Phase stabilization over a 3 km optical link with sub-picosecond precision for the AWAKE experiment

D. Barrientos, J. Molendijk

***Abstract*–The Advanced Wakefield Experiment (AWAKE) aims at studying the proton-driven plasma wakefield acceleration technique for the first time. The testing facility, currently being built at CERN, uses the proton beam at a momentum of 400 GeV/c from the Super Proton Synchrotron (SPS) to accelerate an electron beam to the GeV scale over 10 m of plasma. In order to achieve such an acceleration gradient, the reference signal of the low-level RF (LLRF) system controlling the proton beam must keep in-phase with the reference signal used to generate the electron beam and plasma (laser). Even though the SPS LLRF system is located about 3 km away from the laser and electron beam electronics, the phase drift between the three references has been specified to be in the sub-picosecond range. In order to cope with the experiment requirements, we have developed a custom VME board and a digital control system embedded in a FPGA to compensate for the phase drift between the reference signals at both ends of the optical links. In this work, we present the results of the study developed to analyze the expected phase drift, the selected method to compensate it and the performance of the first prototypes of the board. The use of a very precise phase detector and digitally controlled delay lines, both in the level of tens of femtoseconds allow tracking the phase drifts and compensate for them with a very high precision. Measurements of the achieved precision in the developed module have shown to be in the sub-picosecond range, as demanded by the experiment requirements.**

I. Introduction

THE AWAKE experiment [1-2] is a proof-of-concept facility for proton-driven plasma wakefield acceleration currently being commissioned at CERN. The proton bunches extracted from the SPS machine with a momentum of 400 GeV/c will serve as a drive beam in the 10 m rubidium plasma cell, in order to accelerate an electron (witness) beam up to the GeV scale in a few meters. Therefore, the expected performance corresponds to an accelerating gradient about 3 orders of magnitude higher than RF cavities currently being used.

The LLRF system for AWAKE is in charge of the synchronization of the high intensity laser pulses generating the plasma, the electron and the proton beams. The reference signals for the laser and electron beam are generated in the AWAKE laser room, at the former CNGS (CERN Neutrinos to Gran Sasso) facility, whereas the reference signals for the LLRF system of the SPS accelerator (proton beam) are used at the SPS beam control system (Faraday Cage), about 3 km away from the laser room.

According to the specifications in the design report [1], the synchronization of the laser and electron beam and the proton beam must be on the picosecond level. The synchronization procedure is similar to the mechanism used to transfer the beam from the SPS to the LHC (Large Hadron Collider) accelerator [3-4]. When the beam momentum arrives to flat-top (at the nominal value of 400 GeV/c), a re-phasing procedure takes place to lock the RF frequency, phase and extraction timing of SPS to the external reference signals from AWAKE. The selected frequency to be transmitted is about 400.8 MHz, twice the frequency of the main RF system in the SPS. After re-phasing, SPS and AWAKE references are synchronous and the transfer procedure can take place.

In this work, we present a system to stabilize for phase drifts larger than 1 ps over a 3 km optical link used to synchronize the proton beam to the laser pulse and electron beam in the AWAKE experiment. The motivation for the development of a new VME module to cope with the phase drifts due to changes in environmental conditions of the long fiber links is presented in Section II. Also in this section, a study of the jitter contribution of commercial SFP (small form-factor pluggable) transceivers is included. Section III describes the concept used for the stabilization of the phase drift in the developed module with a complete description of the techniques adopted for noise reduction in the phase detection and calibration of non-linearity in delay lines. Finally, some performance figures of the developed module are presented in Section IV.

II. Design Considerations

A. Expected phase drift

In order to study the phase drift of a reference signal over long optical fibers, a continuously logging system has been set up. A 200 MHz reference signal from a signal generator (Rhode & Schwarz SMC100A) is compared to a copy of the same signal sent through a fiber loop of 6 km length similar to the link used by the AWAKE experiment. The frequency of the reference signal for the logging system was chosen to be 200 MHz as the selected frequency to be transmitted in the long optical link was decided to be about 400.8 MHz in a later stage of the project. Then, the phase of both signals is compared by a custom module using an Analog Devices AD8302 phase detector. The phase difference between the inputs is logged twice per minute and stored in a central database. Similarly, the external temperature is also recorded and stored once per minute.

During analysis, recorded phase data are averaged within the same minute, to have an equal number of temperature and phase samples. Then, samples are also filtered using a moving average within the previous 24 hours, as both temperature and

Manuscript received June 3, 2016.

D. Barrientos and J. Molendijk are with CERN, CH-1211 Geneva 23, Switzerland (e-mail: diego.barrientos@cern.ch).

phase differences have shown to be repetitive in a day-by-day basis. Finally, phase differences and temperatures are correlated using a linear fit.

The data taken over a period of about 10 months with a temperature excursion of about 30 degrees have been fit linearly, obtaining an r-squared value of 0.76, as shown in Fig. 1. With the fit values, the maximum expected temperature excursion of 60 degrees Celsius would correspond to a change of about 429.8 degrees at 200 MHz of the reference signal after the 6 km fiber loop, corresponding to 5.97 ns. This drift, being more than three orders of magnitude above the specification, makes necessary the use of an actively compensated optical link.

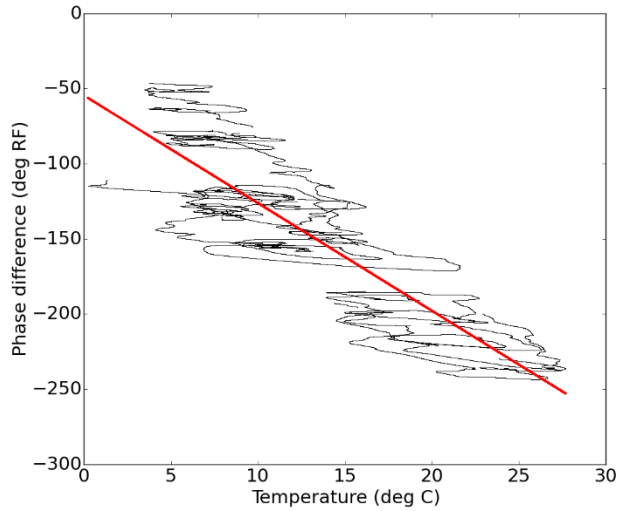

Fig. 1. Correlation between external temperature and phase difference. Black points correspond to experimental data and the red line is a linear fit of the data.

B. *SFP phase noise contribution*

At the first stages in the design of the new module, a complete study of the phase noise introduced by the electrical to optical transceivers has been carried out. Commercial off-the-shelf SFP transceivers have been chosen due to their variety and availability in the market and their performance figures for high-speed optical communications. For these tests, the phase noise of the reference signal from the generator and the signal after being sent through a double conversion in a SFP module (electrical to optical and back to electrical) have been compared. A four-channel module has been developed to convert the 50-ohm single-ended output signal of the signal generator to a 100-ohm differential signal that feeds the optical transceiver via a 1:1 balun transformer. Then, a three meters fiber loopback injects the optical signal in the receiving side of the transceiver. The received differential signal is converted back to 50-ohm (single-ended) by means of the same technique. The phase noise of the signals, shown in Fig. 2, is measured using a signal source analyzer (Keysight E5052B).

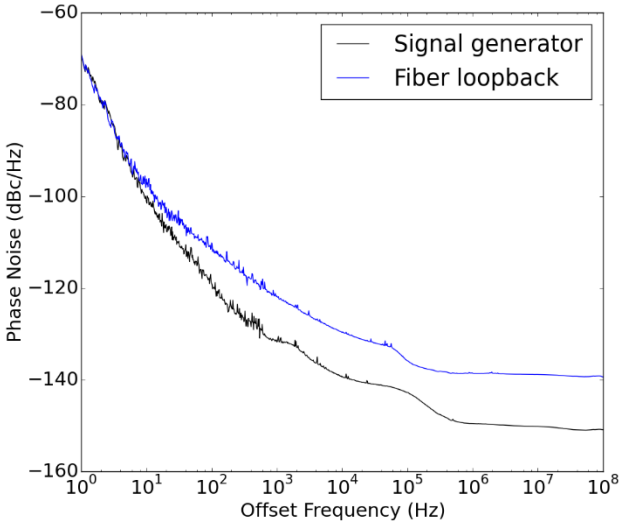

Fig. 2. Phase noise contribution of an electrical to optical SFP transceiver.

In the figure, the phase noise at different offset frequencies from the 400 MHz carrier signal is presented. Integrating the phase noise density over the full range of frequencies (1 Hz – 100 MHz) the jitter contributions of the two signals can be obtained [5].

The computed jitter contribution of the signal generator is 0.21 ps, whereas the jitter of the signal from the optical loopback is 0.64 ps. Thus, the contribution of the electrical to optical and back to electrical conversions and associated circuitry can be computed with a quadratic subtraction of the two quantities. The introduced jitter is 0.60 ps, which is below 1 ps. Tests with other types of transceivers have shown a similar performance, therefore the SFP transceiver solution has been adopted for the design.

III. Phase Drift Compensation Concept

As discussed in Section I, the requirements of the AWAKE experiment include the compensation of the phase drift between the signals in the AWAKE laser room (SPS point 4) and the SPS Faraday Cage (SPS point 3). The 400.8 MHz signal, used as a reference to lock the RF phase of SPS, is the most critical one regarding phase drift. However, other pulsed signals at frequencies of 9.97 Hz and 8.68 kHz used for beam synchronization and extraction, are also transmitted between the two sites, although their phase only has to be stable within a half 400.8 MHz signal period. The expected phase drifts may be caused mainly by external factors, such as temperature changes, affecting the off-the-shelf optical fibers used for signal transmission. Considering these constraints, a new VME module has been developed and built, in order to measure and compensate for the phase drifts of the signals transmitted between the SPS points 4 and 3. A schematic view of the concept is presented in Fig. 3.

978-1-5090-2014-0/16/$31.00 ©2016 European Union

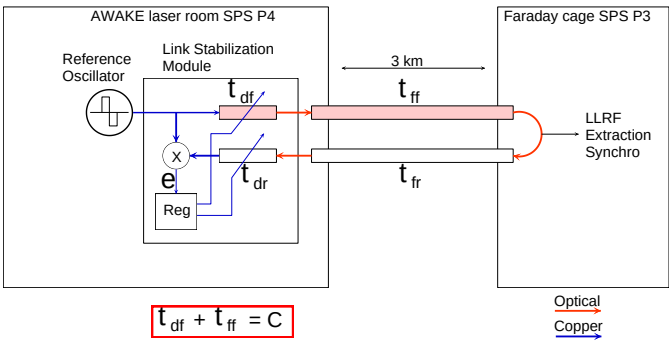

Fig. 3. Schematic view of the concept used for phase drift compensation between the AWAKE laser room (SPS point 4) and SPS Faraday Cage (SPS point 3). The reference signal is synchronous with the laser and electron beam reference signals. Red lines correspond to optical links whereas blue lines represent electrical signals.

As highlighted in the figure, the main function of the link stabilization module is to keep the propagation time of the reference signal between two sites constant. With this purpose, a precise phase detector, described in the following section, measures the phase error (ε) between copies of the reference signal sent to the optical link and reflected back by a 1:2 optical splitter placed at the SPS point 3 (on the right side of the figure). Then, a fully digital regulation loop acts on both electrical delay lines τ_{df} and τ_{dr} to compensate for the detected phase error.

The total delay incurred by the reference signal during the round trip can be expressed as

$$\tau_{total} = \tau_{df} + \tau_{ff} + \tau_{fr} + \tau_{dr} \quad (1)$$

Assuming that changes in the propagation time due to external factors in the optical forward path will be the same as in the return path as both fibers are within the same fiber bunch, the delay error δ can be compensated by modifying the delay lines by Δ. Therefore, the new propagation time becomes

$$\begin{aligned} \overline{\tau_{total}} &= \tau_{df} + \Delta + \tau_{ff} + \delta + \tau_{fr} + \delta + \tau_{dr} + \Delta; \\ \overline{\tau_{total}} &= \tau_{total} + 2\delta + 2\Delta; \\ \varepsilon &= 2(\delta + \Delta) \end{aligned} \quad (2)$$

The digital regulation loop is in charge of bringing the error signal to zero, such that the detected phase remains constant. As a consequence, the sum of δ and Δ become zero, and the propagation time of the reference signal from SPS point 4 to point 3 becomes constant:

$$\tau_{df} + \Delta + \tau_{ff} + \delta = \tau_{df} + \tau_{ff} = C \quad (3)$$

The reference signal used by the regulation loop is the 400.8 MHz signal that requires the highest precision in the phase drift compensation. Additionally, the propagation delay of the other transmitted signals is controlled in open-loop by changing the delay in the forward path by the same amount (Δ) as for the main reference signal.

A. Phase detection

The principle of operation of the link stabilization module relies on a very precise phase detection between the two copies of the reference signal. In order to achieve a high precision, the entire electronics design of the critical path, starting from the optical transceivers, is executed using differential signaling (CML, LVPECL). This differential approach largely suppresses phase shifts caused by logic level shifts. Additionally, the digitization of the two polarities of the differential signal at the output of the phase detector allows to reduce the induced noise due to voltage drifts in the digitization stage, as shown in Section IV.

The implementation of the phase detector is composed by an exclusive-or gate and two low-pass filters, as depicted in Fig. 4.

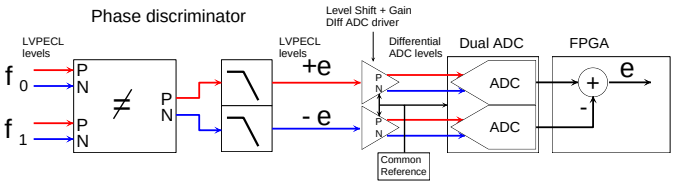

Fig. 4. High-precision phase detector set-up to suppress common-mode noise. Red and blue lines correspond to the two polarities of the differential lines. Both polarities are digitized independently to reduce induced power supply noise.

Both differential lines at the output of the exclusive-or gate are individually filtered with higher order low-pass filter with a cutoff frequency of about 16 kHz to reject the RF components. The remaining DC components of these signals correspond to the positive and negative phase difference between the two inputs, as shown in Fig. 4. Each of these opposite DC levels are fed into two low-noise Analog to Digital Converter (ADC) drivers. The offsets and gains of these differential ADC drivers are configured such that the two ADCs exploit their full range. A common, very low-noise voltage reference has been used for both ADC drivers and ADCs, which are also in a single chip, to largely improve the noise canceling. The two 16-bit values produced by the ADCs are continuously and simultaneously read-out by the Field Programmable Gate Array (FPGA) at a maximum sampling frequency of 714.3 kHz. When a new value is digitized, the difference between the digital values of the two ADCs is computed internally in the FPGA, and is used by the regulation loop as explained previously. This numerical subtraction will effectively cancel out any exclusive-or logic level or reference voltage drift. The theoretical maximum sensitivity of the phase detector for a 400.8 MHz signal is thus 9.5 fs/LSB.

B. Delay control

In order to accurately compensate the phase drift of the signals over the long optical link, differential LVPECL delay lines, controllable by a coarse digital step delay interface and an analog fine tuning delay, have been selected. The range of the coarse delay is 10 ns with step increments of 10 ps, which is enough to cover the maximum expected phase drift; estimated to be below 6 ns (see Section II.A). As detailed at the end of next section, the accurate control of the propagation delay takes place during a relatively short period of time, which ensures a small phase drift excursion. Therefore, the regulation loop will change the coarse delay, if necessary, only

during idle periods in the SPS accelerator, avoiding reference signal glitches due to these changes.

The fine delay control has a range of 30 ps. The analog signal for the delay line is generated by an 18-bit Digital to Analog Converter (DAC) controlled by the FPGA, which allows quasi-continuous delay adjustment. As described in Fig. 2, the phase drift compensation system uses two delay lines for the actual phase drift compensation, while a third delay line was added to create a 90° phase shift of the exclusive-or reference signal, such that the phase discriminator zero occurs when the signals to be compared are in phase.

The only concern about the selected delay line performance during the design phase has been non-linearity effects, i.e. deviations from the ideal value for the coarse delay. To cope with this effect and to enable calibrating the 90° phase shift delay line, two calibration modes, which are largely described in the next section, have been defined.

C. *Calibration modes*

Due to the non-monotonicity and possible drift of the delay lines, some extra components have been introduced in the link stabilization module, as shown in Fig. 5.

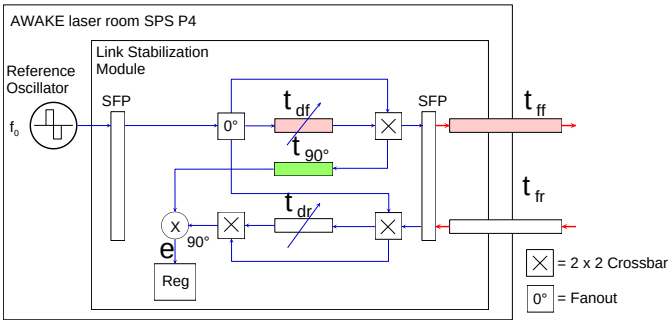

Fig. 5. Schematic representation of the circuitry used for switching between calibration and operation modes.

Digitally controlled crossbars and an extra delay line have been added in order to calibrate the phase detector null and linearize delay lines (τ_{df} and τ_{dr}) used in operation. The two calibration modes are called *cal90deg* and *calLocal*.

In *cal90deg* mode, crossbars are positioned in such a way that both τ_{df} and τ_{dr} delay lines are bypassed. Thus, one of the inputs of the phase detector comes directly from the input reference signal whereas the second one passes through the τ_{90} delay line, which is also fed by the input reference signal. In this mode, the regulation loop searches for a zero error value by changing τ_{90}. The phase detector null is then achieved when the two RF signals to be compared are in phase.

In *calLocal* mode, all delay lines are used, but the compared signals at the phase detector are local. The first branch of the phase detector takes the input reference signal and delays it by τ_{df} and τ_{90}, whereas the second branch goes from the input reference and passes the τ_{dr} delay line. Here, the regulation loop varies either τ_{df} or τ_{dr} in order to have the same propagation time in both branches, even if the coarse delay values are different due to non-linear effects. This calibration mode will be used to guarantee proper tracking of the τ_{df} and τ_{dr} delay lines.

Finally, in operation mode, the first branch of the phase detector only includes the τ_{90} delay line and the second branch has the four components included in Equation (1). To avoid the non-linear coarse steps in the delay lines, only fine delay tuning is used in operation mode (30 ps range) and both delay lines τ_{df} and τ_{dr} are used, as previously explained.

The beam takes approximately 4.3 s to accelerate from 26 GeV/c to 400 GeV/c in the SPS machine. Depending on where the proton beam should be delivered, this process may take place about once every 30 s. Thus, when the beam is not present in the SPS machine and the fine tuning delay of any of the delay lines is approaching the range limits, the link stabilization module may go through the two calibration modes to re-calibrate the phase detector null and lock in a different coarse step value.

IV. Module Performance

After designing and manufacturing the first prototypes of the link stabilization module, some tests have been carried out to validate the design and characterize its performance.

A. *Phase detector noise*

To validate the phase detector design, a 400 MHz signal has been injected at the module input, after being converted from electrical to optical by using the module described in Section II.B. The module has been put in *cal90deg* mode and the delay of the τ_{90} delay line has been kept constant in order to have the minimum phase noise between inputs of the phase detector.

In these conditions, an acquisition of the digitized traces of the raw ADC values and the phase values internally computed in the FPGA has been carried out. The noise analysis of the data is presented in Fig. 6.

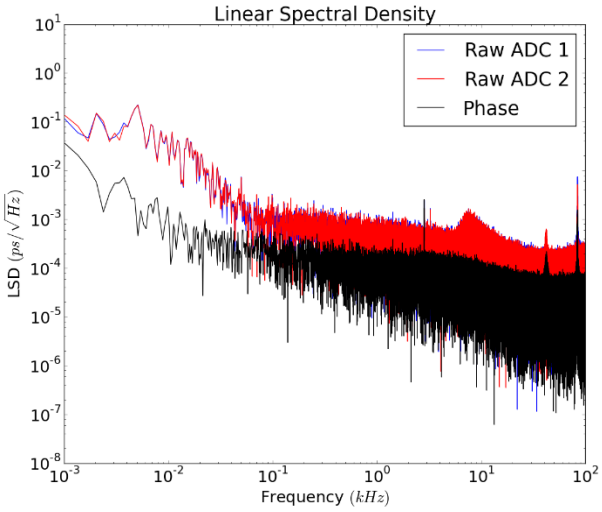

Fig. 6. Linear spectral density of the phase detector expressed in time units. The blue and red curves correspond to the raw digitized data from the ADCs whereas the black curve is the phase value computed internally in the FPGA.

Sampled traces at 714.3 kHz during 2.93 s show different noise figures when comparing the raw digitized values and the internally computed phase value, as depicted in the figure. The noise floor (average) computed between 0 and 100 kHz is 32.6 $fs_{RMS}/\sqrt{Hz}$ for the raw ADC data in both channels, while its value is 6.35 $fs_{RMS}/\sqrt{Hz}$ for the computed phase. These results validate the method used for the phase detector described in Section III.A, as it improves the noise floor by a factor 5.

B. Closed-loop performance

The regulation loop embedded in the FPGA is implemented by a digital Proportional-Integral (PI) controller with configurable rate and gain. In order to test the performance of the loop, the experimental set-up presented in Fig. 7 has been used.

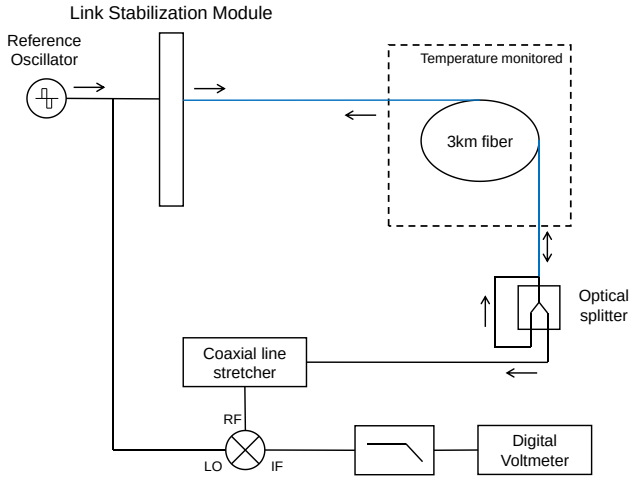

Fig. 7. Test-bench set-up used for measuring the performance of the module in closed-loop. The blue lines correspond to double optical fibers using the same plastic cover, as the ones between SPS points 3 and 4.

As shown in the figure, the 400 MHz reference signal coming from a reference oscillator goes to the link stabilization module, which is set to close the regulation loop after performing the two calibrations, as in operation mode. An RF mixer is used as an independent phase verification detector. One input of the mixer is connected to the output of the reference oscillator and the other one, through an adjustable coaxial phase shifter, to the signal coming from the optical splitter. In this way, the final configuration between SPS points 3 and 4 is reproduced in the laboratory, by using a 3 km dual optical fiber that is placed in a temperature monitored room. The adjustable coaxial line stretcher is used to null the mixer output voltage once the link stabilization module is locked. A precise Digital Voltmeter (DVM) is used to read out and store the DC level at the mixer output.

Firstly, the gain of the mixer is obtained by injecting the same signal with a known frequency to both inputs and changing the length at the line stretcher. Secondly, a set of data is acquired when keeping the temperature of the fiber constant while changing the temperature of the room with the rest of equipment to obtain a correction coefficient for temperature changes in the room where the link stabilization module operates. Finally, data from the DVM and link stabilization module are stored using the same interval basis while closing the regulation loop of the module using a set point that depends on the temperature of the module itself, using a feed-forward correction with the value computed in the previous point.

In this configuration, data from the DVM has been recorded for about 8 hours while tracking a temperature excursion in the fiber of about 0.25 degrees Celsius. The analyzed data have been fitted using a Gaussian curve, as presented in Fig. 8.

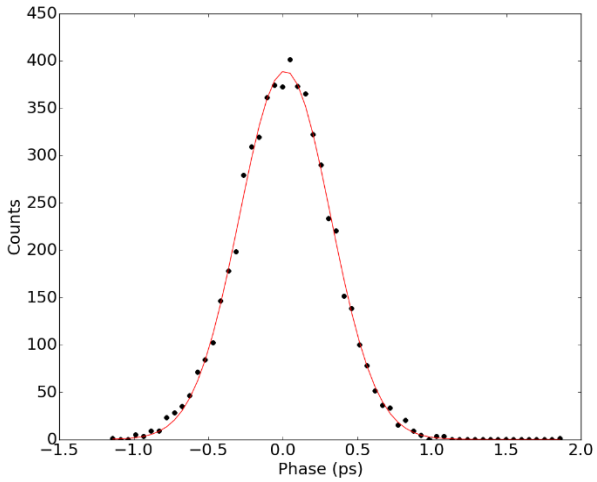

Fig. 8. Histogram of phase values obtained using the mixer. Black points correspond to experimental data and the red curve is a Gaussian fit of the data.

The Full Width at Half Maximum (FWHM) of the Gaussian fit presented in the figure has a value of 0.72 ps. Considering that either in the laboratory or in the logging system described in Section II.A, the observed drift when the temperature changes about 0.25 degrees Celsius is in the order of 50 ps, this preliminary value of the measured drift using an independent system shows very promising results for achieving a sub-picosecond precision in the link stabilization module.

V. Conclusions

A complete description of the proposed solution for the synchronization of the reference signals between the SPS points 3 and 4 for the AWAKE experiment has been presented. The developed module has demonstrated a very good performance in the tests completed in the laboratory with a preliminary measured phase drift of 0.72 ps with a temperature excursion of 0.25 degrees Celsius. In the following months, the module will be commissioned and integrated in the LLRF control system of the experiment. According to the obtained results, it is expected that the module will completely fulfill the requirements of the experiment, which will be verified with the first commissioning beams for AWAKE in summer 2016.

Acknowledgment

Authors would like to thank T. Levens for the phase detector module development used in the logging system and

the platform used for developing the link stabilization module, T. Bohl for the initial set-up of the logging system and phase noise measurements, H. Damerau for his help on the paper review and the AWAKE collaboration for supporting this work.

REFERENCES

[1] AWAKE Collaboration, "A proton-driven plasma wakefield acceleration experiment at CERN", CERN-SPSC-2013-013, SPSC-TDR-003, CERN Geneva Switzerland, 2013.

[2] E. Gschwendtner et al., "AWAKE, The Advanced Proton Driven Plasma Acceleration Experiment at CERN", Nucl. Instrum. Methods Phys. Res., A, doi: 10.1016/j.nima.2016.02.026, 2016.

[3] P. Baudrenghien, T. Linnecar, D. Stellfeld and U. Wehrle, "SPS Beams for LHC: RF Beam Control to Minimize Rephasing in the SPS". CERN-SL-98-027-RF, 6th European Particle Accelerator Conference, Stockholm, Sweden, 1998.

[4] P. Baudrenghien. "Beam Control for Protons and Ions". 9th LEP-SPS Performance.Workshop, Chamonix, France, 1999.

[5] W. Kester, "Converting Oscillator Phase Noise to Time Jitter". Technical Note, Analog Devices, 2008.