Available online at www.sciencedirect.com

SciVerse ScienceDirect

Physics Procedia 37 (2012) 2063 – 2071

Physics Procedia

TIPP 2011 - Technology and Instrumentation in Particle Physics 2011

The FONT5 bunch-by-bunch position and angle feedback system at ATF2

R. J. Apsimon[a], D. R. Bett[a], P. N. Burrows[a], G. B. Christian[a,*], B. Constance[a,b], M. R. Davis[a], A. Gerbershagen[a,b], C. Perry[a], J. Resta-Lopez[c]

[a] *John Adams Institute for Accelerator Science, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford, United Kingdom*
[b] *CERN, Geneva, Switzerland*
[c] *Instituto de Fisica Corpuscular (IFIC), Centro Mixto Universidad de Valencia-CSIC, Valencia, Spain*

Abstract

The FONT5 upstream beam-based feedback system at ATF2 is designed to correct the position and angle jitter at the entrance to the ATF2 final-focus system, and also to demonstrate a prototype intra-train feedback system for the International Linear Collider interaction point. We discuss the hardware, from stripline BPMs to kickers, and RF and digital signal processing, as well as presenting results from the latest beam tests at ATF2.

Keywords: Linear colliders, beam-based feedback

1. Introduction

The motivation for the development of beam-based intra-train feedback systems described in this article lies in demonstrating the feasibility, and technical prototyping, of collision optimisation feedback systems for the interaction point (IP) of a future linear collider, such as the International Linear Collider (ILC) or Compact Linear Collider (CLIC). In the case of ILC, the baseline design for a centre-of-mass energy of 500 GeV envisages approximately 20 km of superconducting-RF based linear accelerator (linac). The design luminosity for such a collider can be expressed as

$$L = f_{rep} n_b \frac{N^2}{4\pi\sigma_x\sigma_y} H_D, \tag{1}$$

where f_{rep} is the pulse repetition frequency, n_b is the number of bunches per pulse, N the bunch intensity, and σ_x and σ_y the horizontal and vertical beam spot-size at the IP respectively. H_D is the luminosity enhancement, or pinch, factor, which is approximately 2 for e^+e^-, due to the self focussing of the electron and positron bunches. The machine repetition rate and current are governed by the choice of RF in the linac, and so in order to maximise the luminosity, the only remaining parameters which can be optimised are the horizontal and vertical IP spot-sizes. To achieve the

*Corresponding author
Email address: g.christian1@physics.ox.ac.uk (G. B. Christian)

1875-3892 doi:10.1016/j.phpro.2012.03.758

nominal luminosity requirement for precision physics at the ILC of 2 x $10^{34} cm^{-2} s^{-1}$, spot sizes of ∼5 nm vertically, and ∼ 500 nm horizontally, are required [1]. The more stringent requirement in one dimension is necessary to mitigate against the effects of energy loss through bremsstrahlung radiation in the field of the opposing bunch, and the choice of the vertical plane is due to emittance preservation considerations. The choice of superconducting-RF acceleration for the ILC limits the repetition rate to 5 Hz, at which frequency a pulse-to-pulse feedback system would be ineffective to correct for dynamic beam misalignments due to ground motion and facilities noise. In the case of CLIC, the repetition rate is greater but the tolerance on the spot size and stability of the IP is much reduced [2]. For either design, maintaining collisions would not be possible without mitigation for dynamic misalignments operating within the duration of the bunch train.

The envisaged IP feedback system (IPFB) operating on short, intra-train timescales, would be one of several beam-based feedback and feed-forward systems in the machine, alongside systems designed for orbit correction on pulse-to-pulse, and longer, timescales. A schematic of the IPFB is shown in Fig. 1. The system works by virtue of the strong beam-beam kick which is imparted on a bunch in the field of the opposing bunch, when the two bunches arrive at the IP with a non-zero offset. Within a certain range of IP offset, defined to be the capture range of the system, the angular deflection varies approximately linearly with offset. The capture range of the system is about ∼100 nm for ILC, and it a task of the slower orbit correction systems to deliver the beam to within the capture range of the IPFB. The large beam-beam deflection angle at the IP translates to a large position offset downstream, which is easily measured with a beam position monitor (BPM). The IPFB system works by measuring this vertical position offset in a stripline BPM in the outgoing beam-line from the IP, and via a fast signal processor, amplifier and kicker, correcting the opposing beam incoming to the IP. One critical parameter of such a system is the latency, or the time it takes for the system to respond to a change in its inputs. This includes both the delay in the electronics and cables and beam time-of-flight around the IP. The system is located as close as possible to the IP, and the measurement and correction are made on opposing beams to ensure the proximity of the BPM and kicker, in order to minimise latency. The feedback system acts to correct the measured position to a given set-point. The first correction will be made after one latency period, and the delay loop is necessary to maintain the correction for subsequent latency periods by providing 'memory' of the sum of all previous corrections. Several prototype systems for the IPFB have been designed and tested on beam-lines, as part of the Feedback on Nanosecond Timescales (FONT) project. Previous and current generations of prototypes are summarised in the following sections.

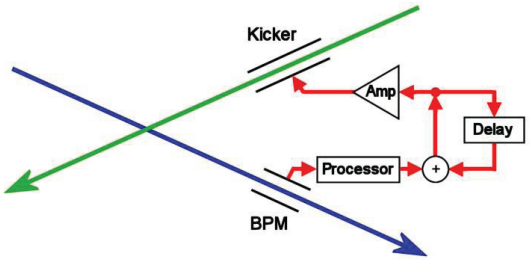

Figure 1: Schematic of the interaction point feedback system.

1.1. Summary of previous FONT feedback prototypes

The original motivation behind the FONT programme was to demonstrate an IPFB system for the 'warm' linear collider designs, such as the Next Linear Collider (NLC) and the Japan Linear Collider (JLC), based on room-temperature RF acceleration. From the point of view of collision optimisation, the 'warm' machines present a greater challenge in that the bunch train is very short, ∼270 ns, and so the IPFB must be very low latency in order to recover a significant amount of luminosity. The first three FONT prototypes were aimed at demonstrating an ultra-low latency correction. The first two, FONT1 [3] and FONT2 [4], were installed on the NLC Test Accelerator (NLCTA) at SLAC from 2001–2004, and FONT3 [5] was installed at the KEK Accelerator Test Facility (ATF) from 2004–2005. The NLCTA ran at 65 MeV, with a 170 ns train length and 87 ps bunch spacing, and FONT1 achieved a latency of 67 ns,

and FONT2 54 ns. Following this, from 2004–2005, FONT3 was installed at the KEK Accelerator Test Facility (ATF), which provided a 1.3 GeV beam, with 56 ns train duration and 2.8 ns bunch spacing. The aim was to take advantage of the ∼GeV energy beam, in that approximately the same kick is required to make a micron-scale correction at 1 GeV, as a nanometre-scale correction at 1 TeV, and to demonstrate a latency as low as 20 ns in order to observe almost three latency periods in the 56 ns train. The actual latency achieved was 23 ns, which although slightly longer than the goal, is a very significant result and still relevant for the IPFB for CLIC, where a train duration of 156 ns is currently envisaged [2].

Following the recommendation of the International Technology Recommendation Panel for the adoption of superconducting RF cavities for the ILC, the FONT programme turned to feedback demonstrations more specific to the ILC time structure. For the ILC, the nominal train duration is almost 1 ms and the bunch spacing ∼180–500 ns [1], relaxing the latency requirement and allowing for bunch-by-bunch feedback. The longer bunch spacing allows for digital signal processing or combined analog-digital techniques to be employed. Such a digital system would facilitate a greater amount of processing, for example the running of sophisticated algorithms or integrating signals from other systems, such as a luminosity monitor, to recover an even great proportion of the luminosity. FONT4 [6, 7, 8, 9], from 2005–2009, was a demonstration of a digital feedback system with ILC-like bunch spacings. This was installed at ATF with a three bunch beam with a bunch spacing that could be varied between ∼140–154 ns. The latency aim for FONT4 was to be less than the bunch spacing in order to achieve bunch-by-bunch feedback, and was measured to be ∼140 ns with the minimum digital processing, and increasing to 148 ns with the inclusion of real-time charge normalisation. FONT5, from 2009 to present, is a demonstration of a bunch-by-bunch 'two-phase' position and angle feedback system and is described in more detail in the following sections.

1.2. ATF2 summary

The ATF2 project is a scaled down mock up, in the upgraded ATF extraction line, of the final focus scheme for ILC and CLIC. The goals are, firstly, to demonstrate a 37 nm vertical spot-size at the optical focal point of the system (called the IP in analogy to a collider interaction point); and secondly, to demonstrate nanometre-level position stability at the virtual IP. Several articles provide more information on ATF2, its goals and current status (see for example articles by White, Yan, Yamaguchi, and [10]).

Fig. 2 shows a schematic of layout of the ATF. A 1.3 GeV s-band linac is used to accelerate the beam, which is then damped in the damping ring (DR) and finally extracted into the extraction line. The FONT5 feedback installation is located close to the start of the extraction line, upstream of the ATF2 final focus beam line. This system will contribute towards meeting the second ATF2 goal, by providing bunch-by-bunch stabilisation at the entrance to the final focus system. The first ATF2 goal is being pursued with single bunch beam, but for the second goal, beam-based feedbacks are considered essential to achieve the level of stability required and therefore a bunch train is implied - the stabilisation being demonstrated on later bunches in the train. The ATF can currently extract up to three bunches from the damping ring with an ILC-like bunch spacing (maximum 154 ns for three bunches), however with a new fast extraction kicker scheme it could provide up to 60 bunches. This fast extraction kicker programme is a separate R&D programme to ATF2, essential for demonstrating the feasibility of the ILC damping ring scheme, see [11] for details.

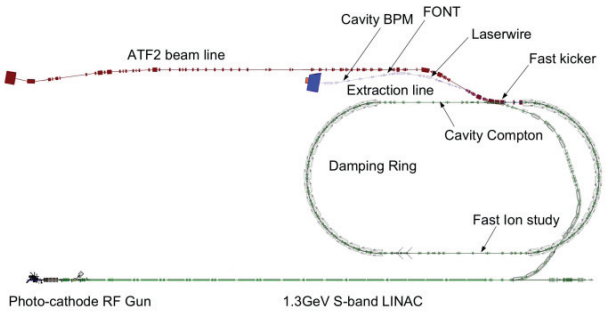

Figure 2: Layout of the ATF facility and ATF2 beam-line.

2. The FONT5 System

The nanometre position stability requirement at the ATF2 IP translates to ∼1 μm stability at the entrance of the ATF2 final focus. For this reason the FONT5 BPM resolution target was set to be 1 μm. The FONT5 system is a bunch-by-bunch position and angle feedback system; a schematic of the beam-line in the region is shown in Fig. 3. The system consists of three stripline BPMs (on movers), P1, P2, and P3 (dark blue in Fig. 3), and two stripline kickers, K1, and K2 (green). The components in light blue are quadrupole magnets, part of the optical lattice of the extraction line. The system is two-phase in that there are two discrete loops, P2-K1 and P3-K2. In the ideal case the two loops would be orthogonal in phase advance. The first loop would correct the position at P2, which equivalently would correct the angle at P3; the second loop would correct the position at P3, equivalent to the angle at P2. In this way, both position and angle components of the beam jitter will have been corrected at either P2 or P3, and therefore also at an arbitrary phase downstream. In reality it is not possible to ensure $\pi/2$ phase advance between the two loops, as well as impose a $\pi/2$ phase advance between the pairs of kickers and BPMs within each loop, but in this case, both loops can be coupled together, in that each kicker drive signal is formed from a linear combination of both P2 and P3 measurements. The relative contribution of each measurement to either loop is determined by the degree of coupling between the loops.

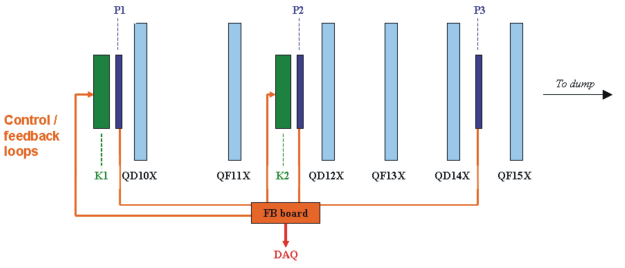

Figure 3: Layout of the FONT5 system on the ATF extraction line.

3. FONT hardware

The signal processing in the feedback loop consists of three main stages: an analogue front-end processor, to downmix the raw stripline signals to baseband, and produce a sum and difference signal, proportional to the charge, and position and charge respectively; an FPGA-based digital processor board, responsible for the sampling of the analogue bunch signals and acting as a feedback controller; and an amplifier to provide the required drive to the kicker. The analogue front-end signal processors process the signals from the opposing strips of the BPM, which exhibit a peak in their power spectrum at 625 MHz, and downmix them to below 100 MHz. A 180° hybrid coupler is used to obtain a difference signal from the raw stripline signals, and a sum signal is formed using a resistive combiner. The resulting signals are then bandpass filtered, and mixed with a local oscillator at 714 MHz sourced from the machine and hence phase locked to the beam, and finally low-pass filtered to reject the high frequency output component of the mixer. This produces sum and difference output pulses with widths of ∼5 ns, FWHM. The latency of the analogue front-end processor is measured to be ∼10 ns.

These signal are then sampled using the FONT5 digital signal processor. A block diagram of the FONT5 board is shown in Fig. 4 This is a 9-channel digitiser and feedback controller, with two kicker drive outputs. Three channels are used for each BPM, nominally horizontal difference, vertical difference, and sum for each. The board is centred around a Xilinx Virtex-5 [12] FPGA, and uses low-latency, fast, ADCs, both of which are clocked at 357 MHz, with a source synchronised to the machine RF. The ADCs used are Texas Instruments ADS5474 [13], 14-bit, monolithic pipeline ADCs, with a latency of 3.5 clock cycles, and the DACs are 14-bit, 210 MSPS, Analog Devices AD9744 [14], with a minimum latency of 3 ns. The digital processor is responsible for synchronising to the machine timing, governing the sampling of the analogue waveforms from the BPM processors, and setting the correct gains for feedback operation. Given the relatively fast sampling rate, the FONT5 board is also used for the data acquisition from the system, and

commands and data are sent serially via a UART over RS232. Other tasks of the FPGA, which can be done more easily within the FPGA logic than would be possible with a purely analogue system, include real-time charge normalisation, where the difference signal is divided by the sum signal in real-time to immunise against the effects of charge variation which would otherwise appear as a position variation; finite impulse response filtering to correct for the droop in the kicker drive signal in the amplifier over time; and static offset removal, where a systematic 'banana' shape to the train profile can be removed alongside the operation of the feedback.

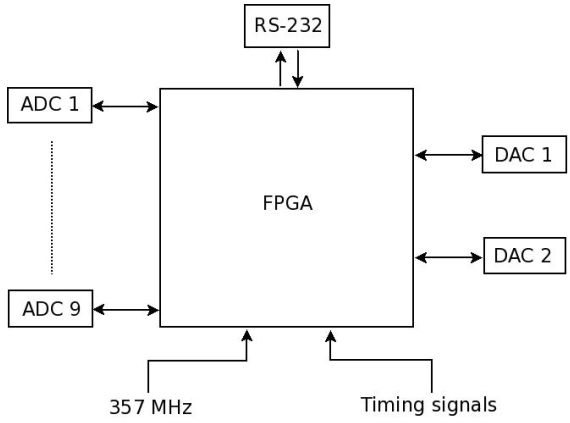

Figure 4: Block diagram of the FONT5 signal processing board.

The final stage of the feedback system hardware are the kickers and their drive amplifiers. Two stripline kickers are used in the feedback system, originating from SLAC, and three custom amplifier units have been manufactured by UK industry [15], one for each kicker and one spare. The amplifiers were specified for pulsed operation, with an operational period of up to 10 μs at a repetition rate of 10 Hz, with an actual duty factor of 0.01 % being used at ATF. The amplifier/kicker bandwidth is ∼30 MHz, with a corresponding 40 ns settling time to 90 % of full kick, and output current of up to ±30 A. A maximum beam deflection of 100 μrad was measured, consistent with the design of the kicker and measurements of the amplifier.

4. Beam-line Test Results

4.1. System Latency

The latency is measured with a special mode of the FPGA firmware, where a constant DAC value is set, providing a constant drive signal to the amplifier, and hence a static kick will be produced on the beam one latency period after the measurement of the first bunch. Provided the latency is less than the bunch spacing then the effect of the kick will be observed in the position of bunch 2, and the kick can be delayed in time until the effect of the kick can no longer be seen in the bunch 2 position, effectively mapping out the leading edge of the kicker pulse with the beam position. Data was recorded with interleaved kicked and un-kicked beam for each delay setting used, to mitigate against slow beam drifts, and averaged at each setting to reduce the effect of beam jitter on the measurement. Fig. 5 shows the average difference between kicked and un-kicked position as a function of additional delay applied, for the loop P3-K2, from April 2010. The system latency is defined as the point where 90 % of the full scale deflection is seen in the kicked beam. For the P3-K2 loop this occurs at a delay setting of approximately 20 ns, which, for a bunch spacing of 151.2 ns, corresponds to a latency of approximately 129 ns. The corresponding value for the P2-K1 loop was measured to be ∼133 ns, and for the most critical path for coupled loop operation, P3-K1, this value increased to ∼142 ns, due to the extra beam time of flight and signal return time. A breakdown of the contributions to the single loop latency is given in Table 1.

4.2. Feedback Performance

The feedback algorithm employed measures the first bunch and attempts to nullify the position of subsequent bunches. The formula for bunch-to-bunch feedback is given by Eq. 2, where y'_n is the position of bunch n with

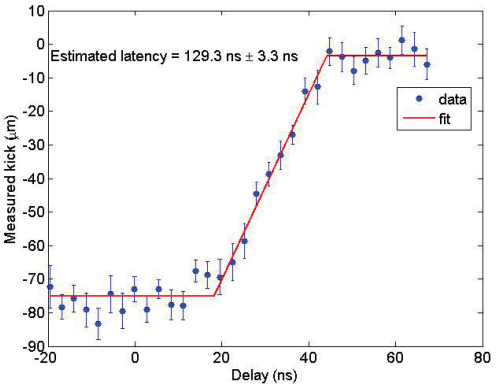

Figure 5: Average difference between kicked and un-kicked positions for bunch 2 at P3, as a function of additional delay applied to the constant amplifier drive. This data was for the P3–K2 loop, with a bunch spacing of 151.2 ns.

Time of flight kicker–BPM (3.5 m)	12 ns
Signal return time BPM–kicker	32 ns
Irreducible latency	*44 ns*
BPM processor	10 ns
ADC (3.5 cycles of 357 MHz)	10 ns
FPGA signal processing (8 cycles of 357 MHz)	22 ns
DAC (1 cycles of 357 MHz + 1 ns)	4 ns
FPGA I/O	3 ns
Amplifier	35 ns
Kicker fill time	3 ns
Electronics latency	*87 ns*
Total latency	**131 ns**

Table 1: Single loop latency component breakdown.

the feedback on, and y_n an y_{n-1} are the measured position of bunch n and bunch *n-1* with the feedback system not operating. As it is the corrected position that is measured for the second and subsequent bunches, and not the incoming position that the bunches would have with the system switched off, the algorithm uses the corrected positions as in Eq. 3. The corrections made to earlier bunches in the train are accumulated on a 'delay loop' register in the FPGA, which forms the memory of the cumulative corrections made to the bunch-train. In the FPGA firmware, the position of each bunch is calculated and multiplied by the appropriate gain factor, and added to the cumulative correction from the delay loop. This signal is then output on the DAC to kick the next bunch in the train, and the delay loop is updated. The feedback algorithm assumes the position of bunch n and bunch *n-1* are well correlated. In order to make the required micron level corrections to a beam jitter of ~3 μm would require measured correlations of 94 % and above. For reference, a 50 % correlation would result in the system breaking even, and for anything lower the feedback system would increase the jitter in the beam position.

$$y'_n = y_n - y_{n-1} \quad (2)$$

$$= y_n - (y'_{n-1} + y'_{n-2} + \ldots) \quad (3)$$

Fig. 6 shows the performance of the system in correcting offsets introduced into the bunch train with an upstream dipole corrector magnet. The mean position of each bunch is plotted, averaged over 50 machine pulses. It can be seen that for every corrector setting the first bunch remains at the initial train offset, as would be expected given that

bunch 1 is only measured but not corrected, and that the second and third bunches converge to an offset near to the nominal *zero* position as measured by the processor. This type of plot shows the system working *on average* but is also useful to verify the correct gain setting for the feedback, as the correct gain setting is that for which the average position will converge for all values of the initial offset. For this data static bunch-to-bunch offsets were removed from the positions of bunches 2 and 3, to correct for a 'banana' shape along the train as measured with the feedback system switched off. The fact that the mean position of bunch 3 is not as close to zero as for bunch 2, indicates that some residual banana still exists in the train after offset removal.

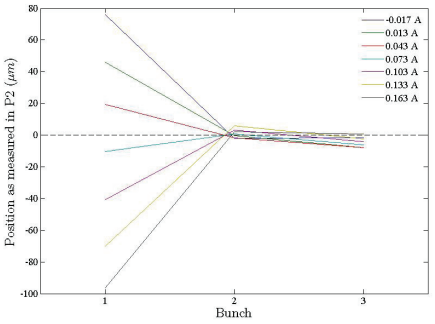

Figure 6: Feedback performance on correcting a series of initial train offsets. Mean position is plotted for three bunches, averaged over 50 machine pulses.

Although Fig. 6 shows the operation of the feedback in nullifying position offsets introduced to the train on average, it conceals the r.m.s. spread in the measured position, which is given by the quadratic sum of the BPM processor resolution and the actual position jitter of the beam. For operation at ATF2 the most important figure of merit for the system is the reduction in the beam jitter, for correlated bunch-to-bunch jitter. Figure 7 shows distributions of the measured beam position of the three bunches with both the feedback system switched off (blue), and on (red), for a dataset where the measured bunch-to-bunch jitter was particularly highly correlated. The incoming (feedback off) bunch-to-bunch correlation were measured to be 98 % for bunch 1 – bunch 2 and 89 % for bunch2 – bunch3. The data shown is for operation of the coupled loop system, and the feedback is interleaved, in that the feedback is operating only every other machine pulse, to mitigate against skewing the results if the incoming beam jitter changes as a function of time. In the case of bunch 1, the position distributions have the same spread, or jitter, of 2.1 μm with both the feedback on and off, as would be expected. For the second and third bunch, however, the operation of the system reduces the jitter, in the case of bunch 2 a factor of five from 2.1 μm down to 0.4 μm, and down to 1.1 μm in the case of bunch 3. The corrected jitters observed are as would be expected given the measured bunch-to-bunch correlations, using Eq. 2 and the standard formula for propagation of correlated errors. The level of jitter correction achieved for bunch 2 with the feedback operating was unexpected, as it was previously believed that the BPM processor resolution was ∼1–2 μm, as measured using the system of the three BPMs, where the position in two of the BPMs is used to predict that in the third. The resolution is then defined as the standard deviation of the residuals from the difference between predicted and measured positions, although the calculation assumes that the resolution will be the same for each processor and gives an 'average' result. The 400 nm position spread measured at P2 indicates that the resolution, at least for the processor at P2, must be below this value in order to be able to measure it, and furthermore, must be at least $\sqrt{2}$ less than this, as the best correction the system could possibly achieve is a factor $\sqrt{2}$ greater than the BPM resolution.

Figure 8 shows the corresponding results at BPM P3 for the same coupled loop feedback run as Fig. 7. As before, bunch 1 is uncorrected and the r.m.s. position spread measured to be 3.4 μm with the feedback off and 3.2 μm with the feedback operating. For the second and third bunch this jitter is reduced from 3.3 μm to 1.8 μm for bunch 2, and 3.3 μm to 1.8 μm for bunch 3. The corresponding bunch-to-bunch correlations were 87 % and 84 % for bunch 1 – bunch 2 and bunch 2 – bunch 3 respectively, again accounting for the correction factors observed. The feedback performance at P3 implies a BPM processor resolution of 1.1 μm or less.

The discrepancy between the feedback results and the measured BPM resolution with the three BPM method is

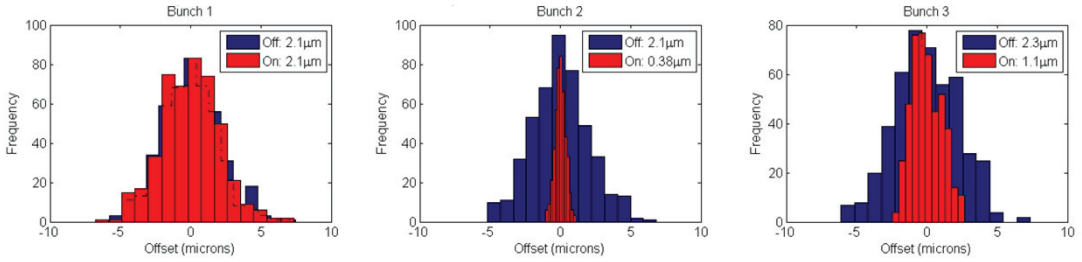

Figure 7: Position distributions for the three bunches at P2 showing the reduction in measured beam jitter with coupled feedback operation, with interleaved feedback off (blue) and feedback on (red).

believed to be due to intrinsic differences in the BPM processors, resulting in differing sensitivity to jitter in the local oscillator for the processors. The largest effect is thought to be due to path length imbalances to the RF hybrid. It was previously assumed that the input cables from the BPM striplines to the processors should be well matched in length, but the actual path length differences will be unique for every processor and hybrid device. It appears from the feedback data that the processor at P2 is particularly well optimised, with the processor at P3 less so, and the processor at P1 being the worst, and that the three BPM resolution method gives an average value across the three processors. An additional effect, having subsequently been observed, is pickup of noise into the ADC clocking network on the digital board which may manifest itself as timing jitter on the ADC clocks, uncorrelated across channels. The present aim is to be able to reproduce the performance at P2 with all three BPM processors, and all processors are currently being optimised, and timing jitter mitigated against, and the performance checked in upcoming tests on the ATF beam-line. Even with a set of very well optimised processors, the system performance will still be determined by the degree of bunch-to-bunch correlation. Unfortunately such correlations as measured above were seen only very rarely, with more typical measurements being around the 50–60 % level, and the mechanism responsible for the lack of correlations is still to be understood.

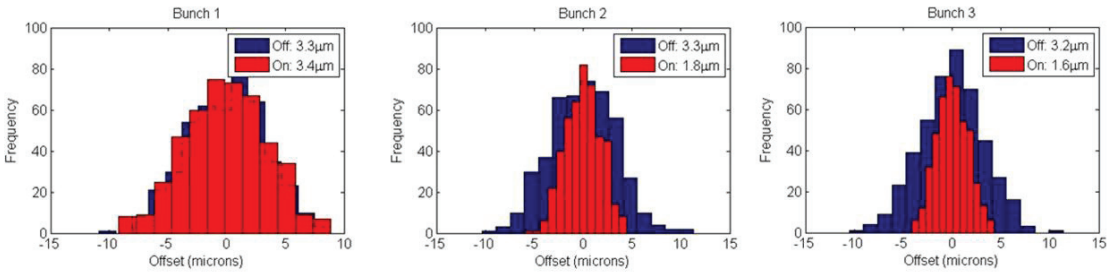

Figure 8: Position distributions for the three bunches at P3 showing the reduction in measured beam jitter with coupled feedback operation, with interleaved feedback off (blue) and feedback on (red).

5. Conclusions

The two main challenges for bunch-to-bunch feedback systems which have been discussed are resolution and latency. Results from the FONT5 bunch-by-bunch feedback system at ATF2 demonstrate a sub-micron correction on the beam jitter, indicating a level of resolution below ∼300 nm, at least for one BPM processor in the system. This performance has still to be demonstrated across all three BPMs in the system, and with the standard three-BPM resolution method, and attempts to understand and optimise the processor performance are currently ongoing. This level of performance would surpass the 1 μm goal set for the FONT5 feedback system at ATF2, and would be more

than adequate for the IP feedback system at the ILC or CLIC, where the beam-beam deflection and large lever-arm from the IP result in a resolution of a few microns being adequate.

Very low latencies, down to 23 ns, have previously been obtained with purely analogue prototype feedback systems FONT1, FONT2, and FONT3, although these were not bunch-to-bunch systems. An analogue system such as FONT3 would be suitable for CLIC where the total train length is 156 ns and the bunch spacing 500 ps. For machines with longer bunch spacing, such as the ILC, more processing can be done, and in a simpler manner with a digital feedback system, although at the expense of increasing the latency. For the current generation digital feedback prototype FONT5, a latency of ∼130 ns has been measured for single-phase, or uncoupled, correction (i.e. correction of position at one phase in the lattice, or with orthogonal position and angle correction). For coupled two-phase correction (i.e. non-orthogonal position and angle correction) this value increases to ∼140 ns, due to the extra path length from the final BPM to the first kicker. These measured latencies are well within the specification for bunch-to-bunch feedback at ILC, where the bunch spacing will be ∼180–500 ns.

The remaining goals for the FONT5 position and angle feedback system are, firstly, to demonstrate the same level of performance of jitter reduction at both feedback control BPMs, as has previously been observed at BPM P2, taking into account the measured bunch-to-bunch correlation. A secondary goal would be to then demonstrate the same level of performance using BPMs external to the feedback loops, and even, if possible, a corresponding reduction of beam jitter at the ATF2 virtual IP.

References

[1] ILC Global Design Effort and World Wide Study, ILC Reference Design Report, http://www.linearcollider.org/about/Publications/Reference-Design-Report
[2] F. Tecker (Ed.) et al., CLIC 2008 parameters, CLIC-NOTE-764
[3] P.N. Burrows, Feedback on Nanosecond Timescales (FONT): Results from First Beam Tests at the NLCTA at SLAC, Proceedings of the 2003 Particle Accelerator Conference, Portland, Oregon, p.687
[4] P.N. Burrows et al., Nanosecond-timescale Intra-bunch-train Feedback for the Linear Collider: Results of the FONT2 Run, Proceedings of EPAC 2004, Lucerne, Switzerland, p.785
[5] P.N. Burrows et al., Tests of the FONT3 Linear Collider Intra-Train Beam Feedback System at ATF, Proceedings of 2005 Particle Accelerator Conference, Knoxville, Tennessee, p.1359
[6] P.N. Burrows et al., Design of the ILC Prototype FONT4 Digital Intra-train Beam-based Feedback System, Proceedings of EPAC, Edinburgh, Scotland, 2006, p.849
[7] P.N. Burrows et al., The FONT4 ILC Intra-train Beam-based Digital Feedback System Prototype, Proceedings of PAC07, Albuquerque, New Mexico, p.416
[8] R. Apsimon et al., Design and Performance of a Prototype Digital Feedback System for the International Linear Collider Interaction Point, Proceedings of EPAC08, Genoa, Italy, p.3245
[9] R. Apsimon et al., Beam Test Results with the FONT4 ILC Prototype Intra-Train Beam Feedback System, Proceedings of PAC09, Vancouver, BC, Canada, p.2676
[10] P. Bambade et al. (ATF Collaboration), Present status and first results of the final focus beam line at the KEK Accelerator Test Facility, Phys. Rev. ST Accel. Beams 13, 042801 (2010)
[11] T. Naito et al., Multibunch beam extraction using the strip-line kicker at the KEK Accelerator Test Facility, Phys. Rev. ST Accel. Beams 14, 051002 (2011)
[12] Xilinx, Inc., http://www.xilinx.com/products/virtex5/index.htm
[13] Texas Instruments, Inc., http://www.ti.com
[14] Analog Devices, Inc., http://www.analog.com
[15] TMD Technologies Limited, http://www.tmd.co.uk