Achieving Picosecond-Level Phase Stability in Timing Distribution Systems With Xilinx Ultrascale Transceivers

Eduardo Mendes, Sophie Baron, *Member, IEEE*, Csaba Soos, Jan Troska, and Paolo Novellini

***Abstract*— This article discusses the challenges posed on the field-programmable gate array (FPGA) transceivers in terms of phase-determinism requirements for timing distribution at the Large Hadron Collider (LHC) experiments. Having a fixed phase after startups is a major requirement, and the typical phase variations observed in the order of tens of picoseconds after startups while using the state-of-the-art design techniques are no longer sufficient. Each limitation observed in the transmitter and receiver paths of the high-speed transceivers embedded in the Xilinx Ultrascale FPGA family is further investigated and solutions are proposed. Tests in hardware using Xilinx FPGA evaluation boards are presented. In addition to a higher phase determinism, the techniques presented make it possible to fine-tune the skew of a link with a picosecond resolution, greatly simplifying clock-domain crossing inside the FPGAs and providing better short-term stability for the FPGA-recovered clock in a high-speed link.**

***Index Terms*— Fast timing, field-programmable gate arrays (FPGAs), high-energy physics instrumentation, optical links, timing circuits.**

I. Introduction

In THE Large Hadron Collider (LHC), a timing signal derived from the radio frequency driving the particle beams is transmitted to the four LHC experiments, allowing them to be synchronized to the particle bunches circulating in the accelerator rings [1]. This signal is called the bunch clock and has a frequency of approximately 40.079 MHz. When it reaches the four experiments, it has to be further distributed to tens of thousands of endpoint nodes which are located in a harsh radiation environment. The bunch clock recovered at the end points is required to have a fixed and deterministic phase with respect to the beam of particles.

We refer to a fixed and deterministic phase relation when the time offset between two nodes (phase between two clocks or a clock and the beam) does not change neither over time nor across multiple system initializations. In the LHC experiments, a physics beam calibration is performed before a data-taking run to compensate for time offsets between different end points of an experiment. Therefore, the absolute phase of each node is theoretically not important. However, from an operational point-of-view, it is highly desirable to avoid performing such a calibration every time a node in the system needs to be restarted as this could create chaotic operating conditions. This leads to a major requirement of a timing distribution system at the LHC experiments: the relative time offsets between different nodes shall not change across multiple system restarts/initializations.

The system responsible for delivering the bunch clock within the LHC experiments is called the timing, trigger, and control (TTC) system. The original implementation of this system, deployed in the 2000s, is based on 80-Mbit/s unidirectional point-to-multipoint optical links. The transmitter, located in the back-end zone (where no radiation is present), is based on discrete electronic circuits, and the receiver, located in the front-end zone (where radiations are present), is a custom-made application-specific integrated circuit (ASIC) featuring a fixed and deterministic phase relation between the output and the input. The TTC system is based on nonflexible discrete components that have now become obsolete. In addition, it can no longer cope with data-rate requirements of future upgrades of the LHC experiments.

Since then, the TTC implementations inside the experiments have evolved. A notable case is the Compact Muon Solenoid (CMS) experiment which had its TTC system upgraded in 2014, giving birth to the timing and control distribution system (TCDS) [2]. The main components of the TCDS are based on the Xilinx Kintex-7 field-programmable gate array (FPGA) family.

In the LHC phase-1 upgrades (2019-2020), the Large Hadron Collider Beauty (LHCb) experiment and A Large Ion Collider Experiment (ALICE) are undertaking a major upgrade of their TTC system by replacing the legacy TTC circuits. Both experiments use FPGA-based boards in the back end to transmit the TTC signals to a new radiation-hard custom-made ASIC, the GigaBitTransceiver (GBTx) [3]. In particular, the ALICE experiment has been chosen to adopt a TTC system based on passive optical networks (TTC-PONs) [4] for the distribution of TTC signals within the back-end zone. Some of the key components of the new ALICE trigger system implementation are the central trigger processor (CTP) and the local trigger unit (LTU) containing a Xilinx Kintex Ultrascale FPGA.

In the LHC phase-2 upgrades (2024), the increased particle accelerator performance (luminosity) will pose new challenges caused by the larger number of concurrent collisions occurring

Manuscript received October 7, 2019; revised December 5, 2019 and January 8, 2020; accepted January 14, 2020. Date of publication January 23, 2020; date of current version March 13, 2020.

Eduardo Mendes, Sophie Baron, Csaba Soos, and Jan Troska are with the Experimental Physics Department, CERN, 1211 Geneva, Switzerland (e-mail: eduardo.brandao.de.souza.mendes@cern.ch).

Paolo Novellini is with Xilinx, 20153 Milan, Italy.

Color versions of one or more of the figures in this article are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TNS.2020.2968112

during the same Bunch Crossing. In order to distinguish between different events, the CMS experiment has approved the construction of a high-resolution timing detector [5]. A Toroidal LHC ApparatuS (ATLAS) experiment has also proposed the construction of a timing detector to cope with the increased luminosity [6]. Both experiments foresee a major TTC upgrade during the LHC phase-2 upgrades and will have detectors which will require a timing precision of the order of picoseconds. The highly stringent timing requirements are posed on the short-term stability of the front-end recovered clock but it is desirable to mitigate nondeterministic timing instability sources related to system restarts or long-term environmental variations.

In general, the baseline for future TTC upgrade architectures of the different experiments is based on high-speed serial optical links to transmit timing. Inside the back-end zone, the clock is transported by point-to-multipoint (based on PON) or point-to-point optical link architectures. The clock distribution to the front end is made through point-to-point links and the radiation-hard transceivers GBTx or low-power gigabit transceiver (lpGBT) [7] deliver the clock to end points, the front-end chips. Modern FPGA devices equipped with high-speed transceivers seem to be a natural choice for the back-end zone. The Xilinx FPGA Ultrascale architecture seems to be a popular choice among the experiments and it is the subject of this article.

When using the FPGA transceivers for timing distribution, a careful design has to be carried out in order to achieve a fixed and deterministic phase after a full transceiver startup (a reset of all its internal blocks). Several studies on fixed-phase transceiver implementation exist in the literature [8], [9]. However, those were performed for legacy FPGA families. In addition, the measurements carried out do not allow to distinguish phase variations of the order of picoseconds due to their low resolution.

II. Limitations on the FPGA Transceivers for High-Precision Timing

The focus of the existing fixed-phase transceiver designs has been to ensure that no unit interval (UI) jumps occur between the transmitter reference clock and the receiver-recovered clock. For a 10-Gbit/s link, no phase jumps bigger than multiples of 100 ps between resets are present. Those are generally related to frequency multiplication and subsequent division and are discussed in [8]. Current solutions to implement a UI fixed-phase transceiver are well known in the high-energy physics community and are reviewed in Sections III-A and IV-A.

With the challenging phase determinism requirements of the phase-2 upgrades of the LHC experiments, sub-UI phase variations of the order of picoseconds may become important and must be understood.

In this article, we characterized different devices of the Xilinx Ultrascale FPGA family. Phase variations in the order of tens of picoseconds were measured in the transmitter and receiver paths using the state-of-the-art techniques presented below. Such sub-UI phase variations can be attributed to transceiver building blocks which add variable delays that are optimized automatically by the transceiver hard IP at every reset. Techniques to control this optimization process to yield sub-UI fixed-phase transceivers for the Xilinx Ultrascale family have been developed and are described in this article. Furthermore, the solutions discussed here can be generalized in order to simplify clock-domain crossing (CDC) problems in FPGAs and make it possible to fine-tune the skew of any data transmission link.

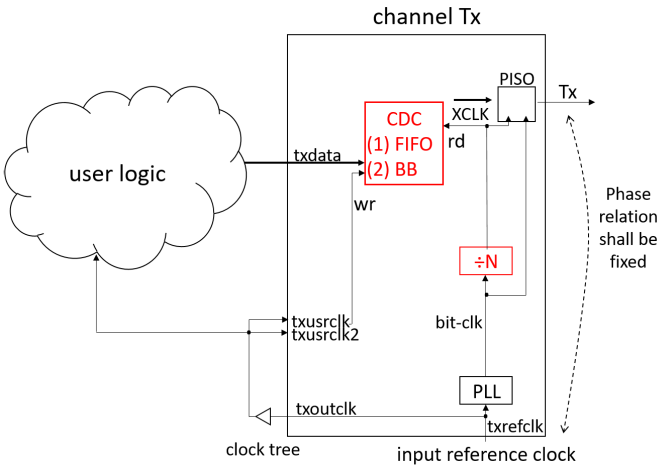

Fig. 1. Ultrascale transmitter architecture (simplified). PISO refers to parallel-in to serial-out register.

III. Transmitter Path

A. State of the Art

A simplified block diagram of the transmitter path of a Xilinx FPGA Ultrascale transceiver is shown in Fig. 1. For the transceiver configurations concerned by this article, txusrclk and txusrclk2 are the same signals and are derived from the input reference clock (txrefclk). The phase-locked loop (PLL) is shown inside the channel for illustration purposes, and all techniques shown in this article can equally work with a channel PLL (CPLL) or with a quad PLL (QPLL).

In the transmitter path, the bit clock coming from the PLL is divided down by N (XCLK) to clock the parallel-in to serial-out (PISO) register. This divider is not synchronized with the input reference clock (txrefclk) which can lead to different phases between txrefclk and Tx serial data after reset. In addition, the way the CDC between the transmitter user logic (txusrclk) and the parallel serializer clock domain (XCLK) is handled can add additional latency uncertainty.

Two transmitter configurations are available in the modern Ultrascale/Ultrascale+ GTH [10] and GTY [11] devices with respect to the way the CDC is handled.

1) *Elastic buffer (FIFO):* The CDC is dealt with by a first-in-first-out memory.
2) *Buffer-bypass (BB):* The CDC is register-based and the phase between clocks is controlled using advanced internal features.

Enabling the FIFO on the transmitter side does not solve the divider phase uncertainty issue mentioned above and

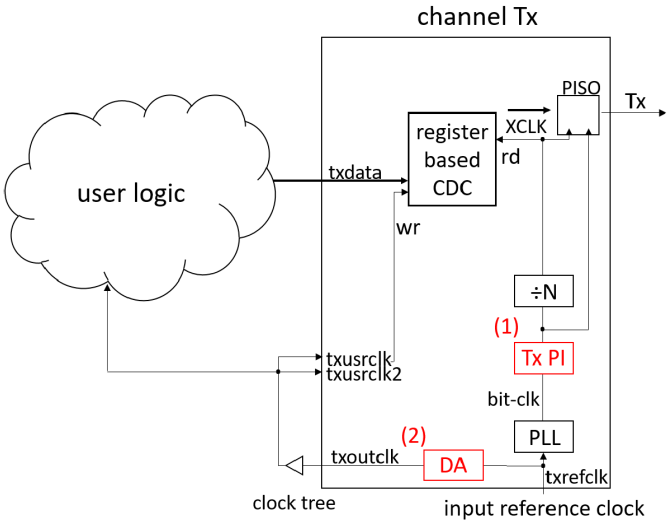

Fig. 2. Ultrascale transmitter architecture in the BB mode (simplified).

it can potentially add additional latency uncertainty due to the synchronization between the FIFO write and read pointers. Therefore, the BB configuration has been preferred for applications requiring minimal phase variation [4], [8], [12]. We observed that this technique can no longer cope with the phase determinism requirements of the LHC experiments' phase-2 upgrades and this will be discussed in Section III-B.

B. Limitations on Standard Transceiver Configuration for High-Precision Timing

In order to understand the additional limitations of the BB configuration that prevent higher stability, it is worth giving some details on the transceiver operation in this mode. The technique employed relies upon two advanced internal features shown in Fig. 2.

1) *Transmitter phase interpolator (Tx PI):* This block is a very fine phase shifter whose architecture is detailed in [13]. In the BB mode, it is used to perform a fine initial phase alignment to ensure proper phase between clocks (txusrclk and XCLK) for the register-based CDC.
2) *Delay aligner (DA):* Performs an on-the-fly compensation for phase variations in the clock-tree of txoutclk due to voltage and temperature variations.

We found out that phase variations in the order of tens of picoseconds after startup are present when the BB technique is used. Our hypothesis to explain this behavior is that every time the transceiver is initialized, the Tx PI-based fine clock alignment is performed and the Tx PI position chosen might be different. These results are shown in Section III-D.

Even though the observed phase variations are minimal and negligible for the large majority of typical link applications, it is desirable to mitigate this effect for the LHC phase-2 experiment upgrades. Furthermore, the BB mode prevents the control of the Tx PI block that can be potentially useful for applications requiring fine-phase tuning of a link. An alternative technique for implementing a fixed-phase

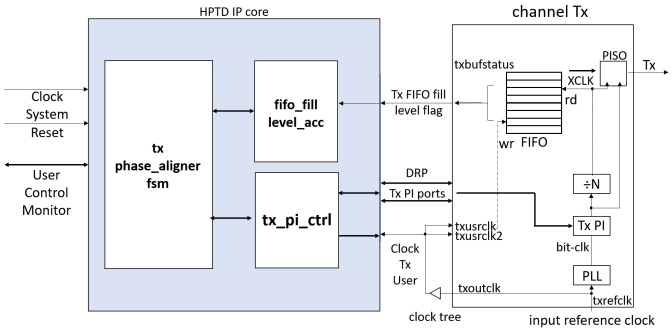

Fig. 3. HPTD IP core architecture.

Ultrascale FPGA transmitter which leads to a fully deterministic design is discussed in Section III-C.

C. High Precision Timing Distribution (HPTD) IP Core

The innovative technique employed here operates in the elastic-buffer configuration of Fig. 1. In this mode, a flag indicating whether the FIFO is more than half full is available (txbufstatus[0]) as shown in Fig. 3. This feature can be used as an early-late phase detector. Shifting the phase setting of the Tx PI block until the point where a transition in this flag is observed allows to always have the same phase between the write and read pointers of the FIFO. This translates into a deterministic phase between txrefclk and the Tx serial data. In principle, this technique is similar to the initial fine alignment performed in the BB mode but the full control of the alignment procedure allows us to implement more sophisticated algorithms. This technique is implemented as a protocol-agnostic soft VHDL core (HPTD IP core) publicly available in [14]. The architecture is composed of three blocks as depicted in Fig. 3.

1) *tx_pi_ctrl:* Controller block for Tx PI responsible for shifting the phase of the FIFO read pointer.
2) *fifo_fill_level_acc:* Accumulator for the txbufstatus[0] flag indicating whether the FIFO is half full. It acts as an early/late phase detector and gives the phase difference between the FIFO write and read pointer.
3) *tx_phase_aligner_fsm:* Finite-state machine responsible for the alignment algorithm. Shifts the phase of the FIFO read pointer (controlling **tx_pi_ctrl**) until a phase detector (**fifo_fill_level_acc**) change is detected.

The full IP core is very light in terms of occupancy, and the total usage for implementation on a Kintex Ultrascale FPGA is 135 Lookup Tables (LUTs) and 172 registers. Moreover, it does not make use of any resource shared within a group of four transceivers, which is called a "quad" in the Xilinx transceiver architecture, making multichannel implementations straightforward.

The alignment algorithm uses the average of the FIFO filling level flag in order to have a high-resolution phase detection. The algorithm can also work with a static configuration setting where the position found for the Tx PI for the first initialization is chosen again for further resets, completely eliminating the

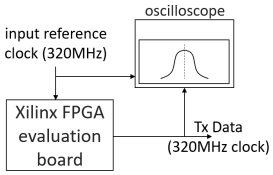

Fig. 4. Block diagram of the transmitter test setup.

phase uncertainty present in the BB technique and removing the uncertainty coming from the clock-tree. The static mode is recommended when the boards will not be subject to large temperature excursions. The results reported in this article were performed using the static mode.

D. Test Results

1) Transmitter Only: The setup used to test this technique is shown in Fig. 4. The same test was executed for two different Xilinx FPGA evaluation boards: a ZCU102 [15] containing a Zynq-Ultrascale+ SoC with the transceiver implemented in a GTH-Ultrascale+ and a KCU116 [16] containing a Kintex-Ultrascale+ FPGA with the transceiver implemented in a GTY-Ultrascale+. The FPGA transceiver receives a reference clock of 320 MHz and it is configured at 10.24 Gbit/s (lpGBT data-rate). In order to ease the measurements, the FPGA transmits a clock pattern at 320 MHz as a data stream. Both transceivers implemented use QPLL, which has a low-jitter *LC*-tank based VCO.

Test measurements are performed with a DSA91204A oscilloscope from KeySight [17]. An edge-to-edge time measurement between the reference clock fed to the transceiver and the Tx data is performed. The sampling rate is set to 10 GS/s with an acquisition window of 2 ms. For each phase measurement, around 1.5 MS are accumulated and the phase value analyzed here is the mean of the distribution obtained. All signals are ac-coupled to the oscilloscope and an absolute threshold level of 0 V is used. The temperature of the laboratory was not controlled during the tests but the variations were smaller than 5 °C while the measurements were performed. After each measurement, a reset of the FPGA transceiver is performed and a new measurement starts.

The two techniques previously discussed were tested using this setup for the two different FPGA evaluation boards. They were tested using the same measurement equipment. The results can be observed in Figs. 5 and 6 for the GTH Ultrascale+ and GTY Ultrascale+, respectively. Each histogram represents approximately 1000 measurement points. For the simple BB technique, we can observe discrete steps of multiples of 1.5 ps for different reset numbers (shown in red). Our hypothesis is that they correspond to different Tx PI positions as previously discussed. In comparison, the technique proposed in this article always presents a deterministic phase after a startup (shown in blue). In summary, previous state-of-the-art techniques provide a standard deviation of the phase after resets of around 4.0 ps, whereas our proposed technique based on the HPTD IP core has a standard deviation of around 0.4 ps. This corresponds to an improvement of a factor of 10 in the phase stability after resets for an environment with a relatively stable temperature.

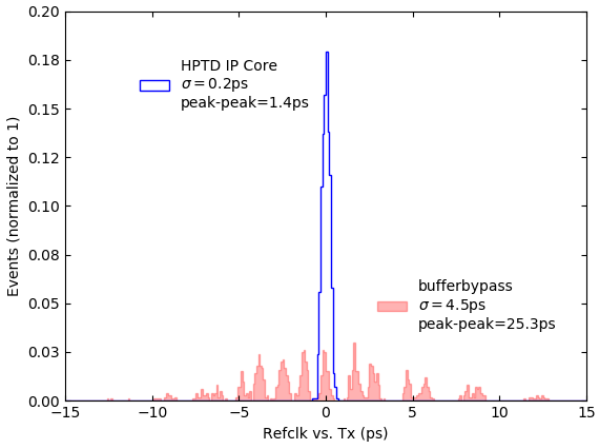

Fig. 5. Transmitter phase results for GTH Ultrascale+. Total number of events is 1000.

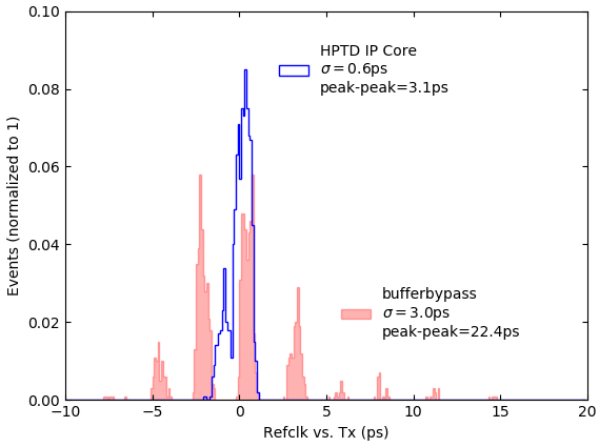

Fig. 6. Transmitter phase results for GTY Ultrascale+. Total number of events is 1000.

2) System: The radiation-hard lpGBT ASIC was not yet available when these tests were performed. Therefore, system tests were performed with its predecessor, the GBTx chip. The setup used for the system tests is shown in Fig. 7. A Xilinx FPGA evaluation board (KCU105 [18]) was used to emulate a back-end board implementing the GBT-FPGA core, a soft IP core developed at CERN to implement the back-end counterpart of the GBTx ASIC. This board contains a Kintex Ultrascale FPGA, and the transceiver implemented is a GTH Ultrascale. The transceiver is implemented following the GBT-FPGA example design which uses a CPLL, containing a ring-oscillator-based VCO. The versatile link demo board (VLDB) [19] was used to emulate a front-end card (containing a GBTx).

The high-speed optical link has a data rate of 4.8 Gbit/s. A pseudorandom bit sequence of length $2^7 - 1$ (PRBS-7) pattern is transmitted using the GBTx encoding scheme implemented in the GBT-FPGA core. The FPGA evaluation board is equipped with a commercial AFBR-709SMZ optical transceiver from FoxConn. The radiation-hard versatile transceiver (VTRx) [20] is used for VLDB. A 2-m-long multimode optical fiber was used between the back-end and front-end optical

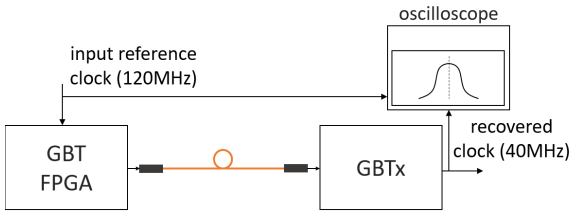

Fig. 7. Block diagram of the transmitter system test setup.

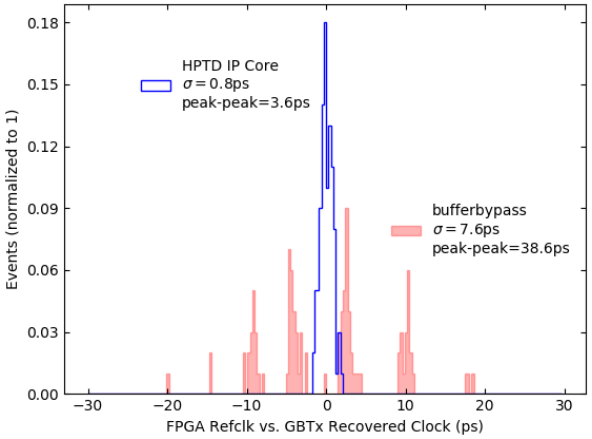

Fig. 8. Transmitter phase results for GTH Ultrascale-system tests with GBTx ASIC. Total number of events is 100.

transceivers. The phase measurement method and conditions are similar to the previous test. After each measurement, a full reset of the chain (reset of FPGA followed by GBTx reconfiguration) is performed before the next one.

To confirm our hypothesis, the two techniques previously discussed were tested using this setup, using the 40-MHz recovered clock from GBTx. It is important to note that GBTx itself is designed to have a fully deterministic phase. The results can be observed in Fig. 8. Each histogram represents around 100 resets. As expected, for the BB technique, discrete steps are observed. It is important to note that the steps are bigger than in the previous test. This reinforces our hypothesis as the Tx PI step is bigger for this transceiver configuration. We can observe that the technique implemented in the HPTD IP core always presents a deterministic phase after a startup.

IV. Receiver Path

A. State of the Art

A simplified block diagram of the receiver path of a Xilinx FPGA Ultrascale receiver is shown in Fig. 9. For the configurations used in this article, rxusrclk and rxusrclk2 have the same signal. The clock and data are recovered from the Rx serial stream by the clock-and-data recovery circuitry (CDR) after equalization. The high-speed clock coming from the CDR is divided down to clock the serial-in to parallel-out (SIPO) register. This divided clock (XCLK) can lock in any edge of the high-speed clock and, therefore, the phase information with respect to the Rx serial stream header is lost.

This first difficulty was overcome in [8]. The technique proposed consists in shifting the recovered parallel clock

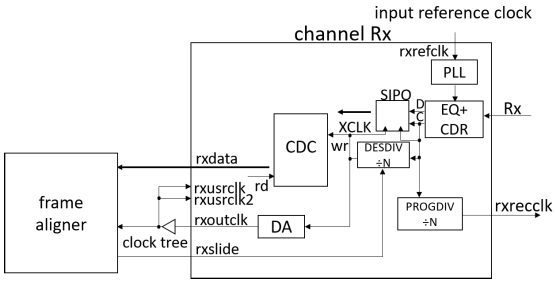

Fig. 9. Ultrascale receiver architecture (simplified). SIPO refers to serial-in to parallel-out register.

(XCLK) until it is aligned with the Rx serial stream header using the rxslide feature. This scheme is adopted in a variety of projects at CERN [4], [12].

Another technique, also proposed in [8], is the roulette approach where a reset is applied to the transceiver until it reaches a position locked to the serial stream header. The trade-off between these two approaches is complexity versus locking time.

Both techniques can provide a UI-level fixed-phase receiver design but an additional limitation was observed in our tests for the design of the sub-UI fixed-phase receiver.

B. Limitations on Standard Transceiver Configuration for High-Precision Timing

During system tests with multiple FPGAs, we have observed an additional limitation in the receiver path for high-precision timing purposes. We believe this limitation to be related to the adaptive blocks present in the receiver equalizer. Those blocks are controlled on-the-fly by slow feedback loops while left in the adaptive mode. It seems that for different startups, the steady state of the loop can be slightly different. This can have an impact on the phase of the FPGA-recovered clock which is not the same for different startups as it will be shown in Section IV-D. In this article, we propose a technique which can be used to remove this additional uncertainty in Section IV-C.

C. Rx Equalizer Frozen Algorithm

The receiver equalizer has several parameters which are adapted on-the-fly by slow feedback loops, such as the automatic gain control, the high- and low-frequency gains, and the offset cancellation. These parameters can be left in the adaptive mode (the feedback loops are closed) or in the frozen mode (parameters are overridden with user-defined values). In order to override the parameters with optimal values, the following algorithm is proposed.

1) Leave the transceiver equalizer in the adaptive mode for all concerned parameters (parameter_ovrden=0).
2) Reset transceiver.
3) Read the adapted values via the digital monitor (dmonitor) port. More information can be found in [10] and [11].
4) Write previously read values through the dynamic reconfiguration port (DRP). More information can be found in [10] and [11].

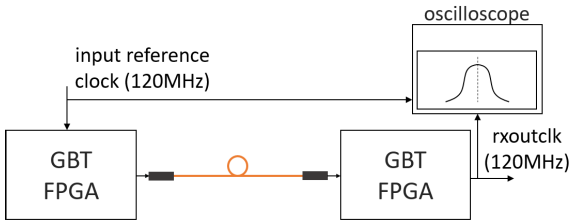

Fig. 10. Block diagram of the receiver test setup.

5) Enable override of parameters (parameter_ovrden = 1).
6) Reset transceiver.

This algorithm is performed for the first startup and the adapted values are used for further resets. On-the-fly adaptation of those parameters is mainly important when systems are subject to large temperature variations. For the well-controlled environment present in the LHC experiments, we expect that overriding the equalizer parameters will not have a major impact on the link quality. Besides, if needed, a new calibration can be applied between the accelerator and experiment running periods.

D. Test Results

The setup used to test this technique is shown in Fig. 10. For this test, two Xilinx FPGA evaluation boards (KCU105) were used in the setup to implement the transmitter and the receiver in different FPGAs. The phase measurement method and test conditions were similar to the ones described in Section III-D. After each measurement, a reset of the second FPGA in the chain was performed and a measurement starts.

The receiver with a frozen equalizer was tested using this setup and compared with a receiver having a fully adaptive equalizer. Both receivers are implemented using a discrete feedback equalizer (DFE). The results can be observed in Fig. 11. We can observe discrete jumps after startup when the equalizer is left in the adaptive mode and a fully deterministic behavior with the Rx equalizer frozen algorithm. Each histogram represents around 200 measurement points. The phase standard deviation is around 6.5 ps for the fully adaptive equalizer with respect to 1.5 ps when the Rx equalizer frozen algorithm is used. When the equalizer used the low power mode (LPM), which is a continuous-time linear equalizer (CTLE), no phase jumps were observed after reset even if the equalizer was left in the adaptive mode. We use the frozen DFE equalizer for the rest of the tests presented in this article.

V. Cascaded FPGA Links

A. Traditional Cascading Challenges

We refer to a cascaded FPGA link when a clock is recovered in an FPGA and it is used forward to drive another high-speed link in a cascaded fashion as shown in Fig. 12. A recovered clock from a high-speed serial link can be noisy due to potential low signal-to-noise ratio (SNR), intersymbol interference (ISI), and signal reflections. The FPGA itself can also have a major impact on the quality of the recovered clock due to fabric switching activity and interference with the voltage-supply source. Therefore, the quality of the FPGA-recovered

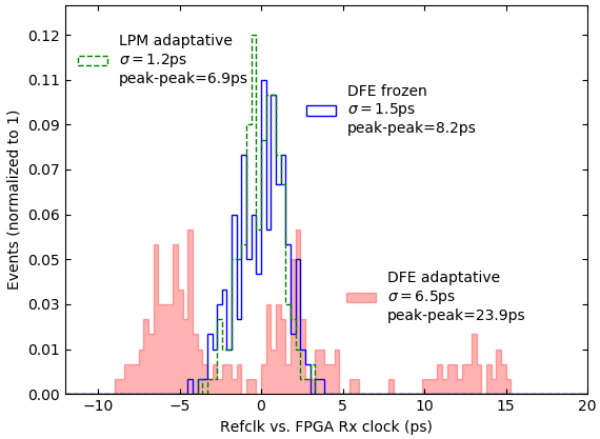

Fig. 11. Receiver phase results for GTH Ultrascale-system tests with two FPGAs. Total number of events is 200.

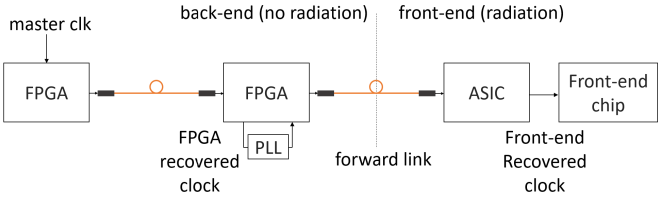

Fig. 12. Example of a cascaded FPGA link.

clock can be highly dependent on the design loaded in the FPGA.

In order to meet the stringent transceiver phase-noise mask requirements [21] for the forward link, it is necessary to clean the FPGA-recovered clock with an external PLL. This adds an additional difficulty for the FPGA designer. The addition of an external path with unknown delay variability is hard to be properly modeled by static timing analysis tools. Therefore, the traditional technique is to rely on a fixed-phase mesochronous CDC technique in order to transfer the data received forward. A possibility for such a CDC, adopted in [4] and [12], is to shift the external PLL phase in order to perform a scan of the metastability-free window.

A new feature of the Ultrascale architecture allows us to have a cleaner FPGA-recovered clock if one targets the ultimate optimization of the timing links and its difficulties are discussed in V-B. In Section V-C, we overcome those difficulties with a novel CDC scheme requiring no shift capability of the external PLL.

B. New Recovered Clock Possibility in the Ultrascale Architecture

In the Ultrascale and Ultrascale+ architectures, it is possible to output a clock directly from the CDR (rxrecclk) to the transceiver reference clock pins without going through the fabric [10], [11]. The path for this clock is shown in Fig. 9. In principle, being directly extracted from the CDR, this clock is not affected by the FPGA-related effects mentioned previously, and its quality does not depend on the design implemented in the FPGA. This is an extremely attractive option for next-generation timing links.

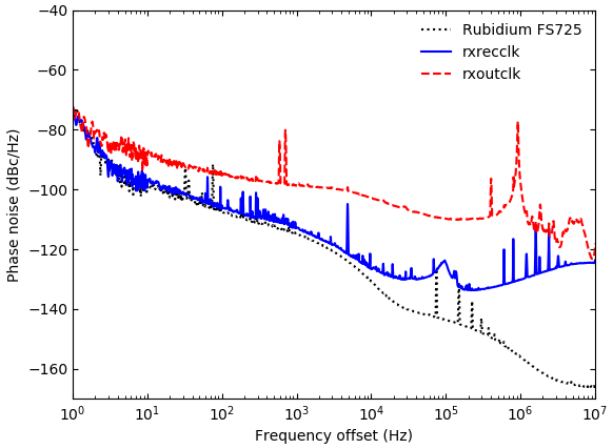

Fig. 13. FPGA-recovered clock quality. Rxrecclk does not pass through the FPGA fabric and, therefore, it exhibits a higher quality than rxoutclk.

TABLE I
INTEGRATED PHASE-NOISE VALUES

clock	1 Hz-1 kHz	1 kHz-10 MHz
Rubidium FS725 (reference)	0.32 ps	0.17 ps
rxreclk	0.33 ps	3.31 ps
rxoutclk	1.27 ps	44.23 ps

Indeed, this can be observed in the phase-noise plots of Fig. 13. These were obtained with the FSWP8 phase-noise analyzer from Rohde and Schwarz [22] for the setup of Fig. 10 having a rubidium clock generator [23] as reference. For rxoutclk, when compared to rxrecclk, we can observe additional white noise potentially coming from the clock tree in the FPGA fabric. A peak at around 1 MHz is very pronounced in rxoutclk. It potentially comes from switching activity of dc–dc converters (which can be related to the board design itself). Other peaks of smaller amplitude are also observed, which were identified to come from a potential crosstalk between the transmitter and receiver reference clocks which have a plesiochronous relation. The time interval error (TIE) jitter values measured with an oscilloscope are approximately 6-ps rms and 50-ps rms for rxrecclk and rxoutclk, respectively.

Typically, an external PLL is used to clean the FPGA-recovered clock in a high-speed optical link as shown in Fig. 12. Therefore, it is important to take into account its effect in order to properly analyze (Fig. 13). We use a bandwidth of around 1 kHz as a reference. Table I shows the integrated phase-noise values for the potential external PLL in-band of 1 Hz–1 KHz and out-band of 1 KHz–10 MHz. Most of the noise present in rxoutclk would be cleaned by such a PLL but we still gain around 1 ps for the in-band of rxrecclk.

The difficulty of using the rxrecclk clock for fixed-phase links is that the divider generating this clock (PROGDIV in Fig. 9) is not synchronized with the receiver deserializer divider (DESDIV in Fig. 9) which yields a phase uncertainty which can be up to one rxoutclk period.

C. Direct Cascading

The techniques developed in Sections III-C, IV-C, and V-C can be used in turn to a highly efficient CDC. The scheme proposed consists in using rxoutclk of the FPGA receiver as txusrclk/txusrclk2 for the forward link FPGA transmitter as depicted in Fig. 14 while using rxrecclk as the reference clock. By using such a technique, CDC is safely handled at the FIFO level of the forward link transceiver and the phase uncertainty of rxrecclk mentioned above is resolved at the transmitted data.

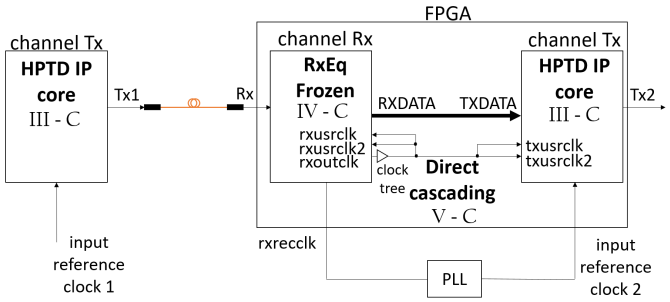

Fig. 14. Scheme for the cascaded link.

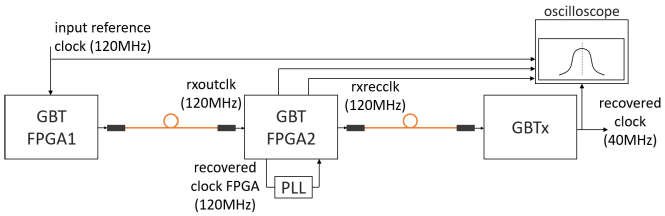

Fig. 15. Block diagram of the cascaded test setup.

D. Test Results

The setup used to test the cascaded link is shown in Fig. 15. For this test, two Xilinx FPGA evaluation boards (KCU105) and a VLDB were used. The PLL used is Si5344 [24] from Silicon Labs configured in a zero-delay mode. The phase measurement method and test conditions are similar to the ones described in Section III-D. A reset of the full chain (reset first FPGA, reset second FPGA, reset PLL, and reset GBTx) is performed and then five phase measurements are performed, taking approximately 5 min.

The measurement lasts for around three days; therefore, major temperature variations were observed in the laboratory which are not typical inside the experiment. In order to show the effect of temperature on the phase, we also analyze the temperature measured by the system monitor [25] in the two FPGAs. A total of 1000 resets were issued during this time period, corresponding to 5000 measurement points.

The proposed fully optimized transceiver design in the cascaded mode was tested using this setup and compared with the traditional implementation. For the optimized design, the recovered clock used is rxrecclk, and for the traditional implementation, rxoutclk is used.

We can observe in Fig. 16 (performed for the optimized link design) that the phase of rxrecclk exhibits two UI jumps spanning the full range of one rxoutclk period, ≈8.3 ns for the 120-MHz frequency. It is interesting to see the phase values over time with a modulo UI operation as shown in Fig. 17. It can be observed that the phase variations follow

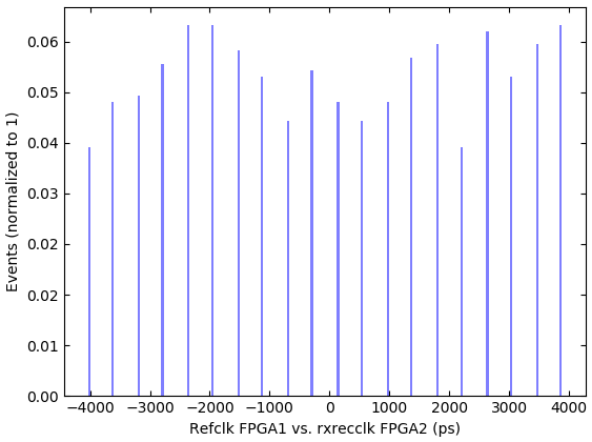

Fig. 16. FPGA2 rxrecclk exhibits two UI phase jumps after startup. Total number of events is 5000.

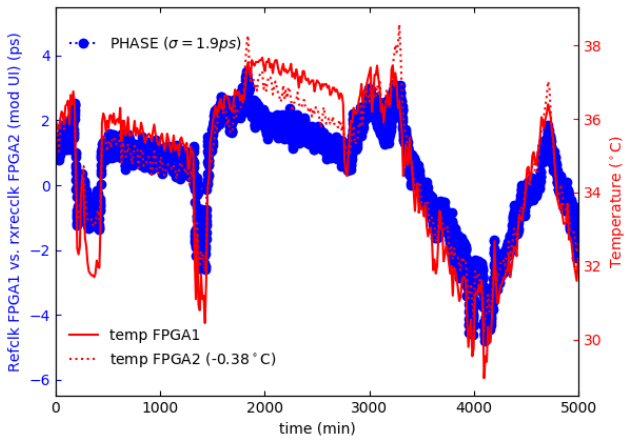

Fig. 17. FPGA2 rxrecclk after modulo UI mathematical operation.

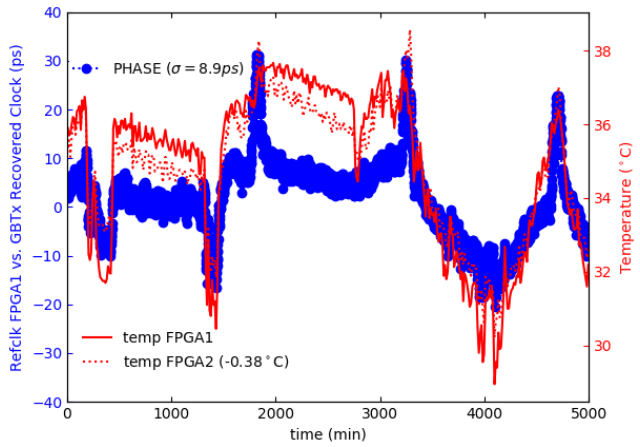

Fig. 18. Results of the cascaded link using the techniques discussed in this article.

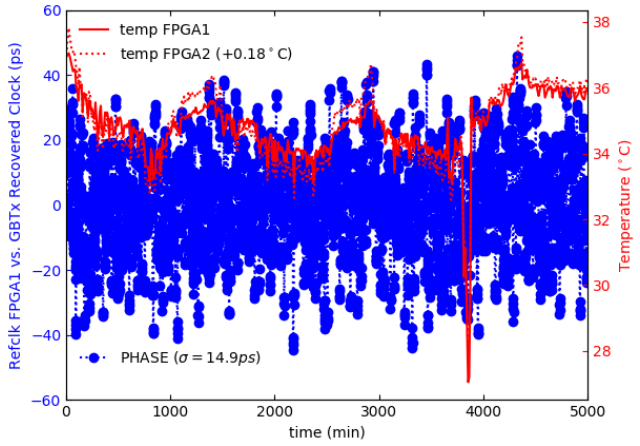

Fig. 19. Results of the cascaded link using traditional techniques.

the temperature variation well. The temperature dependence of the delay was observed to be around 1 ps/°C, which includes an FPGA transmitter, the optical transceivers, short optical fibers, and the FPGA receiver.

The results of the GBTx-recovered clock using the techniques described in this article can be observed in Fig. 18. The phase uncertainty of rxrecclk is resolved at the FPGA transmitter level of the forward link and no discrete jumps can be observed between points. The phase follows the temperature variations well, and the temperature dependence of the delay is around 10 ps/°C.

The results of the traditional cascaded link can be observed in Fig. 19. The correlation with temperature is less clear as the discrete phase variations after startup dominate the results. For an environment with a stabilized temperature, the results are expected to be even better for the proposed techniques.

VI. Scope of This Article

This article was developed for the Ultrascale/Ultrascale+ architectures as the majority of our use-cases adopt those families. The sub-UI phase variations described here are related to specific transceiver blocks such as the transmitter phase alignment (see Section III-B) or the equalizer adaptation (see Section IV-B) which are reinitialized following each transceiver reset. Our approach in this article was to characterize the possible sources of sub-UI phase variations in the aforementioned FPGA architectures and to tackle each of the observed challenges with transceiver-specific solutions that do not require external hardware.

If there is a need for such phase stability on other transceiver types, those should be carefully characterized in a similar fashion to determine appropriate mitigation techniques. A truly transceiver-agnostic implementation may necessitate the use of external components to provide the required feedback/control, which can be potentially impractical on complex board designs featuring hundreds of transceivers.

VII. Conclusion

New techniques for implementing picosecond-precise transmission using the transceivers embedded in the Xilinx Ultrascale FPGA family have been demonstrated and characterized in this article. The proposed techniques reach picosecond resolution, which is an improvement over the prior art of one order of magnitude.

The proposed techniques can bring CDC simplification inside the FPGAs without requiring a phase-shifting capability. They can also provide us with better short-term stability for the recovered clock in an FPGA-based high-speed link. These two features will be a tremendous asset to upgrade the very demanding high-resolution timing detectors of the LHC.

For timing distribution links located in a heterogeneous environment where long-term environment variations (such as temperature and voltage) may eventually dominate the total timing stability, the techniques presented here will be useful for the implementation of timing-compensated schemes. The HPTD IP core allows applications requiring fine skew tuning of a link to reach a picosecond resolution, paving the way for precise timing compensated links implementation. It enables the compensation of all individual links in an FPGA without external components and only a little amount of logic in the FPGA fabric.

REFERENCES

[1] S. Baron, T. Mastoridis, J. Troska, and P. Baudrenghien, "Jitter impact on clock distribution in LHC experiments," *J. Inst.*, vol. 7, no. 12, p. 12023, Dec. 2012.

[2] J. Hegeman *et al.*, "The CMS timing and control distribution system," in *Proc. IEEE Nucl. Sci. Symp. Med. Imag. Conf. (NSS/MIC)*, Oct. 2015, pp. 1–3.

[3] P. Leitao *et al.*, "Test bench development for the radiation hard GBTX ASIC," *J. Inst.*, vol. 10, no. 01, p. 1038, Jan. 2015.

[4] E. Mendes, S. Baron, D. Kolotouros, C. Soos, and F. Vasey, "The 10G TTC-PON: Challenges, solutions and performance," *J. Inst.*, vol. 12, no. 02, p. 2041, Feb. 2017.

[5] L. Gray and T. T. de Fatis. (Nov. 2017). *Technical Proposal for a MIP Timing Detector in the CMS Experiment Phase 2 Upgrade CERN-LHCC. LHCC-P-009*. [Online]. Available: https://cds.cern.ch/record/2296612/files/LHCC-P-009.pdf

[6] C. Allaire, "A high-granularity timing detector in ATLAS: Performance at the HL-LHC," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 924, pp. 355–359, Apr. 2019.

[7] P. Moreira. (2018). *LpGBT Specifications CERN*. [Online]. Available: https://espace.cern.ch/GBT-Project/LpGBT/Specifications/LpGbtxSpecifica%tions.pdf

[8] R. Giordano and A. Aloisio, "Fixed-latency, multi-gigabit serial links with Xilinx FPGAs," *IEEE Trans. Nucl. Sci.*, vol. 58, no. 1, pp. 194–201, Feb. 2011.

[9] I. Papakonstantinou *et al.*, "A fully bidirectional optical network with latency monitoring capability for the distribution of timing-trigger and control signals in high-energy physics experiments," *IEEE Trans. Nucl. Sci.*, vol. 58, no. 4, pp. 1628–1640, Aug. 2011.

[10] *UltraScale Architecture GTH Transceivers User Guide*, document UG576 (v1.5.1), Xilinx, Aug. 2018. [Online]. Available: https://www.xilinx.com/support/documentation/user_guides/ug576-ultrasca%le-gth-transceivers.pdf

[11] *UltraScale Architecture GTY Transceivers User Guide*, document UG578 (v1.3), Xilinx, Sep. 2017. [Online]. Available: https://www.xilinx.com/support/documentation/user_guides/ug578-ultrasca%le-gty-transceivers.pdf

[12] M. B. Marin *et al.*, "The GBT-FPGA core: Features and challenges," *J. Inst.*, vol. 10, no. 03, p. 3021, Mar. 2015.

[13] D. Taylor, M. Klein, and V. Vendramini, *All Digital VCXO Replacement for Gigabit Transceiver Applications (UltraScale FPGAs)*, document XAPP1241 (v1.0), Xilinx, Aug. 2015. [Online]. Available: https://www.xilinx.com/support/documentation/application_notes/xapp1241%-vcxo.pdf

[14] E. Mendes and S. Baron. (Nov. 2018). *Transmitter Phase Aligner for Xilinx Transceivers. CERN. HPTD IP core (rev 0_1_0)*. [Online]. Available: https://gitlab.cern.ch/HPTD/tx_phase_aligner

[15] *ZCU102 Evaluation Board User Guide*, document UG1182, Xilinx, v1.5, Jan. 2019. [Online]. Available: https://www.xilinx.com/support/documentation/boards_and_kits/zcu102/ug1%182-zcu102-eval-bd.pdf

[16] *KCU116 Evaluation Board User Guide*, document UG1239 (v1.2), Xilinx, Sep. 2018. [Online]. Available: https://www.xilinx.com/support/documentation/boards_and_kits/kcu116/ug1%239-kcu116-eval-bd.pdf

[17] K. Technologies. (Jan. 2019). Keysight Infiniium Oscilloscopes. [Online]. Available: https://www.keysight.com/upload/cmc_upload/All/Infiniium_prog_guide.pdf

[18] *KCU105 Evaluation Board User Guide* document UG917 v1.9, Xilinx, May 2018. [Online]. Available: https://www.xilinx.com/support/documentation/boards_and_kits/kcu105/ug9%17-kcu105-eval-bd.pdf

[19] R. M. Lesma *et al.*, "The versatile link demo board (VLDB)," *J. Inst.*, vol. 12, no. 02, p. 2020, Feb. 2017.

[20] J. Troska *et al.*, "Versatile transceiver developments," *J. Inst.*, vol. 6, no. 01, p. 1089, Jan. 2011.

[21] *Kintex UltraScale FPGAs Data Sheet: DC and AC Switching Characteristics Product Specification*, document DS892 v1.17, Xilinx, Jan. 2019. [Online]. Available: https://www.xilinx.com/support/documentation/data_sheets/ds892-kintex-u%ltrascale-data-sheet.pdf

[22] S. Rohde. *FSWP Phase Noise Analyzer Specifications*, document Version 07.00, Aug. 2018. [Online]. Available: https://scdn.rohde-schwarz.com/ur/pws/dl_downloads/dl_common_library/dl%_brochures_and_datasheets/pdf_1/service_support_30/FSWP_dat-sw_en_3607-2090-22%_v0700_120dpi.pdf

[23] S. R. Systems. *FS725 Rubidium Frequency Standard Operation and Service Manual*, document Version 1.3, Dec. 2015. [Online]. Available: https://www.thinksrs.com/downloads/pdfs/manuals/FS725m.pdf

[24] Silicon Labs, *Si5345/44/42 Rev D Data Sheet*, document Rev. 1.2, Sep. 2018. [Online]. Available: https://www.silabs.com/documents/public/data-sheets/Si5345-44-42-D-DataSheet.pdf

[25] *UltraScale Architecture System Monitor*, document UG580 v1.9.1, Xilinx, Feb. 2019. [Online]. Available: https://www.xilinx.com/support/documentation/user_guides/ug580-ultrasca% le-sysmon.pdf