means that several convenient features that are available in LANs, as for finding devices, do not work anymore. For setting up the connections properly, the network masks and gateways in the sub-networks and the IP addresses of the involved devices need to be known.

The test itself was performed almost identical to the one over the microwave link shown above.
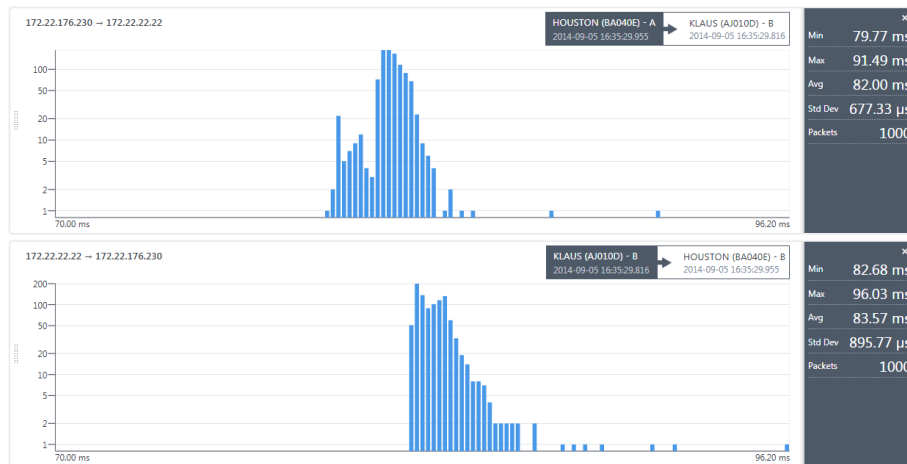
Figure 8: One-way propagation delays between Texas and Austria

The measured propagation delays are in the range of 80 ms to 100 ms. To show the distribution in more detail, the graphs are biased at 70 ms and the distributions are much narrower than the figure suggests. The standard deviation is less than one percent of the mean value in both directions.

The beeline distance between the locations is 8654 km. With the typical figure for the speed of wave propagation on electrical wires or in optical fibers (about 2/3 of the speed of light in the vacuum), the propagation time for this distance equals 43 ms. As the links will not follow the beeline and the actual way for the signals will be longer, at least 50 ms can be attributed to the propagation time in the media alone, leaving about 30 ms to 50 ms for the processing and forwarding the in the involved communication equipment.

8. CONCLUSION

The performance of a power utility communication network is crucial when it conveys mission critical real time information. Theoretical examinations may be useful for the design of such a network, but they are never sufficient for a final assessment of the performance. With modern measurement equipment, such propagation time measurements are easily performed in LANs with and without redundancy as well as with WANs.

BIBLIOGRAPHY

[1] Steinhauser, F.: From LAN to WAN – assessing the communication network for protection, automation, and control. PACWorld conference 2014, Zagreb
[2] Steinhauser, F.: Propagation and Interaction of Ethernet Packets with IEC 61850 Sampled Values in Power utility Communication Networks. DPSP 2014, Copenhagen
[3] Steinhauser, F.: Interaction of Ethernet Traffic in Power Utility Communication Networks. PACWorld conference 2013, Dublin
[4] Ingram, D., et. al.: Direct Evaluation of IEC 61850-9-2 Process Bus Network Performance. IEEE Transactions on Smart Grid, Vol. 3, No. 4, December 2012
[5] IEC 61850-90-4: Communication Networks and Systems for Power Utility Automation. Part 90-4: Network Engineering Guidelines, Technical Report, 2012
[6] IEC 62439-3: Industrial communication networks - High availability automation networks - Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR), 2012
[7] Steinhauser, F.: Measuring the Performance of GOOSE Communication - Assessing IEC 61850 Real Time Messaging, CIGRÉ SEAPAC 2011, Sydney
[8] Steinhauser, F., Schossig, T.: Coexistence of SCADA-Communication and Process-Level Real-Time-Communication in Substation Networks. PACWorld conference 2010, Dublin