Chapter 2

FONT Fast Feedback Systems

The IP fast offset-correction feedback as described in the Reference Design Report for the ILC [11] is being developed under the heading of FONT (Feedback on Nanosecond Timescales). Since the formation of the FONT project, the work has progressed from super-fast analogue electronics suitable for a warm machine to a digital processor suitable for a machine with superconducting technology (such as the ILC). Testing was completed for the analogue electronics at NLCTA (at SLAC) [45] and ATF (at KEK) [46]. The digital electronic tests are ongoing at ATF and, in 2008, ATF2 [47].

The tests are performed on a single beam, not two colliding beams. This is due to the fact that no available test beams have collision points and the R&D nature of the FONT tests requires test facilities rather than physics facilities as regular accesses to the beamline are needed to adjust the electronics. The principle of beam position measurement and correction can be demonstrated with a single beam as illustrated in figure 2.1. An offset ($\sim$100μm) can be given to the test beam from an upstream dipole magnet that kicks the beam off axis. The BPM instrumented with FONT electronics can measure this offset and produce a position signal to be amplified and sent to a kicker positioned between the dipole magnet and the BPM. The kicker kicks the beam such that the position of later bunches is zero in the BPM assuming the gain on the amplifier is set correctly so the correct kick is given. This demonstrates the ability to perform feedback allowing for measurements of its latency and accuracy.

Once the correction is successfully applied, the BPM reads no position offset and so no signal is passed to the amplifier from the BPM processor and the position correction is lost. This is unless part of the signal from the amplifier is delayed and applied to later bunches, thus retaining the correction even when no position offset is seen in the BPM. This delay loop can also tested in the scenario illustrated in figure 2.1.

2.1 FONT1

The first generation of FONT was designed to meet the challenges of the proposed warm machine, the Next Linear Collider (NLC) [48]. This represented the most challenging IP

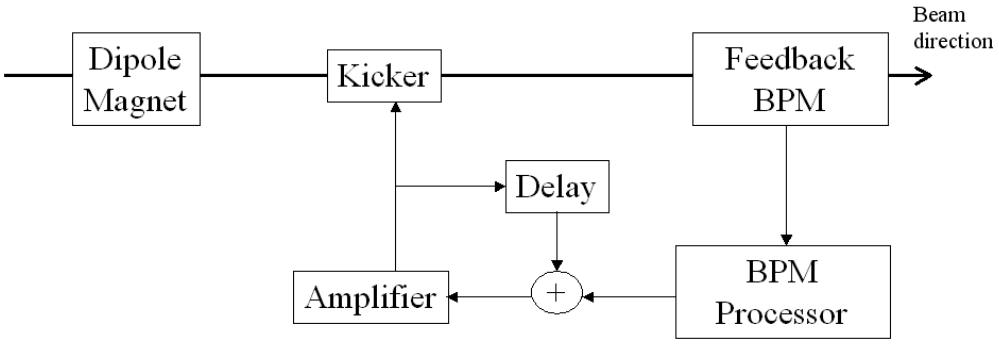

Figure 2.1: Schematic of FONT tests on a single beam showing the three necessary beamline components (the upstream dipole magnet, the kicker and the BPM for feedback) and the components of the feedback system (the BPM processor, the amplifier and the delay loop).

Parameter	NLCTA
Beam energy (GeV)	0.062
Electrons per bunch	1×10^8
Bunches per train	$\sim$2000
Bunch spacing (ps)	88
Train length (ns)	177
σ_x (μm)	500
σ_y (μm)	1000

Table 2.1: Beam parameters at NLCTA [49]

feedback for a potential future linear collider at the time due to the short bunch trains of a mere 270 ns.

Testing for the feedback electronics took place at the NLC Test Accelerator (NLCTA) at SLAC. Here the bunch trains were 177 ns long (see NLCTA parameters in table 2.1). To demonstrate the principle of FONT, it was necessary to show a feedback correction applied to the beam and then retained due to the action of the delay loop. Therefore the goal was to produce a feedback system with a latency[1] of a few tens of nanoseconds. It was predicted to be 67 ns as shown in table 2.2. Most of the latency was irreducible and came from the physical setup at NLCTA. These are the "time of flight" which is the time for the bunch to travel from the kicker to the BPM, and the "signal return delay" which is the time for the signal from the BPM to travel through the cables upstream to the kicker. The only way to reduce this is to move the BPM and the kicker closer together as was done for FONT2. The rest of the latency is from the BPM processor, the feedback circuit (where the delayed position signal was added to the current position signal), the rise-time of the amplifier and the fill-time of the kicker.

The experimental setup was as described in figure 2.1 with 4.19 m between the kicker and the feedback BPM [50]. The FONT BPM was a button BPM as described in section 1.5.1. The BPM processor worked by taking the difference between the signals from the top and

[1]The latency is defined here as the time between the position measurement of the first bunch and the position measurement of the first correction.

		Time (ns)
Irreducible	Time of flight of bunches	14
	Cable delay	18.5
Reducible	Processor	5
	Feedback circuit	10.5
	Pre-Amplifier	5
	Kicker Amplifier	12
	Kicker fill-time	2
	Total	67

Table 2.2: Latency prediction for FONT1 [50].

bottom electrodes and then dividing by the sum (for charge normalisation). However, inverting the sum took a lot of time since it was digitised, inverted and then reproduced as an analogue signal using an Arbitrary Waveform Generator (AWG). Therefore the inverted sum signal was applied to the difference signal for the next pulse.

The upstream dipole magnet was used to give the beam an offset from zero. Five beam positions were used as shown in figure 2.2(a). The red and blue lines show the position signal for the entire bunch train at the largest position offsets given to the beam, the magenta and cyan lines show bunch trains at small position offsets and the green line shows a bunch train close to the centre of the beam pipe (zero position offset). The separation between the bunches was 88 ps which is too fast for the processor to resolve.

The feedback was turned on (without the delay loop) and the measured beam position from the start of the train was used to correct later bunches as shown in figure 2.2(b). From around 140 ns, the beams start to show the effects of a kick that brings them to the same position. This position is not zero due to a static shape to the bunch train that is evident in figure 2.2(a) but the offset beams do all converge on one trajectory demonstrating successful feedback and position correction. The beams return to their original offsets around 210 ns as the correction is lost; the BPM is reading the "corrected" bunches that are near zero and so no longer provides the position signal required.

In figure 2.2(c) the delay loop is on and the correction is retained until the end of the bunch train.

The latency was 75 ns measured as the time between the start of the pulse and the start of the correction and 70 ns measured as the time during which the correction acts (from the data in figure 2.2(b), with the delay loop off) [50].

2.2 FONT2

The second generation of FONT aimed to improve the latency to 53 ns to observe more correction periods. This was achieved by moving the kicker closer to the feedback BPM, from 4 m as in FONT1 to a distance of 2 m. As shown in table 2.3, the expected lower

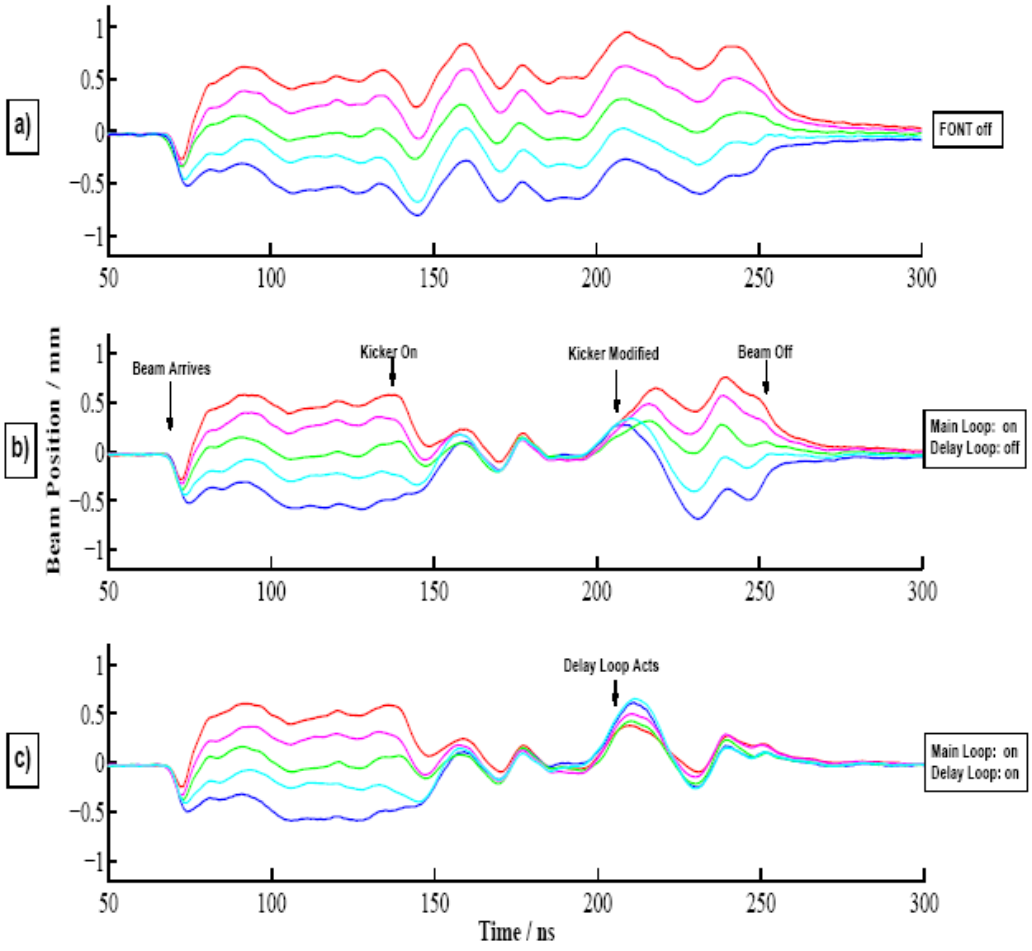

Figure 2.2: Feedback operation for FONT1 showing (a) without feedback, (b) feedback main loop on and (c) both the feedback main loop and delay loop on [49]. Each line is the average of ten pulses, showing the entire bunch train at five different position offsets.

		Time (ns)
Irreducible	Time of flight of bunches	6
	Signal return delay	10
Reducible	Processor	18
	Feedback circuit	4
	Amplifier	12
	Kicker fill-time	3
	Total	53

Table 2.3: Latency prediction for FONT2 [49].

latency comes from the lower "irreducible" part of the latency estimate.

FONT2 also included the introduction of independent witness BPMs downstream from the feedback BPM. The three BPMs were instrumented with the same design of processor and calibrated against known NLCTA beamline BPMs.

The processor designed for FONT2 shifted the 11.424 GHz button BPM response to 400 MHz by mixing it with a 11.024 GHz reference (known as "downmixing"). The sum signal was split and one output signal was passed through a limiting amplifier, producing a square wave reference at 400 MHz. Using this, the sum and difference signals were downmixed to frequencies near DC ("baseband"). Instead of normalising a pulse using the sum signal from the previous pulse, a new scheme was introduced using logarithmic amplifiers in the processor which avoided the long process of digitising and inverting the sum signal. Logarithmic amplifiers were used with the signal from the bottom electrode being subtracted from the signal from the top giving a normalised position output. The resolution for the BPM processor was calculated using two of the instrumented BPMs to predict the position in the third from a consideration of the geometry. The residual was defined as the difference between the predicted position in the third BPM and the actual position in the third BPM from the calibrated output signals. The standard deviation of the residuals was found. This contained information about the resolutions of the three BPM processors. With the assumption that the noise in each BPM was uncorrelated and the resolution for each BPM was the same, the resolution was calculated from the standard deviation of the residuals to be 10.9 μm [49].

The amplifier for FONT2 was improved on the amplifier used in FONT1 by being more compact and therefore more appropriate for use in the space-restricted ILC beamline. It had the same drive as the FONT1 amplifier. However, since the distance between the kicker and the feedback BPM was greatly reduced, more kick was needed to bring the beam to zero offset. For that reason, two kickers were used.

The static bunch position profile visible in figure 2.2 was removed using an AWG. This "beam flattener" generated a signal that was the negative of position variations around the mean position of the bunches in the train as calculated over many pulses. The signal was added to the signal from the BPM processor before the amplifier but after the delay loop.

For the feedback demonstration, the beam was given eight offsets using an upstream dipole magnet. This is shown in figure 2.3(a). The signal from the feedback BPM was split

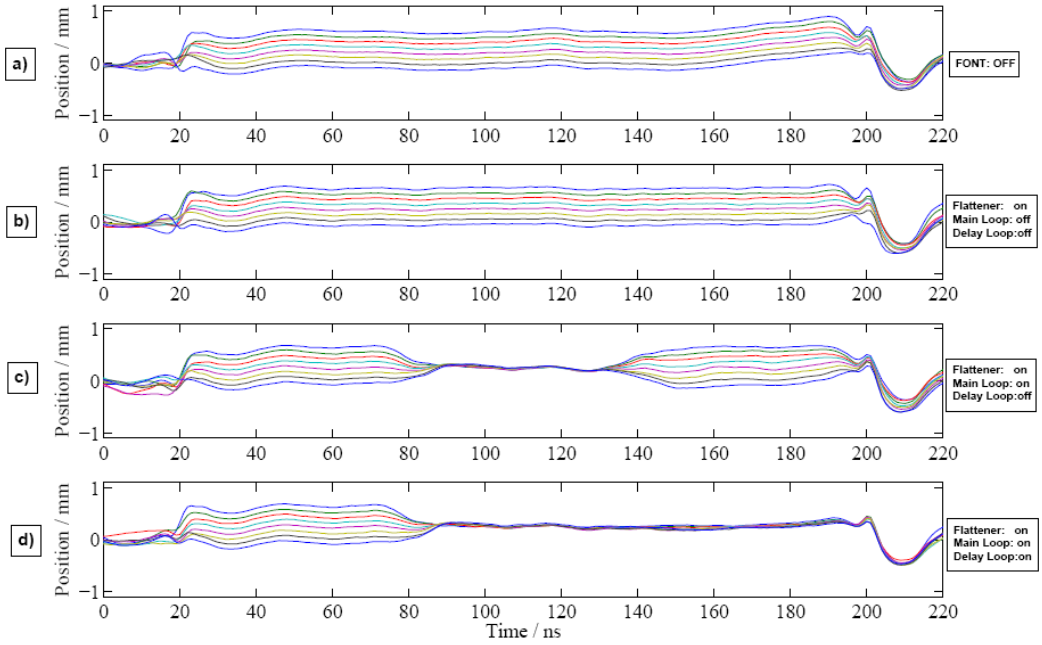

Figure 2.3: Feedback operation for FONT2 showing (a) without feedback, (b) no feedback, beam flattener on, (c) feedback main loop on plus beam flattener on and (c) the feedback main loop, the delay loop and the beam flattener all on [49]. Each line shows the bunch train at a different position offset.

to in-tunnel electronics for the feedback and out-of-tunnel electronics to be recorded on the FONT DAQ [51]. These position signals are from the out-of-tunnel electronics and show a positive "zero" offset compared to the in-tunnel electronics, the central magenta and cyan curves being around zero in the in-tunnel electronics.

Figure 2.3(b) shows the bunch train for the eight offsets and the beam flattener turned on. It successfully improved the beam profile, removing the upwards slope in position in the last 40 ns of each bunch train.

The feedback was turned on with the delay loop off. Figure 2.3(c) shows each offset beam being kicked around 80 ns to a common position, zero as measured in the feedback electronics. The latency of the feedback was found to be 57 ns, measured as the time between the onset of the kicker and the return to the original offset. After one latency period, the correction is lost as the feedback BPM sees a beam at zero position.

The delay loop was turned on showing that the correction to zero was retained after the second latency period, shown in figure 2.3(d).

This feedback test demonstrated that with beams initially spread in a range of positions spanning ∼0.7 mm, the feedback could kick them all to a common position with a spread of ∼0.05 mm. That is a 14:1 correction [49].

Parameter	ATF
Beam energy E (GeV)	1.28
ΔE/E	0.08%
Electrons per bunch	$(0.2 - 1.0) \times 10^{10}$ single bunch
	$(0.3 - 0.5) \times 10^{10}$ multi-bunch
Bunches per train	20
Bunch spacing (ns)	2.8
Train length (ns)	56
σ_x (μm)	70
σ_y (μm)	7
σ_z (mm)	8

Table 2.4: Beam parameters at ATF [52].

		Time (ns)
Irreducible	Time of flight of bunches	3
	Signal return delay	5
Reducible	Processor	5
	Amplifier	5
	Kicker fill-time	2
	Total	20

Table 2.5: Latency prediction for FONT3 [49].

2.3 FONT3

2.3.1 Moving to ATF

After two successful stages at NLCTA, FONT moved tests to the Accelerator Test Facility (ATF) at KEK.

As shown in table 2.4, the beam energy is 1.28 GeV at ATF. A feedback system that can correct at the micron level at ATF can also correct at the nanometre level at the ILC where the beams are up to 250 GeV (or 500 GeV with the 1 TeV centre of mass upgrade). Therefore, demonstrating micron-level feedback at ATF is a test of a system that can be taken to the ILC.

The beam at ATF is small and flat (shown in table 2.4) and has smaller beta functions than the beam at NLCTA. This results in smaller beam jitter and the chance at a much better correction.

However, beam tests at ATF put pressure on the feedback to be even faster as the entire bunch train at ATF was only 56 ns long. To get more than two latency periods and clearly demonstrate both the feedback and the delay loop behaviour, the new goal for latency was 20 ns, the budget for which is shown in table 2.5.

A similar setup was used at ATF as at NLCTA. The layout of the ATF extraction line

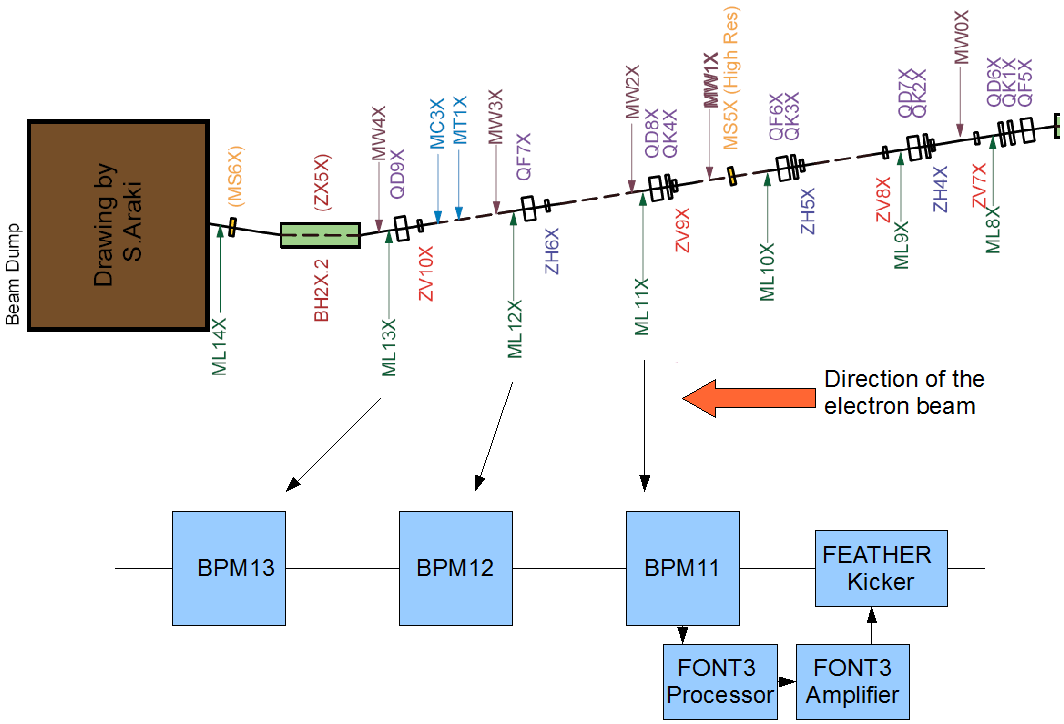

Figure 2.4: Top: ATF extraction line [53]. The prefix ML refers to stripline BPMs, QD are defocussing quadrupoles, QF are focussing quadrupoles, QK are skew quadrupoles, MW are wire scanners, ZH and ZV are correctors (for x and y respectively), MT and MC are current monitors and BH bend magnets. Bottom: Block diagram showing the components used for FONT3.

and the beamline components used for the FONT3 tests is shown in figure 2.4. There were upstream corrector dipole magnets to put offsets in the beam (ZV7X and ZV8X), a kicker (1 m upstream of ML11X) and a feedback BPM (ML11X). Two further BPMs (ML12X and ML13X) acted as witnesses downstream of the feedback loop. The BPMs used were stripline BPMs of length 12 cm which has its peak response at 625 MHz. The FONT3 processors were used with the y pickoffs on BPMs ML11X, ML12X and ML13X, called BPM11, BPM12 and BPM13 from this point forward.

2.3.2 FONT3 Processor

The processing scheme was returned to a simple difference calculation between the top and bottom BPM pickoffs. To meet the latency goal, there was no charge normalisation. With the goal of micron-level correction after two corrections, a 10% variation in charge could be tolerated. The charge profile for the bunch train at ATF was flat enough for this. The occasional bunch train with greater than 10% charge variation was removed from the data in analysis. There was a unique processor on the x pickoffs of BPM13 that outputted the sum signal for monitoring the charge and performing offline charge normalisation.

The FONT3 processor (see figure 2.5) began with two signals: the signal from the top strip and the signal from the bottom strip. At ATF, the striplines are shorted at the upstream

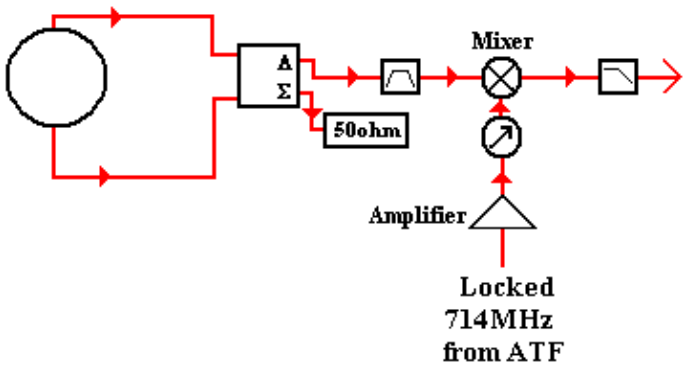

Figure 2.5: A block diagram of the FONT3 processor [49]. It takes signals from the top and bottom strips of a stripline BPM (left of diagram) and passes them through a difference hybrid (Δ). Then filters and mixers are used to produce the output to the amplifier.

end with the pickoff at the downstream end as described in section 1.5.2. At the pickoff, the BPM processor is attached. Care was used to make sure the processor was as close to the pickoffs as possible in the case of BPM11, the feedback BPM, and the cables used were the shortest possible to minimise the latency of the full feedback system.

One signal was subtracted from the other using a 180° hybrid (except for the summing processor which used a hybrid without the phase change). Exact subtraction of one signal from another when they are only picoseconds in duration requires the signal paths in the hybrids, and the cables between the striplines and the hybrids, to be well matched. However, the slightest difference between the signals in time produces a common mode signal when they are subtracted in the hybrid and in reality there is always some common mode component.

After the hybrid, the signal was downmixed to baseband to get an output that can be used to drive the kicker. It was first passed through a band-pass filter centred on 714 MHz and then passed through a mixer. The mixer required a reference signal (or local oscillator, LO) as an input. In this case, it was 714 MHz from the ATF central clock, locked to the beam. The mixer took the reference and the input from the early stages of the BPM processor (the RF signal) which after going through the bandpass filter had frequencies centred around 714 MHz. It multiplied the two signals resulting in a signal with low frequencies (under 200 MHz) and a high frequency signal around 2×714 MHz as in equation 2.1.

$$\cos(\omega_1 t)\cos(\omega_2 t) = \frac{1}{2}\left(\cos(\omega_1 + \omega_2) + \cos(\omega_1 - \omega_2)\right) \tag{2.1}$$

The LO and RF must be in phase with each other. If the phase difference between the two inputs is zero, the output is purely dependent on the beam position. If it is exactly 90° out of phase, the output is the beam position independent common mode. Most likely when the processor is first assembled, the phase difference would be some arbitrary amount. In this case, the output from the mixer depends both on position and on the phase difference. The correct phase needs to be set for the correct operation.

To adjust the relative phase between the two inputs into the mixer, the signal after the band pass filter was split two ways and sent to two mixers. The LO was also split two ways,

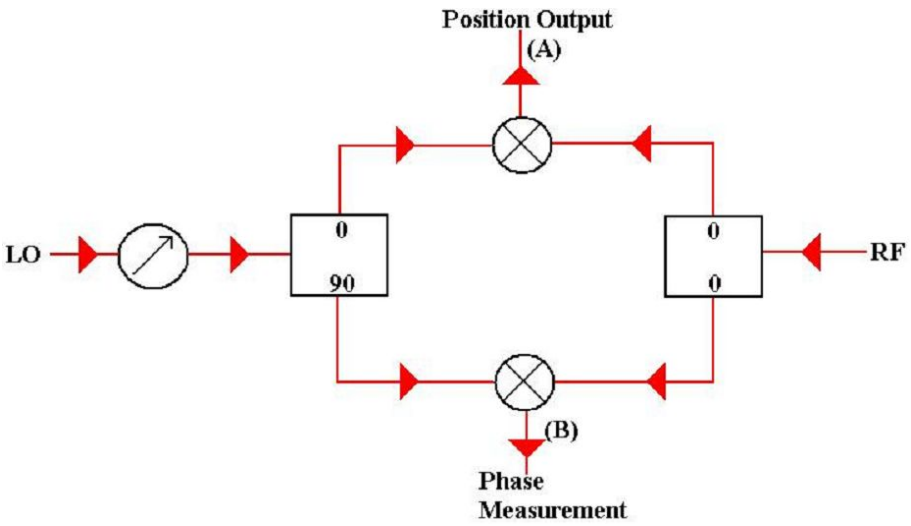

Figure 2.6: A block diagram showing the setup used to correctly phase the LO with the RF [49].

one LO path being longer than the other to introduce a 90° phase shift shown in figure 2.6. The phase of the LO was adjusted before the split and the output from the mixer with the 90°-shifted LO watched on an oscilloscope. The adjustment continued while the beam was moved by the upstream dipole magnet until the output became purely position independent. Since this common mode signal is 90° out of phase from the position signal, this method brought the LO that was not 90°-shifted to the correct phase to be mixed with the RF signal, thereby producing the purely beam position dependent output.

To remove the higher frequencies (around 2×714 MHz), and any 714 MHz leakage, from the baseband signal, the final component in the processor was a low-pass filter.

The LO phase was adjusted by variable phase shifters that could be controlled with a 0-15V power supply. They adjusted the phase of the LO from 0 to 180°.

Any jitter on the LO phase would be translated to the position signal as amplitude jitter.

This filtering scheme, using a band-pass filter before the mixer, was preferred over missing out the band-pass filter stage and just using a low pass filter to select the baseband signal as there was a strong 357 MHz signal at ATF with multi-bunch mode as this was the bunching frequency. Just a low-pass filter after the mixer would have worked if the filter had a sharp cutoff to remove the 357 MHz signal. However, a sharp cutoff compromises on speed so to keep latency low, the band pass filter was introduced to reduce the level of 357 MHz signal before the mixer.

Processor latency

The processor was modeled in SystemView [54], an RF simulation package. According to the simulation, the latency for this processor design was 4.1 ns, not including the physical size of the processor and the time for the signal to travel [49].

In practice, the latency was measured to be 4.6 ± 0.3 ns which is consistent with the predicted value assuming a signal path of ∼10 cm [49].

Processor calibration and resolution

The System View simulation of the processor predicted a resolution of 0.24 μm.

The three instrumented BPMs were calibrated using an upstream dipole magnet. Using the transfer matrix of the dipole magnet and the distance between the dipole and each BPM, the offset given to the beam by the dipole in the BPM could be calculated. The current in the dipole magnet was changed in steps, moving the beam through approximately 2 mm in each BPM.

Due to timing jitter of the oscilloscope trigger and scope noise when taking a single data point, the maximum point of the processor output was not used for calibration. Instead, normalised difference signal Y (equation 2.2) was calculated from the area under the difference signal ($V_{\text{Difference}}$) and the area under the sum signal (V_{Sum}).

$$Y = \frac{\int V_{\text{Difference}}\, dt}{\int V_{\text{Sum}}\, dt} \tag{2.2}$$

The resolution was calculated using the backslash (or "left matrix divide") operator in Matlab. This operator finds `x = A\b` using a least squares method to solve the equation `Ax = b`. Here `A` is a matrix that includes the positions of the beam in two of the three instrumented BPMs. The vector `x` uses the two BPMs to predict the third where `x` is calculated using the measured third BPM positions in the vector `b`. The resolution of the third BPM is then found by calculating the residuals, that is the difference between the predicted position and the actual position. The standard deviation of the residuals, is given as the resolution.

It was possible to run the beam at ATF in single-bunch or multi-bunch mode. The resolutions in single-bunch mode were found to be 24.5 μm, 12.2 μm and 34.3 μm (without much attention to the correct LO phasing) [49]. In multi-bunch, the presence of the 357 MHz bunching frequency put more power into 357 MHz and its harmonics (including 714 MHz). This boost in power at 714 MHz improved the resolution of the processors. The best resolution numbers achieved were 2.3 μm, 4.1 μm and 4.9 μm [49]. Shown in figure 2.7 and figure 2.8 are calibration and resolution data for one measurement of the FONT3 processor. The performance of the FONT3 processor was often not consistent between beam tests. The reliability issue was addressed in FONT4.

As shown in figure 2.7, the processors gave a linear response to the beam position over a range of ∼2 mm. The charge per bunch was approximately 0.5×10^9 electrons when this data was taken. The linearity depends on the input to the mixer [49], which cannot be above 0.25 V, and therefore the range scales with bunch charge.

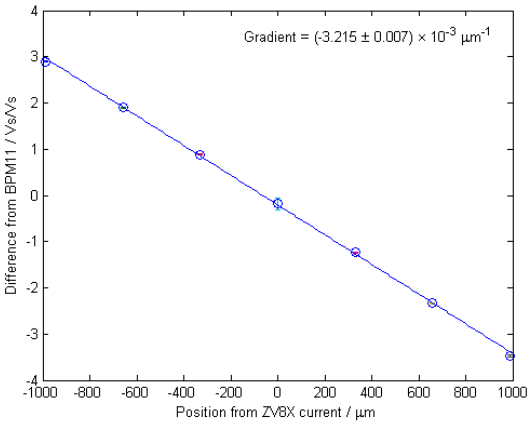

(a) BPM11 calibration

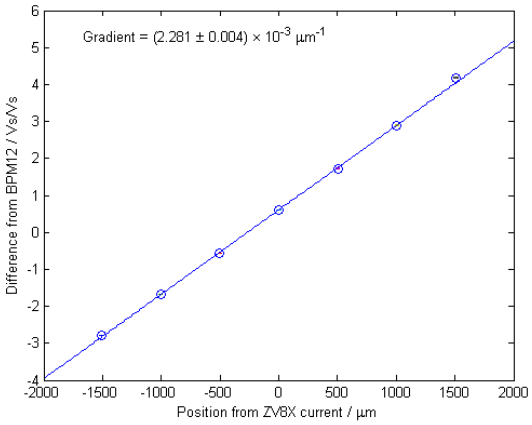

(b) BPM12 calibration

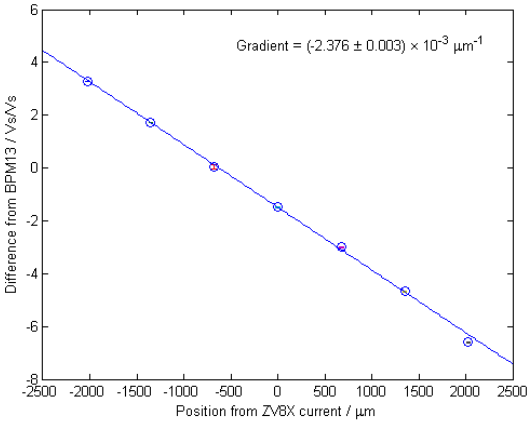

(c) BPM13 calibration

Figure 2.7: Calibration curves for three BPMs instrumented with the FONT3 analogue processor. Y (equation 2.2) versus position is shown.

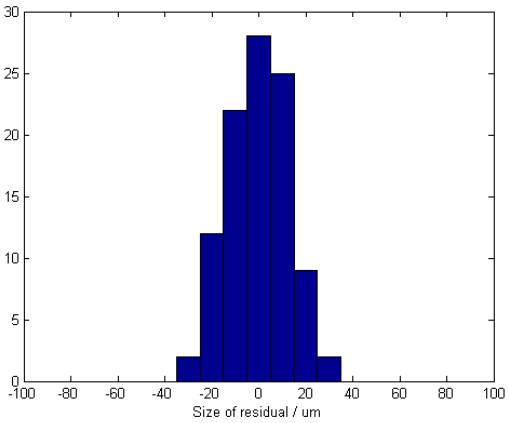

(a) BPM11 residuals showing a resolution of 12.8 μm.

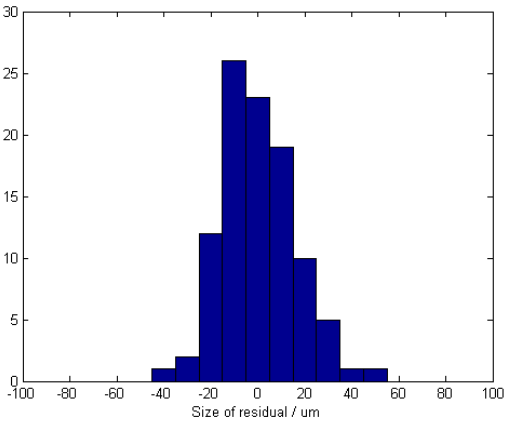

(b) BPM12 residuals showing a resolution of 15.8 μm.

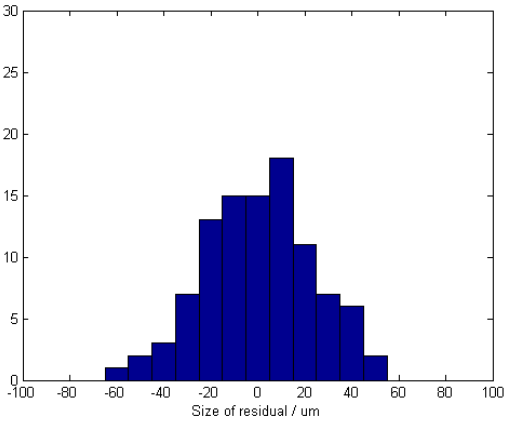

(c) BPM13 residuals showing a resolution of 23.4 μm.

Figure 2.8: The distribution of the residuals for three BPMs instrumented with the FONT3 analogue processor.

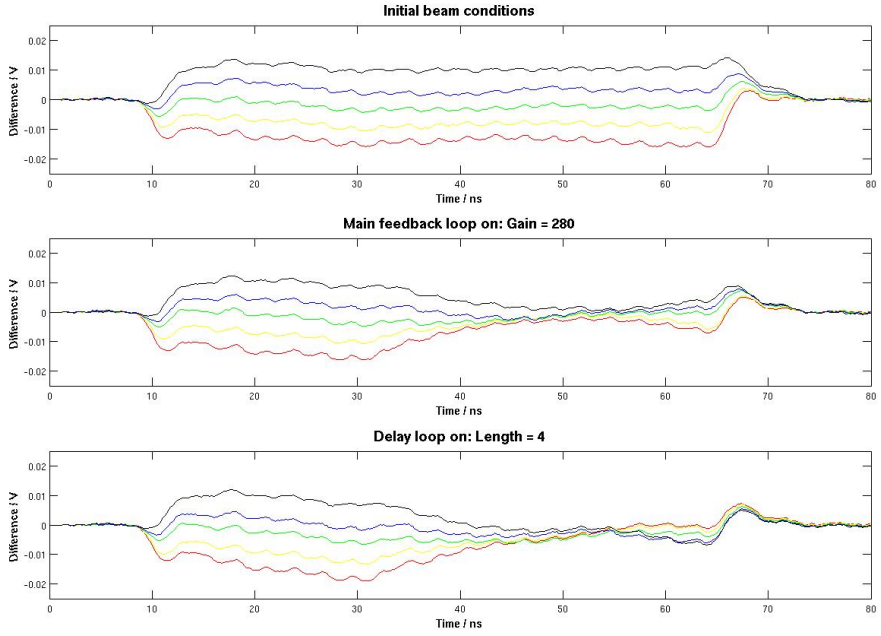

Figure 2.9: Feedback operation for FONT3 showing the signal from the difference processor on BPM11. (a) No feedback. (b) Feedback main loop on. (c) Feedback main loop plus the delay loop on [55]. Each lines shows the bunch train at a different initial position offset.

2.3.3 FONT3 Feedback Results

To perform the feedback, this processor was attached to all three stripline BPMs. The processor output from the feedback BPM was split two ways: half the signal went to the amplifier and half came out along long heliax cables to be recorded on the FONT DAQ. The lengths of cables used between the feedback BPM pickoffs and the processor, the processor and the amplifier and the amplifier and the kicker were kept to a minimum to reduce latency. There was no such latency concern with the witness BPMs so the signals from the pickoffs were carried along long heliax cables to the processors and the FONT DAQ.

A 56 ns bunch train was used, comprising 20 bunches 2.8 ns apart. The bunch train produced a signal in the striplines and was processed in the electronics. With a bunch separation of 2.8 ns, the single bunches in the train could be resolved as shown in figure 2.9(a). As with FONT1 and FONT2, an upstream dipole magnet was used to give the bunchtrain a position offset.

The feedback was turned on and the output from the processor was amplified and drove the kicker just 1.2 m upstream of the BPM. The gain on the amplifier was adjustable and chosen to bring the bunchtrain to zero in the feedback BPM. Figure 2.9(b) shows the best guess at setting the main loop gain (value 0.3106 [56]) during the tests (however, subsequent analysis of the data indicated that the gain setting was slightly too low [49]). The offsets within ± 100 μm (in blue, green and yellow) converge at zero around 48 ns in this figure. The largest offsets (in red and black) do not meet for reasons discussed in the next section. As expected with no delay loop, the correction is lost and the bunches begin to return to

Item	Value
Main Gain	0.3650
Latency / Delay length (ns)	24.0 ± 0.5
Delay Gain	0.6
Correction ratio	23:1

Table 2.6: The results from FONT3 feedback, summarised [49].

their original offsets around 55 ns in this figure.

The delay loop was of variable length and variable gain. The settings were chosen so that after the second latency period, the beam stayed corrected for the rest of the bunch train. The feedback results shown in figure 2.9(c) show the best guess at the delay loop settings during the beam tests (length 21.7 ns and gain 0.6). Subsequent analysis suggested that the delay loop length was set too low as the latency of the feedback was longer than this setting. Unlike in figure 2.9(b), the bunches do not return to their original offsets at 55 ns; instead they are given a kick that is too large and the bunches are over-corrected.

The latency of the system was measured to be 24.0 ± 0.5 ns from an analysis of the signals from when the main loop was on [49]. This was the average of two latency measurements: one where the latency was defined as the time between the signal from the very start of the bunchtrain, when it first goes above the noise level, and the point where the position falls below 90% of its maximum, and the second where the latency was defined as the time between when the signal climbed to half its maximum and the time where the signal fell to half its maximum.

Bunch trains $\sim$130 μm offset from each other were corrected to within $\sim$6 μm, a 23:1 correction.

Data were recorded with other settings and subsequent analysis indicated that the ideal settings were as shown in table 2.6.

2.3.4 Simulations of FONT3 feedback

The FONT3 feedback results shown in figure 2.9 required some investigation to be understood properly as the feedback behaviour was not ideal: figure 2.9(b) shows that the larger offsets (in red and black) do not meet and figure 2.9(c) shows that there is overcorrection in the third latency period (beginning $\sim$55 ns in the figure). These were hoped to be better understood through simulating the feedback.

Simulink [38] (a tool for modeling dynamic systems) was used to model the feedback as shown in figure 2.10. Experimental data of the beam's position in y were used as the input. The recorded difference signals from the feedback BPM, 1.2 m downstream from the kicker, were loaded at the "Beam Generator" and, using a transfer matrix, the beam was tracked back to what it would have looked like at the kicker. The kicker imparts an angle change on the beam and over the course of the 1.2 m drift space between the kicker and the BPM, the

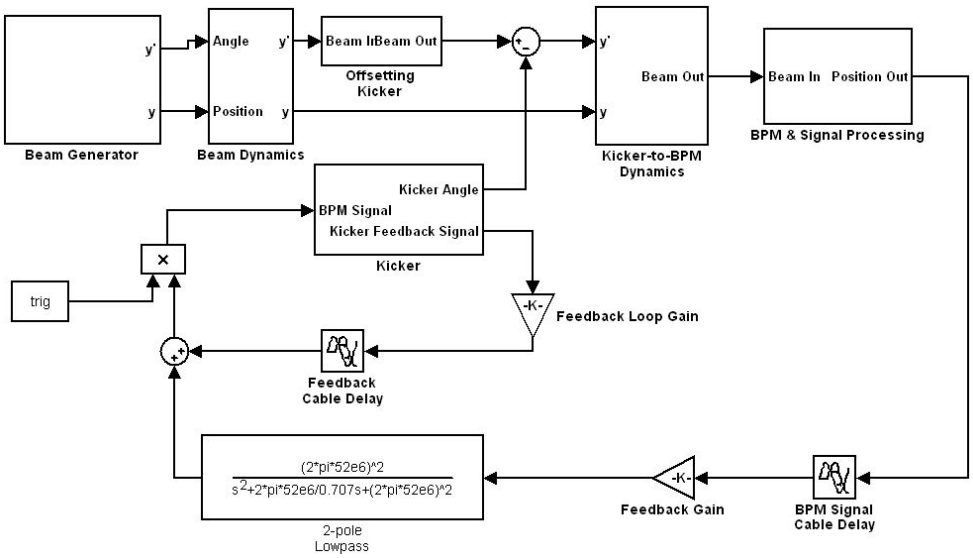

Figure 2.10: The FONT3 feedback as simulated in Simulink [57].

y position changes. As the input was already what is seen as the BPM processor output, it was not necessary to simulate the BPM processor. Instead a delay was introduced equal to that of the latency of the processor. A further delay was needed for the cabling between the processor and the amplifier. The amplifier had a bandwidth of 52 MHz and was modeled by a 2-pole Bessel function [56]. The kicker was modeled simply as an integrator with suitable delay to give a rise-time of 8 ns (which appeared to match the rise-time observed in the run at ATF) with its output affecting the beam angle. As the kicker was modeled this way rather than a full modeling of the dipole magnet and its effect on the charged beam, the amplification setting used in the feedback experiment and the gain set in this simulation bear no resemblance to each other and the gain had to be set by repeatedly running the simulation for different gains and deciding which produced the output that best matched the real feedback data. A feedback loop that takes the amplifier output through a delay equal to the latency of the feedback system, allows the correction be kept for the next latency period. The various delays used in the simulation were chosen to agree with the delays in the feedback tests at ATF.

Settings were found that gave the best feedback as shown in figure 2.11 and then parameters changed to investigate how and why the data differed. The input beams came from runs done at approximate offsets of ± 70 μm, ± 35 μm and 0 μm at the kicker. The simulation took the inputs and produced results as expected: the offset beams were corrected to zero offset at the BPM (see figure 2.11). However there was also a sharp positive signal at the end of the 56 ns bunch train. This signal appears to be from the BPM processor as it is also seen in the input to the simulation. It also appears in the data from ATF (figure 2.9) but it is more enhanced in the simulation possibly due to the amplifier bandwidth. Being after the 56 ns bunch train, it has no affect on the beam and is not a concern.

It was speculated that the amplifier may have a non-linear response for the range of initial offsets that were used. The non-linearity of the amplifier was introduced, modeled using the function shown in figure 2.12.

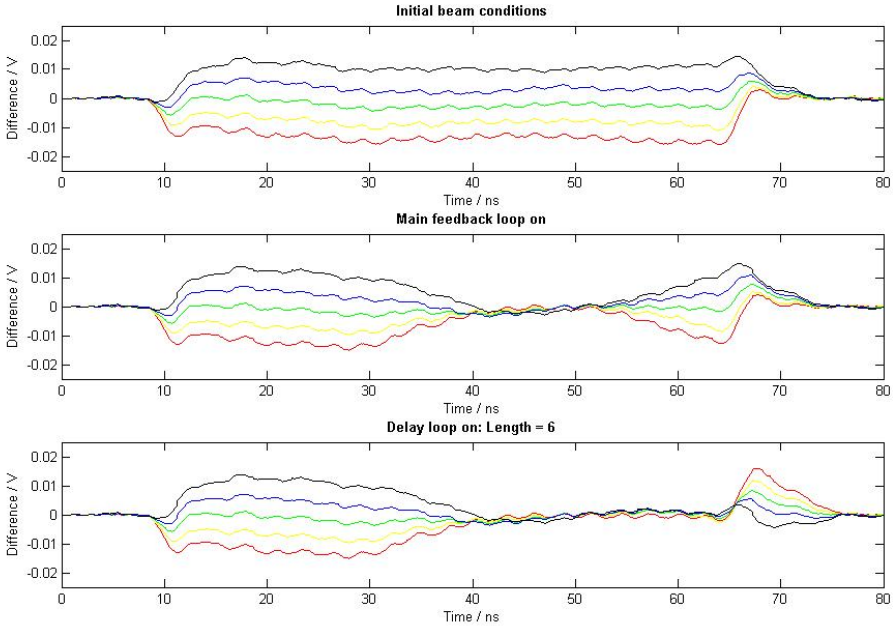

Figure 2.11: Simulation results with the signal from the difference processor versus time. Top: The average input bunch trains to the simulation. Middle: The feedback simulation with main loop on, delay loop off. Bottom: The feedback simulation with the main loop on and the delay loop on. The gain and delay loop length used here were set to produce the best feedback result possible. In each case, the five lines are the averages of 120 bunch trains.

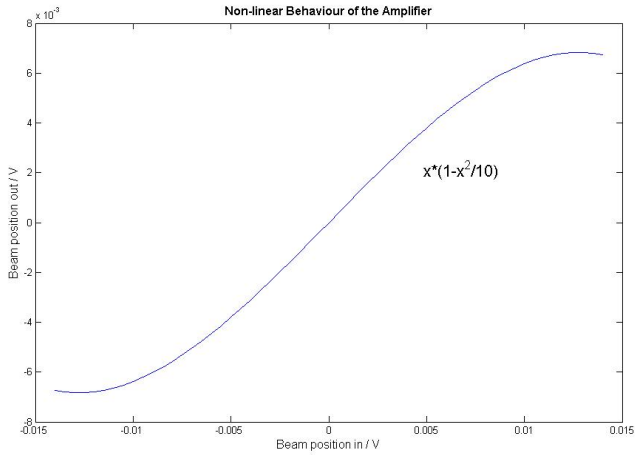

Figure 2.12: A function derived from empirical observation describing the non-linearity of the amplifier [56].

For a direct comparison to high-gain data from ATF, the simulation was run with a higher than ideal main loop gain (equivalent to a gain of 0.4655 in the FONT3 system) and the appropriate delay loop length to match the latency (24 ns). The non-linear response for the amplifier was included in this simulation. Figure 2.13(a) shows the initial beam conditions, data from ATF that was used as input to the simulation. Figure 2.13(b) shows the simulation results with a linear amplifier response. Figure 2.13(c) shows the simulation results with the non-linear amplifier response as shown in figure 2.12. Figure 2.13(d) shows the high-gain data taken during the FONT3 tests at ATF. Introducing the non-linear amplifier created a closer match to the data though it appears that the compression effect was slightly over-estimated.

Figure 2.9(c) showed overcorrection of the beams when the delay loop acted. Simulation results using a delay 2 ns shorter than the system latency (figure 2.14) confirm that there is a greater correction after the second latency period (around 55 ns) than after the first. This is because the kicker sees both the delayed signal and the offset beam signal from the BPM processor, resulting in a signal that the sum of both. The compression effect again appears to be slightly over-estimated.

The simulations demonstrated that overcorrection can occur with the delay loop set 2 ns too short and that the amplifier was non-linear so that the bunchtrains with the greatest offset were not corrected ideally.

2.4 FONT 4

With the decision to have a bunch spacing of hundreds of nanoseconds at the ILC and a train comprising thousands of pulses came the next step of making the FONT feedback a more powerful tool capable of learning and applying more complex algorithms. The extra time between bunches[2] allows the output of the super-fast BPM processor to be digitally sampled and processed by a field-programmable gate array (FPGA), applying algorithms that can predict bunch positions from an analysis of the full train.

For example, see figure 2.15. Here adjustable gains are used rather than one single main loop gain. Simulations for the TESLA design described in section 1.4.2 were used to demonstrate how using multiple gains in the feedback could lead to a lower loss of luminosity than using just one single gain.

The "train flattener" described in section 2.2 (FONT2) also demonstrates the kind of algorithm that can be applied with a digital processor. Any static bunch train profiles can be removed.

The digital processor can also have input signals from other monitors such as the luminosity monitor, using extra data to achieve the best correction possible.

This generation of FONT, called FONT4, used a new PCB version of the FONT3 analogue BPM processing scheme, a digital processor and a new amplifier. It was tested at ATF using the same setup as FONT3 with the digital processor used on the feedback BPM only.

[2]The bunch spacing for the ILC is proposed to be between 153.8 and 461.5 ns, see table 1.2.

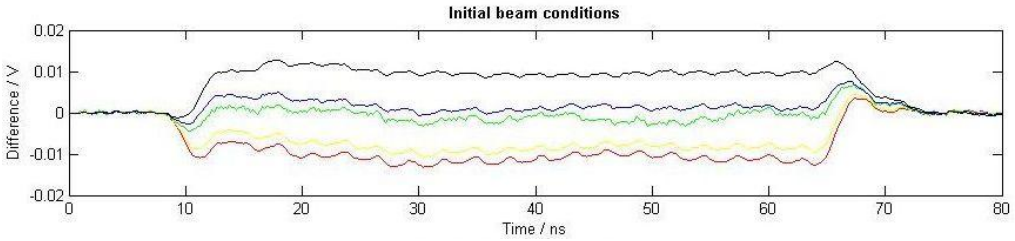

(a) Five initial beam positions at ATF (used as the input to the simulation).

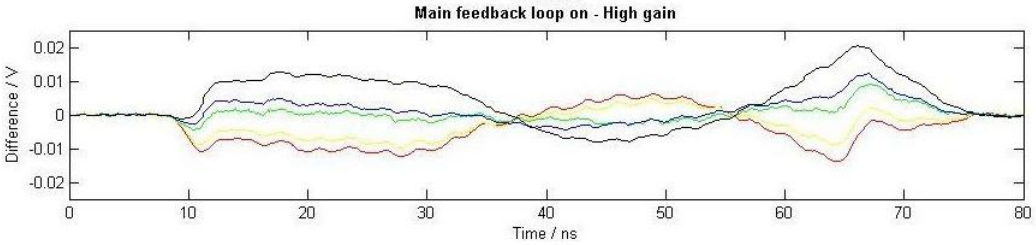

(b) The simulated feedback with the main gain set high and the delay loop off. The amplifier is linear in this simulation.

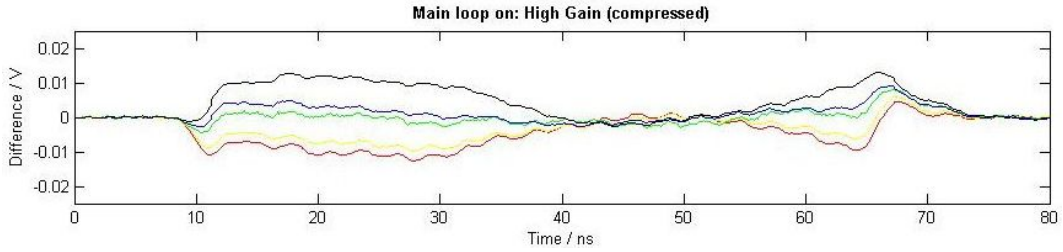

(c) The simulated feedback with a high gain setting, non-linear amplifier and no delay loop.

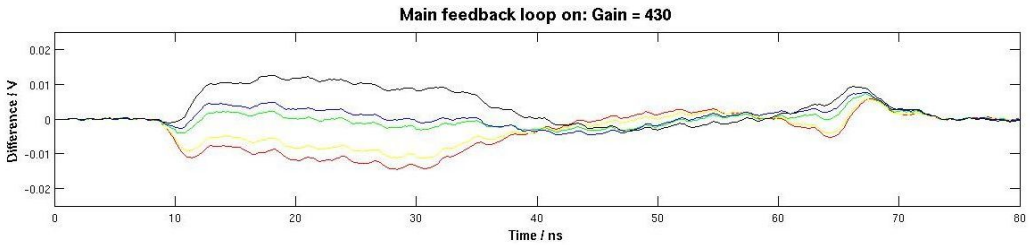

(d) Data taken at ATF with a high gain and no delay loop [55].

Figure 2.13: A comparison of simulated results and real results from ATF.

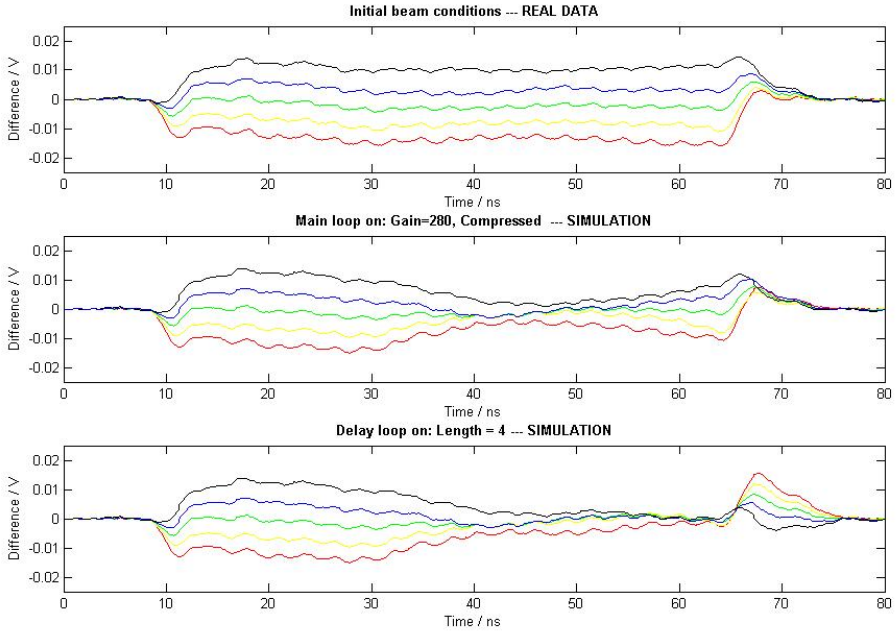

Figure 2.14: Simulation results with the signal from the difference processor versus time. Top: The average input bunch trains to the simulation. Middle: The feedback simulation with main loop on, delay loop off. Bottom: The feedback simulation with the main loop on and the delay loop on. The gain was the ideal value as in figure 2.11. The non-linear model of the amplifier was used. The delay loop length was 22 ns (2 ns short of the system latency). These are settings to model the situation in figure 2.9.

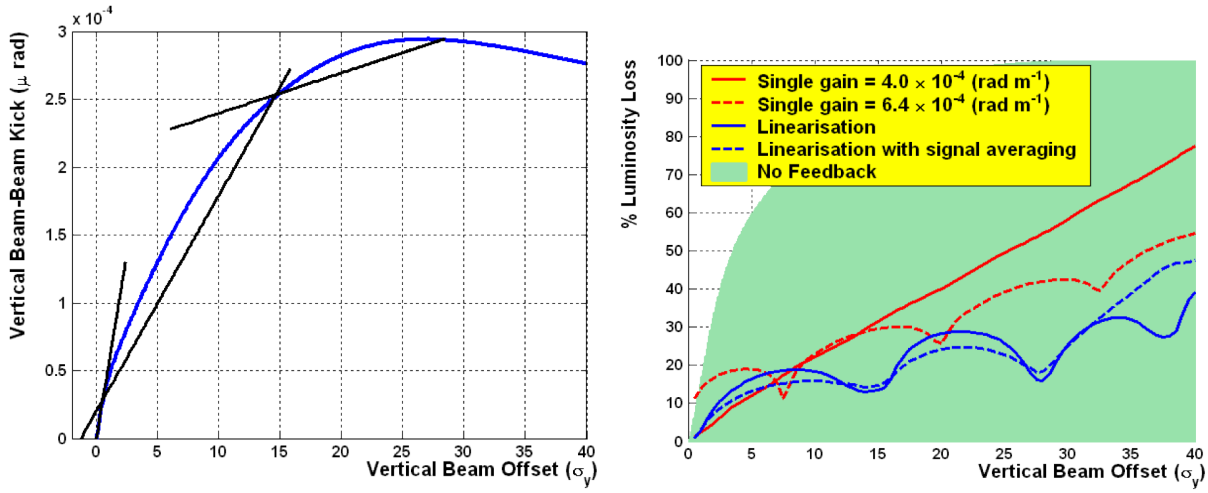

(a) The blue curve shows the kick from the beam-beam interaction at different vertical beam offsets. The black lines are examples of how this curve can be coarsely divided into linear regions.

(b) The luminosity loss in TESLA simulations of beams colliding with an offset. Without feedback, luminosity loss is shown by the green area. With feedback with a single gain setting, the luminosity loss is the area under the red lines (with the dashed red line showing a higher gain setting than the solid red line). The blue lines show feedback with multiple gains that roughly linearise the beam-beam kick relationship with offset.

Figure 2.15: Simulation results of having multiple gains depending on the offset compared to one single gain [58].

Latency is not as tight a constraint as for warm machines but it is desirable to make as many corrections as possible to ensure the best correction. Therefore the goal of FONT is to operate on a bunch to bunch basis. The smallest bunch spacing at the ILC is 153.8 ns therefore the latency budget for the FONT4 system (including time of flight and signal times) needs to be at most 150 ns. The application of sophisticated algorithms would increase the latency further.

ATF was able to provide three bunches with up to 154 ns bunch spacing. It was decided to set the FONT4 feedback latency to 140 ns as shown in table 2.7 in case it was necessary to use the smaller bunch spacing at ATF for reasons of beam quality in the ATF extraction line.

2.4.1 Analogue Processor

The FONT3 analogue processor was modified to become the front-end of the FONT4 processor. The ADC required suitably broad sum and difference signals to make sampling at the peak reliable and accurate. The width of the signals from the FONT3 processors was approximately 3.2 ns FWHM which was too narrow for sampling accurately. The bandpass and lowpass filters in the FONT3 processor were modified to give the processor a smaller bandwidth and broaden the output signal aiming for ∼10 ns. In FONT3, the lowpass filter was 170 MHz and the bandpass filter 500-900 MHz. For the simulation of a possible FONT4

		Time (ns)
Irreducible	Time of flight of bunches	4
	Signal return delay	10
Reducible	Analogue Processor	7
	Digital Processor	68
	Amplifier rise time	40
	Kicker fill time	3
	Total	132

Table 2.7: Latency prediction for FONT4 [59].

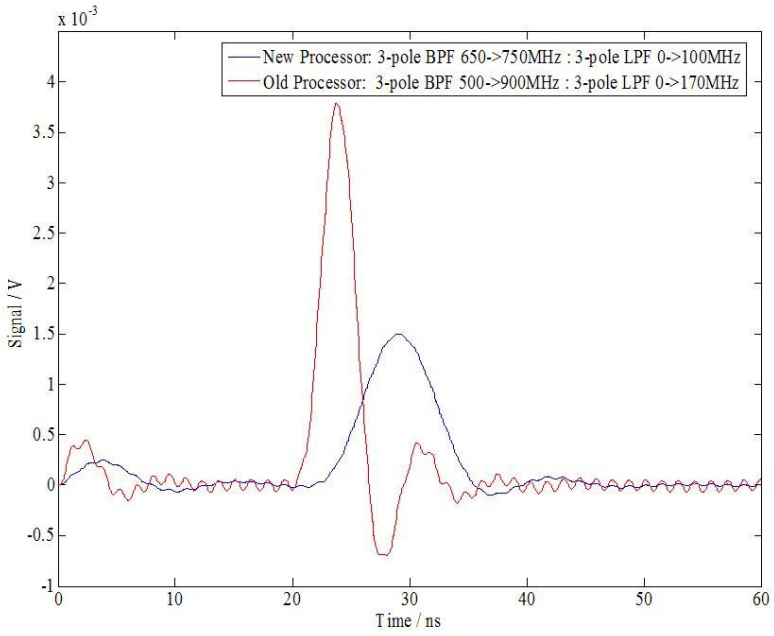

Figure 2.16: A System View simulation of the FONT3 processor output (red) and the FONT4 processor output with the revised filter values (blue) [60].

processor, the filter cutoff frequencies were assumed to be 100 MHz for the lowpass filter and 650-750 MHz for the bandpass filter. Figure 2.16 shows the simulated pulse from the FONT3 processor compared with the simulated pulse for this FONT4 analogue processor with the revised filtering scheme. The broadened pulse comes at the expense of amplitude and also latency. The actual filter scheme that was decided on for FONT4 analogue processor had a 90 MHz lowpass filter and a bandpass filter with ±130 MHz about 714 MHz.

The modified FONT3 design (figure 2.17) was mounted on a PCB and held in a rigid case. This was both more convenient (as it resulted in less set-up than the connectorised FONT3 processor) and reliable (the repeated disconnection and connection of components in the FONT3 processors created slightly different behaviour each time the processors were used). The FONT4 analogue processors also gave both sum and difference as outputs for the same pair of BPM pickoffs as shown in the schematic in figure 2.18.

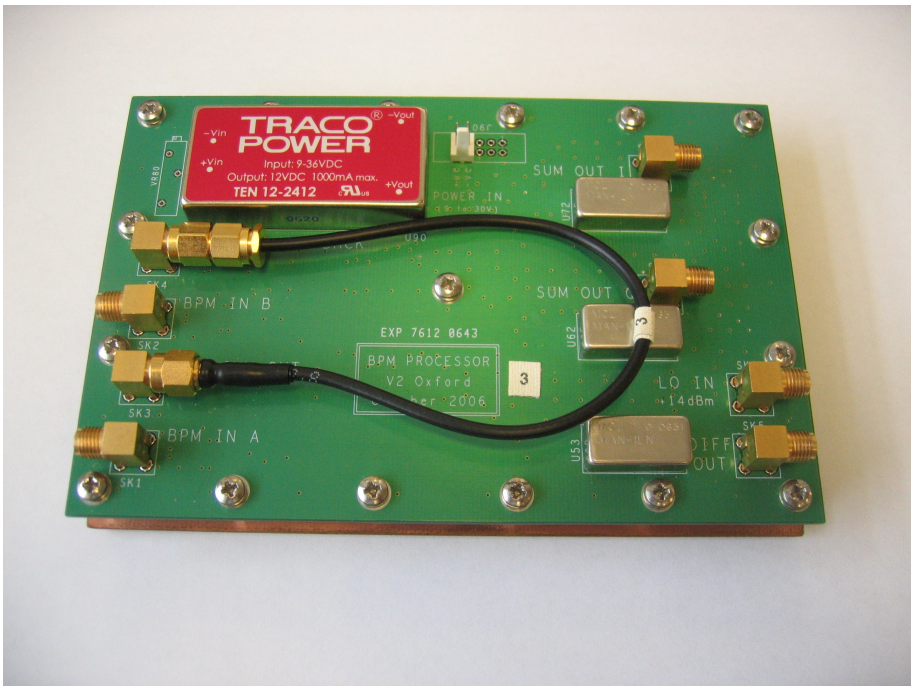

Figure 2.17: The FONT4 analogue processor [61].

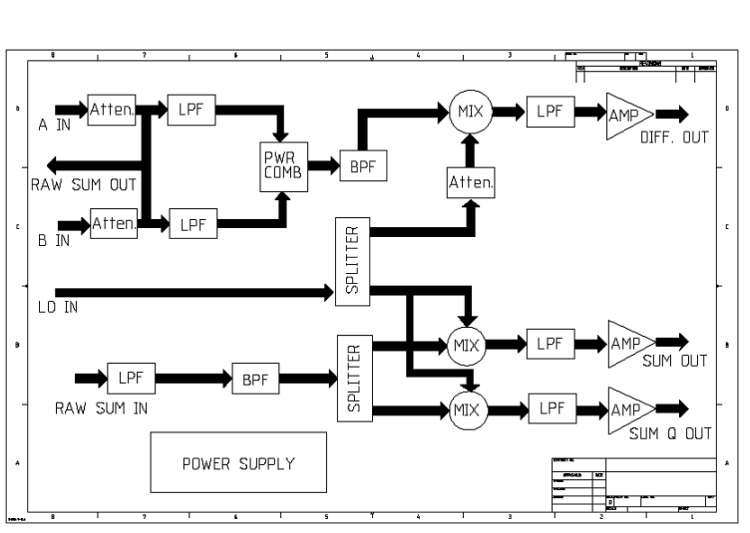

Figure 2.18: Block diagram of the FONT4 analogue processor components [61]. This is essentially the same design as the FONT3 processor (see figure 2.5) except with both the difference and the sum as output and different values for the filters.

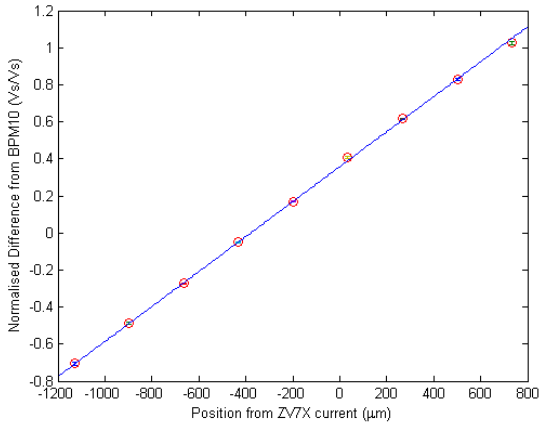

(a) BPM10 calibration

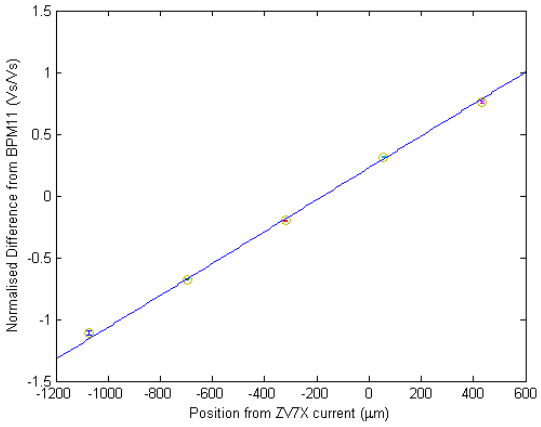

(b) BPM11 calibration

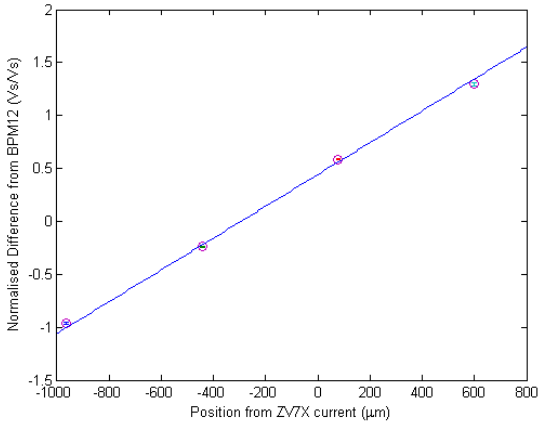

(c) BPM12 calibration

Figure 2.19: Y (equation 2.2) versus position. Calibration curves are shown for three BPMs instrumented with the FONT4 analogue processor.

The FONT4 analogue processors were calibrated using an upstream dipole magnet of known kick, exactly the same as for the FONT3 processors. The calibrations are shown in figure 2.17. As with the FONT3 processors, the resolutions were found using two BPMs to predict the third and looking at the distribution of the residuals (shown in figure 2.20). The resolutions were found to be 13.0, 7.9 and 9.9 μm.

This shows an improvement on the resolutions for the FONT3 processors (for single bunch operation), despite the smaller bandwidth of the FONT4 processor, due to the improvements in setting up the processors and phasing the LO. However, jitter on the LO of approximately 4° existed [62] and this made the resolution worse. Further improvements can still be made to reduce the common mode signal and increase the accuracy of the subtraction in the hybrid. There is ongoing work to improve the resolution of the FONT4 analogue processor.

The latency of the FONT4 analogue processor was measured using test bench equipment to generate a short pulse like half a stripline bipolar shown in figure 2.21. The latency was measured to be 10.4 ns from the peak of the input to the peak of the output. That is ∼4 ns longer than the FONT3 processor which was measured at ATF to be 6.2 ns from the peak of

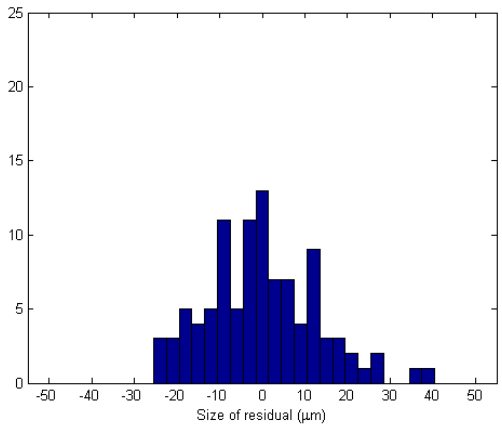

(a) BPM10 residuals

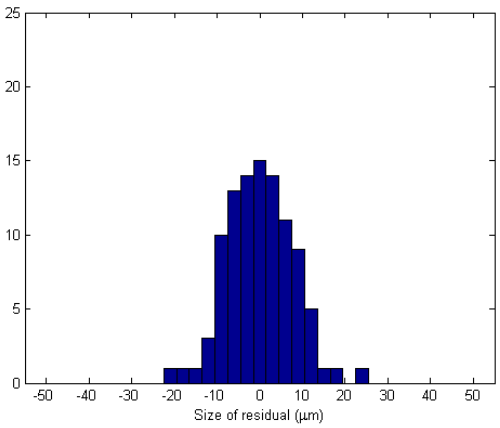

(b) BPM11 residuals

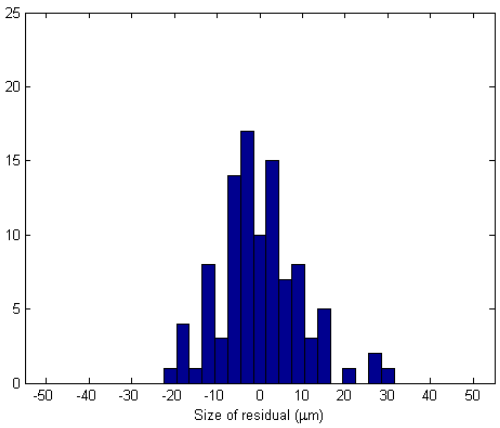

(c) BPM12 residuals

Figure 2.20: The distribution of the residuals for three BPMs instrumented with the FONT4 analogue processor.

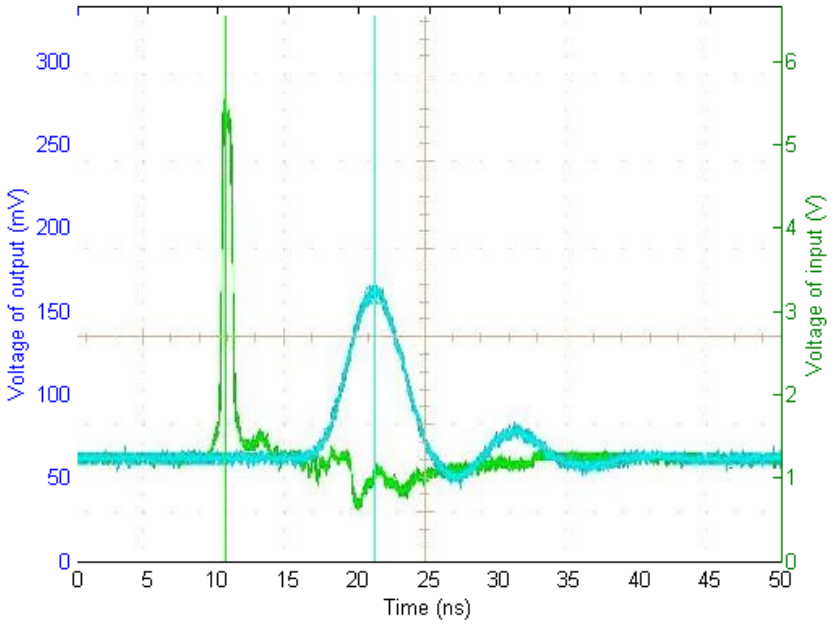

Figure 2.21: Measurement of the FONT4 analogue processor latency showing the input signal (green) and the processor output (blue) [63].

the positive-going part of the input signal to the peak of the processor output. However, the latency measurement for FONT4 included an extra low pass filter on the PCB, introduced to filter out any oscillations from the on-board amplifier which was never actually used in the tests. The low pass filter will be removed in the future FONT4 analogue processors.

The width of the FONT4 analogue processor output was 4.4 ns FWHM and approximately 10 ns full-width.

2.4.2 Digital Processor

The output from the analogue BPM processor was digitised by a 14 bit ADC capable of 105 MSPS (mega samples per second) and then processed by a Xilinx Virtex4 FPGA [64]. As the analogue processor signals were ∼10 ns wide at the maximum, this meant that only one sample per bunch was possible. It was important to make this sample at the peak of the bunch signal for the greatest sensitivity. This required a complex arrangement involving Block RAM embedded on the FPGA [65].

Sampling at the peaks of the bunches

Available at ATF was a 357 MHz (the ATF bucket spacing frequency) signal locked to the beam. This was used to clock the Block RAM (BRAM) which outputted enable bitstreams[3] for the ADC. The bitstreams were preloaded patterns for 16 different bunch spacings (from

[3]These sequences of 1s and 0s would instruct the ADC to sample or not sample.

120.4 ns to 162.4 ns in 2.8 ns steps) that mixed 357/4 MHz and 357/5 MHz clocks. By mixing the two clocks, sampling at the exact spacing between the bunches was possible whereas with only one clock the bunch spacing would need to be a multiple of the clock's time period. A bunch spacing of 154 ns, for example, with an enable bitstream based on the 357/4 MHz clock, would not have the peaks of the bunches sampled because there is a non-integer number of samples between the peaks. Based on two clocks, it is possible to have an integer number of samples in the 154 ns and therefore the sampling is in phase with the bunches. The slightly faster sampling (357/4 MHz) was used at the bunches whereas 357/5 MHz was used between the bunches.

To time the BRAM output such that the enable bitstream coincides with the bunch signals, a pre-beam trigger selected the ring-clock cycle on which the bunches were extracted. Finding the peak of the signal from the first bunch involved incrementing the address of the BRAM in 2.8 ns steps (from the 357 MHz clock) with respect to the 2.16 MHz ring frequency. Further fine tuning in case the peak of the bunch signal fell between the 2.8 ns steps was achieved using silicon delay lines on the chip.

FPGA Processing

Two methods of signal processing were developed: without charge normalisation and with charge normalisation.

Without charge normalisation was the simplest method. The difference signal from the analogue processor was sampled as above and then multiplied by a single gain number.

Charge normalisation used lookup tables to perform the normalisation and multiplication by a gain. The sum signal from the digitised analogue processor signal was inverted and multiplied by a gain via a lookup table and then this was multiplied by the digitised difference signal from the analogue processor.

The final signal after the logic was then passed to a DAC (which used the enables from the block RAM). The analogue output signal was input to the amplifier.

The delay loop is incorporated into the FPGA logic and retains the correction from one bunch to the next.

Latency

The ADC, DAC and FPGA processing were measured for latency as a function of the clocking frequency. The results, as shown in figure 2.22, suggest a latency of 70 ns at 89.25 MHz (357/4 MHz). This is the minimum latency for the digital board as the latency test does not involve processing and uses only 357/4 MHz as the system clock. It is estimated that 25 ns of the latency is from the FPGA processing, 3 ns from input and output to the FPGA and the rest from the ADC and DAC [66].

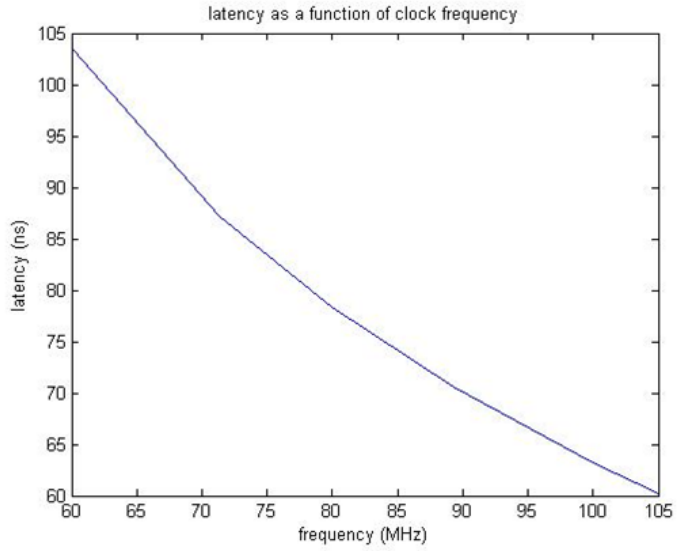

Figure 2.22: The latency of the FONT digital processor versus the frequency of the ADC and DAC [65].

2.4.3 Amplifier

The outline specification for a fast amplifier for FONT was done in Oxford and two units were manufactured by TMD Technologies [67]. It provides ±30 A to the adjustable gap kicker.

The amplifier rise time was specified to be 35 ns to 90% peak output. The specified pulse length was 10 μs to allow for use with the proposed new extraction scheme at ATF [68] which will have up to 60 bunches.

2.4.4 Method

A mode of producing three bunches at large bunch spacings was used (not to be confused with the multibunch trains with the small 2.8 ns spacing used in FONT3). Three RF buckets were filled in the damping ring with spacings a multiple of 2.8 ns. Typically, two spacings were used: 140 ns (2.8×50) and 154 ns (2.8×55). All three trains were extracted on one (∼300 ns long) kicker pulse. The charge was between 0.3 and 0.6 $\times 10^{10}$ per bunch (a little lower than the bunch charge of 1-2 $\times 10^{10}$ at the ILC).

The same setup was used as for FONT3 (described in section 2.3). The first bunch was measured in BPM11 and the FONT correction was applied to bunches 2 and 3 with the upstream kicker. Analogue processors were used on BPMs 12 and 13 to act as witness to the correction. Five different initial position offsets were given to the beams using an upstream corrector (ZV7X). The experiment was performed with and without charge normalisation. Data were recorded using the on-chip ILA (Integrated Logic Analyser) core for the feedback BPM (BPM11) and the oscilloscopes as with FONT3 for the witness BPMs.

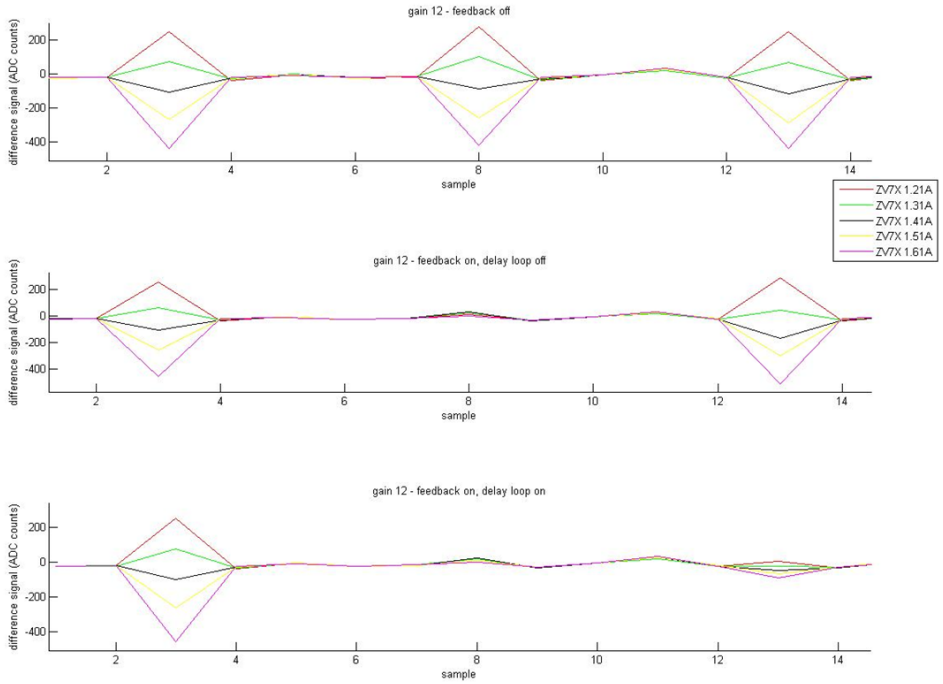

Figure 2.23: ADC counts versus sampling number. Top: the three bunches at their five initial offset positions without feedback. Middle: feedback on and delay loop off. Bottom: the feedback and delay loop on [69]. Shown are averages of 11 bunch trains.

2.4.5 Results

Figure 2.23 shows the FONT4 feedback work without charge normalisation. The five different position offsets given to the bunches using corrector ZV7X are shown as lines of different colours. As there was only one sampling point per bunch, there is not an indication of the bunch shape. Without the delay loop on, the correction is not kept for the third bunch. This shows a 7:1 correction in bunch 2.

Figure 2.24 shows the FONT4 feedback work with charge normalisation. Here, only three different position offsets were used.

Feedback with charge normalisation shows greater success in bringing the bunches to zero y offset on average. Shot by shot however, it appeared worse due to the added random error of the resolution of the charge measurement.

In both methods of signal processing, bunch 3 was not corrected as well as bunch 2. The gain was optimised for bunch 2 and the bunches did not have the same initial position offset due to variation in the amplitude of the extraction kicker pulse. This meant that bunch 3 was under-corrected if it had a larger initial offset position or over-corrected if it had a lower initial offset position. This type of behaviour where different bunches require slightly different gains can be addressed in the FPGA logic in future tests.

The latency of the FONT4 system without charge normalisation was measured by adding

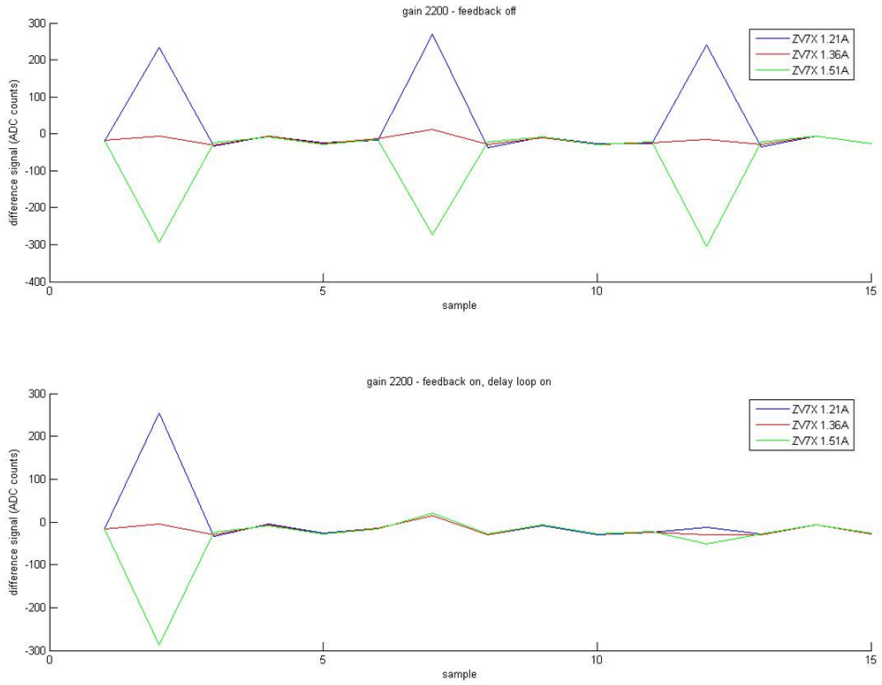

Figure 2.24: ADC counts versus sampling number. Top: the three bunches at their three initial offset positions without feedback. Bottom: the feedback and delay loop on [69]. Shown are averages of 11 bunch trains.

delay to the DAC output. The bunch spacing was set to 154 ns (the maximum available) and the feedback turned on. As long as the latency of the FONT4 system plus the additional delay was less than the bunch spacing, the second bunch was still kicked. The delay was increased until the second bunch was no longer kicked. The position of the second bunch is shown in figure 2.25 versus the additional delay. With no additional delay, the bunch is kicked to a large offset. With 40 ns of additional delay, the bunch was not kicked at all. Between 0 and 40 ns, the kick experienced by bunch 2 gradually falls as the amplifier has not reached its peak output. 90% of the full kick corresponds to approximately 15 ns of additional delay. With a bunch spacing of 154 ns, the 15 ns of slack indicates a latency for the FONT4 system of ∼139 ns.

Charge normalisation is expected to add 8.4 ns to the latency [66]. This brings the basic FONT4 feedback system latency to approximately 148 ns. This is below the smallest bunch spacing in the ILC accelerator parameters.

2.5 Summary

The FONT systems have been developed from low-latency analogue processors to digital processors, demonstrating their ability to correct beam offsets on test beams at NLCTA and ATF.

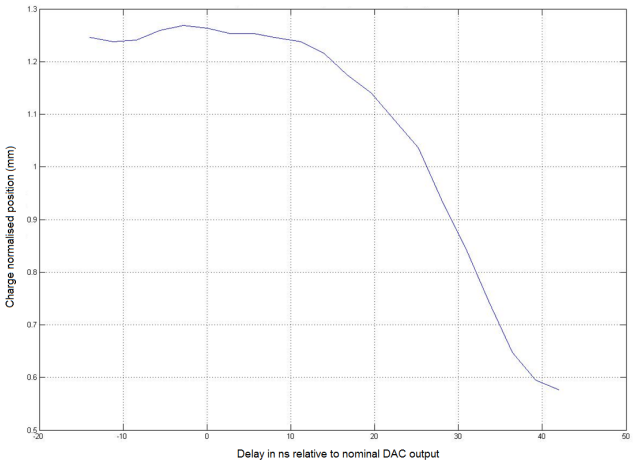

Figure 2.25: The position output of the digital FONT4 processor with delay added to the DAC output to measure the latency [69].

FONT1 successfully demonstrated how beams with different offsets could be corrected to the same trajectory within 70 ns and the correction held for the duration of the train. The correction ratio was 10:1 and a static position profile was present in each pulse that prevented all the bunches from being kicked to zero.

FONT2 improved on both latency, reducing the time to 57 ns, and effectiveness of the correction utilising a digital algorithm to remove static position offsets from the bunchtrain. The correction ratio was 14:1. FONT2 also used a solid state amplifier which was compact and a better design for the ILC.

FONT3 saw a great improvement in latency down to 24 ns and a better correction of 23:1 despite a non-linear response from the amplifier reducing its range of correction, as confirmed through simulation. These beam tests on a 1.3 GeV beam corrected the beam to 6 μm demonstrating that nanometre stability of a TeV was possible with the FONT3 feedback system.

The decision to build the ILC with superconducting cavities with large bunch spacings of 153.8 to 461.5 ns reduced the pressure on low-latency electronics allowing the use of digitisation making the feedback a more powerful tool with complex algorithms and more inputs. These digital FONT4 tests are ongoing. Current tests at ATF show full feedback without charge normalisation with a latency of 140 ns. Feedback with charge normalisation was also a success.

The future goals for FONT at ATF (and ATF2) include improving the resolution of the processor such that the natural jitter of the beam can be stabilised. The resolution of the FONT4 analogue processor, based on the FONT3 analogue processor, is $\sim$10 μm and is required to be $\sim$1 μm for micron stabilisation of the ATF beam.

Essentially, the BPM processor electronics, amplifier and delay loop have been tested without the need for an ILC-like beam. However, the conditions that the BPM is in at the test beam environments are very clean whereas a real Interaction Region (IR) is full

of backgrounds. Therefore the feedback studies were done in the absence of backgrounds. The remainder of this thesis addresses the concern of the operation of the BPM in high backgrounds.