ISPCS 2008 – International IEEE Symposium on Precision Clock Synchronization for Measurement, Control and Communication
Ann Arbor, Michigan, September 22–26, 2008

Spider Transparent Clock

John C. Eidson[1], Andrew Fernandez[1], Bruce Hamilton[1], Jad Naous[2], Dieter Vook[1]
[1]Agilent Technologies, 5301 Stevens Creek Boulevard, Santa Clara, CA
[2]Stanford University, Stanford, CA

Abstract *– This paper discusses a device, the spider transparent clock, which can be used to retrofit existing bridges and routers to allow them to deliver highly accurate time using the IEEE 1588 protocol. The spider transparent clock corrects for the internal queuing jitter and asymmetry introduced by these bridges and routers.*

Keywords *– Transparent Clock, IEEE 1588, Jitter, Legacy Networks, High Accuracy*

I. INTRODUCTION

IEEE 1588 is a protocol that provides a way for devices on a LAN to maintain synchronized time with errors as low as a few nanoseconds [1]. However, to attain this accuracy requires hardware support at both the end devices and at intermediate bridges or routers.

Without IEEE 1588-capable bridges or routers[1], the indeterminate packet residence time due to queuing in the bridges and routers introduces jitter and asymmetry in the link delay and clock offset measurements. These link delay and offset errors can increase timing errors in the slave clocks to tens of microseconds or more [2].

For many applications, timing errors this large are unacceptable. For example, transmission of professional uncompressed video requires better than 1μs absolute synchronization between devices, and phase jitter filterable down to 100ps [3]. Some 2G and 3G wireless base stations require phase offset below 10μs [4]. Military and test and measurement applications require time accuracy below 10ns in critical applications [5].

Replacing an Ethernet bridge or router with an IEEE 1588 specified boundary or transparent clock can cost tens of thousands of dollars. This paper discusses a device, termed a spider, which allows a few ports of an ordinary bridge or router to act as a transparent clock. The spider should be particularly attractive in networks where only a few of the intermediate or end devices need highly accurate time.

II. SPIDER TRANSPARENT CLOCK

A spider is a network device that intercepts and modifies IEEE 1588 packets arriving and leaving selected bridge or router ports. Like a transparent clock, the spider measures a packet's residence time within the bridge or router and then modifies the egress packet's correction field as specified in clause 11.5 of IEEE 1588-2008. The previously indeterminate residence times of the packets in the bridge are now known, and jitter and asymmetry in the link delay measurement is greatly reduced.

The basic architecture of the spider transparent clock is shown in Figure 1. Shown is a 4-port spider that makes ports 1–4 of the bridge function like an IEEE 1588 end-to-end transparent clock with the attendant reduction in timing jitter and asymmetry. Ports 5–8 function like an ordinary bridge and will introduce timing jitter and asymmetry due to the uncorrected queue delay variations.

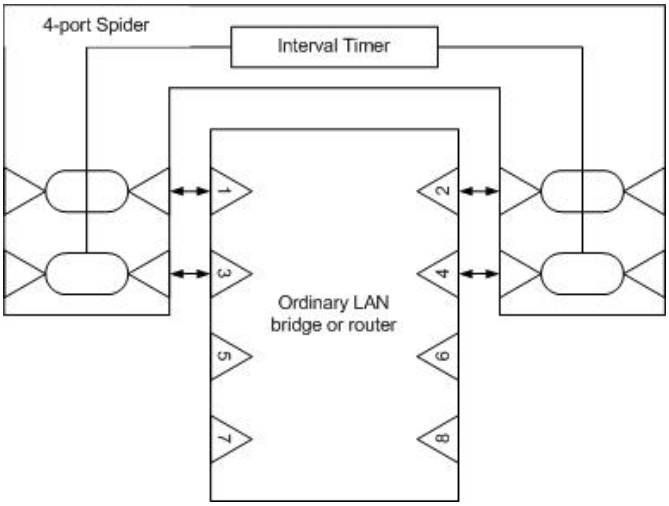

Fig. 1. Spider transparent clock block diagram

A network designer can place spiders at selected ports on the bridges and routers in a network to create highly accurate timing paths without replacing existing infrastructure or requiring control of network traffic loads. Figure 2 illustrates the use of the spider to design a highly accurate timing path between several devices that is independent of the cross traffic in the network.

For example, if the traffic source in Figure 2 generates variable traffic directed at traffic sink-2 and traffic sink-3, this traffic will share links H and J with the IEEE 1588 event messages passing between the grandmaster and ordinary clocks. In the absence of the spiders, the resulting queuing delays would degrade the synchronization at all ordinary clocks. However the presence of the spiders on the ports common to the path between the grandmaster and ordinary clocks 2 and 3 eliminates the effects of this queuing in the bridges. Ordinary clock-1 will not be degraded by queuing in

[1] IEEE-1588 capable bridges and routers are called either boundary or transparent clocks depending on their design [1].

bridge-1 but will be degraded by queuing in bridge-2 since port-7 on bridge-2 is not protected by the spider.

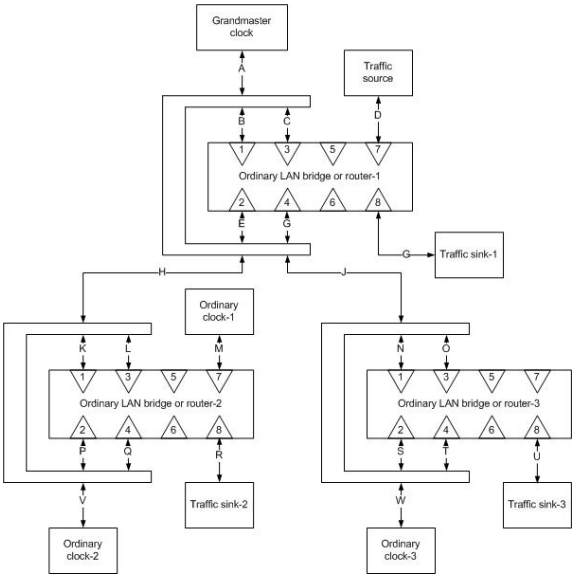

Fig. 2. Network Using Spiders

III. INTERNAL ARCHITECTURE

The internal architecture of each spider port is illustrated in Figure 3.

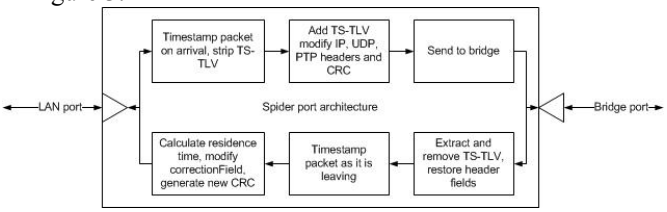

Fig. 3 Spider port internal architecture

The spider operates by measuring the actual residence time of IEEE 1588 packets as they traverse the bridge or router. The timing information is transported from ingress to egress in a special IEEE 1588 TLV extension appended to IEEE 1588 event packets. This extension is called a TS-TLV and is an IEEE 1588 organization specific TLV [1]. The TS-TLV is not recognized by, and thus is ignored by all IEEE 1588 devices except spiders designed by the vendor defining the TS-TLV [1]. All spider ports share access to a common timer for generating ingress and egress timestamps. It is important that all ingress and egress ports have equal latency paths to the common spider timer to avoid the introduction of asymmetry by the spider itself.

An IEEE 1588 event packet entering the spider from the LAN is timestamped at the ingress port. Any existing TS-TLV is removed. Next a TS-TLV is added to the packet with the ingress timestamp inserted by the spider hardware into the timestamp field. The IP, UDP, PTP headers and CRC of the packet are modified to create a legal Ethernet packet. The modified fields are indicated in bold and underlined font in Figure 4. The packet is then passed to the bridge where it is forwarded according to the addressing rules of the network, just like any other data packet. At the egress port of the spider the ingress timestamp is noted, the TS-TLV is removed and the IP, UDP, and PTP headers restored. The event packet is timestamped on egress and the residence time computed and added to the correctionField. Finally a new CRC is generated. Since the residence time is measured for all IEEE 1588 event packets, i.e. Sync and Delay_Req, the spider also corrects for any internal asymmetry in the up and downlink paths through the bridge or router.

An IEEE 1588 event packet is 24 octets longer between the ingress and egress ports of the spider and during transit of the bridge than it was at ingress to the spider. This additional length is accommodated preferably by using octets from the interframe gap, if available, while preserving the minimum 12 octet interframe gap required by the Ethernet standards. If the network loading does not permit this, some packets must be dropped, or the spider must include a queuing mechanism on the bridge side of ingress ports.

All of these operations can be done in an ASIC or FPGA without requiring a microprocessor. If syntonization of the spider clock to the grandmaster is needed, then the ingress port receiving grandmaster originated IEEE 1588 Sync events provides the Sync event packet contents and ingress timestamp to the spider's common timer to syntonize the timer rate to that of the grandmaster.

As discussed in IEEE 1588, syntonization reduces errors in the measurement of residence times that result from the difference between the definition of the second in the spider and the grandmaster. This is important for high accuracy systems. For example, a residence time of 1 ms and a difference of 0.01% in the definition of the second between the spider and grandmaster results in an error in the measured residence time of 100 ns. Spiders can support multiple IEEE 1588 domains using the same syntonization techniques discussed in IEEE 1588-2008 for transparent clocks [1].

The format of the IEEE 1588 event packets delivered by the spider to the bridge ingress port and received by the spider from bridge egress ports is illustrated in Figure 4.

Destination MAC (48)						Source MAC HI (16)	
Source MAC LO (32)				Ethertype (16)		IPV (4) IHL (4)	TOS (8)
IP Total Length (16)		ID (16)		Flags & Fragmentation (16)		TTL (8)	IP Protocol (8)
IP Header Checksum (16)		IP Source (32)				IP Destination HI (16)	
IP Destination LO (16)		UDP Source Port (16)		UDP Destination Port (16)		**UDP Length (16)**	
UDP Checksum (16)		tran (4)	1588 type (4)	Reser (4)	Ver (4)	**1588 Message Length (16)**	Domain (8)	Reserved (8)
1588 Flags (16)		correctionField HI (48)					
correctionField LO (16)		Reserved (32)				sourcePortIdentity HI (16)	
sourcePortIdentity LO (64)							
sequenceId (16)		control (8)		LMI (8)		originTimestamp HI (32)	
originTimestamp LO (48)						tlvType (16)	
Tlv_length (16)		organizationId (24)				organizationSubType (24)	
Spider UUID (48)						Spider RX Timestamp HI (16)	
Spider RX Timestamp LO (48)							

Fig. 4 TS-TLV format inside the bridge

The TS-TLV is the last 24 octets shown in Figure 4 outlined by the heavy lines. The Spider RX Timestamp fields contain the ingress timestamp. The UUID field contains an

identifier unique to each spider which allows the egress port to ignore TS-TLVs that did not originate from an ingress port on the same spider.

IV. SYSTEM ISSUES

The spider transparent clock is designed to create highly accurate timing paths through a subset of the ports on an ordinary bridge or router as illustrated in Figure 2. In Figure 2 the high accuracy path includes ports 1–4 of each of the bridges shown in the figure. Thus ordinary clocks 2 and 3 have a high accuracy path to the grandmaster clock.

Ordinary clock-1 is attached to port-7 which is not serviced by the spider and will therefore suffer from jitter introduced by bridge-2 but not from jitter introduced by bridge-1. The spider attached to port-1 of bridge-2 will append a TS-TLV to IEEE 1588 event packets entering the bridge. Since port-7 is not serviced by the spider, the TS-TLV will be appended to IEEE 1588 event packets received by ordinary clock-1. Since ordinary clock-1 ignores the TS-TLV and the remainder of the packet is both legal and reflects correct IEEE 1588 information, clock-1 can correctly synchronize to the grandmaster, albeit with degraded time accuracy.

The spider is intended for use in network topologies where the paths from the grandmaster always enter a bridge or router via a spider port. However, even in networks where this is not the case, see for example Figure 5, the spiders process the IEEE 1588 event packets without protocol error.

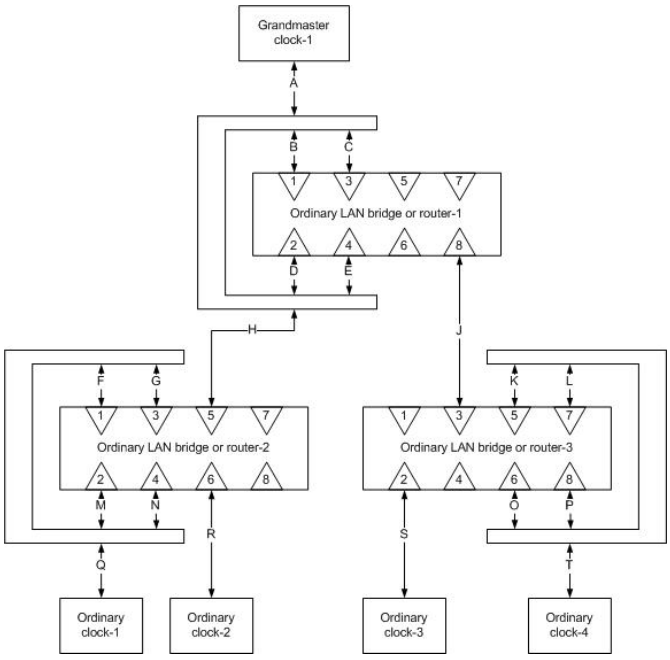

Fig. 5 Non-optimum Network

For example, in Figure 5 the event packets entering port-5 of bridge-2 do not have a TS-TLV appended. However, these packets have been corrected by the bridge-1 spider for residence time jitter and asymmetry in bridge-1. Ordinary clock-1 will receive event packets via the bridge-2 spider. The egress spider port between path M and Q will not find a TS-TLV appended to downlink event packets and will simply pass the packets to clock-1. On the uplink, the ingress port between paths H and D will strip the TS-TLV as described earlier. Therefore both clocks 1 and 2 will experience the full jitter introduced by bridge-2 but no jitter from bridge-1.

Similarly, in Figure 5 the event packets entering port-3 of bridge-3 do contain a TS-TLV appended by the spider ingress port of bridge-1. The egress spider port between path P and T will ignore this TS-TLV based on the Spider UUID field and will simply pass the packet to clock-4. Both clocks 3 and 4 will be degraded by the jitter introduced by both bridge 1 and 3 since the event packets do not pass through both ingress and egress ports of the same spider.

Normally only a single spider is used with each bridge or router. However it is acceptable to have a spider span multiple bridges or routers as illustrated in Figure 6.

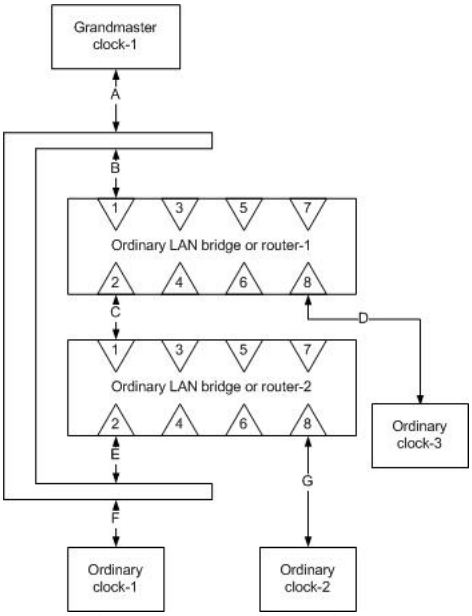

Fig.6 Spider Correcting Multiple Bridges or Routers

In Figure 6, the path between ordinary clock-1 and the grandmaster is fully corrected for jitter and asymmetry introduced by both bridges. However, both clocks 2 and 3 will be degraded by the jitter in the bridges since the event packets do not pass through both the ingress and egress ports of the same spider.

The design of the spider in Figure 6 should have all ingress and egress ports in close proximity to permit these ports to have equal low latency paths to the common timer in the spider. Therefore for network topologies such as shown in Figure 6, the connections between the spider and the bridges, e.g. paths B and E, will require longer cables than used in Figure 2. However the latency in these cables appears to the spider as part of the latency of the bridges and will be included in the measured residence time.

V. PERFORMANCE MEASUREMENTS

A two-port spider similar to Figure 1 was prototyped using a NetFPGA board configured for 100-BaseT Ethernet [6]. The design is such that these boards can be daisy chained to create spiders supporting more than two bridge ports. Performance measurements were made using the network design shown in Figure 7.

The traffic source broadcasts traffic to all bridge ports. The synchronization between the grandmaster and the ordinary clock was observed by measuring the time difference between 1 PPS signals generated by the clocks. The 1 PPS signals are coincident with the rollover of the fractional-seconds portion of the IEEE 1588 clock in each device. The LSB of the spider is 8 ns and the LSB of the IEEE 1588 clocks is 10 ns.

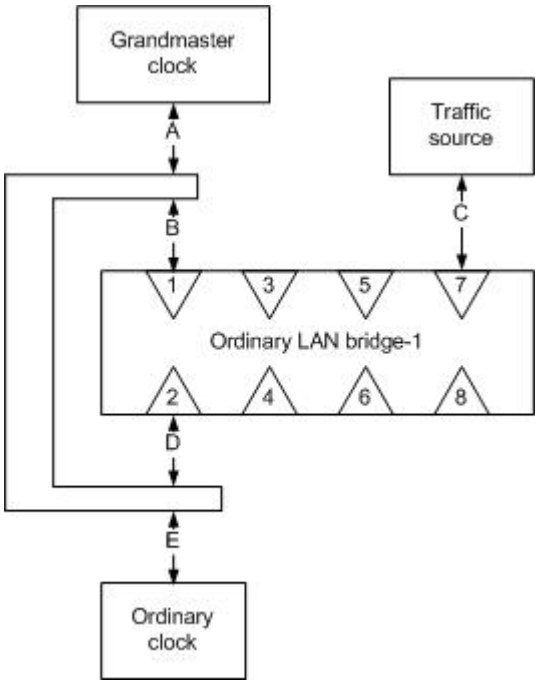

Fig. 7 Experimental topology

As a baseline, the synchronization of the two clocks connected directly using a crossover cable with no intervening bridge is illustrated in Figure 8. For this measurement the two clocks shared a common instrument grade oscillator to eliminate all but quantization errors in the two clocks. In the remainder of the measurements the grandmaster and ordinary clocks each had an instrument grade oscillator. The oscillator in the spider is a low grade oscillator typical of those used in a PC. The measurements shown in Figures 10 – 11 are limited by jitter introduced by the oscillator in the spider, errors due to the servo characteristics of the slave clock, and the effects of the intervening bridge and spider.

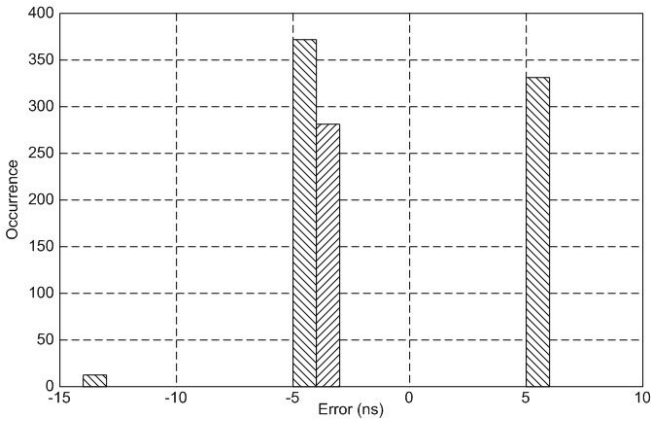

Fig. 8 Synchronization using a crossover cable

Next the degradation due to the bridge loading by the traffic generator was measured with the results shown in Figure 9. In this case the two clocks were connected directly to the bridge with no spider present. The data in Figure 9 is therefore a good indication of the jitter introduced by the loaded bridge. The bulk of the data lies within ±40us but with significant outliers to ±70us.

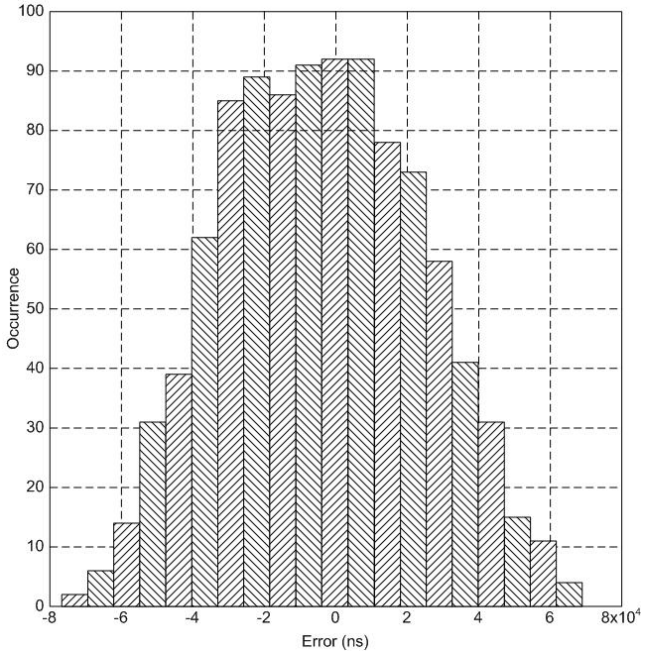

Fig. 9 Synchronization using a loaded bridge (no spider)

The improvement provided by the spider is illustrated in Figure 10. The measurements of Figure 10 were made with an un-syntonized spider. The bulk of this data lies within ±400ns of the mean but with significant outliers to ±800ns.

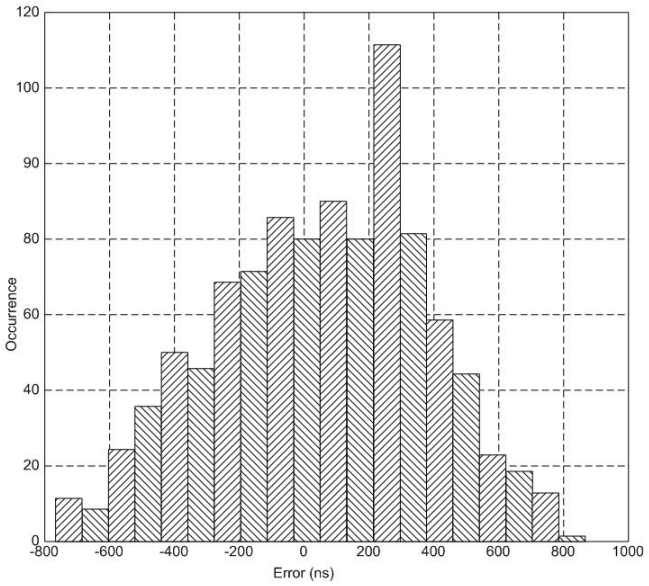

Fig. 10 Synchronization using a loaded un-syntonized spider

The data in Figure 10 represents a significant improvement over the data in Figure 9, illustrating the removal by the spider of jitter introduced by the loaded bridge. The width of the distribution has been reduced approximately 2 orders of magnitude. An offset of the mean of approximately 80ns was hidden in the data of Figure 9 but is visible in Figure 10.

Finally, the performance with a syntonized spider is shown in Figure 11. The bulk of this data lies within ±80ns of the mean with only a few outliers to ±800ns. In addition the shift in the mean has been reduced from approximately 80ns for the un-syntonized spider to less than 25ns for the syntonized spider. If the measurement of the residence times by the spider was perfect, one would expect the width of the distribution in Figure 10 to resemble the data in Figure 11.

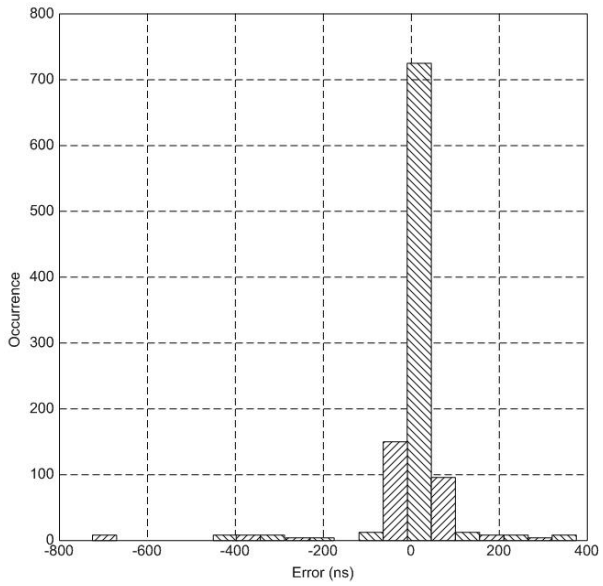

Fig. 11 Synchronization using a loaded syntonized spider

The residual jitter shown in Figure 10 is a result of the difference in the definition of a second in the grandmaster and spider. The ratio of the widths of the distributions in Figures 9 and 10 suggests that the definition of the second differs by approximately one part in 100 between the grandmaster and the spider. The further reduction in Figure 11 indicates the syntonization reduced the difference in the second of the two clocks to better than one part in 500.

The results of these experiments are summarized in Table 1.

Table 1 Summary of performance measurements

Topology	Sync accuracy
Crossover cable	± 10ns
Loaded bridge (no spider)	± 70 000ns
Loaded bridge un-syntonized spider	± 800ns
Loaded bridge syntonized spider	± 80ns

VI. CONCLUSIONS

This paper describes a prototype of a spider transparent clock for retrofitting existing bridges and routers to convert two or more ports into an IEEE 1588 transparent clock. The performance is expected to improve in future versions of the spider that make use of the new PHY level 1588 silicon and better syntonization.

REFERENCES

[1] IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control System- IEEE Std 1588-2008

[2] André Vallat, Dominik Schneuwly, *Modeling Packet Delay in Ethernet and IP Networks,* IEEE Workshop on Synchronization in Telecommunication Systems (WSTS-2008), May 13-15, 2008, Broomfield, Colorado

[3] Michael D. Johas Teener, *Audio/Video Bridging for Home Networks*, Joint ITU-T/IEEE Workshop on Carrier-class Ethernet 31 May – 1 June, 2007 Geneva, Switzerland, available at http://www.itu.int/ITU-T/worksem/cce/programme.html

[4] Geoffrey Garner, *Ethernet Qos, Timing and Synchronization Requirements*, Joint ITU-T/IEEE Workshop on Carrier-class Ethernet 31 May – 1 June, 2007 Geneva, Switzerland, available at http://www.itu.int/ITU-T/worksem/cce/programme.html

[5] Vook, D.; Hamilton, B.; Fernandez, A.; Burch, J.; and Srikantam, V., *An Update on Nanosecond-level Time-Synchronization with IEEE-1588*, Proceedings of 2005 Conference on IEEE-1588 Standard for Clock Synchronization 11 Oct. 2005, available at http://ieee1588.nist.gov/IEEE%201588%20Agenda/Paper-Vook_1588conf_2005.pdf

[6] http://www.netfpga.org/