White Rabbit Multi-Point Time Distribution Network

E.F. Dierikx, Y. Xie
 VSL
 Delft, The Netherlands
 edierikx@vsl.nl

A. Savencu
 OPNT
 Amsterdam, The Netherlands

J. Lopez, J.L. Gutiérrez Rivas
 Seven Solutions
 Granada, Spain

Abstract—A nanosecond-accuracy, scalable, multi-point time- and frequency dissemination has been set up in an existing DWDM optical fiber network spanning approximately 200 km, using the White Rabbit (WR) Precision Time Protocol (PTP). Presented novelties in this paper are: a performance study of 5 cascaded WR switches allowing a wide fan-out; reference delay compensation ensuring that the time received by the user is synchronized with the reference timescale within a few nanosecond offset; experiments with a master clock priority algorithm for redundant operation and experimental results of hold-over performance.

Keywords—Time and Frequency transfer; White Rabbit Precision Time Protocol

I. INTRODUCTION

This work, in the framework of the EMPIR WRITE project, was done for bringing accurate optical fiber synchronization techniques from the experimental stage towards a robust and standardized technology ready for adoption by wide range of (mainly industrial) users. Time and frequency synchronization over long-distances using WR [1] has been successfully demonstrated before [2],[3]. WR may be less accurate and stable than other optical fiber synchronization techniques [4], but has the advantages of excellent scalability to larger multi-user networks, ease of integration in data networks and standardization in IEEE 1588 v2.1 [5]. Until recent, WR equipment was missing some important features like reference delay correction and a best master clock algorithm for implementing redundant links. In addition, work is in progress for implementing a hold-over mechanism to maintain accurate time and frequency in case all fiber links fail.

II. MEASUREMENT SETUPS

For performance evaluation of novelties developed in the WRITE project, a multi-point WR network was set up in an existing DWDM optical fiber network. A schematic diagram is shown in Fig.1 and is indicated as "Link 1". The network consists of five cascaded nodes: one grandmaster (GM), three boundary clocks (BC) and one slave node (SL). The network was set up as a round trip with the first and last WR device in the same location to facilitate the performance evaluation.

This project 17IND14 WRITE has received funding from the EMPIR programme co-financed by the Participating States and from the European Union's Horizon 2020 research and innovation programme.

Fig. 1: Schematic diagram of the WR network set up in the WRITE project (orange), with additional links used for redundancy test (green and purple) and three time interval counters (TIC). GM: grandmaster, SL: slave, BC: boundary clock.

For the connections between the nodes, only two DWDM wavelength channels were used in a single bi-directional fiber, allowing other data traffic in the remaining channels. The management and monitoring of the WR equipment was arranged over the same wavelengths as used for the WR-PTP. The total span of the network is approximately 200 km.

Besides Link 1, Fig.1 shows two addition links, Link 2 and Link 3, that were used for experiments with a master clock priority algorithm.

The hold-over function of the WR switch is not yet ready for field tests. Initial tests were performed in a laboratory set-up shown in Fig. 2.

Fig. 2: Schematic diagram of the test setup for the low-jitter WR switch (WR-LJ) with the hold-over extension board connected to a rubidium oscillator.

The device under test is a modified version of a low-jitter (LJ) WR switch that includes an extension board supporting several types of hold-over oscillators. In this test case a Stanford Research System PRS 10 rubidium oscillator was used as hold-over oscillator. Under regular conditions, the WR-LJ is locked to the WR-GM via the WR protocol and the phase-locked loop in the WR-LJ makes the rubidium oscillator lock in both frequency and phase to the WR-GM. In this experimental version of the WR-LJ the hold-over operation can be manually initiated. The time interval counters TIC1 and TIC2 measure the phase difference of the 1PPS output of the WR-LJ and the 10 MHz from the rubidium oscillator with respect to the reference timescale to which the WR-GM is locked.

III. RESULTS

A. Reference Delay Compensation

After careful measurement of the cable delay between the time reference and the input 1 PPS of the WR-GM, the correction value was inserted in the configuration of the WR-GM. The remaining time offset between the time reference and the WR-GM 1 PPS output as shown in Fig.3 is 0.0 ns. The noise in the measurement is dominated by the time interval counter.

Fig. 3: Measured difference between the reference timescale UTC(VSL) and the 1PPS output from the WR-GM.

B. Network Accuracy Performance

The network has been operational for already more than one year. All WR equipment and WDM multiplexers were calibrated before the installation. The fiber delay asymmetry was also calibrated during the installation. The remaining round-trip offset immediately after installation was less than 0.5 ns. Fig. 4 shows results of the round-trip accuracy performance of this WR-link including five cascaded WR switches, more than 1 year after the initial installation and calibration of the link. The remaining offset is still within 2 ns. There is no obvious evidence of drift of the calibration parameters. The change has mainly been caused by the replacement of the time interval counter during the course of the experiment.

Fig. 4: Measured difference between the WR grandmaster and a WR slave after three intermediate WR boundary clocks, after more than 1 year of operation. The initial calibration of the link has changed less than 2 ns. Gaps in the dataset are caused by interruptions of the time interval counter software.

C. Redundant link operation

With the redundant link set-up is shown in Fig. 1. and WR switch with modified firmware to support redundant link operation, several tests have been performed to validate the behavior of the WR slave connected to the redundant links. For this experiment, a phase and frequency offset generator was installed between the time reference and the WR-GM of Link 1. Three types of tests have been performed: link break tests, phase step tests and frequency offset tests.

Fig. 5 shows results from the link break tests. The 1 PPS output signal from the WR-SL with redundant link support is shown with respect to the time reference. Initially, the WR-SL is locked to Link 1. At some point, Link 1 is interrupted and the WR-SL locks to Link 2 which has been assigned the next priority in the configuration. When Link 2 is also interrupted, the WR-SL locks to Link 3 which has been assigned the lowest priority of the available links. Then when Link 2 is restored, after about 1 minute of stability evaluation, the WR-SL locks back to Link 2. And finally, Link 1 is restored as well, thus again after about 1 minute of stability evaluation the WR-SL locks back to Link 1.

Fig. 5: Results of link break tests in a redundant link operation. The upper part of the graph shows when each of the three links is "on" or "off". The lower part of the graph shows changes in the 1PPS output of the WR-SL locking to the available link with the highest priority.

Then the experiment is repeated with similar response. For visibility, the links have been given intentional phase offsets.

Fig. 6 Shows results of redundant link operation when phase steps occur in the main link. In the configuration file of the redundant slave switch, limits for phase deviation have been set to +100 ns and -100 ns. This means that if the main link (Link1) deviates by more than these limits from the two other links, the slave will lock to the link with the next highest priority, which is Link2 in our experiment. The limits can be chosen by the user but should be reasonable with respect to the level of synchronization of the reference timescale for the three redundant links. In our experiment, Link1 has an intentional offset of approximately -40 ns from Link2 and Link3. In Fig. 6, the upper part of the graph shows the offset of the WR-GM from the reference timescale, UTC(VSL), and the lower part of the graph shows the response of 1PPS output of the redundant slave to the phase steps introduced to the Master of Link1. At 12:30h, a phase step off -60 ns is introduced in the Master of Link1. Initially, the Slave follows this step, but when the Slave's offset with respect to Link2 and Link3 exceeds -100 ns, the Slave locks to Link2. At 12:35h, the phase of the Link1 Master is set back to its initial offset of -40 ns. The Slave takes some time to evaluate the link differences and after about 40 s it restores the lock to Link1. Other steps of -100 ns, +300 ns and -200 ns have been introduced to the Master of Link1, and the response of the Slave is as expected.

Fig. 6: Results of phase offset tests in a redundant link operation. The upper part of the graph shows the phase offset of the Masters with respect to the reference timescale UTC(VSL). The Master of Link3 is not shown because it has the same offset as Link2. The lower part of the graph shows changes in the 1PPS output of the Slave locking to the link with the highest priority that has a phase offset within the given limits.

In the next experiment, frequency offsets have been introduced to the Master of Link1. In this case, the behavior of the redundant switch is similar to the behavior in the previous experiment. That means that the Slave doesn't evaluate the frequency offsets of the Master links, but still evaluates the phase offsets. A frequency offset in the Master of Link1 causes a gradual increase of the phase offset of this Master and the Slave with respect to the two redundant Masters. When the phase offset of Slave exceeds the limit, the Slave will lock to

the Master with the next highest priority. Some experimental results are shown in Fig. 7. At about 12:55h, a relative frequency offset of -1×10^{-9} Hz/Hz is introduced to the Master of Link1. Initially the Slave follows the increasing phase offset of the Master, but when the phase offset exceeds the limit of -100 ns with respect to the two other Masters, the Slave locks to Link2. When the frequency offset of the Master of Link1 is reversed to $+1 \times 10^{-9}$ Hz/Hz, the Slave will wait until the phase offset of Link1 has reached half of the limit, so in this case -50 ns, before locking back to Link1.

Fig. 7: Results of frequency offset tests in a redundant link operation. The upper part of the graph shows the phase offset of the Masters with respect to the reference timescale UTC(VSL). The Master of Link3 is not shown because it has the same offset as Link2. The lower part of the graph shows changes in the 1PPS output of the Slave locking to the link with the highest priority that has a phase offset within the given limits.

D. Hold-over results

In case a WR Slave or Boundary Clock loses all connections to its Masters, as a last resort, it could run in "hold-over" mode. In this case, the time and frequency provided from the Slave to the end user are only based on the stability of the internal oscillator in the Slave. In the current generator of commercially available WR devices, hold-over mode is not well supported. The internal oscillator is typically a crystal oscillator which is not selected for being extremely stable on its own, and, in addition, in case of loss of the Master link, the frequency steering voltage will typically run either to its minimum or maximum setting, giving the oscillator a large frequency offset. To support hold-over operation, two types of development are required: hardware development allowing new types of oscillators to be plugged in, and firmware development supporting hold-over operation.

The hardware support has been realized by means of a printed circuit board that can be mounted in a low-jitter WR switch.

The firmware is under development. So far, for experimental purposes, basic functions have been implemented for adjusting the proportional and integral gain of the PLL and learning buffers have been implemented for storing that last phase error measurements and corresponding steering values

over a specified time with a specified sampling interval. The data stored in these buffers during regular operation is to be used for creating predictions of the frequency steering voltage during hold-over operation. The prediction algorithm may need further improvement. The best results achieved so far are from a hold-over function that repeatedly replays the steering values that are in the buffer. The results from initial tests with a Stanford Research PRS 10 rubidium oscillator are shown in Fig. 8. The buffer size is 256 samples and the sampling rate is 1 sample per second. After at least 5 minutes of learning, the Slave is set to hold-over operation. The lines in Fig. 8 show for four experiments how the phase of the Slave clock drifts off during 4000 s of hold-over operation.

Fig. 8: Preliminary test results of a low-jitter WR switch in hold-over operation with a commercial rubidium oscillator as hold-over oscillator.

Fig. 8, shows that in this limited set of measurements the clock drift is typically less than 50 ns within one hour, which corresponds to an average frequency offset at the level of

1×10^{-11} Hz/Hz. This is a very useful result for several applications including telecommunication, electrical grid operation, time stamping of financial transactions, etc.

IV. CONCLUSIONS

This paper shows that developments in the WRITE project have resulted in WR equipment becoming more suitable for adoption in industrial applications. The delay reference compensation allows that the end user receives a copy of the reference time within an accuracy of less than 1 ns. The redundant link support feature ensures that the end user receives time from an alternative source in case a fiber link is interrupted, and ensures that the time offset is kept within specified limits in case one of the master clocks deviates too much. Preliminary tests of hold-over operation have shown that with a commercial rubidium oscillator, the time offset can be kept within 50 ns over 1 hour.

V. REFERENCES

- [1] White Rabbit project, <https://ohwr.org/project/white-rabbit/wikis/home>
- [2] E. Dierikx et al., "White Rabbit Precision Time Protocol on Long-Distance Fiber Links," in *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 63, no. 7, pp. 945-952, July 2016, doi: 10.1109/TUFFC.2016.2518122.
- [3] E. Dierikx et al., "Optical Fiber Trans-national Time Transfer for Comparing Two UTC Realisations," 2019 Joint Conference of the IEEE International Frequency Control Symposium and European Frequency and Time Forum (EFTF/IFC), Orlando, FL, USA, 2019, pp. 1-2, doi: 10.1109/FCS.2019.8856091.
- [4] P. Krehlik, Ł. Śliwczyński, Ł. Buczek and J. Kołodziej, "Fiber-Optic UTC(k) Timescale Distribution With Automated Link Delay Cancellation," in *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 66, no. 1, pp. 163-169, Jan. 2019, doi: 10.1109/TUFFC.2018.2878319.
- [5] IEEE 1588-2019 - IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems