

Measurement of Switching Latency in High Data Rate Ethernet Networks

Tomas Hegr¹, Miroslav Voznak², Milos Kozak¹, Leos Bohac¹

¹*Department of Telecommunications at Faculty of electrical engineering at Czech Technical University in Prague,*

Technicka 2, 166 27 Prague, Czech Republic

²*Department of Telecommunications at Faculty of electrical engineering VSB-Technical University of Ostrava,*

17. listopadu 15, 708 33 Ostrava, Czech Republic

tomas.hegr@fel.cvut.cz

Abstract—The paper deals with a methodology of switching latency measurement in switched Ethernet networks. The switching latency is parameter necessary for simulation and design of low-latency networks that are often intended for real-time control inherent to many industrial applications. The proposed measurement methodology provides a simple way of switching the latency determination and vendor quoted latency values verification directly at the physical layer. Numerous experimental measurements were carried out to support the arguments in this paper and to demonstrate the usability of the proposed methodology. All results are presented and analysed up to 10GBase-R Ethernet including OpenFlow switches.

Index Terms—Switching latency, measurement methodology, OpenFlow.

I. INTRODUCTION

The Ethernet is one of the most progressive transmission technologies at the data link layer nowadays. Ethernet proved to be a suitable technology in most infrastructure levels, from LANs to carrier WANs. One of the most important parameters is a maximum transmission delay in high-demanding areas as data centres or substations in smart grid [1]. This parameter is even more indispensable when designing real-time control in industrial networking where Ethernet boldly put down roots.

One example of the challenging deployment is a Substation Automation (SA) system as described in standard IEC 61850 [2]. This standard requires that a data network has to ensure the transmission delay less than 3 ms for sampled values and control messages.

The SA system is only one of many application examples of the real-time communication which shows how important it is to know all network parameters precisely when designing a network topology. Incomplete data in this sense is the switching latency and its progress for various frame lengths. Although vendors indicate switching latency, it is mostly defined for 64 B frames only and under undefined

conditions. For this reason, we decided to design a new measurement methodology allowing verification of information provided by vendors.

The key objective was to develop and verify a measurement methodology that enables to determine the switching fabric latency for data rates up to 10 Gbps without specialized instruments. Although the methodology was firstly published in [3], the methodology proposed in this paper is enhanced and was extensively verified on switches supporting Ethernet speeds up to 10GBase-R including OpenFlow switches.

Since the OpenFlow (OF) appears to be a promising technology, it was decided to include OF switches in the tests. The OF protocol described in [4] is a part of an emerging Software-Defined Networking (SDN) concept which is progressively getting into the production environment. SDN allows to control data flows in the network by a controller more effectively than it is possible in traditional distributed networks. The SDN approach enables to systematically compute and implement optimal flow paths and thus determine transmission resources necessary for the real-time traffic.

The paper is organized as follows. Section II presents related works and standards. Sections III and IV present measurement limits and the measurement methodology. The last Section V describes a number of experimental measurements carried out in our laboratory with the aim to verify and demonstrate the methodology's applicability.

The expanded uncertainty of measurement was in most cases up to 10 % relative to the estimated switching latency. In this paper, we consider the real-time traffic as specific traffic with high demands to the latency and jitter transferred via regular switched Ethernet network, not the deterministic real-time Ethernet.

II. RELATED WORKS

The perennial weakness of the latency measurements is the source-receive time synchronization. Published works dealing with the switching latency measurement use different approaches. The first option is external synchronization through dedicated wiring using timing signals, e.g. IRIG-B

Manuscript received October 30, 2014; accepted January 16, 2015.

This research was funded by Students grants at Czech Technical University in Prague SGS13/200/OHK3/3T/13, SGS12/186/OHK3/3T/13 and partially the grant of SGS reg. no. SP2015/82 conducted at VSB-Technical University of Ostrava.

(Inter-Range Instrumentation Group mod B) or 1 PPS (Pulse Per Second), or time synchronization protocols, e.g. Network Time Protocol (NTP) suggested by Loeser *et al.* in [5]. Even though NTP is suitable for many applications, it is recommended to implement Precision Time Protocol (PTP) defined in IEEE 1588 in order to achieve high synchronization accuracy [6].

The second option is to use specialized internally synchronized card providing a high-precise frame timestamping and measuring frames in loopback as is suggested by Ingram *et al.* in [7]. Both approaches require additional specialized hardware which depends on the used transmission technology or synchronization protocol.

The next option is to measure latency directly by means of special data frames called CFrames, as suggested in [8]. The authors of this paper suggest using a special CFrame flow forwarded through the internal switching fabric and to measure latency between the ingress and the egress port directly at a switch backplane. Since the integrated box design prevents accessing the switching fabric, the proposed methodology views the measured switch as a black box. This approach takes into account delays caused by internal processes and the resulting value is then more meaningful.

The methodology design relies on RFCs by IETF and on fundamental standard for switched Ethernet networks IEEE 802.3:2012 [9]. The elementary description of the switching latency is based on RFC 1242 of 1991, which defines the latency of the store-and-forward devices [10]. According to the recommendation, the switching latency, or in other words processing time of the passing frame, is defined as a time interval starting when the last bit of the input frame reaches the input port and ending when the first bit of the output frame is seen on the output port. This method is typically called Last In First Out (LIFO).

Further documents related to the measurement methodology include RFC 2544 of 1999 [11]. A wide range of specialized measuring instruments implement this recommendation as a basis. This document defines, *inter alia*, the time intervals necessary between individual readings and also frame lengths needed for measurements.

Ultimately, the root document for an evaluation of the measurement accuracy is technical report Evaluation of Measurement Data - Guide to the Expression of Uncertainty in Measurement by Joint Committee for Guides in Metrology [12]. It specifies a calculation of measurement uncertainty and its handling.

III. SWITCH ARCHITECTURE AND MEASUREMENT LIMITS

Generally, the switch can be seen from different perspectives. From the hardware point of view the switch is generally composed of line cards, CPU, various memory structures storing Forwarding Information Base (FIB) and the switching fabric. Most fabrics are usually implemented in form of an Application Specific Integrated Circuit (ASIC). This arrangement is shown in Fig. 1. All components are connected by an internal bus situated on the switch backplane. The line card contains at least one interface for signal processing at the Physical layer (PHY) and Medium Access Control (MAC). It also contains a local

FIB and the fabric ASIC if the line card serves more ports. The architecture of modular and large enterprise switch is different both in terms of backplane design and line card construction. These switches are usually equipped by additional CPUs and memories.

The switch can also be viewed from the frame processing and memory utilization perspective. A significant amount of current switches uses some kind of shared memory with distinct data structures. Architecture called Combined Input and Output Queuing with Virtual Output Queuing (VOQ) [13] is frequently applied in order to reach efficient utilization of resources and the best delay to throughput ratio. In this case, the incoming frames are arranged into a shared memory dedicated to the appropriate output port queues VQO. Once a frame is processed, the frame is forwarded to the output queue of the destination port. This prevents the head of the queue blocking.

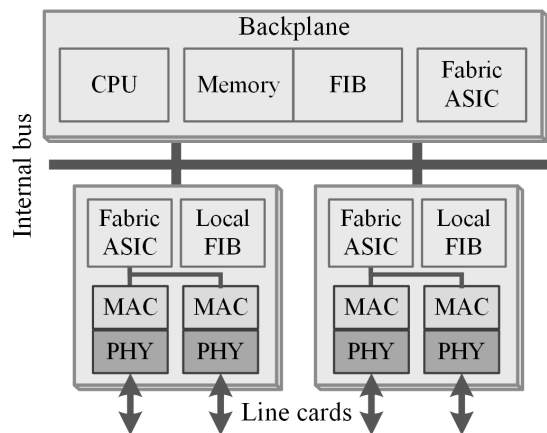


Fig. 1. Physical arrangement of components in a common Ethernet switch.

Accordingly, the overall processing time of the frame transmission between an input and output port is composed of several independent delays. The minimum measurable switching latency in the commonly used architecture can be estimated by (1)

$$t_{sw} = t_{iq} + t_{sf} + t_{oq} + 2t_{lc}, \quad (1)$$

where t_{sw} stands for the total switching latency, t_{lc} represents the line card delay, i.e. the processing time of the frame passing between layers and the time needed to transfer the frame via the internal bus to the switch backplane, t_{sf} is the switching fabric delay itself, t_{iq} is the input queue delay (e.g. VOQ) and t_{oq} represents the output queue delay. The line card delay does not involve an input buffering delay eliminated by the LIFO measurement approach.

The use of memories and their arrangement can vary considerably. A general switch determines the output port from the destination MAC address via the FIB stored in a Content-Addressable Memory (CAM) whose cells support binary states only. The Ternary CAM (TCAM) was introduced to overcome such a limitation. TCAM provides a third state representing the “do-not-care” value. This state allows using wildcards during the look up process in FIB, and it also allows defining Access Lists (ACL) without the

need to store them for each individual address. Although TCAM is very effective in matching, the cost and size of its implementation is high as one TCAM cell consists of 16 transistors [14]. Due to this reason, vendors often implement FIB through hash tables [15] for all types of lookups including ACL. It is expected that the forwarding processed by CPU, i.e. not in TCAM, will be considerably slower.

IV. MEASUREMENT METHODOLOGY

The measurement methodology is based on the LIFO method. It advantageously uses Manchester encoding at 10Base-T channel. This means that the channel is not burdened by any broadcasting in the rest state between transmissions and consequently it is possible to unambiguously identify the passing test frame. Other Ethernet types at higher data rates keep uninterrupted signal broadcast on the transmission channel to preserve the sender-receiver synchronization. Thus, it is not possible to determine the head and tail of the passing test frame at the physical layer without decoding the signal. This type of measurement is challenging and requires high performance packet analyser. We decided to use a two-channel oscilloscope commonly available on technical workplaces.

The test traffic consists of Internet Control Message Protocol (ICMP) packets. It is generated by a sender using the ping application. This application is sufficient for measuring purpose because it allows setting the packet length and time spacing between individual packets. All unnecessary switch services generating unsolicited traffic or consuming switch performance must be disabled at the switch otherwise it would not be possible to unambiguously identify test packets. The unwanted traffic additionally causes a queue filling which influences and distorts measured data. Ultimately, it is necessary to set up static ARP entries at both pingging sides avoiding Address Resolution Protocol (ARP).

The original methodology was intended for measuring the switching latency between 10Base-T Ethernet ports only, but the measurement steps remained similar. The time difference measurement is carried out on the oscilloscope which is connected directly to the transmission medium at the physical layer by active differential probes. Where possible, it is necessary to deactivate the Automatic MDI/MDI-X (Medium Dependent Interface) feature, i.e. pair swapping, at the measured ports. The measurement is usually carried out on the TD+ and TD- pair before and after the switch, i.e. in the sender-to-receiver direction.

Readings are made with respect to RFC 2544 in series of different frame lengths (64 B, 128 B, 256 B, 512 B, 1024 B, 1280 B, 1518 B). The number of repetitions must be at least 20 times with the reported value being the average of the recorded values as required by RFC 2544. Naturally, the higher number of repetitions, the lower the statistical error. The threshold voltage level is based on the resistance of the used probe. The set of the ports determined for measuring is extensively described by RFC 2889.

With the goal to measure higher data rates, it was necessary to extend the wiring diagram and methodology steps due to the aforementioned synchronization

broadcasting. The enhancement depicted in Fig. 2 reside in extending the original schematic by two auxiliary devices keeping 10Base-T Ethernet on input and output ports.

At first, it is necessary to measure the characteristic delay between auxiliary switches without the evaluated switch and subsequently to create a correction table. The measurement of characteristics is made using the same procedure as described above for all frame lengths and the examined data rates. It is recommended to take far more than 20 readings to reduce the correction uncertainty in further applications. Once the correction table is drawn up, it is possible to connect the evaluated switch between those auxiliary ones and repeat all measurements.




Fig. 2. Schematic for high-speed Ethernet scenarios. SWAUX 1 and 2 are auxiliary switches and SWMEAS is the examined one.

In the extended methodology, it is necessary to cleanse the results obtained from the measurements performed by means of the correction table. While the correction table consists of the arithmetic mean delay for all frame lengths and the examined data rates obtained by the pre-measured series between the auxiliary devices, the correction itself must be expanded to include the input buffering delay and signal propagation delay at a newly created network segment. This new segment is located between an auxiliary switch and the evaluated one. The delay value cannot be included in the pre-measured characteristics so it must be calculated. Both delays can be estimated very accurately as the input buffering delay behaves clearly linearly and the cable propagation delay remains constant. Whereas the input buffering produces a significant additional delay and must be considered, the signal propagation delay is almost negligible.

Although the measurement itself is carried out at the physical layer, it is possible to use a net bit rate (also referred as data rate) to estimate the frame input buffering delay. This assumption can be made as the frame is equipped with the preamble, Start Frame Delimiter (SFD) and Check Sequence (CRC) at the MAC layer. These frame fields are encoded together with the rest of the frame. They are explicitly mentioned because they are not usually provided to higher layers such as MAC addresses or EtherType. The length of all these fields must be taken into account in the correction expression. The arithmetic mean value for a given frame length is computed as shown in (2).

$$\bar{t}_{sw} = \frac{1}{N} \sum_{i=1}^N t_{mes_i} - \bar{t}_{aux} - \frac{l_{hf} + l_{pt}}{R} - t_{sp}, \quad (2)$$

where t_{aux} is the mean delay of auxiliary switches taken from correction table [s], l_{hf} is the header length with preamble, SFD and CRC (208 bits) [bit], l_{pt} is the length of

the ping test frame [bit], R is the net bit rate [bit/s] and finally the optional t_{sp} signal propagation delay [s]. The signal propagation delay can be evaluated by (3), where l_c is the cable length, c represents the speed of light and NVP stands for the Nominal Velocity of Propagation. NVP expresses the speed with which electrical signals travel in the cable relative to the speed of light in vacuum

$$t_{sp} = \frac{l_c}{NVP \times c}. \quad (3)$$

The subsequent part of the measurement methodology is to determine the measurement accuracy. The overall measurement accuracy is given by the expanded standard uncertainty covering both A and B type. The standard A type uncertainty characterizes the dispersion of the measured values. For the first measurement methodology the A type uncertainty can be estimated as the experimental standard deviation of the mean as shown in (4). It quantifies how well t_{sw} approximates the expected mean value

$$u_A(t_{sw}) = \frac{s(t_{sw})}{\sqrt{N}} = \sqrt{\frac{1}{N(N-1)} \sum_{i=1}^N (t_{mes_i} - \bar{t}_{sw})^2}. \quad (4)$$

As the extended measurement is compounded of two measurement, the standard A type uncertainty of the measured values must be expanded by the uncertainty of the correction measurements. This combined uncertainty can be evaluated as the sum of squares of the particular uncertainties for scenarios with/without inserted evaluated switch as shown in (5)

$$u_A(t_{swC}) = \sqrt{u_A^2(t_{sw}) + u_A^2(t_{aux})}. \quad (5)$$

The combined standard measurement uncertainty then can be determined by (6), where $u_B(t_{sw})$ corresponds to a standard B type uncertainty primarily caused by the specific measuring instrument characteristics. It is commonly estimated on the basis of oscilloscope parameters such as the sampling rate, resolution, skew delay, etc. Finally, it is necessary to multiply the value of the combined uncertainty $u_C(t_{sw})$ by the coverage factor $k_t = 2$ to obtain the expanded uncertainty and achieve 95 % confidence level as shown in (7):

$$u_C(t_{sw}) = \sqrt{u_A^2(t_{swx}) + u_B^2(t_{sw})}, \quad (6)$$

$$U = k_t u_C(t_{sw}). \quad (7)$$

V. ANALYSIS OF EXPERIMENTAL MEASUREMENTS

In contrary to the previous experimental measurements, the objective was to test the methodology at 10GBase-R Ethernet including OF switches.

Measurements were realized on Tektronix DPO4032 oscilloscope with a maximum sampling frequency 2.5 GS/s. This sampling frequency is sufficient as the 100 MS/s is the minimum. The oscilloscope supports the external network

connection so readings were automated using Python and PyVISA library [16].

This automated approach significantly increases the reading resolution that has an impact on the standard B type uncertainty. While the lowest measured switching latency was about one microsecond the measurement resolution was in nanoseconds. The B type uncertainty was for experimental measurements estimated to 60 ns based on the used instruments. Moreover, the standard A type uncertainty was also decreased since the process automation enables to take more readings within the same time range.

Several thousand readings for dozens of switch-data rate combinations were taken. All measurements were made in one direction between random ports or between ports supporting the desired data rate. This procedure was chosen because randomly realized measurements showed that the measurement direction or selected port pairs do not differ significantly in values obtained. The correction characteristics between the auxiliary switches had a clear linear progression.

In most cases, the achieved expanded uncertainty for automated measurements was up to 8 % relative to the estimated mean value. This is primarily due to more precise readings and the number of readings increased to 50 samples. This is an improvement to manual measurements where the expanded uncertainty mostly fluctuates between 10 % and 15 %. In some cases when switching latency is around 1 μ s, the expanded uncertainty relative to given mean can reach up to 30 % in peak. This is caused by the enlargement of the sampling window especially for large frames since the inserted new segment adds a significant buffering delay.

The correction characteristic of delay between auxiliary switches had a linear progression in all variants as shown in Table I. It was used a linear regression to estimate correction characteristics. The linearity is confirmed by the coefficient of determination R^2 which reaches nearly 1 for all data rates.

TABLE I. CORRECTION FUNCTIONS.

Ethernet	Correction linear function	R2[-]
10Base-T	$Y=7.995E-7x+3.344E-5$	1.0000
100Base-TX	$y=7.978E-8x+1.667E-5$	1.0000
1000Base-T	$y=7.803E-9x+1.555E-5$	0.9996
10GBase-R	$y=9.024E-10x+2.402E-5$	0.9970

A. Enterprise Switches

Switches supporting 10GBase-R Ethernet at least on up-link ports were designated as enterprise switches. In our case, these include switches with SFP+ (Small Form-factor Pluggable) transceiver or with an older version of XFP transceiver. Dell 5524 was used as the auxiliary switch (SWAUX) split into two VLANs (Virtual LAN) meaning two auxiliary switches as described in the second methodology. Although this approach does not follow the original idea, it proved to be fully applicable.

Measurement results presented in Table I show a high stability of the expanded uncertainty for given switches at 50 readings. The uncertainty is slightly above 0.1 μ s in all cases which means up to 6 % relative to the estimated latency. The only exception is Dell S4810 where the switching latency

falls down below 1 μs and the expanded uncertainty reaches up to 15 %. In principle, the absolute values can be affected by SFP+ transceivers or by the fact that only up-link ports were available on the tested switches.

TABLE II. SWITCHING LATENCIES OF 10GBASE-R SWITCHES.

Switch	Frame length [B]				
	64	256	512	1024	1518
	Switching latency $\pm U$ [μs]				
Dell 5524	2.05 0.11	2.21 0.11	2.18 0.11	2.07 0.11	2.14 0.12
Dell S4810	0.82 0.12	0.88 0.11	0.94 0.12	0.88 0.12	0.85 0.13
Cisco Catalyst 3750x	4.27 0.11	4.73 0.10	4.91 0.12	5.49 0.12	5.94 0.13
Foundry Edgelron 8x10G	3.27 0.12	3.74 0.11	4.06 0.12	4.76 0.11	5.51 0.11
HP 5406zl	2.07 0.12	2.20 0.11	2.37 0.11	2.85 0.12	3.25 0.11
HP 3800E	1.97 0.11	2.17 0.11	2.28 0.12	2.52 0.11	2.87 0.12

Measured latencies are visualized in Fig. 3. Results indicate possible differences in the switch architecture. While most latencies record a slow linear increase, the latency for Dell switches remains almost constant. This behaviour suggest that most likely there is no additional frame transfer between line cards and backplane. The remaining lines demonstrate an opposite development.




Fig. 3. Switching latency dependant on the frame length for 10GbBase-R.

Absolute values for particular switches are surprisingly high in comparison with the lower data rates. This indicates a convergence toward the real switching latency. The phenomenon is illustrated in Fig. 4, where are three switches supporting data rates from 10 Mbps to 10 Gbps.




Fig. 4. Dependency of switching latency on data rate at 64B length frames.

B. OpenFlow Switches

The OF protocol covers the lower part of the SDN architecture and represents an interface between a logically

centralized controller and controlled switches. OF enables uploading of forwarding instructions into the switch forwarding table. Consequently, any traffic passing through the switch must match some uploaded rule to take an action. The matching rule consists of header fields from L4 to L2 and a physical input port, in OF terminology referred to as tuples. All tuples or their parts can be wildcarded [4].

Matching rules were designed only for the destination MAC address as is common with L2 switches. Since the ARP is eliminated by static records on both client sides, it is necessary to upload just two matching rules to the examined switches. All other tuples are wildcarded. To perform the measurement, we chose a Floodlight controller and its tool Static Flow Entry Pusher (SFEP) [17]. All forwarding modules in the controller were deactivated to prevent any unwanted matching rules being generated. SFEP is built as a controller module and its interface is accessible via JSON (JavaScript Object Notation) and the controller web interface. Such an approach enables to setup a time-unlimited matching rule in both directions for test packets.

Four hardware switches supporting OF and one server PC with the OF service were evaluated. While three switches from HP and Dell have truly integrated OF support in the firmware, the fourth RouterBoard has the OF support in form of the additional software package. The last examined switch was clearly software-based Open vSwitch running on a small server built on dual-core Atom at 1.4 Ghz processor with 2GB RAM and running stripped Debian Wheezy as operating system. The server had two integrated network interface cards up to 1000Base-T Ethernet. Table III provides an overview of results.

TABLE III. SWITCHING LATENCIES OF OPENFLOW SWITCHES.

Switch	Mode	Frame length [B]				
		64	256	512	1024	1518
		Switching latency $\pm U$ [μs]				
Dell S4810	10 Gbps	1.00 0.12	0.98 0.11	0.92 0.11	0.76 0.13	0.76 0.12
HP5406zl		261.90 3.38	244.66 5.17	272.41 8.23	292.00 7.07	272.31 6.94
HP3800E		169.08 1.51	182.43 3.48	170.37 3.57	188.56 4.75	207.00 6.29
Dell S4810	1 Gbps	2.18 0.15	1.89 0.13	2.04 0.14	1.98 0.13	2.03 0.13
Open vSwitch		33.09 1.45	36.85 1.22	37.39 1.27	41.77 0.98	46.69 1.01
RB2011LS-IN		15.69 0.59	20.06 0.21	25.22 0.78	35.09 0.21	46.77 2.08
HP5406zl	100 Mbps	271.35 5.92	249.70 4.73	277.98 7.18	294.30 6.91	299.27 9.58
HP3800E		179.51 6.88	181.05 3.31	170.67 7.21	197.96 6.10	208.02 6.95
Open vSwitch		233.45 6.32	209.93 5.74	163.46 5.84	88.13 3.98	48.24 0.96
RB2011LS-IN	10 Mbps	16.57 0.16	20.89 0.53	26.56 1.09	35.55 0.24	45.08 0.61
HP5406zl		262.28 4.08	232.26 10.48	240.73 9.96	281.57 9.75	301.49 7.35
HP3800E		160.45 0.76	165.26 6.72	196.91 3.12	197.28 4.89	213.56 4.91

The estimated mean switching latency on all switches is considerably higher than on common L2 switches, with one exception. These high values are produced by the frame processing and matching rule evaluation via the software

way, thus by CPU. It is expected that the switching latency will grow significantly during the higher switch load since the CPU performance must be split among other ports.

Even if the Open vSwitch creates its own flow rules derived from OF matching rules and applies them as accurate as possible for a particular traffic, it shows great disproportion between results for different data rates. This can be caused either by non-optimized NIC drivers or the internal process scheduler. In case of HP switches, it is apparent that the port data rate has no significant effect on the overall switching latency unless the CPU is powerful enough. Unfortunately, we were not able to strictly redirect matching rule processing to hardware in the HP boxes.

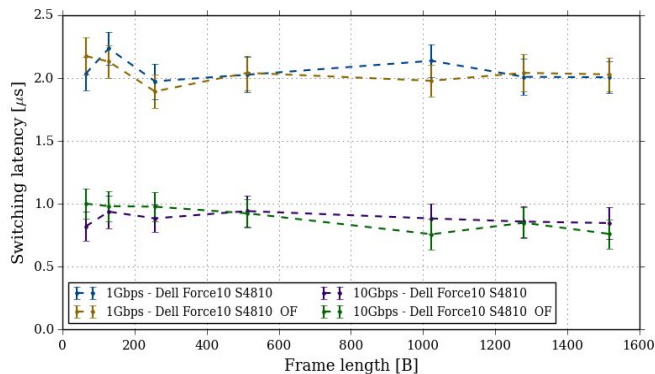


Fig. 7. Switching latency on Dell S4810 for the OF matches traffic and non-OF switching mode.

The only exception among the evaluated switches is the one from Dell which shows latencies oscillating around 2μ s at 1 Gbps and even below 1μ s at 10 Gbps with the expanded uncertainty close to 0.1μ s, as shown in Fig. 7. We noted that OF and non-OF switching latencies are nearly identical. This switch is intended for data centres and is proclaimed to be ultra-low-latency. The great latency stability is the consequence of the dedicated CAM block to the OF process. All OF rules are internally processed as ACLs and thus probably highly optimized. Although only one switch gives sufficient values, it shows that the OF could be implemented even in demanding low-latency networks.

VI. CONCLUSIONS

Even though vendors publish switching latencies for their devices these values are mostly limited only to 64 B frames. Moreover, these latencies are obtained under unspecified conditions. This may not be precise enough in high demanding installations. We propose a measurement methodology that enables to determine the switching latency by commonly available tools. This is very handy for network engineers because they can verify their design with it. The measurement methodology was proved even for high data rates as 10GBase-R by a reasonable expanded uncertainty of the measurement. This uncertainty was up to 15 % relative to the obtained values in case of automated readings. The proposed methodology is applicable even for other transmission means than Ethernet without significant

modifications.

Moreover, this paper presents a range of experimental results over different switch categories. These values can be advantageously utilized for example in simulations giving a possibility to create detailed data network models. This also applies for OpenFlow switches which are not yet broadly researched. In the OpenFlow part, the method of performance comparison was suggested. The results indicate that OpenFlow has a potential to be deployed even in demanding low-latency networks.

REFERENCES

- [1] L. Cepa, Z. Kocur, Z. Muller, "Migration of the IT Technologies to the Smart Grids", *Elektronika ir Elektrotechnika*, vol. 7, no. 123, pp. 123–128, 2012. [Online]. Available: <http://dx.doi.org/10.5755/j01.eee.123.7.2390>
- [2] IEC 61850-9-2:2011: Communication networks and systems for power utility automation - Part 9-2: Specific communication service mapping (SCSM) - Sampled values over ISO/IEC 8802-3, "International Electrotechnical Commission" Std., Rev. 2.0, 2011.
- [3] T. Hegr, L. Bohac, Z. Kocur, M. Voznak, P. Chlumsky, "Methodology of the direct measurement of the switching latency", *Przeglad Elektrotechniczny*, vol. 89, no. 7, pp. 59–63, 2013.
- [4] ONF, *OpenFlow Switch Specification 1.4.0*, Open Networking Foundation Std., October 2013. [Online]. Available: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/openflow-spec-v1.4.0.pdf#0pt>
- [5] J. Loeser, H. Haertig, "Low-latency hard real-time communication over switched ethernet", in *Proc. 16th Euromicro Conf. Real-Time Systems, (ECRTS 2004)*, 2004, pp. 13–22. [Online]. Available: <http://dx.doi.org/10.1109/emrts.2004.1310992>
- [6] M. Pravda, P. Lafata, J. Vodrazka, "Precision clock synchronization protocol and its implementation into laboratory ethernet network", in *33rd Int. Conf. Telecommunication and Signal Processing*, Vienna, Austria, 2010, pp. 286–291.
- [7] D. Ingram, P. Schaub, R. Taylor, D. Campbell, "Performance analysis of iec 61850 sampled value process bus networks", *IEEE Trans. Industrial Informatics*, vol. 9, no. 3, pp. 1445–1454, 2013. [Online]. Available: <http://dx.doi.org/10.1109/TII.2012.2228874>
- [8] A. Poursepanj. (2003, November) Benchmarks rate switch-fabric performance. CommsDesign. [Online]. Available: <http://m.eet.com/media/1095854/feat1-dec03.pdf>
- [9] IEEE Standard for Ethernet - Section 1, IEEE Std., 2012.
- [10] S. Bradner, "Benchmarking terminology for network interconnection devices", RFC 1242 (Informational), Internet Engineering Task Force, Jul. 1991, updated by RFC 6201. [Online]. Available: <http://www.ietf.org/rfc/rfc1242.txt>
- [11] S. Bradner, J. McQuaid, "Benchmarking methodology for network interconnect devices", RFC 2544 (Informational), Internet Engineering Task Force, Mar. 1999, updated by RFCs 6201, 6815. [Online]. Available: <http://www.ietf.org/rfc/rfc2544.txt>
- [12] JCGM, "Jcgm 100: Evaluation of measurement data - guide to the expression of uncertainty in measurement", Joint Committee for Guides in Metrology, Tech. Rep., 2008. [Online]. Available: <http://goo.gl/ryF5ka>
- [13] H. Chao, *High performance switches and routers*. Hoboken, N.J: Wiley-Interscience, 2007. [Online]. Available: <http://dx.doi.org/10.1002/0470113952>
- [14] A. R. Patwary, B. M. Geuskens, S. L. Lu. "Low-power ternary content addressable memory (tcam) array for network applications", *Int. Conf. Communications, Circuits and Systems, (ICCCAS 2009)*, pp. 322–325. [Online]. Available: <http://dx.doi.org/10.1109/icccas.2009.5250516>
- [15] D. Warren. Switch with adaptive address lookup hashing scheme, February 10 2004. US Patent 6,690,667.
- [16] H. Grecco, "Python VISA library", 2013. [Online]. Available: <https://github.com/hgrecco/pyvisa>
- [17] Floodlight, "Static Flow Pusher API", 2013. [Online]. Available: <http://goo.gl/fMDfaI>