PERFORMANCE OF IEEE 1588 IN LARGE-SCALE NETWORKS

Georg Gaderer, Nataša Simanić, Patrick Loschmidt, and Bojan Ćorić
Institute for Integrated Sensor Systems, Austrian Academy of Sciences
Viktor Kaplan Straße 2, A-2400 Wiener Neustadt, Austria
{*Georg.Gaderer, Natasa.Simanic, Patrick.Loschmidt, Bojan.Coric*}@*oeaw.ac.at*

Abstract

For industrial application purposes, the performance of the already well-established IEEE 1588 protocol represents a significant step forward with respect to NTP. The high precision and the determinism offered by this protocol are two of the main reasons for considering its deployment. In large-scale computer networks, the Network Time Protocol forms the counterpart as a protocol suite for precision clock synchronization within the Internet. The particular strength of this protocol lies in the well-engineered control algorithms and its wide distribution. This paper investigates the possibility of using IEEE 1588 for clock synchronization over the Internet. For this sake, experiments of synchronizing clocks between two distant locations are presented. The experimental setup employs a reference clock steered to GPS time and a slave clock synchronized to the reference via the Internet. Finally, the output of the slave is compared to the output of a GPS timing receiver. In order to have a competitive comparison with other protocols, the PTP clocks are controlled by a servo from an industrial application. This is mandatory, as IEEE 1588 does not specify any control algorithms. The results show that IEEE 1588, which allows for high synchronization rates, can help in solving the problem of synchronizing the nodes.

INTRODUCTION

Synchronization of clocks in computer networks obtained a new application with the wide and high broadband availability of the Internet. As in the Internet, devices are usually connected to the global network; it is practically possible to synchronize the clocks of these devices. This world-spanning network also has the advantage that, regardless to its location, a device can always request time service from a known server, which again can be positioned practically anywhere on the world.

However, as the packet-oriented structure of the Internet was never designed to provide any means of guaranteed delivery times, QoS, or even jitter restrictions, the distortion of the synchronization cannot be determined in a formal way. It is obvious that this uncertainty heavily depends on the distance in terms of router and link hops between server and a client, i.e., one would expect less jitter if the server is connected directly in the local network, than when it is hundreds of kilometers away. To determine the best protocol for such purposes, several candidates have to be discussed.

This paper analyzes the applicability of the IEEE 1588 protocol (which was originally designed for local area networks) to synchronize computer clocks over large distances. For this matter, the results of an experiment are presented where two remote clocks were synchronized, one residing in Wiener Neustadt, Austria, and the other one in Melbourne, Australia.

PRECISION TIME PROTOCOL (IEEE 1588)

The IEEE 1588 standard **[1]** defines a set of rules to exchange timing information and measure transport delay within computer networks. It was developed with a specific application area in mind, test and measurement, where simple deployment and a strict hierarchical approach are of great interest. While certain features ease set up in LANs, e.g. automatic master discovery or multicast communication, other beneficial properties regarding usage in WANs were added later or are still not defined. Among others, unicast communication, necessary to directly address a single node in the Internet, was added in version 2 of the standard, while the actual synchronization algorithm, as well as timestamp filtering, are still free to be implemented by the software designers. Compared to NTP, there is no defined control loop, and methods that help to detect faulty or inaccurate servers are not standardized. Furthermore, PTP uses explicit delay measurements at a lower rate than the Sync messages, which is a significant difference from NTP. The latter always determines (at a lower rate than PTP) offset and delay in a single measurement. The influence of this distinction between the two protocols is not yet known. This paper takes an experimental approach to quantify the properties of a large-distance (Internet-based) communication channel and to evaluate their influence on the achievable performance of PTP.

TEST SETUP

To be able to compare the outputs of the remote clocks, one needs to have a globally available time base. For this sake, GPS was chosen, as the jitter between two GPS receivers is, in the worst case, a few hundred nanoseconds, which is far below the jitter of the synchronization over the Internet. The test setup depicted in the first figure consists of two stations, where one is configured as an IEEE 1588 master (Austria) and the other as a slave (Australia). The Austrian master uses GPS to steer its local clock using the 1 PPS output of the receiver and a serial correction via the NMEA protocol. The measurements show that the accuracy is within a few nanoseconds with respect to GPS time. The second station is equipped with a LANTIME M600 GPS **[2]** receiver. However, this station does not synchronize to GPS time, but to IEEE 1588 time over the Internet. Finally, the PTP time of the slave is compared to the 1 PPS timestamp of GPS. This comparison can be treated as the absolute offset between the master and slave, due to the fact that the master is almost ideally synchronized to GPS.

The two clocks and the IEEE 1588 hardware used for this experiment are the commercially available Syn1588 network cards from Oregano Systems. They not only feature an IEEE 1588 hardware timestamper, but also a 1 PPS output and event timestamper interfaces for highly accurate comparison to the reference time. These cards are built as standard PCI interface cards, allowing integration into standard PCs or into industrial grade fanless systems, as in the presented application **0**.

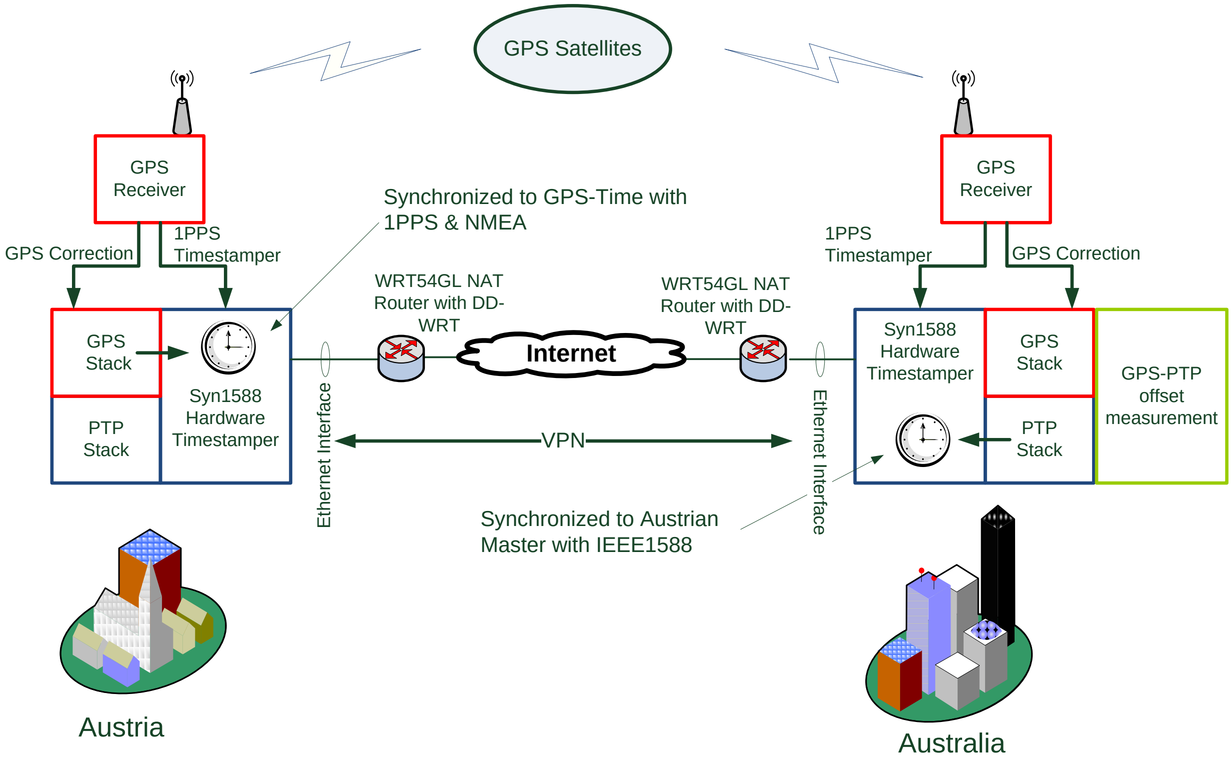

Setup of the test environment.

An important detail of the setup is the interconnection between the two PTP nodes. As the IEEE 1588 standard primarily foresees synchronization based on multicast, the experiment was performed over a VPN in a tunnel to have both nodes (master and slave) in one multicast domain. Alternatively, also unicast (available since PTP version 2) could be used, but this hinders some benefits of PTP, like automatic master discovery. For the sake of simplicity, a COTS router, namely a Linksys WRT54GL, with the open source DD-WRT firmware image, were used in the setup.

RAW TIMESTAMP MEASUREMENTS

As a first step, to get an idea of the behavior of the network in between Austria and Australia, the raw error of the timestamps can be recorded and analyzed. For this analysis, two kinds of measurements can be considered:

- The actual timestamps, carrying the send-time at the master
- The result of the delay measurement, which is, according to IEEE 1588, a combination of a delay-request timestamp and an ordinary synchronization message.

For the first type, the synchronization message offset was logged for 10^5 s (or 83 hours) and is presented in the figure below. The diagram shows two interesting observations: while most of the timestamps are well below 0.2 seconds, a relatively high number of observations show offsets of up to 0.8 seconds.

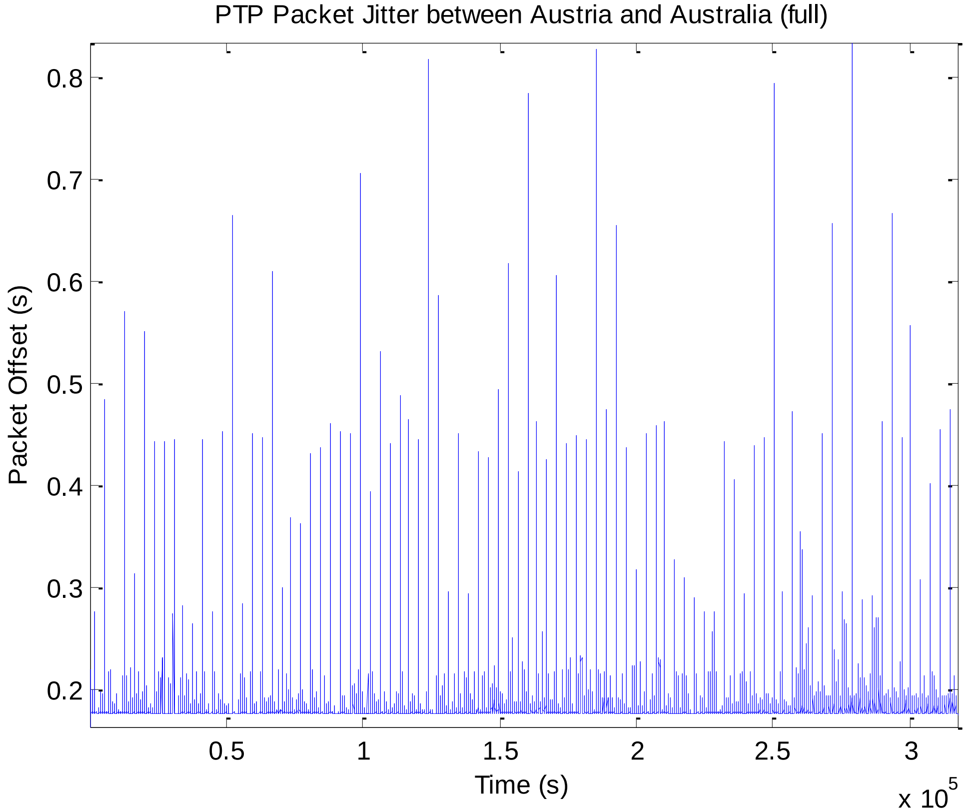

Offset between the actual time and the timestamp carried in the synchronization message at the slave.

The relatively long observation time of the figure above can mislead in terms of the distribution of these outliers. One could interpret the peaks to be regular and, thus, to have a deterministic source. However, a zoom into the timescale (shown in the zoom in the next figure) reveals that the spikes are not regular, neither in terms of their occurrence nor in their values. This figure also shows that the delay is most of the time far below the estimated 200 ms mentioned above, as it is obviously well below 20 ms. This fact and the distribution of the measurements are manifested even more in detail in the histogram on the next page. This graph shows the distribution of the measurements. It can be seen that the distribution is limited in terms of a lower boundary and can be estimated with an Erlang distribution.

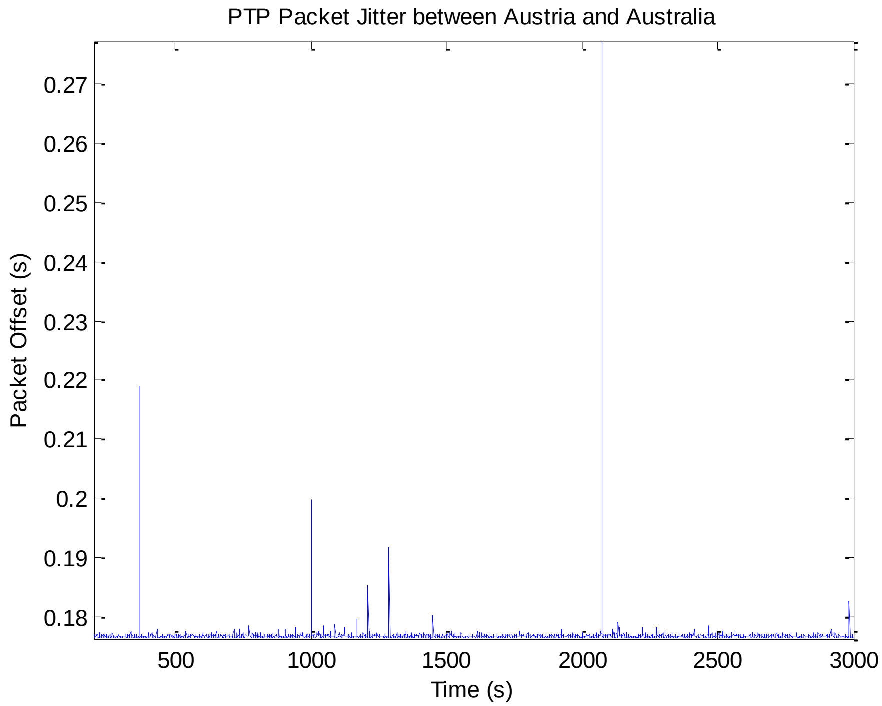

Zoom of the previous figure of packet delay offsets. Samples were taken once a second.

This lower boundary can be descriptively explained by the fact that packets cannot arrive earlier than the speed of light and the minimum processing delay at each hop. The asymmetry to longer packet offset times is then, in contrast to this, caused by processing delay, which can happen theoretically at each hop throughout the network **[4]**.

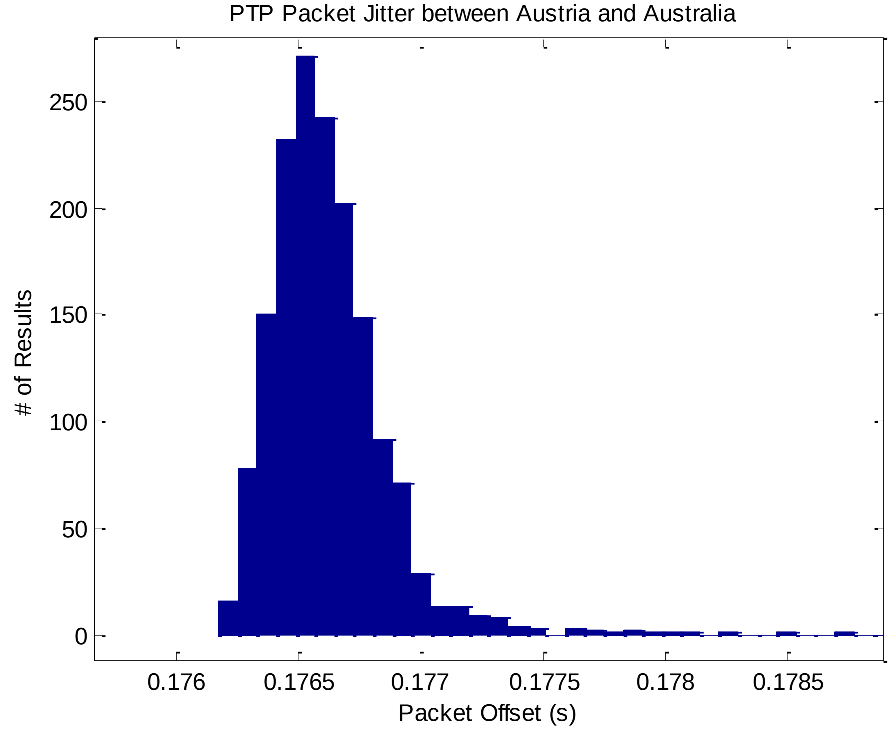

Histogram of the packet offset.

From these measurements, several conclusions can be drawn:

- The timestamps without transport delay correction are accurate to a range of some tenths of a second with an additional jitter of a few milliseconds.
- The distribution of the jitter cannot be approximated with a normal distribution; instead, proper filtering has to be applied (for example, outlier filtering). This filter can be, for example, nonlinear **[1]**.
- In order to smooth the delay measurement error, the timestamps should also be filtered with a low-pass filter.

CONTROLLER DESIGN

The conclusion of the previous chapter can be taken as a starting point for the design of a controller to steer the clock. For the design of a proper controller, it is probably most important to remove the outliers. For the proposed system, the outlier deletion has to be done for two kinds of measurements: the synchronization offset and the delay measurements. The delay measurements are partly done by the synchronization messages and partly by specialized messages called Delay Request messages. Those are initiated towards the master. In principle, these messages are the same as Sync messages.

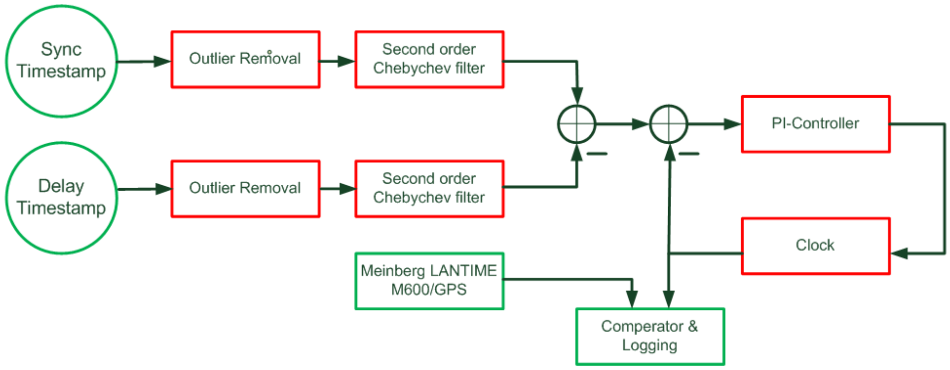

Controller structure.

The outlier-filtering algorithm is the following: First a sliding window of 100 samples is defined. This window determines a mean value and standard deviation for all values by sliding over the historical data. Using this, a hull curve is composed by the moving average and standard deviation of the window. Finally, all data which are above this hull curve are removed. This is applied subsequently to all data.

The outlier-free sampled data are then, as shown in the picture above, filtered with a second- order Chebychev filter for several reasons. On one hand, the data, even without the outliers, are still too noisy; on the other hand, the high dynamics of the measurements, which are done per default once a second, can be removed considering the short-term stability of the oscillator. As the Oregano Systems network cards are equipped with a standard 50 ppm crystal oscillator, a narrowband filter can be used to remove the noise. Finally, the filtered synchronization messages and delay requests can be taken as the input of the PI-controller to steer the clock. The effects of this nonlinear filter are presented in the graph on the next page, showing original and resulting data.

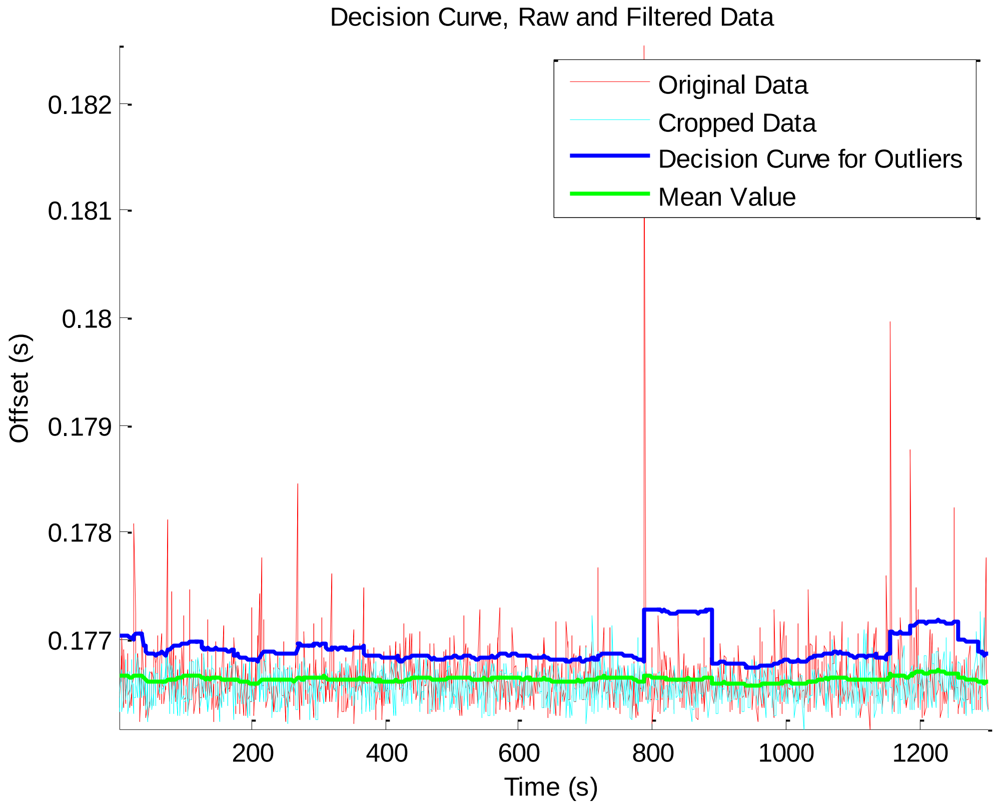

Results of outlier filtering.

RESULTS

Taking the scheme described above, synchronization between Austria and Australia has been established. The final graph shows the delay error, the synchronization message offset, and the overall time error. The time error shows a relatively smooth, low jittering behavior with accuracy below 1 ms. The figure also depicts that the delay can vary significantly. For example, at around 4000 s, the delay jumps about half a millisecond, which is not likely in LANs. Explanations for the change of the delay behavior are, of course, difficult to provide; in some cases during the experiment, it could be noticed that they were caused by sudden changes of the number of hops between the master and the slave. As the connection in between the two endpoints was not deterministic for the experiment, an exact explanation, however, cannot be given.

The analysis of the delay request timestamps with respect to the timely related synchronization messages can provide a second result: As even the time on the network for two of such closely related messages is typically not the same, the residual offset of around 30 µs (notably throughout the experiment and not in that picture) indicates a systematic error. It results from the directional measurement of the delay, which obviously shows the timely asymmetry of the channel. It is interesting to mention that the error (e.g., range of 10 µs) also changed during the duration of the investigation.

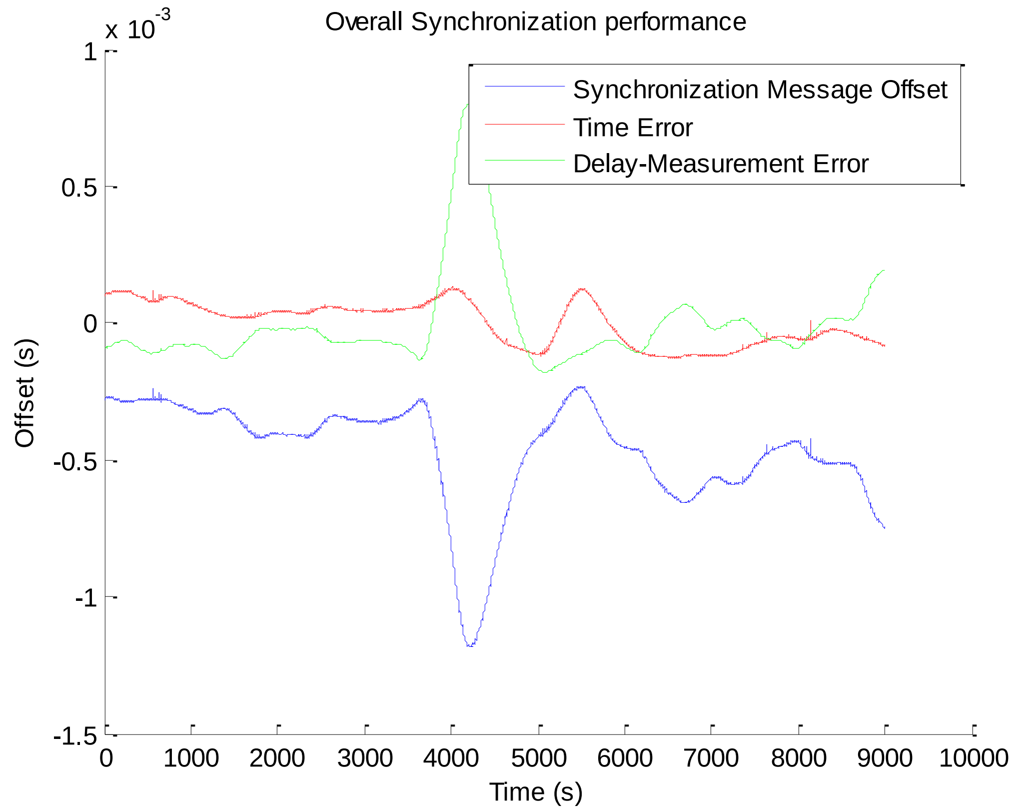

Synchronization performance.

CONCLUSION AND FURTHER WORK

In conclusion, it can be said that synchronization using IEEE 1588 over long-haul networks is reasonable in terms of the resulting accuracy. One problem is the definition of proper filter algorithms in order to cope with outlier detection caused by (in terms of delay) a noisy channel, as this is not specified in the PTP standard. In addition to that, control algorithms and data pre-filtering are also not defined. An important issue for the seamless integration of PTP over large-distance networks is the establishment of a VPN. Only in that case it is possible to benefit from a single PTP domain and, e.g., master failover. Finally, with such a system it was possible to reach a synchronization accuracy between Austria and Australia of a 32.5 µs mean and a 9.58 µs standard deviation over an observation period of 83 hours. Future work on this topic will include additional analysis of the jitter contribution to measurements, for example within Europe, from one office to another or to other continents like America. Finally, as a next step, more sophisticated synchronization algorithms will be deployed, such as the one proposed by Hadžić and Morgan **0**.

ACKNOWLEDGMENTS

The authors wish to thank Julien Ridoux, the University of Melbourne, Oregano Systems, and Meinberg for their vital support conducting the measurements needed for this paper.

REFERENCES

[1] "*IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems,"* IEEE Std 1588-2008 (Revision of IEEE Std 1588-2002), 269 pp.

[2] H. Gerstung, 2008, "*Synchronizing PTPv1 and PVPv2 clients with one common time source,"* in Proceedings of the 2008 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control and Communication (ISPCS), 24-26 September 2008, Ann Arbor, Michigan, USA (IEEE), pp. 1-6.

[3] Oregano Systems, "*Syn1588 PCI NIC Brief Datasheet*," available at: *http://www.oregano.at/pdf/syn1588_pci_nic_briefds.pdf*, accessed 20 December 2010.

[4] J. Burch, K. Green, J. Nakulski, and D. Vook, 2009, *"Verifying the Performance of Transparent Clocks in PTP Systems,"* in Proceedings of the 2009 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control, and Communication (ISPCS), 14-16 October 2009, Brescia, Italy (IEEE), 6 pp.

[5] I. Hadžić and D. R. Morgan, 2009, *"On Packet Selection Criteria for Clock Recovery,"* in Proceedings of the 2009 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control, and Communication (ISPCS), 14-16 October 2009, Brescia, Italy (IEEE), 6 pp.