

15. Testing of devices, distributed functions and systems

Performance Measurements for IEC 61850 IEDs and Systems

Fred Steinhauser, Thomas Schossig, Andreas Klien, Stephan Geiger
OMICRON electronics, Austria

GOOSE Performance

A common starting point for discussions about performance is the meanwhile famous figure 16 in IEC 61850-5. It outlines the issue as it shows that there are several times involved in a transfer of information from one device to another. But as simple as the figure is, it is not straightforward to implement a test that uses the indicated signal flow and delivers the times shown.

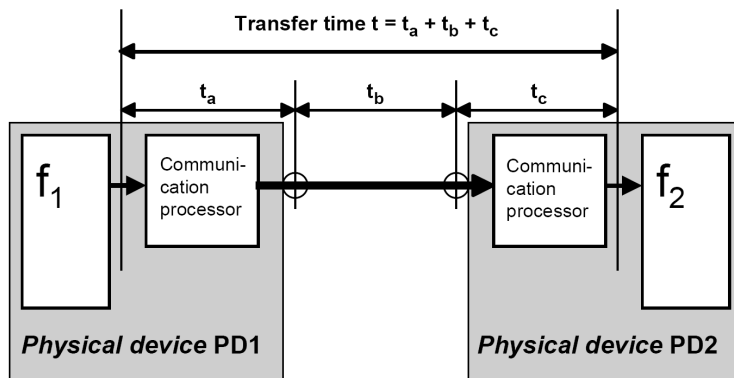


Fig. 1 Transfer time figure from IEC 61850-5

IEC 61850-10 proposes tests with combined measurements with physical I/Os. It is not guaranteed that these physical I/Os are always provided and the test setup and instrumentation becomes more complicated. And if these physical I/Os are provided, they should be implemented with minimal delays (solid state outputs), as the errors incurred by normal relay contacts would be too large for a useful measurement.

The Testing Subcommittee of the UCA International Users Group (UCA IUG) currently works on test procedures that will allow the assessment of an IED's GOOSE performance.

When discussing communication performance ratings for IEDs, it is often commented that this may not be useful as these figures do not guarantee a system performance, which is really important in practical applications. While this is true, it might be argued why these performance figures are questioned when it comes to communication and not in other cases.

A comparable case is with rated pick up times for distance relays. As well, these ratings do not guarantee the performance of a protection scheme, since several other factors (e.g. the performance of the signal channel) have significant effects. Nevertheless, no protection engineer would accept if this rating is waived, as it gives important hints for the applicability of the relay.

This is similar with the communication performance of IEDs. Such ratings are important for assessing the applicability of an IED for an intended system functions. If an IED already consumes most or all of the time allowed for the system function, it can be seen right away that it is not suited for this purpose.

Round-Trip Test

The round-trip test is also sometimes called a "ping-pong" test. A stimulus is sent to a device under test (DUT). The DUT is set up to reply with a response as fast as possible. The time between sending the stimulus and receiving the response is the round-trip time t_{RT} .

The test procedures currently under development at the UCA IUG are widely based on round-trip tests. The scenario in the figure from IEC 61850-5 has to be adapted for round-trip tests. Compared to Fig. 1, the round-trip scenario contains two transfers. The t_a corresponds to $t_{out,TS}$ and $t_{out,DUT}$. The t_c corresponds to $t_{in,DUT}$ and $t_{in,TS}$, and t_b corresponds to t_{Net} .

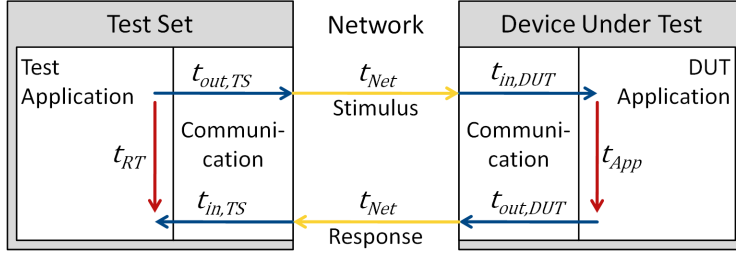


Fig. 2 Times in a round-trip test

The test method is a compromise. On the one hand it is easy to implement, but it also has some shortcomings. The actual measurement t_{RT} is the total of seven individual times, thus introducing some averaging, and some assumptions have to be made to come to useful results.

The assumption which is easiest to implement is to neglect t_{Net} . Such scenarios can be created in a test setup by a direct point-to-point link between test set and DUT or by only using one high performance switch in between. In such cases, the t_{Net} can be assumed to be one order of magnitude smaller than the other times involved.

Another assumption is the symmetry or (isotropy) of t_{in} and t_{out} of a device. If this cannot be assumed, the test evaluation has to stop at the sum of both (equations (3) and (4)).

The application time t_{App} is explicitly mentioned to account for processing delays in the DUT, such as cycle times of a programmable logic. Most protective relays utilize optimized signal paths for important protection signals (e.g. trip signals), while the construction of a receive/send feedback for such a test would require the use of a programmable logic. These programmable logics typically have scan cycles that essentially contribute to in the test result. If a vendor can make a specification of the cycle time, it can be accordingly deducted from the round-trip times to evaluate the (essentially smaller) communication delays, which will then be effective for the protection applications.

By neglecting t_{Net} , the averages of the times total to the average of the round-trip time t_{RT} . The variances (squares of the standard deviations) are similarly related.

$$\bar{t}_{RT} = \bar{t}_{out,TS} + \bar{t}_{in,DUT} + \bar{t}_{App} + \bar{t}_{out,DUT} + \bar{t}_{in,TS} \quad (1)$$

$$\sigma_{RT}^2 = \sigma_{out,TS}^2 + \sigma_{in,DUT}^2 + \sigma_{App}^2 + \sigma_{out,DUT}^2 + \sigma_{in,TS}^2 \quad (2)$$

This implies the assumption that the individual times are not correlated and no covariances have to be taken into account. Given that t_{App} and σ_{App} are known, as well as the data of the test set, we get (3) and (4) for the DUT.

$$\bar{t}_{in,DUT} + \bar{t}_{out,DUT} = \bar{t}_{RT} - (\bar{t}_{in,TS} + \bar{t}_{out,TS} + \bar{t}_{App}) \quad (3)$$

$$\sigma_{in,DUT}^2 + \sigma_{out,DUT}^2 = \sigma_{RT}^2 - (\sigma_{in,TS}^2 + \sigma_{out,TS}^2 + \sigma_{App}^2) \quad (4)$$

Isotropy means that the average sending and receiving times are equal and that the individual times follow the same distribution as expressed in (5) and (6).

$$\bar{t}_{in,TS} = \bar{t}_{out,TS} = \bar{t}_{TS} \quad \text{and} \quad (5)$$

$$\sigma_{in,TS}^2 = \sigma_{out,TS}^2 = \sigma_{TS}^2$$

$$\bar{t}_{in,DUT} = \bar{t}_{out,DUT} = \bar{t}_{DUT} \quad \text{and} \quad (6)$$

$$\sigma_{in,DUT}^2 = \sigma_{out,DUT}^2 = \sigma_{DUT}^2$$

If isotropy applies, the results become to (7) and (8).

$$\bar{t}_{DUT} = \frac{\bar{t}_{RT} - \bar{t}_{App} - \bar{t}_{TS}}{2} \quad (7)$$

With these values evaluated for the DUT, a rating for t_{DUT} , restrained with a certain probability can be derived. For example, to cover 99.99%, an interval of about $\pm 4\sigma_{DUT}$ (still assuming a standard distribution) would be needed. Thus, a rating could look like (9).

$$\sigma_{DUT} = \sqrt{\frac{\sigma_{RT}^2 - \sigma_{App}^2 - \sigma_{TS}^2}{2}} \quad (8)$$

$$t_{DUT} < (\bar{t}_{DUT} + 4\sigma_{DUT}) \quad (9)$$

for 99.99% of all cases

Calibrating the test set

The test as described above can as well be used to evaluate the performance of the test set itself. When using a second test set as the DUT, $\bar{t}_{DUT} = \bar{t}_{TS}$ and $\sigma_{DUT} = \sigma_{TS}$ apply. The results simplify to (10) and (11).

$$\bar{t}_{TS} = \frac{\bar{t}_{RT}}{4} \quad (10)$$

$$\sigma_{TS} = \frac{\sigma_{RT}}{2} \quad (11)$$

Rally Test

This may be a rather simple test that can be performed without a test set for generating a stimulus. It is set up with two DUTs, which are configured to form a self exciting loop, running as fast as the IEDs can do. What is required is the possibility to measure how long the "rally" has gone on and a counter for the number of loops that have completed during this time interval.

The result of the test carries heavy averaging, since only the average loop time can be determined. Nevertheless it gives useful indication for the capabilities of the devices, especially when the performance of the other functions of the IEDs under the rally condition can be verified.

Influence of Network Load

In practical installations, other network traffic will be also present on the substation network and it must be expected that this traffic interferes. There are two main issues where the network load might affect the system's performance.

The IEDs have to process packets that arrive at their network ports regardless if they are intended for them or not. Depending on how the filtering of the traffic is implemented, this may cause load for the IED's processor, possibly leading to a situation that it becomes unable to respond to other requests. This condition is also known as denial of service (DoS).

When crossing the network, the packets involved in the application will be forwarded by switches, where they may interfere with the other data packets. By using priority tagging, data packets can be assigned to different priority levels. Important traffic can obtain higher priority and will be favored by the switches. But even packets of higher priority are affected by other traffic to a certain degree.

Definition of Network Load

Although quite common until now, just stating a figure such as a percentage of the nominal network bandwidth or an absolute bandwidth for the network load is not a sufficient specification. Identical figures of this kind can be produced by totally different load scenarios, e.g. by sending long packets at a low rate, or short packets at a high rate (see Fig. 3). It is obvious, that the two scenarios have potentially totally different impact on the network traffic and the connected IEDs.

Long packets will cause the network ports in switches to be occupied for a longer time. Other packets that have to wait for the port to become idle will collect larger through this.

Shorter packets will occur at a high rate and put stress on the IEDs which have to receive the packets even if they are not intended for them.




Fig. 3 Packet size over packet rate
(MLEP = Maximum Length Ethernet Packet)

There are several other questions about the nature of the load traffic that can be asked: Are the packet lengths constant or varying? If they are varying, are the packet lengths evenly distributed or do they follow another distribution?

The same considerations apply to the timing when packets are injected into the network.

Considering the many different variables that may characterize network load, the number of all possible combinations will easily explode and it is impossible to cover all of them.

In the following, two distinct load scenarios, which can be easily reproduced, are used.

Network Load Scenarios

The figure below shows the general setup used to measure the influence of network load on the results of the round-trip measurements. There are two switches connected by a trunk link.

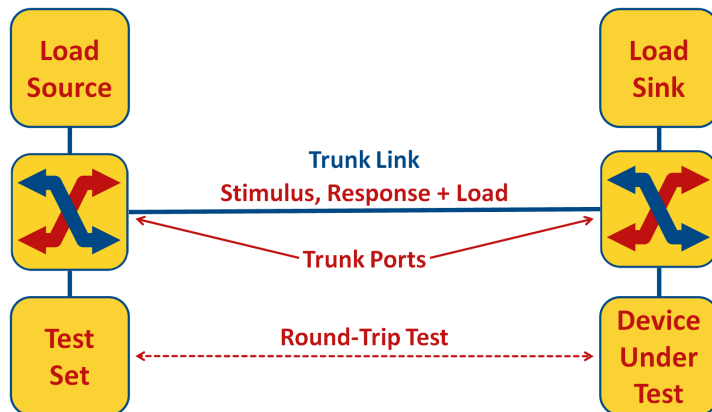


Fig. 4 Test setup for testing with network load

The GOOSEs involved in the test have to cross both switches, as well as the load traffic.

The trunk ports are the points where congestion can be potentially expected. It must be noted that any packet to be sent over the trunk has to wait until the port becomes idle, when another packet is just being sent. Even for a pending high priority packet, the sending of a lower priority packet is not aborted, the priority mechanism becomes only effective for the next packet being sent.

Sampled Values as Network Load

Sampled Values being present on the network is an important case with practical relevance which is expected to be challenging for the network and IEDs. As the Sampled Values are multicast messages, a dedicated load sink is neither defined nor required. The switches will distribute the packets to all ports. In this case, also the test set will be loaded with the Sampled Values, which has to be considered when assessing the measurements.




Fig. 5 Average and standard deviation of t_{DUT} under varying SV load

Sampled Value packets are relatively short (approximately $13\mu s$), so the jitter introduced by congestion in the trunk ports can be expected to be decent. But the packets keep the DUT busy, possibly delaying its response to the incoming GOOSE stimulus.

Fig. 5 shows how the results from the round-trip test are affected by injecting one, two, or three Sampled Values streams as load.

The times for the DUT have been evaluated and the figures have been normalized to \bar{t}_{DUT} for the figure. All tests have been done with 1000 repetitions, as suggested in IEC 61850-10 and the draft of the performance test procedures of the UCA IUG.

ICMP Messages as Network Load

Injecting long ("maximum Ethernet size") packets can be used to simulate file transfers, as they would occur e.g. during a download of a fault record. These packets occupy the trunk ports for a much longer time than the Sampled Values from the test case above, so they will be "in the way" for a longer time. Thus, a higher jitter can be expected from this traffic.

Ping, a standard networking utility, may be used to inject ICMP (Internet Control Message Protocol) packets into a network. The Linux versions take parameters for packet size (-s) and the interval (-i) between two packets. The addressed ("pinged") computer replies with more or less identical messages, thus creating the same network traffic in opposite direction. By this, a reasonably defined network load can be produced.

Nevertheless, the observed effects from a first series of tests were insignificant. At an injection rate of 500 pings per second, there was no relevant effect on the timing of the GOOSE messages. As the links are full duplex, the network traffic and the probabilities for congestion at the trunk ports are kept separate in each direction. Further tests looking more closely into this matter may be useful.

Testing with TAP

A TAP is a device that is inserted into a network link to "sniff" the traffic from the link for capturing and analysis purposes. A TAP is designed to affect the data flow on the link as less as possible. Purely "passive" TAPs have no effect on the timing of the crossing data at all, but also "active" TAPs incur only so small delays that this is still negligible. TAPs used for tests as discussed here have the option to assert accurate time stamps to the captured data. From the timestamps of the captured data, accordingly detailed timing results can be derived.

For the test cases and results presented in Fig. 5, the results have been verified by measurements made with a TAP. The TAP was inserted directly at the DUT, thus delivering just the times elapsed in the DUT itself and eliminating the contribution of the test set to the round-trip time.

Remarks and Outlook

Several assumptions have been made to keep the test cases and evaluations simple and easily understandable. The applied statistics mostly assume a normal distribution, allowing the application of the commonly known definitions for average value and standard distribution. The assumption of statistical independence of the individual times to eliminate the covariances has been already mentioned. All of these assumptions can be questioned in the one or other way.

The methods shown here could be a starting point for further refinements of the test and evaluation methods to waive some of the mentioned simplifications and assumptions. It may be worthwhile to look closer into the actual distributions of the times and using other metrics for characterizing the time properties, such as medians and distinct quantiles, which may be also more robust against disturbances and outliers in the obtained raw data.

Similar issues as with the interference of GOOSE packets with other traffic apply to the propagation of Sampled Values in a process bus. The accumulation of jitter when crossing multiple switches can be measured in a similar way as shown here. An interesting question is under which conditions the jitter reaches values so that packets of a Sampled Values stream catch up to their predecessors.

Conclusions

The performance of modern protection systems depends on the performance of its components, mainly the communication network and IEDs. Timely response of an IED is an interoperability issue, as it determines the applicability of a device in a system function. The system's response cannot be faster than the response of its components, and when designing a system, it must be taken care that the individual responses leave some margin for the system to achieve its ratings.

Therefore, to assess the applicability of an IED, ratings for its communication performance are required. The described methods to assess the performance can be performed with modern, accurate protection test sets and do not require specialized communication test equipment. Even the verification of the performance of the test equipment itself is possible this way.

The results obtained show that the tested devices are rather robust against the loads applied in the described scenarios. Further refined tests may give more insight into distinct scenarios.

References

- [1] Ito, H., Ohashi, K.: Implementation of High Performance IEC 61850 GOOSE Messages. PAC World Magazine Winter 2008, p. 40-46
- [2] Bastigkeit, B., Schossig, T., Steinhauser, F.: Efficient Testing of Modern Protection IEDs. PAC World Magazine Winter 2009, p. 54-59
- [3] Klien, A: Performance Analysis of IEC 61850 GOOSE Applications. Bachelor Thesis. Vienna University of Technology, 2010
- [4] IEEE 802.1q:2003 Virtual Bridged Local Area Networks
- [5] UCA International Users Group, Testing Subcommittee: Test procedures for GOOSE performance according to IEC 61850-5 and IEC 61850-10, draft version 0.3, March 2010

Biographies



Fred Steinhauser was born in Austria in 1960. He studied Electrical Engineering at the Vienna University of Technology, where he obtained his diploma in 1986 and received a Dr. of Technical Sciences in 1991.

In 1998 he joined OMICRON, where he worked on several aspects of testing power system protection. Since 2000 he works as a product manager with a focus on substation communication issues.

Fred Steinhauser is a representative of OMICRON in the UCA International Users Group. As a member of WG10 and WG17 in the TC57 of the IEC he contributes the standard IEC 61850. He is also a member of SC B5 of CIGRÉ.



Thomas Schossig (IEEE) was born in 1970. He received his diploma (master degree) in Electrical Engineering at the Technical University of Ilmenau (Germany) in 1998.

He worked as a project engineer for control systems and as a team leader for protective relaying at VA TECH SAT in Germany from 1998 until 2005.

In 2006 he joined OMICRON as a product manager for substation communication products. Additionally, he is responsible for the IEC 61850 trainings at OMICRON. He is author of several papers regarding IEC 61850 and protection testing and member of standardization working groups.



Andreas Klien was born in 1986. He is studying Computer Engineering at the Vienna University of Technology, where he received his bachelor's degree in 2010. He is currently in the master's program for Computer Sciences.

Since he joined OMICRON in 2005, he works as a software developer mainly for IEC 61850 products.

In the course of his bachelor thesis he studied performance measurement methods for GOOSE protocol implementations and the effects of network load on the communication performance in a substation automation system.



Stephan Geiger was born in Austria in 1987. He graduated from the Technical College of Bregenz with a focus on Electrical Engineering in 2007.

In 2008 he joined OMICRON, where he works in the product management as an application engineer for secondary testing related issues.

During his involvement in the verification of products for testing IEC 61850 features, he has performed extensive test series.

OMICRON is an international company serving the electrical power industry with innovative testing and diagnostic solutions. The application of OMICRON products allows users to assess the condition of the primary and secondary equipment on their systems with complete confidence. Services offered in the area of consulting, commissioning, testing, diagnosis, and training make the product range complete.

Customers in more than 130 countries rely on the company's ability to supply leading edge technology of excellent quality. Broad application knowledge and extraordinary customer support provided by offices in North America, Europe, South and East Asia, Australia, and the Middle East, together with a worldwide network of distributors and representatives, make the company a market leader in its sector.

Americas

OMICRON electronics Corp. USA
12 Greenway Plaza, Suite 1510
Houston, TX 77046, USA
Phone: +1 713 830-4660
+1 800-OMICRON
Fax: +1 713 830-4661
info@omicronusa.com

Asia-Pacific

OMICRON electronics Asia Limited
Suite 2006, 20/F, Tower 2
The Gateway, Harbour City
Kowloon, Hong Kong S.A.R.
Phone: +852 2634 0377
Fax: +852 2634 0390
info@asia.omicron.at

Europe, Middle East, Africa

OMICRON electronics GmbH
Oberes Ried 1
6833 Klaus, Austria
Phone: +43 5523 507-0
Fax: +43 5523 507-999
info@omicron.at