A Fast Feedback System Designed to Maintain Luminosity at a Linear Collider

Stephen Molloy
Department of Physics

Queen Mary College
University of London

Thesis submitted for the Degree of Doctor of Philosophy in the University of London

· 2006 ·

Abstract

An Interaction Point FeedBack (IPFB) system has been designed for intra-pulse correction of beam-beam misalignments at the collision point of a high energy linear collider. This system measures the large beam-beam kick imparted to the outgoing bunch from one beam, and uses this to provide a corrective kick to bring the opposing beam into collision.

A prototype system is developed at the Next Linear Collider Test Accelerator (NLCTA) at the Stanford Linear Accelerator Center (SLAC) to test this scheme. Prototype BPM processors, amplifiers, and kickers where installed on the NLCTA beamline in order to show the effectiveness of such a feedback system, and a total correction delay of 53 ns is demonstrated.

A $\sim$1 μm correction of a 1 GeV beam scales to a 1 nm correction of a 1 TeV beam as may be present at a high energy linear collider, so a prototype demonstration of an IPFB system at a machine of 1 GeV would be a demonstration of the technology to be applied at the collider. Such a system is designed and results from its installation on the beam extracted from the 1.3 GeV damping ring of the Accelerator Test Facility (ATF) at the High Energy Research Laboratory (KEK) in Japan are shown. The 56 ns train length of the ATF beam necessitates reducing the system latency by a factor of $\sim$2 in order that the feedback is observed. To this end, BPM electronics are designed to process stripline outputs in $\leq$5 ns. This processor is demonstrated to have a resolution of $\sim$4 μm, and a signal delay of 4.6 ns. A total correction time of 23 ns is achieved.

And if I close my mind in fear, please pry it open . . .

J. Hetfield

Contents

List of Figures

List of Tables

Acknowledgments

First and foremost, thanks to Lisa Lynch. It was all her idea in the first place.

The list of people I am indebted to for their help over the last three years cannot possibly fit on one page, however I will try to mention some of the main players, and will hope that people who aren't mentioned will forgive me.

Top of this list must be my supervisor, Prof. Philip Burrows, who provided me with a seemingly endless amount of support and encouragement, and gave me the opportunity to work at both SLAC and KEK during a very exciting time in accelerator physics. This experience has proven to be of enormous help to my career, and I am truly indebted to him.

I am extremely grateful to the entire FONT group, without whose help this thesis could never have been written. Thanks to Colin Perry, whose experience and knowledge of electronics provided the backbone of FONT. Also to Joe Frisch for assistance with almost every area of this experiment, Glen White for his knowledge of data aquisition systems and cups of coffee at critical times, Alexander Kalinin for advice on RF systems, and Glenn Christian for his patient help while running the experiment. Many, many more people provided direct help to the FONT group over the years, and I extend my thanks to them, and apologise that the limitations of this page do not permit me to mention them by name.

For all the help they gave to a very naive, young, student, I have to thank the NLCTA team – including, but not limited to, Doug McCormick, Tonee Smith, Janice Nelson, Justin May, Kathleen Ratcliffe, and Keith Jobe. A special mention is reserved for Marc Ross, whose support and assistance has simply been awesome.

Thanks to the team at the ATF in KEK, for their warm welcome and assistance. Special thanks goes to Urakawa-san and Hayano-san for allowing us access to the ATF facility.

Last, but not least, I have to thank Dr. Stepan Boitsov and Simon Våge, whose wit and chaotic humour kept me sane from over 5'000 miles away. Skål til deg!!

Chapter 1

Introduction

1.1 Accelerators

The history of charged particle acceleration has been driven primarily by the study of fundamental particles and their interactions. Throughout the history of these investigations, the characteristic energy of each new machine has tended to increase, and this is mainly due to two considerations.

The first is that the structures to be resolved are extraordinarily small, occasionally below 10^{-15} m, and the probe used to examine these structures must have a spatial resolution less than or equal to this. It is clear that visible radiation, whose wavelength, λ, is of order 10^{-7} m, is inadequate for this task, so an alternative probe must be used.

If particles of energy, E, are used, then their spatial resolution is comparable to the de Broglie wavelength, λ_b,

$$\lambda_b = \frac{hc}{E} \tag{1.1}$$

This implies that higher and higher energy particles must be used to investigate smaller and smaller phenomena.

The second characteristic influencing the design energy of accelerators is that theories of the fundamental interactions have tended to imply the existence of particles of higher and higher mass. The amount of energy required to produce a particle of mass, m, derives from Einstein's fundamental relation,

$$E = mc^2 \tag{1.2}$$

which implies that a higher beam energy must be used to create particles of higher and higher energy.

Table 1.1 shows the rest masses of some particles, and in order to produce these, correspondingly high collision energies must be used.

Particle	Symbol	Rest Mass / MeV	Uncertainty
Electron	e	0.511	0.04 eV
b quark	b	4250	150 MeV
Z boson	Z	91187.6	2.1 MeV
t quark	t	174300	5.1 GeV
Higgs	H^0	>114400	95% Confidence Limit

Table 1.1: Example rest masses of some particles[23]. H^0 is the postulated, uncharged, Standard Model Higgs.

1.2 Colliders

A commonly used method of particle production is in the high energy collision of a particle with its anti-particle, in order that they annihilate and their rest mass and kinetic energy become the rest mass and kinetic energy of new particles.

A simple way to do this would be to collide a positron (e^+) beam with a fixed target known to contain a large number of electrons (e^-) with which they can annihilate. Conservation of momentum implies that the newly created particles will continue to move with the centre-of-mass reference frame, therefore a detector placed downstream of the fixed target would suffice to detect and identify the particles created in the reaction. A disadvantage of this system is that the kinetic energy of the centre-of-mass system in the rest frame of the laboratory is not available for particle reactions, and is thus wasted. It can be shown[27] that the ratio, η, of the energy available for particle reactions, E^*, to the accelerated beam energy, E_1 is,

$$\eta = \frac{E^*}{E_1} = \sqrt{\frac{2}{\gamma}} \tag{1.3}$$

where $\gamma = \frac{E_1}{mc^2}$, so in the case of highly energetic beams this is extremely inefficient.

An alternative scheme is to set up the head-on collision of particle beams of the same energy. In this case the centre-of-mass frame of the system is not moving with respect to the laboratory frame, thus no energy is wasted and the energy available for the production of new particles is maximised. For this reason, it is this method that is typically chosen for modern colliders. A disadvantage of this system is that the momentum vectors of the new particles may be pointed in any direction, and the entire solid angle around the interaction point must be covered by the detector.

1.3 Future Colliders

Present theories of high-energy physics (HEP) strongly imply the existence of, as yet, unobserved particles beyond the reach of present day colliders. These particles – one or more Higgs bosons, supersymmetric partners of Standard Model particles, etc. – are a result of theories designed to fill 'gaps' in the Standard Model (e.g. the Higgs field is postulated to account for the differing masses of the fundamental particles), or to go beyond the Standard Model to a more complete theory of fundamental interactions (e.g. supersymmetry).

Focussing on the Higgs particle for the moment, there are presently a very large number

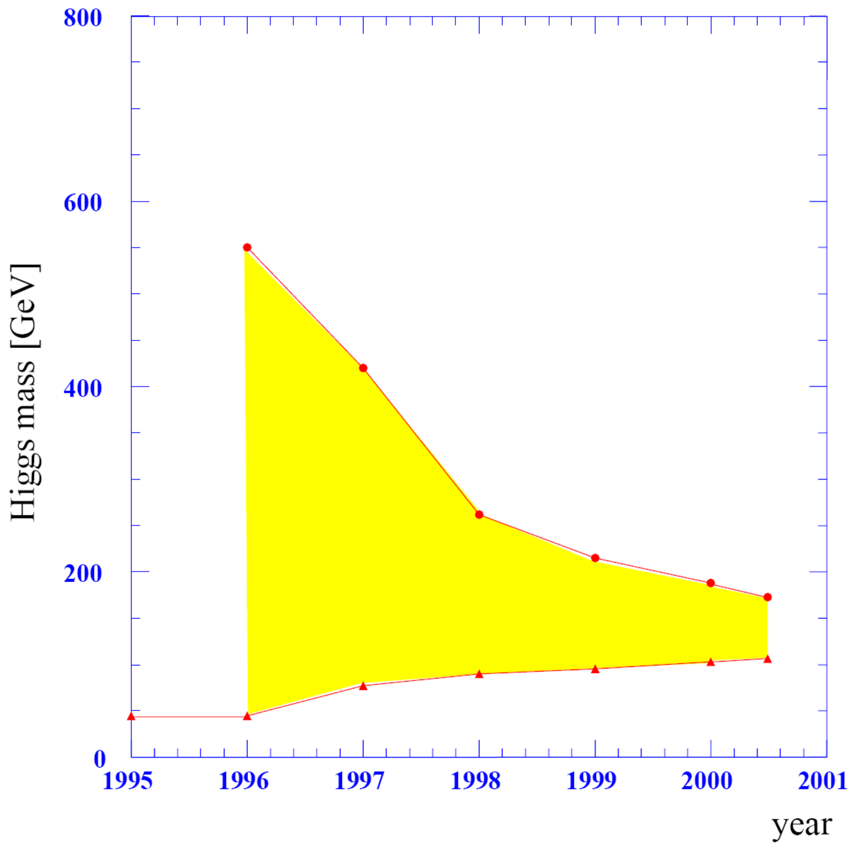

Figure 1.1: Plot showing how the upper and lower 95% confidence limits on the mass of the Standard Model Higgs have evolved over time[1].

of theories (e.g. different versions of supersymmetry) that make predictions about its interactions, properties, and even the number of types of Higgs bosons that exist. Despite the existence of such a wide range of possible properties, some experimental limits have been put on this, as shown in figure 1.1, which shows that, although the confidence limits on the mass of the Standard Model Higgs have been converging, they are still quite widely spaced.

In order to fully constrain this value, as well as to examine the plethora of other theories of fundamental particles and their interactions, it has been agreed by the worldwide HEP community[28] that the next accelerator should be an electron-positron (e^+- e^-) collider with an initial centre-of-mass energy of 500 GeV, later upgradeable to 1 TeV, and that it should be built in parallel with the Large Hadron Collider (LHC)[29] at CERN[30]. Several accelerators were designed for this purpose, with each machine going through extensive testing in order to decide on one final design. In general each of these designs accomplish the goal of achieving 500 GeV e^+- e^- collisions with the same basic components, however, large differences in the operation of each of these components do exist. Since the main thrust of this thesis concentrates on solving a problem that affects one particular machine design more than the others, it will be this design that is focussed on. This is known as the Next Linear Collider (NLC)[2] (the International Linear Collider[31] will be discussed in section 1.5).

1.4 The Next Linear Collider

The NLC is a proposal for a future linear collider, and it has been designed primarily at the Stanford Linear Accelerator Center (SLAC)[32]. It consists of several subsystems[2] to create intense beams of positrons and polarised electrons, accelerate them to very high energies, focus them to very small spots, and collide them in such a way that the results of the collisions can be measured by sensitive particle detectors. Each of these subsystems will be described, and it will be shown how each is necessary in order to achieve the physics goals roughly outlined in section 1.3.

1.4.1 Overview

Figure 1.2 shows a schematic layout of the NLC, showing the main accelerator systems, and the approximate physical size of the machine. It can be seen that the total length of the collider is ∼30 km, and that it is ∼200 m wide. Provision has been made for a second interaction point (IP) at which would be placed a second detector, including its own dedicated beam focussing system and dump, in order that measurements made at one detector can be independently confirmed. The following systems shown in this schematic will be discussed individually,

- Sources
- Damping rings
- Linac
- Beam Delivery System

1.4.2 Sources

The purpose of the sources are to create beams of electrons and positrons that are optimised for clean injection into the machine, and there are several methods of doing this[33].

The method chosen for generation of the electron beam in the NLC design is based on the photoelectric effect, and involves a polarised laser directed at a Gallium Arsenide (GaAs) photocathode. This method allows control of the beam modulation via the timing of the laser system, and also allows control of the polarisation of the electrons via the polarisation of the laser photons.

The laser must be modulated in order to produce trains of 192 bunches separated by 1.4 ns at the machine repetition rate of ∼120 Hz. In order to produce relatively square bunches of charge, the laser beam is modulated with the first and third harmonic of the bunch spacing.

The laser system will produce bunches of charge with the required bunch separation, however the bunches will be ∼700 ps long, which is excessively large for the accelerating RF in the main accelerator (see section 1.4.4), and for the acceleration in the damping rings (see section 1.4.3), so the beam needs to be passed through a 'buncher' in order to longitudinally compress the bunches. Here they are accelerated by RF whose wavelength, λ,

$$\lambda \geq 2\sigma_z \tag{1.4}$$

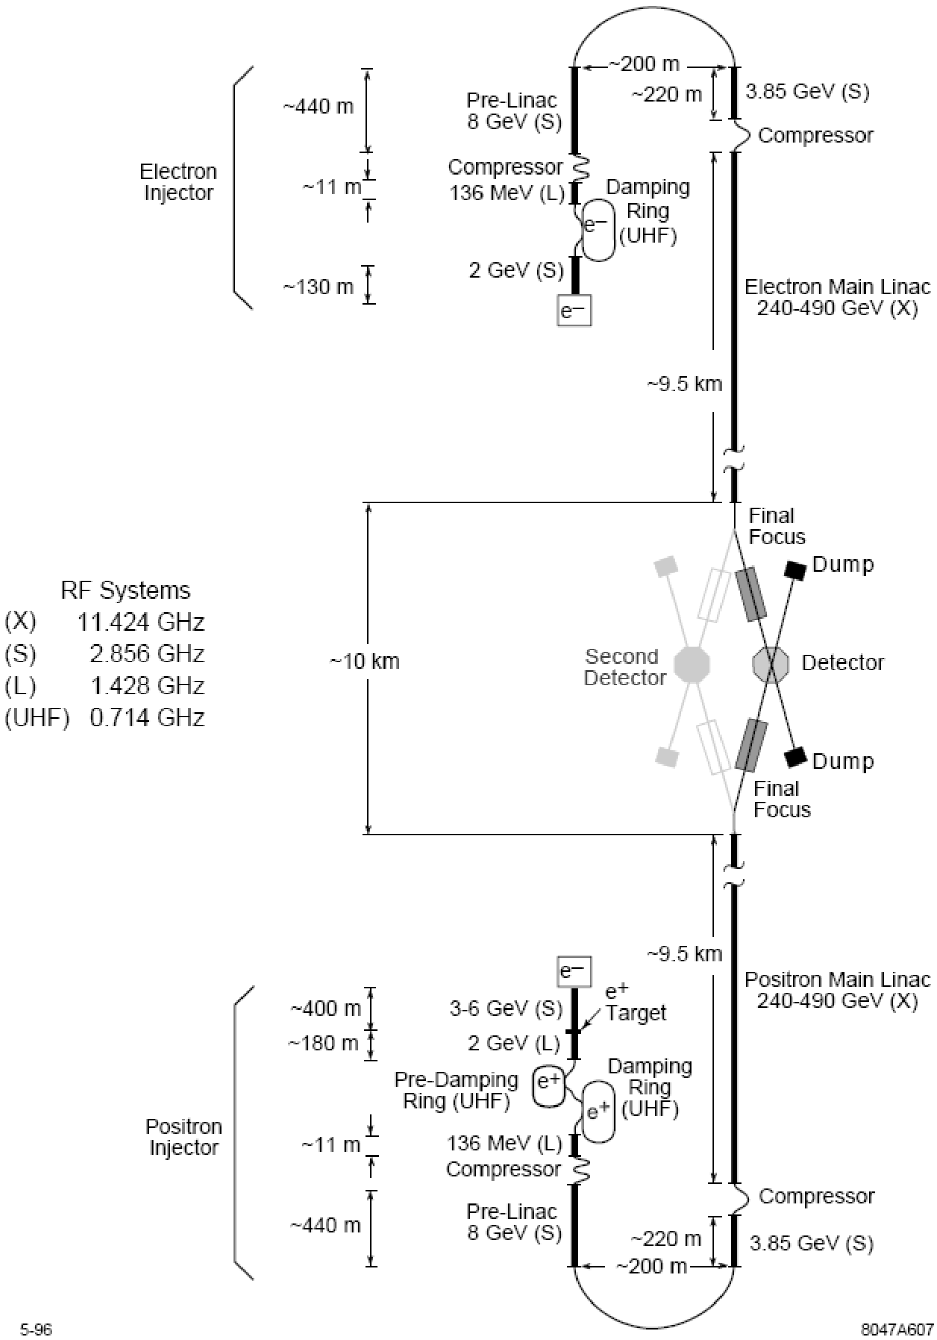

Figure 1.2: Schematic layout of the NLC[2].

where σ_z is the bunch length, and they are accelerated at the phase of the RF where the centre of the bunch coincides with the zero-crossing on the falling edge of the accelerating RF. As the bunches are not yet ultra-relativistic – $\frac{v}{c} \approx 0.75$ – any additional energy they receive results in an increase in speed. Therefore, particles at the head of the bunch will lose energy, thus slowing down, while particles at the tail of the bunch will gain energy, thus speeding up. In this way the particles toward the rear of the bunch will catch up with particles at the head, and the bunch length will shrink.

In the case of the electron beam, it is then possible to transfer the bunches to a 2 GeV accelerator in order to bring them to the right energy for injection into the damping rings.

The positron source is quite different from the electron source. In order to create the positrons, an electron beam is first created using a thermionic gun, which simply heats up the cathode in order to give its electrons enough energy to escape, and be accelerated toward an anode. Electrons produced in this way will have a large spread in the directions of their momentum vectors, so they pass through a solenoidal field during the initial acceleration phase in order to reduce this spread.

These electrons are then accelerated to ∼3–6 GeV and collided with a $W_{75}Re_{25}$ target where they initiate an electromagnetic shower. The positrons produced in this shower are separated from the electrons using a magnetic field, and the positron beam (modulated with the collision frequency of the initial electron beam into the target) is captured and accelerated to 2 GeV for injection into the damping rings.

1.4.3 Damping Rings

The momentum vectors of the beams produced by the sources will be spread to quite a large degree, and this must be reduced in order that the bunches be compressed to the very small spot size necessary for high luminosity collisions. This spread is known as the 'emittance' of the beam, and this concept is illustrated in figure 1.3.

Figure 1.3 shows a measurement of a bunch of particles measured at a particular point in the beamline, and the particles direction, $x' = \frac{dx}{d\vec{s}}$ (where $\vec{s}$ is the unit vector in the direction of motion of the bunch), is plotted against its position, x. A similar plot will exist for the y-plane. This plot was made for a large number of particles, and it can be seen that an ellipse is formed and the emittance is defined as the area of this ellipse (typically given in units of m.rad or mm.mrad). In the case of a more tightly controlled beam, the maximum deviation of x and x' will be smaller, and thus the emittance will be lower. It can be shown[27] that, in a conservative system, the emittance is a constant, and that it will vary as the inverse of the particle energy, therefore, the large energy gain of the particles in the linac (see section 1.4.4) will greatly reduce the emittance. Simulations of the collisions, however, have shown that this does not bring it to a low enough value to achieve the desired luminosity, therefore some other way of reducing it should be found.

If the beam is made to travel through a transverse magnetic field (for example in a circular machine, or the periodically repeating fields in an undulator or wiggler magnet) it will emit synchrotron radiation[34]. This radiation is emitted in the direction of motion of the particle, and thus the particle will lose energy in the direction of its motion reducing the length of its momentum vector. If it is then accelerated, the energy gain is purely

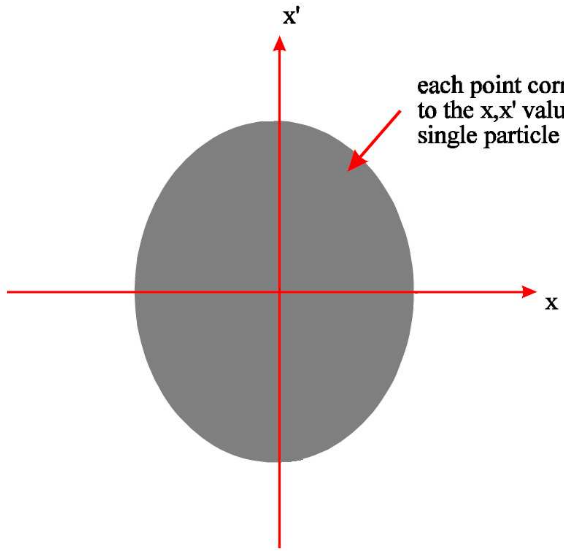

Figure 1.3: Tranverse phase space plot for a large number of particles in a beam. The area of the ellipse is defined as the emittance of the beam[3].

longitudinal, i.e. in the desired direction of motion of the beam. Thus, the length of the momentum vector of the particle has been returned to its original value, but the angle it makes with the direction of motion of the beam is smaller. This is equivalent to reducing the extent of the ellipse shown in figure 1.3 in both x and x', and this will shrink the area of the ellipse, and thus the value of the emittance.

Figure 1.2 shows that both beams travel through a damping ring after their production, and the positron beam is sent through two rings due to the very large initial emittance of the beam emerging from the electromagnetic shower. The normalised emittance[1] of the electron beam is reduced from, $\gamma\epsilon_{x,y} = 1 \times 10^{-4}$ m.rad (where the subscripts refer to the plane of the measurement) to $\gamma\epsilon_x = 3 \times 10^{-6}$ m.rad and $\gamma\epsilon_y = 3 \times 10^{-8}$ m.rad, and the normalised emittance of the positron beam is reduced from $\gamma\epsilon_{x,y} = 0.06$ m.rad to $\gamma\epsilon_{x,y} = 1 \times 10^{-4}$ m.rad by the pre-damping ring, and the second ring is identical to the electron ring.

After the beams have been damped they are sent through various compression stages, in order to shrink the bunch length from $\sim$5 mm to $\sim$100-150 μm. This is important as the long bunches that emerge from the damping rings would not be accepted by the linac, and a large amount of the beam would be lost.

1.4.4 Linac

The linac is arguably the defining subsystem of the entire machine. Many methods of particle acceleration exist, however all present high energy machines use the principle of RF acceleration, i.e. acceleration by an alternating field, as opposed to a constant voltage. Since the method of acceleration places firm restrictions on the quality and structure of the beam entering and leaving the linac, the choice of technology for this subsystem defines choices

[1] The area of the ellipse multiplied by the relativistic γ of the beam in order to eliminate the inverse dependence of the emittance on the beam's energy.

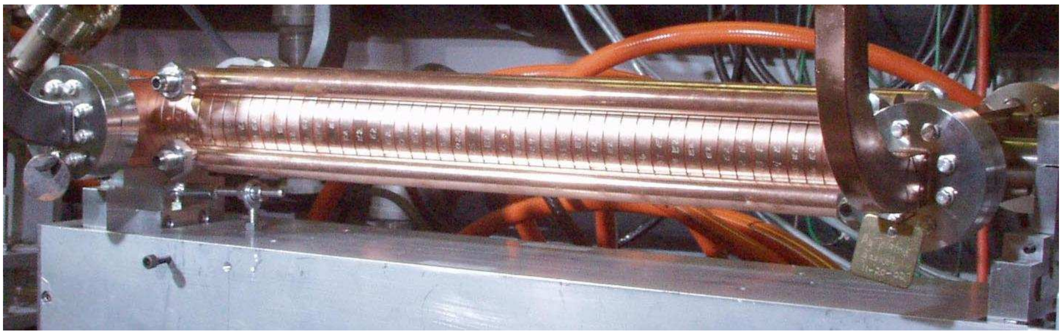

Figure 1.4: Photo of a prototype travelling wave accelerating structure installed on the NLC Test Accelerator beamline[4].

made for almost every other subsystem in the machine. The application of an alternating electric field to a particle bunch allows much higher fields to act on the beam, while avoiding the corona discharge limit of DC[2] acceleration.

RF acceleration works by generating an alternating field, which can be a travelling wave or a standing wave, in a structure through which the beam is sent. Consider the case where the alternating field is a travelling wave whose phase velocity, v_p, is equal to the speed of light, c, and the particles to be accelerated are already ultra-relativistic (i.e. their velocity, $v_e \approx c$). If the particles enter the cavity at the RF phase where the electric field is maximum, then, due to the fact that $v_e = v_p$, the particles will always be at this phase, and will thus receive the maximum energy from the RF.

In practice, the particles are not normally placed at the phase of the maximum field. In order to compensate for correlations between the particle's energy and its position in the bunch, the bunch is placed at a phase calculated to induce approximately the opposite energy/longitudinal-position correlation.

The NLC linac design relies on $\sim$16 km of travelling wave, copper cavities that use an accelerating frequency of 11.424 GHz. This relatively high frequency allows higher fields to be reached before breakdown occurs, thus giving a higher accelerating gradient[3], shorter linac, and reduced power consumption for a given beam energy. Due to fields excited in the cavity by the passage of the charged bunch ('wakefields') that could degrade the quality of the following bunches, sufficient time must be given between each bunch for this field to decay. It is for this reason that bunches are not accelerated at every crest of the 11.424 GHz field. Instead they are placed at every sixteenth crest, thus giving a bunch separation of $\sim$1.4 ns.

Figure 1.4 is a photograph of a prototype travelling wave structure installed for testing at the NLC Test Accelerator in SLAC. The 11.424 GHz accelerating field is injected into the cavity at the upstream end (left hand side of this image), and leaves at the downstream end.

[2]Direct Current. This implies the accelerating voltage does not depend on time.
[3]Energy gain of the particles per unit length of the linac.

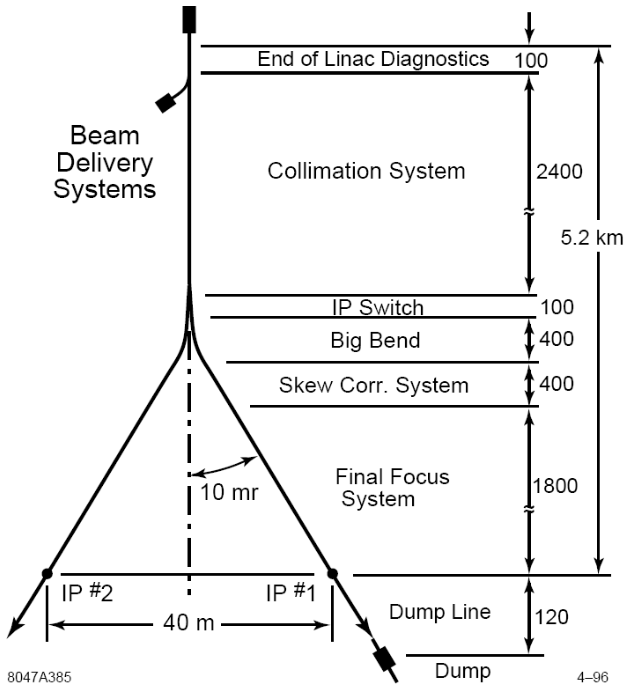

Figure 1.5: Schematic of the NLC Beam Delivery System (BDS) for one beamline[2]. The beam is moving from the top of the image to the bottom, and all regions downstream of the linac are shown.

1.4.5 Beam Delivery System

Figure 1.5 shows a schematic of the NLC BDS for one beamline. At the end of the linac, the beam enters a series of diagnostics to determine the quality of the beam as it enters the BDS. This is extremely important, as, for example, a misteered beam could cause severe damage downstream of this point, including the detector. If a beam error is detected at this point, an emergency dump has been provided for, and this can be seen as the small right turn (from the point of view of the beam) at the top of figure 1.5.

The next section is a series of collimators that scrape off beam particles whose deviation from the nominal bunch position is very large. The bunches are expected to be approximately Gaussian shaped in all three dimensions, but with a 0.1% 'halo' extending very far from the mean position, and it is this halo that the collimation system is designed to remove. As well as particles with large position offsets, the collimation system will remove particles whose trajectory is excessively steep, and whose energy is far from the nominal. In this way the detector is protected from backgrounds that could damage it or degrade the quality of the particle physics data.

Following this is a switch and a bending region that allow the beam to be switched between two interaction points. Complementary detectors will be installed at each of these collision points, which will share the available luminosity.

Before the beams collide, they are demagnified by $\sim$80 horizontally and $\sim$300 vertically in order to achieve the necessary luminosity. This demagnification is the job of the final focus system, which will also include corrective systems for energy-spread induced position deviations, mechanical systems to stabilise the final focussing magnets, and a full suite of beam diagnostics.

Finally, the beams collide and exit via a fully instrumented dump line to the dumps.

1.5 International Linear Collider

For many years the high energy physics community has been united in the belief that an electron-positron collider capable of reaching the TeV range is of the utmost importance in both complementing the physics results from the LHC, and in extending the field of particle physics in its own right. To this end, however, two main[4] technologies with which to accelerate the beam have been developed. One of these is known as the X-band accelerator, and is the technology of choice for the NLC design, as described in section 1.4.4. As explained this uses 11.424 GHz radiation[5] to accelerate the beam through conventional copper cavities. Acceleration through copper cavities is a mature, well known, technology, however, it was a major R&D challenge to obtain the very high (65 MV/m) accelerating gradients.

An alternative technology was used in a proposed machine called the Teraelectronvolt Energy Superconducting Linear Accelerator (TESLA)[1] and this was mainly designed at the Deutsches Elektronen-Synchrotron (DESY)[35]. This proposal was based around the use of superconducting niobium cavities to accelerate the beam. Since these cavities are superconducting, none of the power in the 1.3 GHz (L-band) accelerating RF is lost to resistive heating in the cavity itself. Instead all of the power is available for transfer to the beam, making this technology, in principle, much more efficient.

Figure 1.6 shows the basic layout of TESLA, and it can be seen that the basic components listed in section 1.4.1 are also included. This design also includes a facility known as an x-ray laser that uses the accelerated electrons to create very bright, highly collimated, ultra short pulses of synchrotron light for use in, for example, observing chemical reactions.

The TESLA design requires an accelerating gradient of $\sim$25 MV/m for the beams to reach the collision energy of 500 GeV, so proportionally more of the TESLA design is filled with accelerator structures compared to the NLC design.

Several other differences between NLC and TESLA are also visible, and many of these are driven by the choice of accelerating technology. The most obvious difference is the much longer damping ring in the case of TESLA. Due to the very small power loss in the walls of the superconducting cavities, the beam induced wakefields are much longer lived in the case of TESLA compared to NLC, and this necessitates a much longer inter-bunch separation – $\sim$330 ns for TESLA compared to $\sim$1.4 ns in NLC. This much increased bunch separation leads to much longer bunch trains[6], and, since an entire train must fit inside the damping ring, the damping ring must become longer.

[4] An additional technology (CLIC) relying on acceleration using the energy recovered from a drive beam of electrons has been developed at CERN, however it is generally recognised that this promising technology still needs some time before it reaches maturity.

[5] The frequency spectrum is split into numerous bands, and the IEEE US designates frequencies within the range 8 GHz to 12 GHz as ‘X-band’ radiation. This is why this technology is referred to as the X-band design.

[6] There are many more bunches per train in TESLA (2820) compared to NLC (192) in order to obtain approximately the same luminosity in each machine.

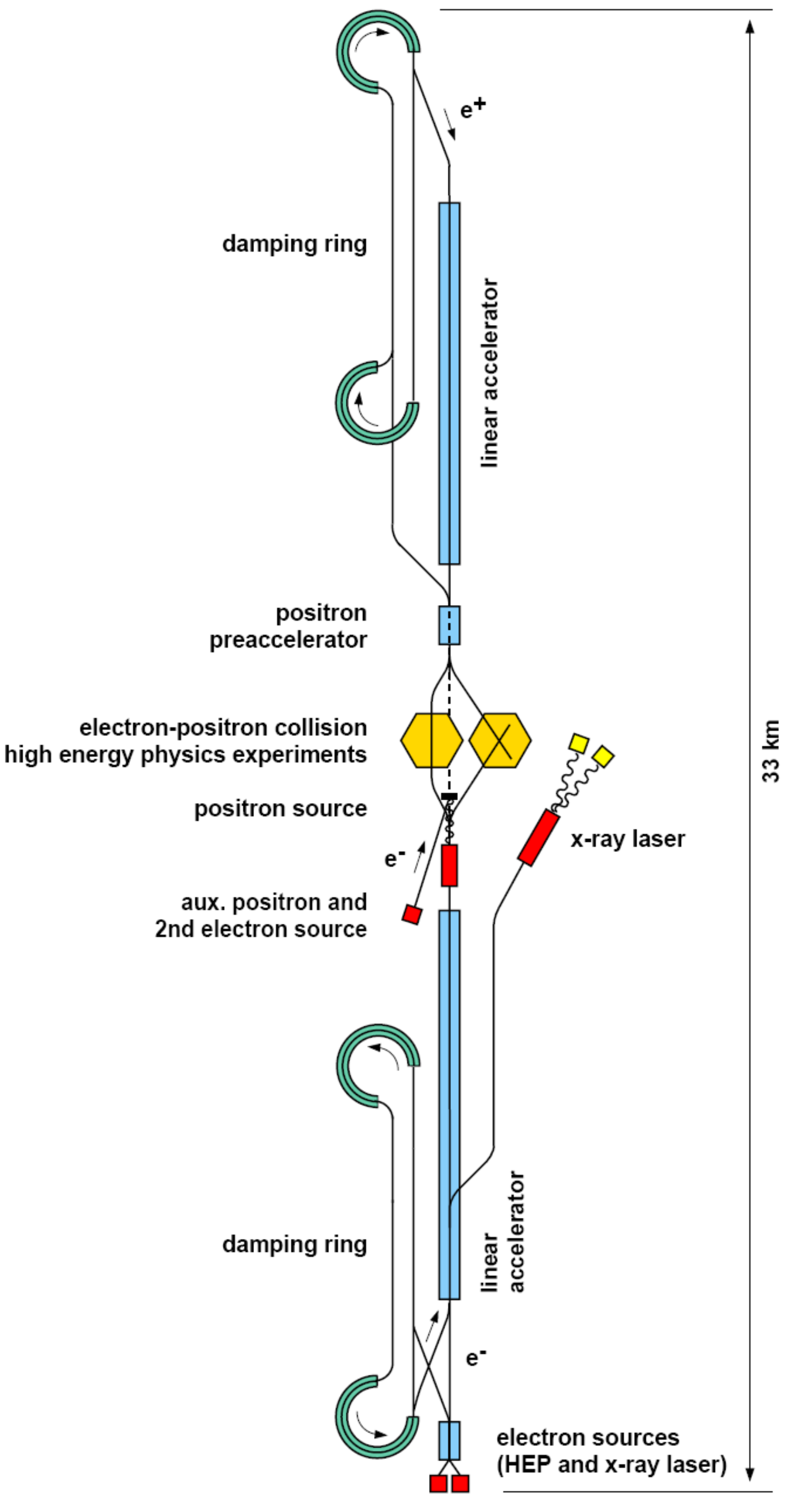

Figure 1.6: Schematic of the TESLA design[5]. It can be seen that, in comparison with figure 1.2, this design incorporates the same basic subsystems as the NLC design, however an x-ray laser facility has been included in the TESLA baseline design.

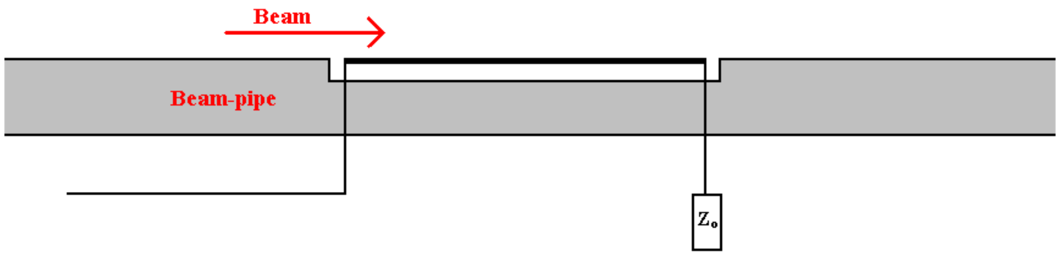

Figure 1.7: Stripline pickup installed on the lower wall of the beam pipe. The downstream end of the stripline is terminated through an impedance, Z_0.

1.5.1 International Technology Review Panel Decision

The International Technology Review Panel (ITRP) was given the task of evaluating both of these technologies – the superconducting 'cold', and the X-band 'warm' design – and giving a recommendation to the international community as to which design to press forward with, and in August 2004 they released their decision to move forward with the cold technology[36]. It was made clear that this decision was not a recommendation to move forward with the TESLA design, but rather was a statement that they felt that the goals of the high-energy community could be best met with the superconducting technology. Thus a new design – the International Linear Collider (ILC) – was created, and it is this design that is currently evolving toward a baseline proposal.

As of early 2006, the ILC design (controlled by the Global Design Effort (GDE)) has reached the stage of releasing a baseline design, and rapid progress is being made toward realising this machine.

1.6 Beam Position Measurement Methods

In designing and building high energy accelerators it is very important to include suitable diagnostics, including methods of measuring the energy, emittance, spot size, and position, and many diagnostic techniques have been built up over many decades (see [37]). This section will concentrate on the main methods of measuring the position of the beam.

There are three main beamline instruments that are used to measure the position of the beam as it moves down the beam pipe[38] – stripline pickups, button pickups, and RF cavities. Each of these will be described in turn, as well as the method of extracting the beam position from their output.

1.6.1 Striplines

Figure 1.7 shows a sketch of a stripline pickup installed in the lower wall of the beam pipe. The upstream end of the stripline is connected to a cable that brings the signal to the processing electronics, and the downstream end is terminated through an impedance, Z_0. It can be seen that the stripline serves as a break in the electrical continuity of the beam pipe wall, and it is this discontinuity that creates the output signal.

In the rest frame of the electron bunch, its field lines are distributed evenly in all radial directions, however, since the bunch is moving at an ultra-relativistic velocity in the laboratory frame, its field lines will be Lorentz–contracted to a very narrow disc perpendicular to the momentum vector of the bunch. Due to this very tight contraction of the field lines, the image current monitored at any point in the beam-pipe wall will, therefore, vary in the same way as for the beam current. The image current will then have the same time-varying profile as the beam, but with the opposite charge. If the image current is integrated around the circumference of the pipe it will be equal in magnitude, but opposite in sign, to the beam charge[7].

As the beam passes over the upstream end of the stripline the field lines 'jump' the discontinuity, coupling with the stripline, and exciting a pulse. This pulse is able to travel in both directions, so it splits in two[7], with one half travelling out to the processing electronics, and the other half travelling down the stripline. Since the signal propagation speed in the stripline is very close to the speed of light, the excited pulse and the beam will reach the end of the stripline simultaneously. At this point the field lines will move across the discontinuity from the stripline to the beam-pipe wall, thereby inducing a pulse of the opposite polarity to the original. Once again this pulse will be able to travel in both directions, so half will travel out through the impedance, Z_0, while the other half travels back down the stripline to the processing electronics. Note that the initial positive pulse that travelled with the beam to the downstream end of the stripline will be cancelled out by the negative pulse, so it is possible to instrument both ends of a stripline to monitor beams travelling in the same pipe, but in opposite directions, (for example, near the interaction point of a collider, where the beam-pipe is shared by the incoming and outgoing beams).

The un-processed output of a stripline will then be a pulse (positive polarity if generated by a negatively charged beam) whose time structure is identical to that of the beam, followed by a similar pulse of the opposite polarity. The time separation, t, between these will be equal to the return trip time along the stripline, i.e.,

$$t = \frac{2l}{c} \tag{1.5}$$

where the stripline has length, l, and the signal travels at the speed of light, c.

This signal is shown in figure 1.8.

Signal Processing – Difference over sum

For a single stripline, the amplitude of this output signal will depend linearly on the charge of the beam, and it will also depend on the proximity of the beam to the stripline as follows[38],

$$V_{peak} \propto \arctan \frac{w}{2\,(R-d)} \tag{1.6}$$

where V_{peak} is the peak amplitude of the output pulse, w is the stripline width, R is the half-separation of the striplines, and d is the distance of the beam from the centre of the

[7] The pulses will still have the same time characteristics as the excited pulse, but will now have half the amplitude.

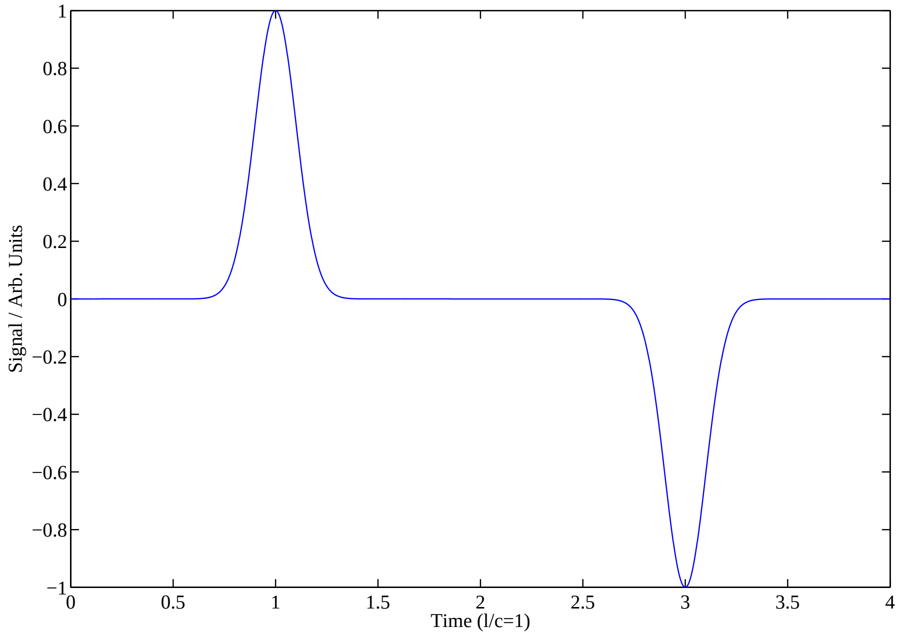

Figure 1.8: Theoretical output of a stripline. The amplitude is plotted in arbitrary units, and the time axis is plotted in units such that $\frac{l}{c} = 1$. The beam is modelled as having a Gaussian charge structure.

beam-pipe. A Taylor expansion of this leads to,

$$V_{peak} \propto \frac{w}{2\,(R-d)} - \frac{w^3}{24\,(R-d)^3} + \frac{w^5}{160\,(R-d)^5} + \cdots \tag{1.7}$$

Since the half-separation of the striplines is typically larger than their width by a factor of 5 or more, and when $d \ll R$, this approximates to the following,

$$V_{peak} \propto \frac{w}{2\,(R-d)} \tag{1.8}$$

i.e. an inverse dependence on the beam's distance from the strip.

With two oppositely placed striplines (e.g. on the top and bottom of the beam-pipe) it is possible to subtract these signals in order to extract a signal whose amplitude is directly proportional to both the charge and the distance of the beam from the centre of the beam-pipe in the plane of the stripline pickups. This can be seen by the following,

$$V_1 = f\,(t)\,Q\rho\frac{w}{2\,(R-d)} \tag{1.9}$$

$$V_2 = f\,(t)\,Q\rho\frac{w}{2\,(R+d)} \tag{1.10}$$

where V_1 and V_2 are the signals output by the top and bottom striplines respectively, $f\,(t)$ is the function that defines the shape of the output pulse, Q is the bunch charge, and ρ is the impedance of the measurement electronics. It can then be shown that

$$V_1 - V_2 = f(t)\,Q\rho w d\left(\frac{1}{R^2 - d^2}\right) \tag{1.11}$$

and since $d \ll R$,

$$V_1 - V_2 \approx f(t)\,Q\rho w d \tag{1.12}$$

The sum of these signals can be shown to be proportional to the charge and not the position,

$$V_1 + V_2 = f(t)\,Q\rho w R \tag{1.13}$$

So if the difference between the signals is normalised by their sum, this will yield a signal that is proportional only to the beam offset from zero.

$$\frac{V_1 - V_2}{V_1 + V_2} = \frac{d}{R} \tag{1.14}$$

This is the so-called 'Difference-over-sum' method.

The difference and sum may be found in software after sampling the output waveform, or may be found with analogue electronics connected directly to the stripline. The division by the sum signal is normally performed in software due to difficulties with analogue division.

Signal Processing – Log ratio

An alternative scheme is to calculate the logarithm of each of the signals from the opposite pickoffs (e.g. using logarithmic amplifiers), and subtract one from the other. This leads to a signal that is proportional to the beam position and is independent of its charge, as follows[38],

$$y \equiv \log(V_1) - \log(V_2) = \log\left(\frac{V_1}{V_2}\right) \tag{1.15}$$

where y is the beam position, and V_1 and V_2 are the amplitudes of the signals from the top and bottom striplines respectively.

As V_1 and V_2 are both dependent on the charge, equation 1.15 shows that taking the difference of the logs is equivalent to normalising by charge. The logarithm of the ratio of the stripline outputs will provide a result that is approximately linear with the beam offset (again with the stipulation that the offset be much less than the stripline dimensions).

Signal Processing – Amplitude Modulation to Phase Modulation (AM/PM)

The methods explained so far involve manipulating the signal amplitudes in order to cleanly extract the beam position. The AM/PM[39] method on the other hand, first converts the amplitude information into phase information, and calculates the beam position from that.

This is often accomplished with the use of a series of splitters and delay lines. The signal from each pickoff is split, with one half of the signal being delayed by 90° of the frequency

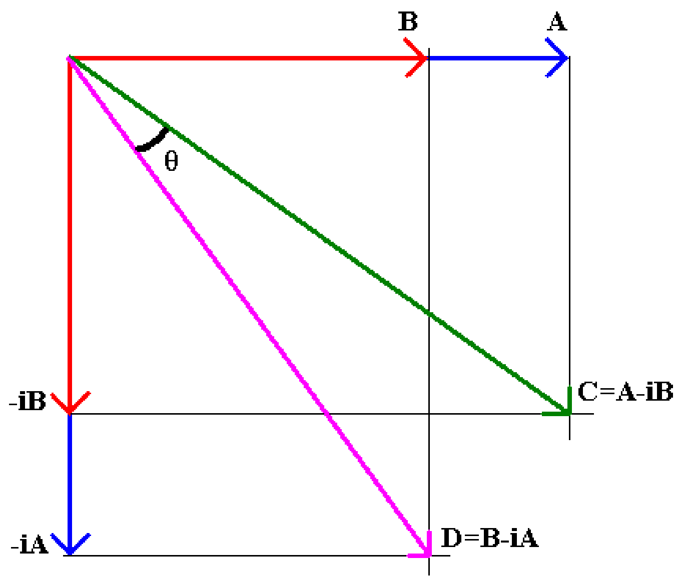

Figure 1.9: Diagram showing the amplitude/phase relationships of the signals in the AM/PM processing scheme.

of interest. This rotated component of each pickoff is then added to the un-rotated half of the signal from the opposing stripline. The result of this operation will be two signals whose relative phase is related to the beam position, and whose amplitude is related to the beam charge. This can be seen from figure 1.9

In figure 1.9 the signals are shown in complex space, where signals are represented by vectors whose lengths are proportional to their amplitude, and whose rotation angles are proportional to their relative phase. Signals A and B represent the output of the top and bottom striplines respectively, and it can be seen that they have the same relative phase but differing amplitudes (the case of a beam closer to the top pickoff is illustrated here). In the AM/PM scheme these signals will be rotated by 90° (illustrated by a multiplication by $-i$ to give $-iA$ and $-iB$) and added to the un-rotated component from the opposing stripline. It can be seen that the results of this ($C \equiv A - iB$ and $D \equiv B - iA$) will have the same amplitude, but differing phases, and that the phase difference will depend on the relative size of A and B, i.e. the position of the beam. It has been shown (for example in [40]) that the position, y, calculated by this method is,

$$y \propto \arctan\left(\frac{A}{B}\right) \tag{1.16}$$

and so will be approximately linear with respect to the beam position close to the centre of the device.

Comparison of Processing Schemes

Figure 1.10 shows the transfer functions of the three processing schemes discussed. It is clear from this plot that the difference over sum scheme is superior in terms of its linearity. The

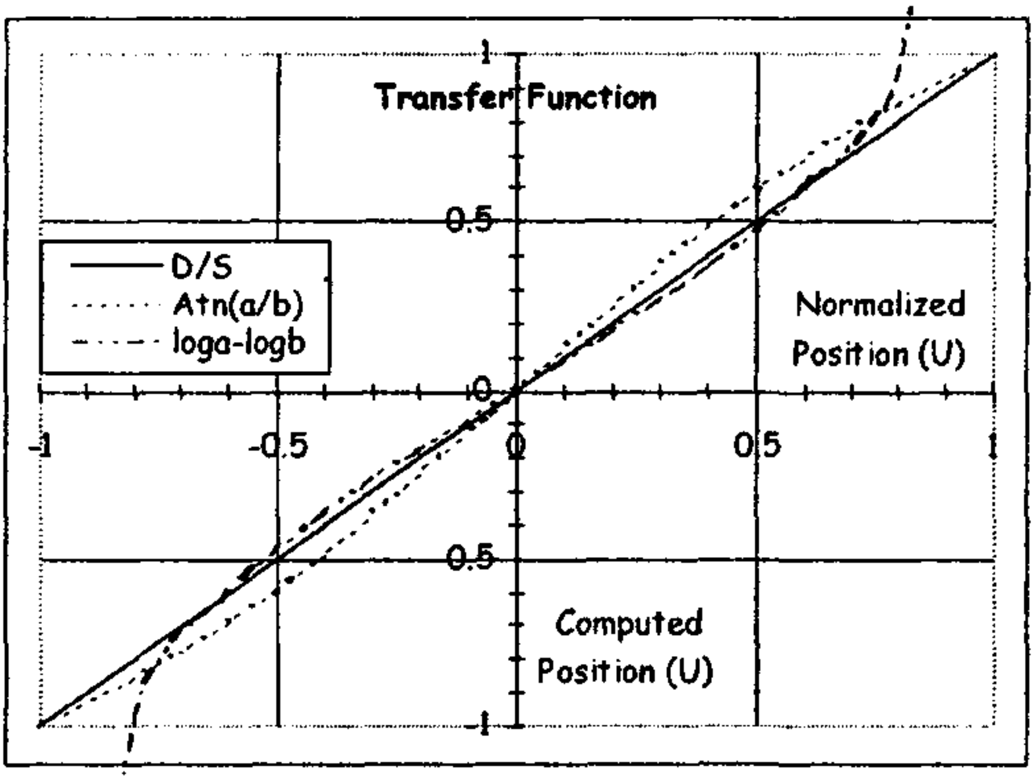

Figure 1.10: Plot of the transfer functions of the three processing schemes discussed[6]; Difference over sum (D/S), AM/PM (Atn(a/b), and log ratio (loga-logb). The axes are in units of the separation of the BPM electrodes.

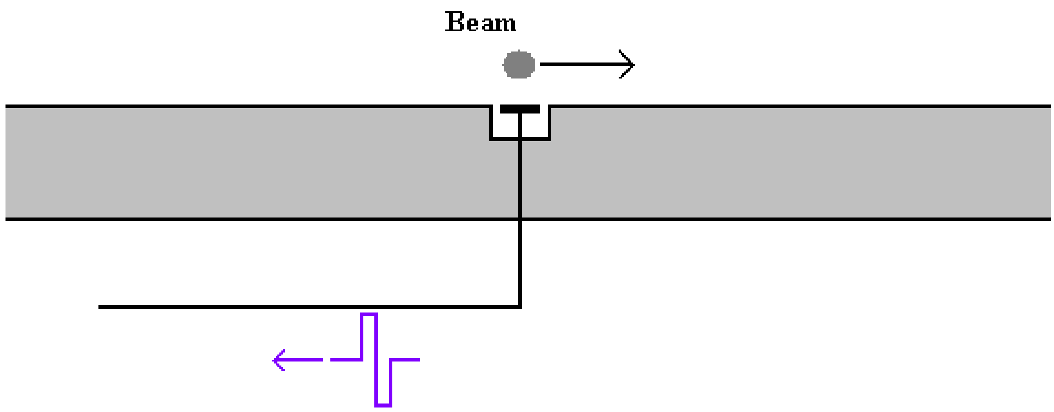

Figure 1.11: Button pickup installed on the lower wall of the beam pipe.

log-ratio scheme becomes very non-linear if the beam moves beyond $\pm 0.75U$, but behaves reasonably well inside these limits. It can be seen that the AM/PM scheme is the worst in terms of linearity, although it does respond well to large beam offsets.

1.6.2 Buttons

So-called 'button' BPMs are very similar in operation to striplines, however the electrode is much reduced in size. Figure 1.11 shows a sketch of a button electrode installed on the lower wall of the beam-pipe. The output signal is produced in the same way as for the stripline, i.e. the field lines inducing a pulse when they jump to and from the electrode. Due to the small length of the button in comparison to the stripline, the output signal will be much shorter, and will, therefore, be a fast doublet whose amplitude depends on the proximity and charge of the beam.

This signal can be processed in the same way as for the stripline output, i.e. difference over sum, log ratio, or AM/PM.

1.6.3 Cavities

The electro-magnetic fields generated by the passage of a bunch of charge through a RF cavity can also be used to measure the position of the bunch. Figure 1.12 shows two EM modes (only the electric field is shown) generated by a bunch of charge passing through the cavity at a slight offset from the centre.

The TM010 mode is known as the monopole mode, and its amplitude is proportional only to the charge of the bunch, and not its position. The TM110 mode (known as the dipole mode) is dependent on both the charge and the position of the beam. The amplitude is only proportional to the absolute magnitude of the displacement, however the phase of this mode changes by π when the beam crosses zero, so the polarity of the offset may be determined from examination of the phase. Since the monopole and dipole modes will have a fixed relationship, a measurement of their relative phases will yield the direction of the offset, and a comparison of their amplitudes will give the magnitude of the offset from zero.

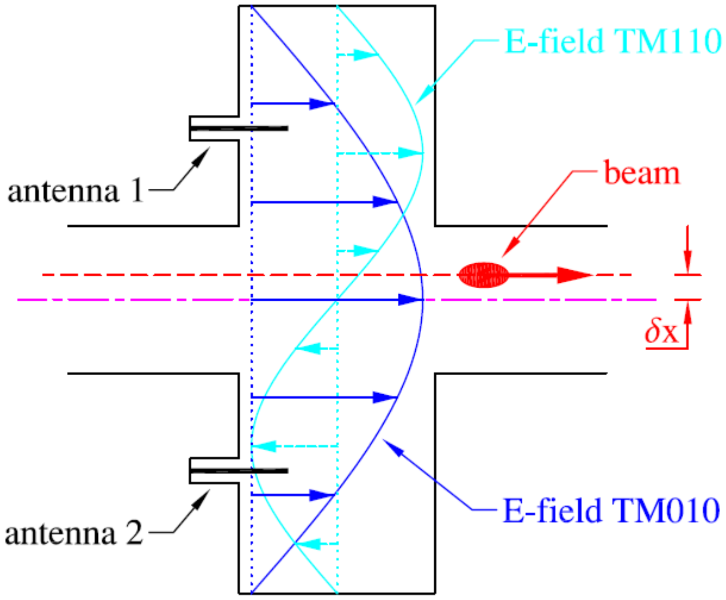

Figure 1.12: Sketch of a RF cavity, with a beam passing through displaced a small distance, δx, from centre. The electric field of the TM010 and TM110 modes generated by the passage of the beam are shown[7].

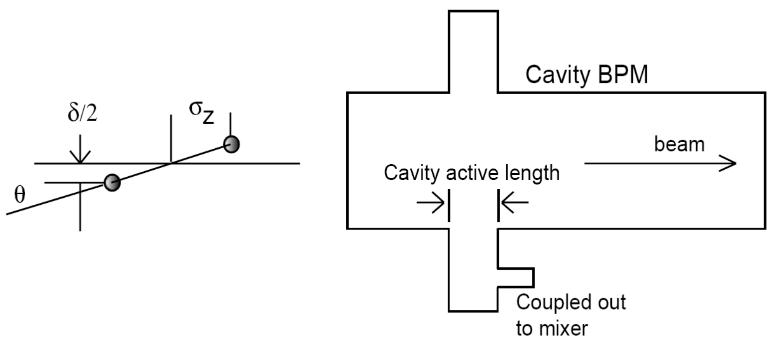

Figure 1.13: Sketch of a cavity with an incoming 'tilted' beam. The head and tail of the bunch are represented by single macro particles[8].

Figure 1.13 shows that cavities also respond to tilted bunches, or bunches with a tilted trajectory. In this figure the tilted bunch is represented as two macro-particles travelling through the centre of the cavity. As the 'head' particle travels through the cavity it will induce a monopole oscillation in proportion to its charge, and a dipole oscillation in proportion to its charge and position (the head is offset from centre by $\delta/2$). The 'tail' particle (whose charge is identical to the head particle) produces an identical monopole mode, and a dipole mode with the same amplitude, but with a relative phase of 180° compared to the head dipole mode. If the bunch length, σ_z is short compared to the cavity active length (as is typically the case), then the addition of two out-of-phase signals will be like a subtraction of two identical signals that have been slightly displaced in time, i.e. the differential. The differential has a relative phase of 90°, which will distinguish it from the position signal. Therefore, careful phase-space monitoring of the output signal allows the beam offset and tilt to be measured simultaneously.

An advantage of cavity BPMs over striplines and buttons is their extremely small resolution limit. Striplines and buttons produce very wide-band signals, so the processing of their signals typically involves filtering out a large portion of the high frequencies. Since such a large amount of the signal must be removed for processing this impacts the signal-to-noise ratio. Cavities, on the other hand, produce very narrow band signals, so it is possible to retain much more of the original signal power during processing. This allows a very high signal-to-noise ratio, and resolutions of $\sim$20 nm have been obtained with these devices[41].

Chapter 2

Luminosity Maintenance Using Beam-Based Feedback

Table 2.1 shows some beam parameters for the NLC design[24] and the Nominal ILC design[25], both for the initial centre-of-mass energy of 500 GeV and the upgrade energy of 1 TeV. In each of these designs the beam is focused to a few nanometres in the vertical, and it is expected that ground motion will cause the final focusing magnets to move by hundreds of nanometres, thereby moving the focal point of the bunch by the same amount.

Figure 2.1 shows the measured ground motion spectra from several sites around the world, and it can be seen that the trend of the ground motion versus frequency, ω, goes as ω^{-4}.

Figure 2.2 shows the integral of the power spectra measured under various condition at the SLC in order to show how the RMS ground motion changes with frequency. Table 2.1 shows that the repetition rate of the NLC is 120 Hz, so figure 2.2 shows that, in the 'noisiest' environment measured here, the RMS motion at this frequency will be at least $\sim$1 nm. The noise present at the bunch frequency of 714 MHz (inverse of the bunch separation time shown in table 2.1) will be many orders of magnitude smaller due to the ω^{-4} dependence shown in figure 2.1. These plots therefore, imply that there will be random noise with an

	NLC		**ILC Nominal**		**Units**
Property	**500GeV**	**1TeV**	**500GeV**	**1TeV**	
Electrons/bunch	0.75	0.75	2.0	2.0	10^{10}
Bunches/train	192	192	2820	2820	
Train Repetition Rate	120	120	5	5	Hz
Bunch Separation	1.4	1.4	307.7	307.7	ns
Train Length	0.269	0.269	867.7	867.7	μs
Horizontal IP Beam Size (σ_x)	243	219	655	554	nm
Vertical IP Beam Size (σ_y)	3.0	2.3	5.7	3.5	nm
Longitudinal IP Beam Size	110	110	300	300	μm
Luminosity	2.0	3.0	2.03	2.82	$10^{34} cm^{-2} s^{-1}$

Table 2.1: Some parameters of the NLC[24] and Nominal ILC[25] design for the linear collider, at both the initial centre-of-mass energy of 500 GeV and the upgrade energy of 1 TeV.

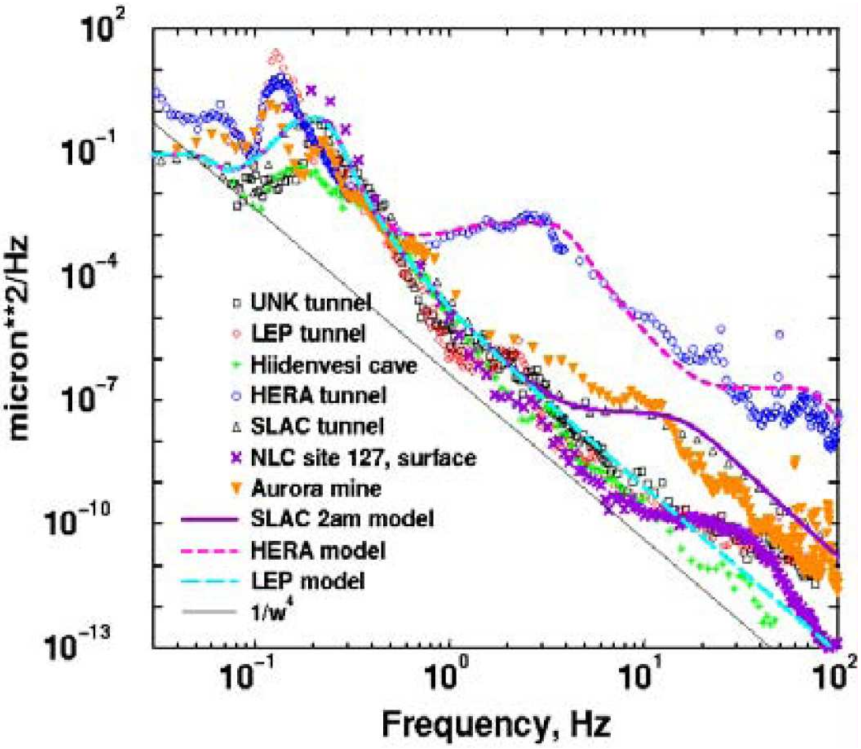

Figure 2.1: Ground motion measurements at various sites[9].

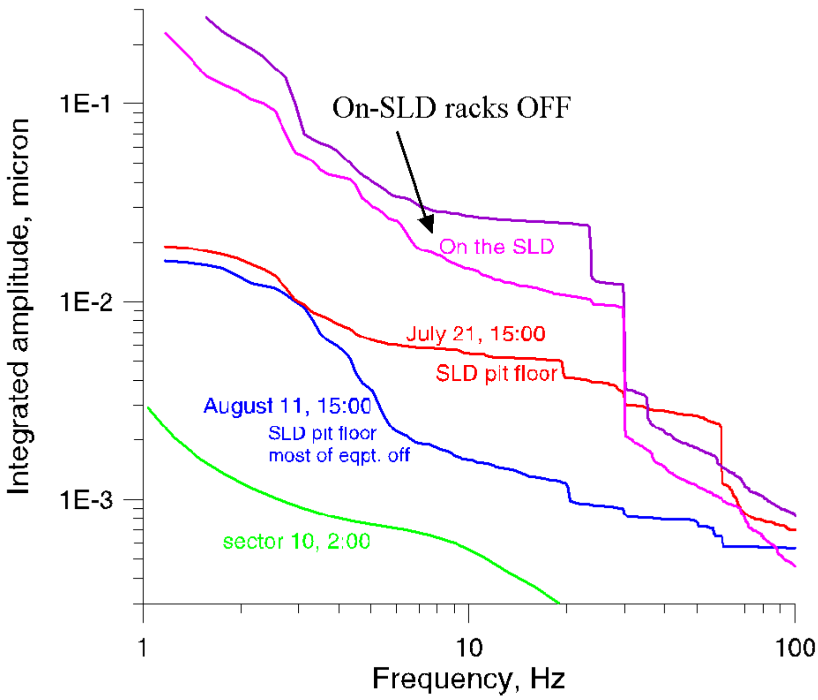

Figure 2.2: Integrated ground motion from a quiet region (sector 10) of the SLC[10] compared with the motion of the SLD[11] in a variety of circumstances[12].

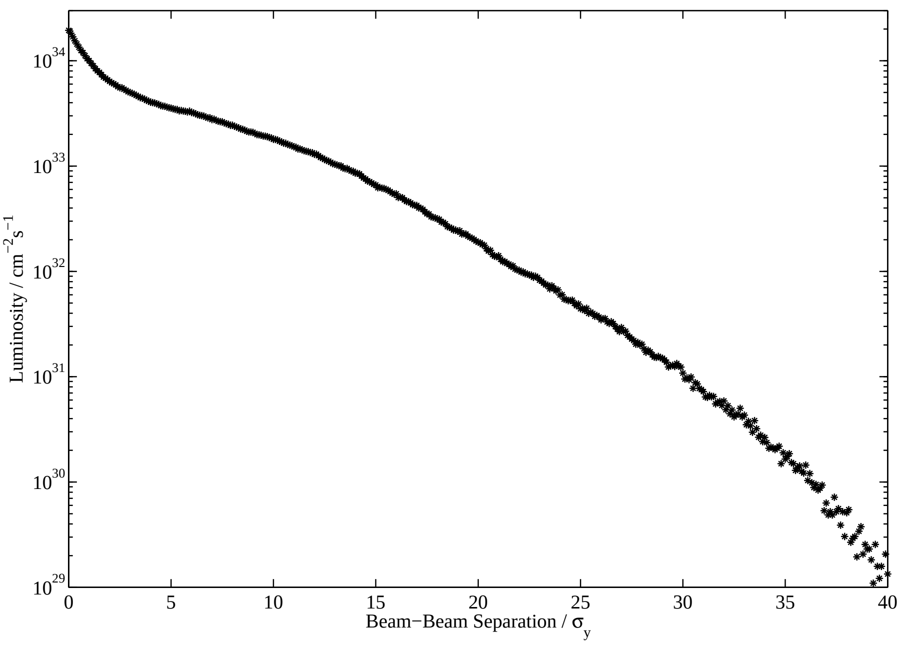

Figure 2.3: Simulation of the dependence of the luminosity on the offset of the colliding bunches[13]. This simulation used the Nominal ILC parameter set at a centre-of-mass energy of 500 GeV (see table 2.1).

amplitude of $\geq$1 nm on the relative positions of the trains, but the positions of the bunches in each train will be strongly correlated. The ILC bunching frequency of $\sim$3.2 MHz also means that the bunch position within the trains will be strongly correlated, however the lower repetition rate of 5 Hz implies a much larger RMS intra-train noise.

GUINEA PIG[42] was used to perform simulations[13] of the colliding beams in the Nominal ILC design in order to find the dependence of the luminosity on the vertical offset of the colliding bunches, and figure 2.3 shows the results of this. The horizontal scale of this plot is in units of the vertical beam size, and it can be seen that the simulated luminosity drops by orders of magnitude with only a small offset. It is, therefore, necessary to implement a correction scheme to maintain a low beam-beam offset, and the Feedback On Nanosecond Timescales project (FONT) was designed for this purpose[43].

2.1 The FONT Concept

An offset between the charged bunches will lead to a large angular kick being given to the outgoing bunches, and GUINEA PIG was used to calculate how this kick scales with the beam-beam offset[13]. The results of this are shown in figure 2.4. It can be seen that even a small offset causes a kick of several hundred micro-radians, leading to a motion of several hundred microns a few metres downstream of the IP. A beam position measurement at this point will, therefore, give an indication of the offset at the IP. The FONT concept is to use

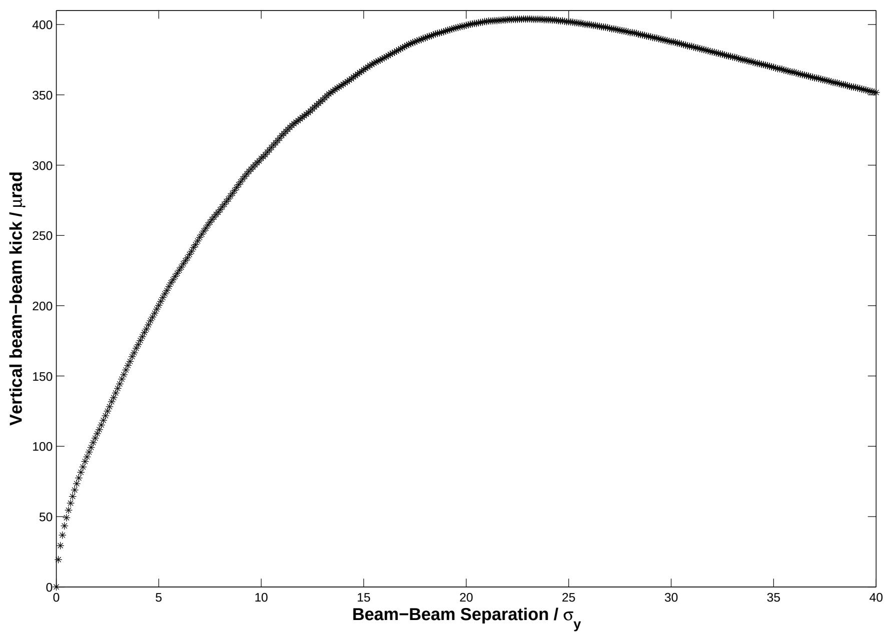

Figure 2.4: Simulation of the of the beam-beam kick due to the offset of the colliding bunches[13]. This simulation used the Nominal ILC parameter set at a centre-of-mass energy of 500 GeV (table 2.1).

this signal to operate a kicker on the incoming beamline of the opposite beam. An amplifier at the kicker input will be set to provide the correct gain to move the opposite beam by the right amount to bring it into collision, as shown in figure 2.5.

Due to the spectrum of the noise (figures 2.1 and 2.2) the position of each new bunch train will be unpredictable, which means that this feedback must operate within each bunch train. It should, therefore, be able to fully correct the beam position within a small fraction of the train length. Table 2.1 shows that the train length of the NLC is ∼270 ns, therefore, such a feedback system for this machine is very challenging. The FONT prototypes were designed to test the ability of performing such a correction on the NLC beam, so much of the effort was put into reducing the time delay (latency) of the system.

Figure 2.5 shows how FONT would be implemented at a high energy collider. It is set up as close to the interaction point (IP) as space allows in order to keep the system latency as low as possible.

If the amplifier gain has been set correctly there will be zero offset after the first correction between the colliding bunches and, therefore, zero kick of the outgoing beam. This will be measured by the FONT BPM as being at the nominal zero, and the time between the beam's initial arrival at the BPM and this zero measurement is known as the latency of the system. The zero signal will then travel around the loop to the kicker, so the kicker will then give a zero kick to the beam after the second latency period. Thus a round-trip delay loop is necessary to preserve the correction.

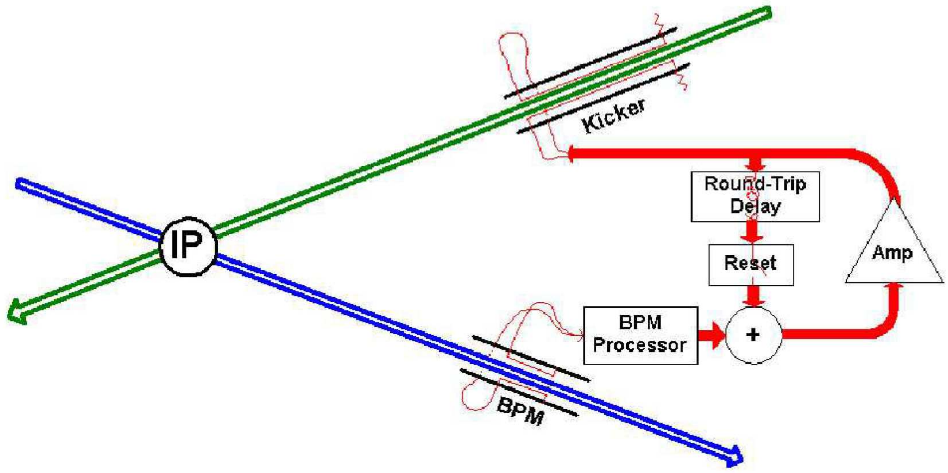

Figure 2.5: Proposed FONT operation at the ILC[14]

Property	**NLC – 500 GeV**	**NLCTA**	**Units**
Electrons/bunch	75	1	$\times 10^8$
Bunches/train	192	$\sim$2000	
Bunch Separation	1400	88	ps
Train Length	269	177	ns

Table 2.2: A comparison of the NLCTA beam with the NLC 500GeV parameter set.

This delay loop is a length of delay line whose length is set so that a signal propagating along it will take exactly one latency to travel its full length. So the signal from the BPM is split in two – one half to fire the kicker, and the other to loop back on itself. In this way the correct signal will continue to be sent to the kicker after the beam has been centred.

Between each bunch train the ground motion induced offset of the beam will have changed randomly, and the correction for the previous pulse will no longer be applicable. For this reason a reset is included to set the delay line signal to zero between each bunch train. In this way the feedback is starting anew with the onset of each new train.

The reset also prevents the feedback loop integrating the noise output by the BPM between each bunch train.

2.2 FONT1

All current colliding beam machines are dedicated to particle physics experiments, and this means that a FONT prototype test at such a machine would be extremely limited in scope. It was, therefore, necessary to perform the test on a single beam machine, and figure 2.6 shows the layout of FONT1 as installed on the NLC Test Accelerator (NLCTA)[44] beamline. Using this layout, it is possible to introduce an offset into the beam with a dipole magnet upstream of the kicker, which can then be measured at the BPM, and corrected by the kicker. So it can be seen that the FONT concept can be tested using a single beam layout.

Table 2.2 shows some beam parameters of the NLC 500 GeV design compared to the

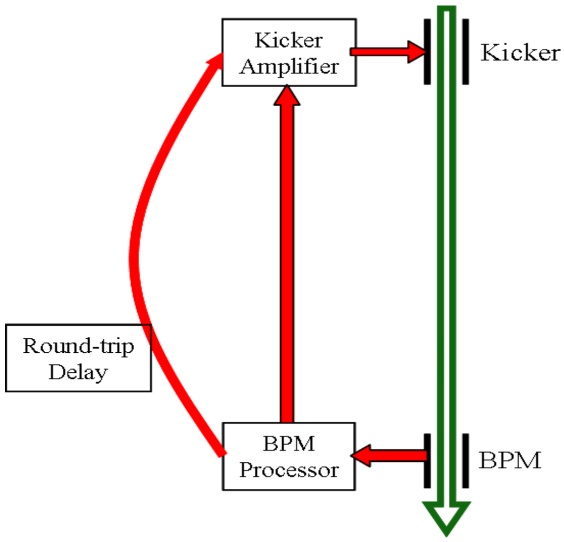

Figure 2.6: FONT1 operation at the NLCTA[15].

NLCTA beam. As previously mentioned, the correction needs to take place within a fraction of the train length, and it can be seen that the NLCTA bunch train is ~65% as long as the NLC beam, which gives enough time to observe several correction periods (the latency of the prototype is predicted in table 2.3). The bunch charge and bunch separation, however, are substantially different from the NLC, and this will require changes in the BPM technology used in the prototype. Since the main aim of the FONT tests was to demonstrate a fast beam correction within a few tens of nanoseconds, the NLCTA can be seen to be more than adequate for this test.

An initial prototype was installed in the NLCTA during 2001 and 2002, and full beam tests were completed in September 2002[15]. As well as the installation of the feedback BPM and kicker, including full processing electronics, feedback circuitry, and data-acquisition system, a dipole magnet was installed for the purpose of introducing a range of beam offsets.

2.2.1 BPM Hardware

See figure 2.7 for a photograph of a fully assembled FONT1 BPM. The BPM is the central object, and the spoolpieces extending to the left and right are narrow pieces of beampipe designed to strongly attenuate the X-band (11.424GHz) radiation generated by upstream accelerating structures. A 'pickoff' (see figure 2.8) tuned to resonate at X-band (for tuning procedure see section 3.3.1) is placed at each of the four remaining faces of the cube, and a current is excited on these by the passing of the beam. The presence of two opposing beam pickups in each plane allows measurement of the beam's position in both x and y, however the prototype experiment aims to correct the position in y only. The current excited on the pickoff is related to both the distance of the beam from the pickoff and the charge of the beam, so the signal must be processed to extract the position[45].

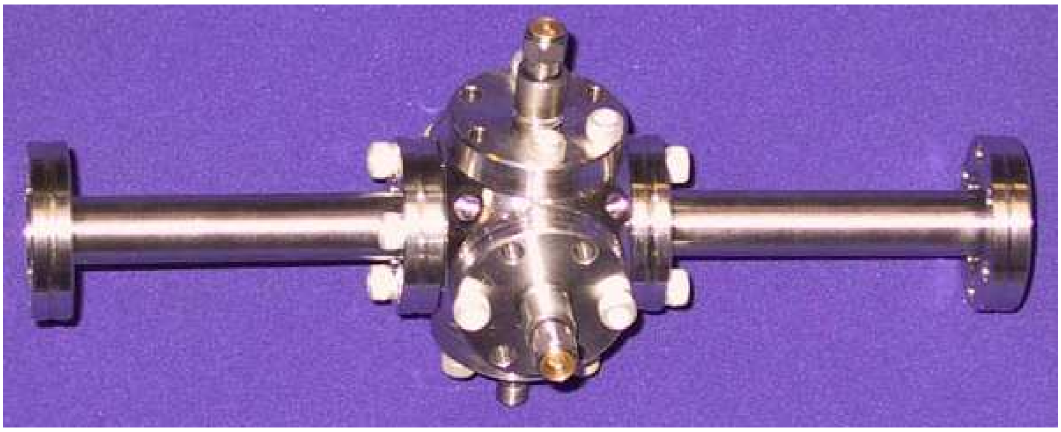

Figure 2.7: A photograph of a fully assembled BPM.[15]

Figure 2.8: A photograph of a 'pickoff' before assembly.[15]

2.2.2 Signal Processing Algorithm

The processing algorithm was based on the "difference over sum" scheme (see section 1.6). Since the analogue division of signals is costly in terms of time it was decided to digitise the sum signal and invert it in software. An Arbitrary Waveform Generator (AWG) can then be used to send this inverted waveform to an analogue multiplying circuit to find its product with the difference signal. Due to the time lag of the processing in software, the inverted signal will be from a previous pulse, necessitating the use of an algorithm to correct for this. The first two terms of a Taylor expansion were used for this purpose[15],

$$\frac{\Delta}{\Sigma} \approx \frac{\Delta}{\Sigma'} - \frac{\Delta(\Sigma - \Sigma')}{\Sigma'^2} = \frac{\Delta}{\Sigma'}(1 - \frac{d\Sigma}{\Sigma'}) \tag{2.1}$$

where Δ is the difference in signal amplitude between the top and bottom pickoffs of the BPM, Σ is their sum, and the position of the beam is calculated as being Δ normalised by Σ. As explained $\frac{1}{\Sigma}$ was calculated in software using Σ for a previous pulse, as indicated by Σ'. $d\Sigma$ is defined as $\Sigma - \Sigma'$.

2.2.3 Kicker Amplifier

In order to clearly observe the effect of the feedback it was necessary to provide a maximum beam kick of ~1 mm, and calculations[15] indicated that the kicker amplifier (see figure 2.6) would have to output ~500 V in order to induce a kick of this size.

Figure 2.9: A photograph of the electronics layout within the FONT1 kicker amplifier.[15]

Component	Expected Latency / ns
Beam time of flight	14
Signal return delay	18
Total irreducible latency	32
BPM Processor	5
Feedback circuit	16
Kicker amplifier	10
Kicker fill time	3
Total reducible latency	34
Total system latency	66

Table 2.3: The expected delay of each component in FONT1[15].

The amplifier design[46] was based on planar triode technology, and it amplified the BPM signal through various stages to bring it up to the high power required. Figure 2.9 is a photograph of the interior of this device in order to show the electronics layout.

2.2.4 Latency

Table 2.3 shows the expected delay associated with each part of the system. The time of flight and signal return delay are considered irreducible since they are a function only of the kicker-BPM separation and the speed of signal propagation in the cables. The remaining latency components are considered reducible as they are functions of the technical design of each component.

In total a latency of 66ns is expected[15].

2.2.5 Results

The results of the first prototype experiment are shown in figure 2.10 (adapted from [15]).

In figure 2.10 the system can be clearly seen to be operating as expected. Each plot shows the beam position against time for five different dipole settings. Figure 2.10a shows the original condition of the beam with no feedback acting in order to illustrate the baseline conditions. The five dipole settings are clearly visible, and it can also be seen that there are position variations along the bunch train that are independent of the dipole setting and that are present on every pulse.

Figure 2.10b shows the effect of turning on the feedback loop without activating the delay loop. After one latency period the correction signal is applied to the beam and it can be seen that the various dipole offsets are kicked to zero. Since the delay loop is inactive, this zero signal is allowed to travel to the kicker and the beam returns to its original position after the second latency.

Figure 2.10c shows the results with the delay loop activated along with the main loop. The correction begins at the end of the first latency period as expected, and it can be seen to continue through to the end of the bunch train.

A system latency of 67ns was measured[15], and this compares well with the prediction from section 2.2.4.

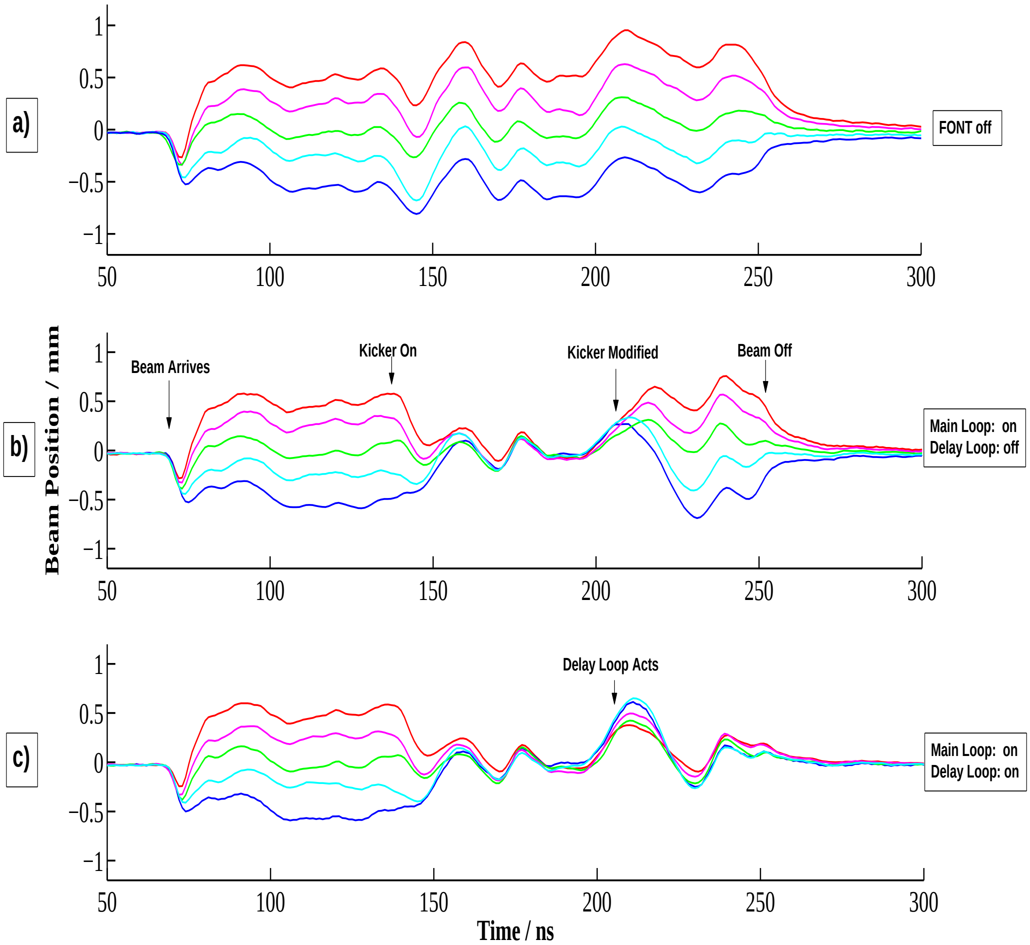

Figure 2.10: Beam position versus time for five beam positions (adapted from [15]); a) Baseline beam position. b) Main loop on, delay loop off. c) Main loop and delay loop on.

2.2.6 Conclusion

FONT1 successfully demonstrated the principle of a fast feedback system to correct the position of a bunch train within a single train, and it was shown that this correction could be performed within the 266ns NLC train length.

Figure 2.10 demonstrated the success of the experiment, but it also shows a specific limitation of this prototype. Since the NLCTA was built to test accelerating RF structures, it was not designed for very high quality beam conditions, and figure 2.10 shows that a static position shape exists on every pulse. This shape is circulated around the feedback loop, and contributes to the structure of the correction signal, making the results harder to interpret. It is this effect that causes the position spread in figure 2.10c just after the delay loop begins to act. For this reason it is desirable to design a system to limit the size of the static position variations.

A further limitation of this demonstration was that the feedback plot shown in figure 2.10 was recorded from the feedback BPM itself, and it would be a stronger demonstration of this concept if the correction were also to be observed on independent witness BPMs.

Also, the kicker amplifier's physical size (see figure 2.9) would prevent its installation in the rather crowded environment of a linear collider interaction point, so to show that this concept is directly applicable to a high energy linear collider it is necessary to scale down the size of the amplifier while maintaining its high output power.

Chapter 3

FONT2

After the success of FONT1 in 2002, several improvements were suggested and put in place to be tested during 2003. These will be discussed now. See figure 3.1 for the layout of FONT2 as compared with FONT1.

3.1 FONT2 Goals

1. Observe >2 corrected periods, compared to the ∼1.5 periods observed in FONT1 (figure 2.10).
2. Install additional BPMs to act as independent witnesses to the position correction.
3. Develop a more elegant BPM processing scheme incorporating real-time charge normalisation.
4. Replace the 'tube' style amplifier used in FONT1 with a solid-state amplifier. This allowed a smaller more compact design[47] that is more suited to installation near the IP of the ILC.
5. Include an algorithm to remove the static position profile observed in FONT1 (figure 2.10).

3.2 Modifications

1. In order to view an additional latency period the feedback BPM and kicker were brought half as close together as before (see figure 3.1). This would halve the latency time due to the beam's time of flight and the signal delay between the BPM and the kicker. The estimated saving due to this was expected to be ∼16 ns (see table 3.1). Latency minimisation was also taken into account during the design and construction of the electronics, and in the arrangement of electronics between the BPM and the kicker. Given the same amplifier power output, the shorter distance between the BPMs and the kicker means that the position change at the feedback BPM due to the kicker will be half as much as in FONT1. Figure 3.1 shows a second kicker with a dedicated amplifier that was added to regain the effective kick.

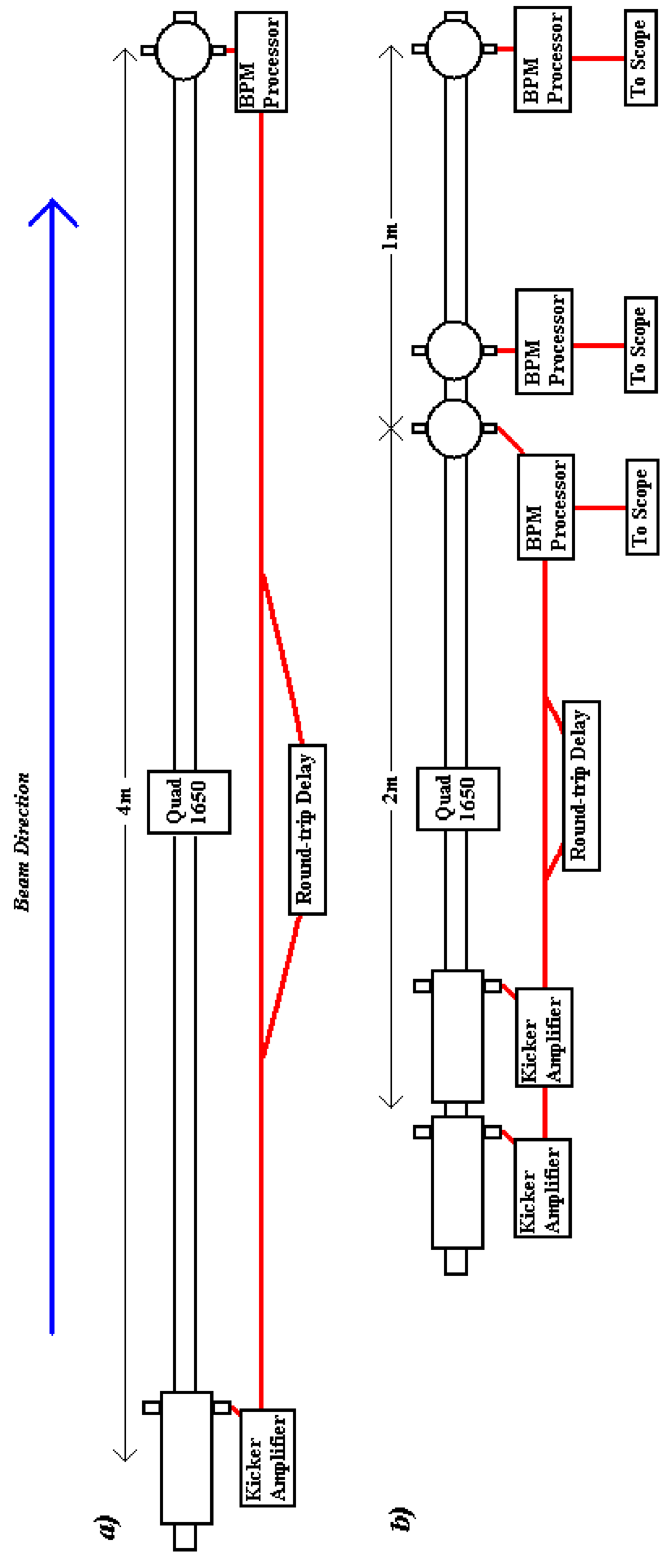

Figure 3.1: Layout of: a) FONT1, and b) FONT2, showing the location of the BPMs and kickers with respect to the fixed beamline component, quadrupole 1650. Also shown is the additional BPMs, kickers, and kicker amplifier.

Latency Component	FONT1 / ns	FONT2 / ns
Beam time of flight	14	6
Signal return delay	18	10
Total irreducible latency	32	16
BPM Processor	5	18
Feedback circuit	16	4
Kicker amplifier	10	12
Kicker fill time	3	3
Total reducible latency	34	37
Total system latency	66	53

Table 3.1: Comparison of the expected delay of each component in FONT1 and FONT2

2. Figure 3.1 shows that two additional BPMs were installed to act as independent witnesses. To minimise the latency the BPM furthest upstream was used for the feedback.

3. A log-ratio BPM processing algorithm (see section 1.6) was implemented so that the position signal would be normalised by the charge from the current pulse, as compared with the charge of a previous pulse as in FONT1. Fast logarithmic amplifiers were used in the processing electronics to obtain a correctly normalised signal in a short period of time. See figure 3.2 for a block diagram of the scheme used[16].

4. Solid state amplifiers were used to apply the drive to the kickers, and these were designed to be physically smaller than the FONT1 amplifiers, but with comparable drive[47].

5. In order to correct the static position variations along the bunch train, a beam flattening algorithm was developed, as shown in equation 3.1.

$$g(t) \equiv -\frac{\sum_{i=1}^{N}\left(f_i(t) - \frac{\int_{t_0}^{t_1} f_i(t)\cdot dt}{(t_1 - t_0)}\right)}{N} \tag{3.1}$$

   where $f_i(t)$ is the function that describes the position variation with time within bunch train i, t_0 and t_1 are defined as the start and end times of the train respectively, and N is the number of measurements. The flattening signal, $g(t)$, is then the negative of the position variations around the mean position of the bunch train, averaged over N pulses.

   $g(t)$ is calculated by digitising the output of the BPM processor for N pulses, and performing the calculation shown in equation 3.1. This signal is then added to the amplifier input using an Arbitrary Waveform Generator (AWG) whose output is timed to coincide with the beam (see figure 3.4).

3.3 The BPMs

3.3.1 Pickoff Tuning

The response of the pickoffs (see figure 2.8) can be maximised by tuning their length to be resonant at the beam bunching frequency, and this procedure will now be described.

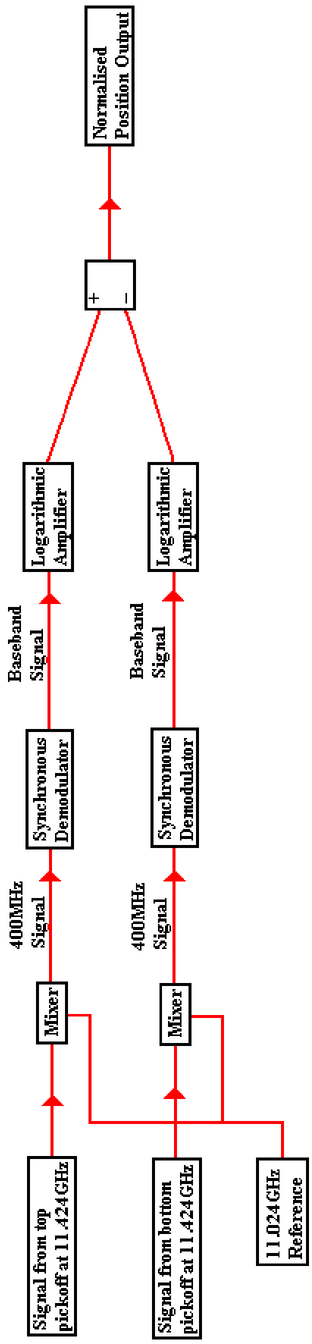

Figure 3.2: Block diagram of BPM processing scheme[16].

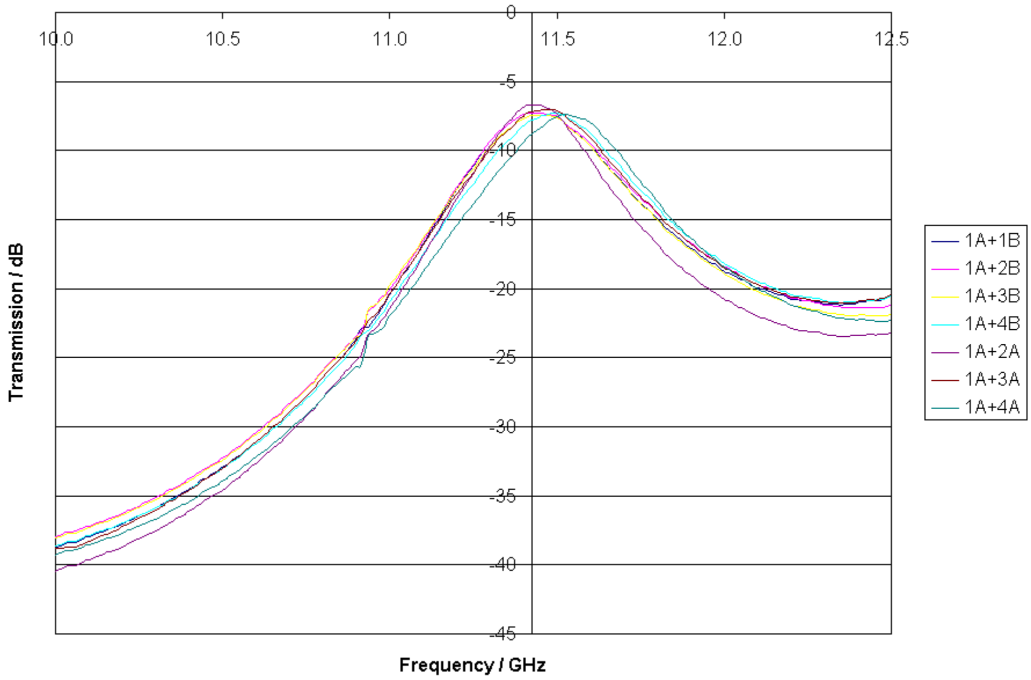

Figure 3.3: Network analyser measurement of the transmission from one pickoff (labelled '1A') to each of the other seven tuned pickoffs. The vertical axis has been drawn at 11.424 GHz to indicate the intended resonant frequency.

The tuning of the BPM pickoffs was performed by slowly filing down the antenna whilst checking its resonant frequency with a network analyser. The analyser was used to input a range of frequencies into the SMA connection of the pickoff and to measure the power received by another pickoff attached to the same BPM cube.

As there were two new BPMs to be installed there were eight pickoffs to tune – two for the horizontal plane and two for the vertical for each BPM. Figure 3.3 shows the results of sending X-band power into one pickoff (labelled '1A'), and being received by each of the other seven. The peaks indicate the pickoff resonant frequency, and one pickoff (4A) appears to be resonant at a higher frequency than required, with a transmission ∼3 dB lower than the others at 11.424 GHz. This pickoff was installed on the horizontal plane of one BPM, since high position resolution would only be necessary for the feedback, which was in the vertical plane[1].

3.3.2 BPM Signal Processing

The BPM signal was processed using the log-ratio scheme (see section 1.6). Figure 3.2 is a block diagram of the processor.

Due to the tuning of the pickoffs the signal emerging from the BPM will contain the majority of its power at ∼11.424 GHz, and this is first translated down to ∼400 MHz using

[1] Horizontal measurements would only be necessary to determine the magnitude of the coupling between the planes. See section 3.5.3

a mixer (see section 4.3.2 for details on mixers and their operation). At its simplest a mixer is a subtraction/addition device, which will add and subtract the frequency of a reference signal (11.024 GHz in this case) from the main input signal. This results in a signal whose main frequency peak is at 400 MHz.

This 400 MHz signal is then passed to a synchronous demodulator to be further mixed to baseband. The synchronous demodulator operates in the same way as the mixer but with a different arrangement of diodes. In this case the reference is a 400 MHz signal, so the output is primarily at baseband.

Next, the signals from both pickoffs are passed through logarithmic amplifiers in order to obtain signals whose levels are proportional to the logarithm of the input amplitude.

The signal from the bottom pickoff is subtracted from the signal from the top in order to obtain the normalised position output.

3.3.3 Experimental Layout

Figure 3.4 shows the experimental layout. In the tunnel, the signal from the feedback BPM was processed through various stages (as described in section 3.3.2) to extract the beam position signal, which was then sent through the feedback delay electronics and amplifiers to the kickers. The raw pickoff signals from the three BPMs (after a splitter in the case of the feedback BPM) were sent outside the tunnel to identical processing electronics. The processed signal was digitised by a scope and transferred via GPIB[48] to the Matlab[19] DAQ system. Matlab was used to calculate the beam flattening signal from the output of BPM 1 (equation 3.1), and this was sent back into the tunnel to be added after the delay electronics.

A fourth scope channel was used to gather charge profiles from BPM 1, while charge information from other BPMs, including information on x position, was sent to VME[49] digitisers.

3.3.4 BPM Calibration

The BPMs were calibrated by moving the beam to several positions with an upstream dipole, and comparing their output with the output of a previously calibrated BPM (see figure 3.5). In this figure the red lines show two alternative beam trajectories through the FONT system (only one FONT BPM has been shown, however the equivalent calculation can be obtained for the other two by analogy) and on to a calibrated NLCTA stripline BPM. It is assumed that the beam trajectories are the same upstream of the dipole magnet, and that the change (represented by θ) is only caused by changing the voltage across the dipole. It is also assumed that the position read out by the NLCTA BPM is correct.

If there is a voltage, V_D, across the dipole, there will be corresponding position readings, y_f and y_n, at the FONT and NLCTA BPMs respectively. At first, however, only a voltage, V_f, will be read out from the FONT BPM, and this must be converted to a position using a coefficient, K_f (units of mm/V). The distances, z_f and z_n, are the distances from the

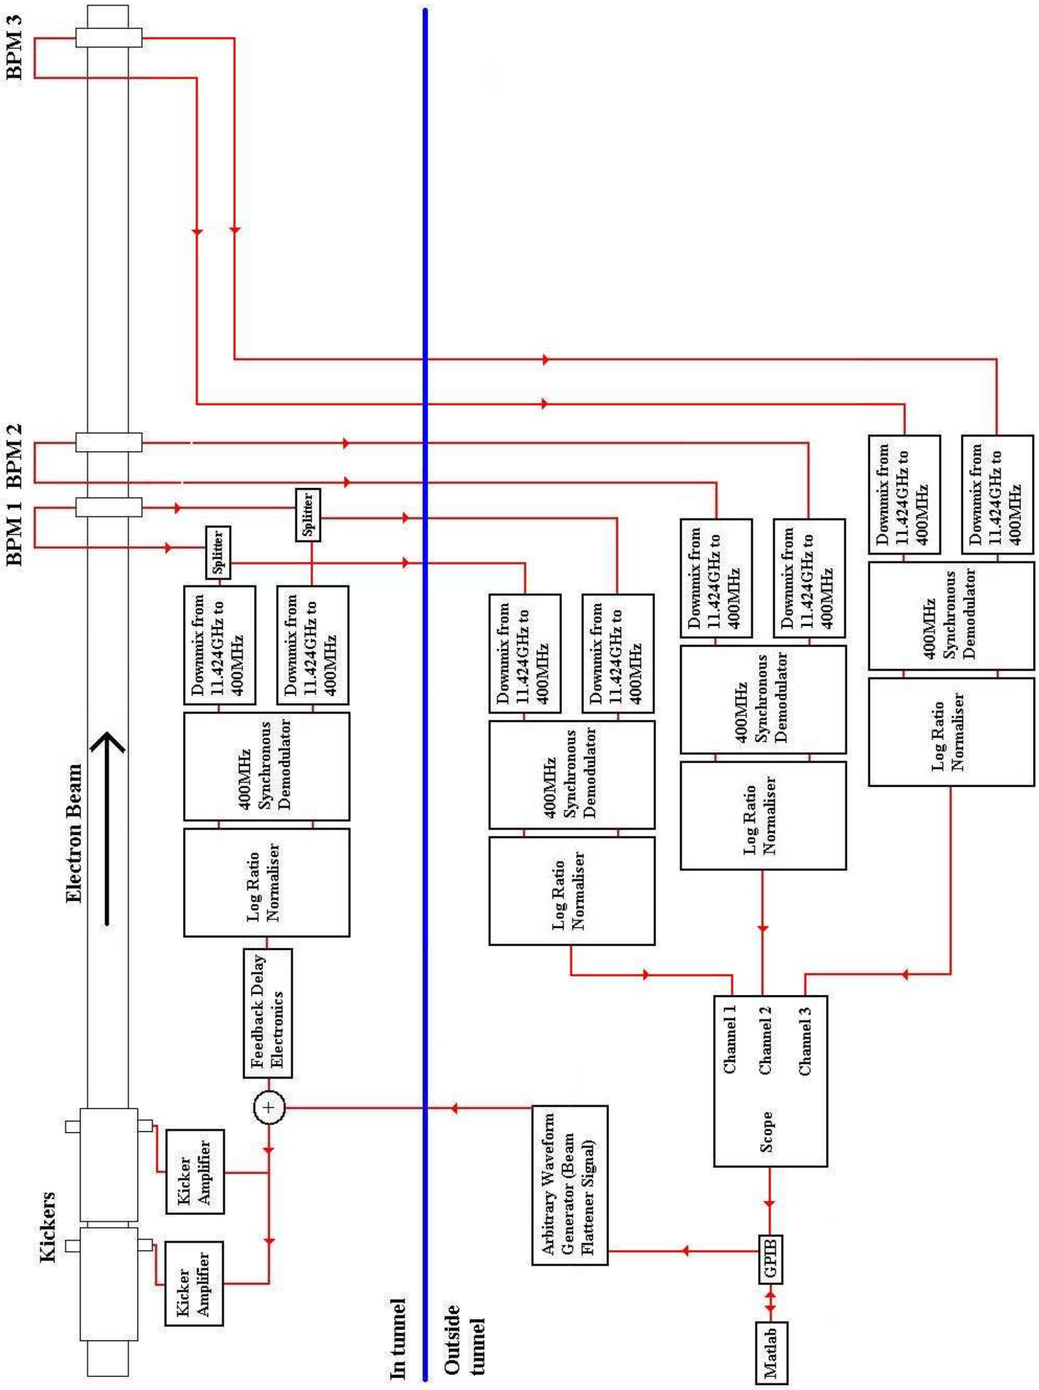

Figure 3.4: Block diagram showing the layout of the apparatus.

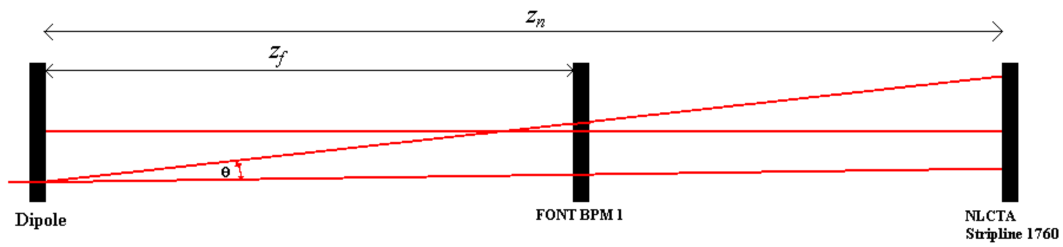

Figure 3.5: Beam trajectory from the dipole magnet through a FONT BPM and then to a NLCTA stripline BPM.

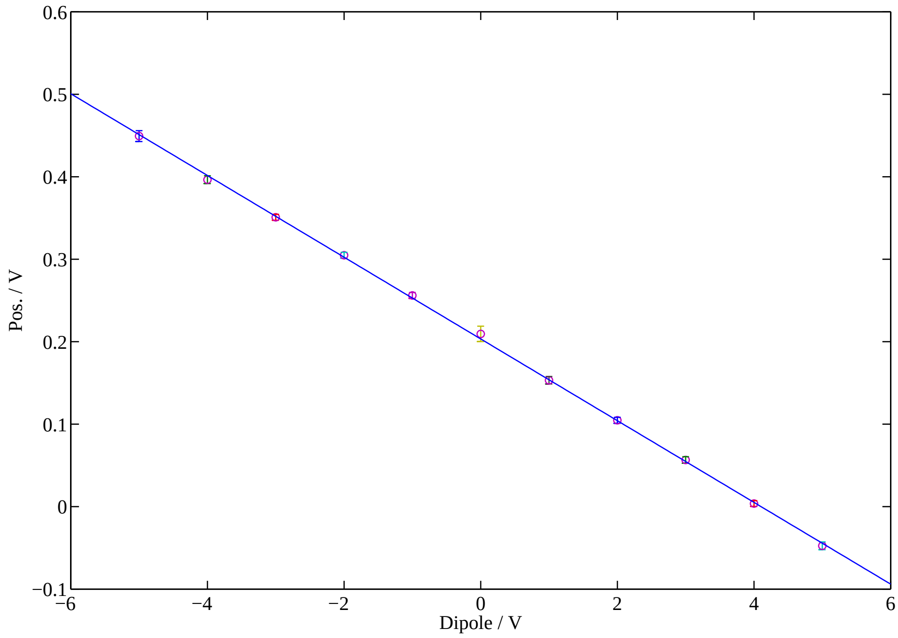

Figure 3.6: Voltage output vs. dipole voltage for BPM 1

dipole to the FONT and NLCTA BPMs respectively.

$$\Delta y_f = \frac{z_f}{z_n} \Delta y_n \tag{3.2}$$

where Δy indicates the change in position caused by a change, ΔV_D, in the dipole voltage.

$$\Delta y_f = K_f \Delta V_f \tag{3.3}$$

where ΔV_f is the change in output voltage of the FONT BPM caused by the change in the dipole voltage.

Equation 3.2 is calculated using similar triangles (see figure 3.5), and equation 3.3 assumes that the voltage output is proportional to the beam position.

The dipole was swept through a range of voltages and the voltage output from the FONT

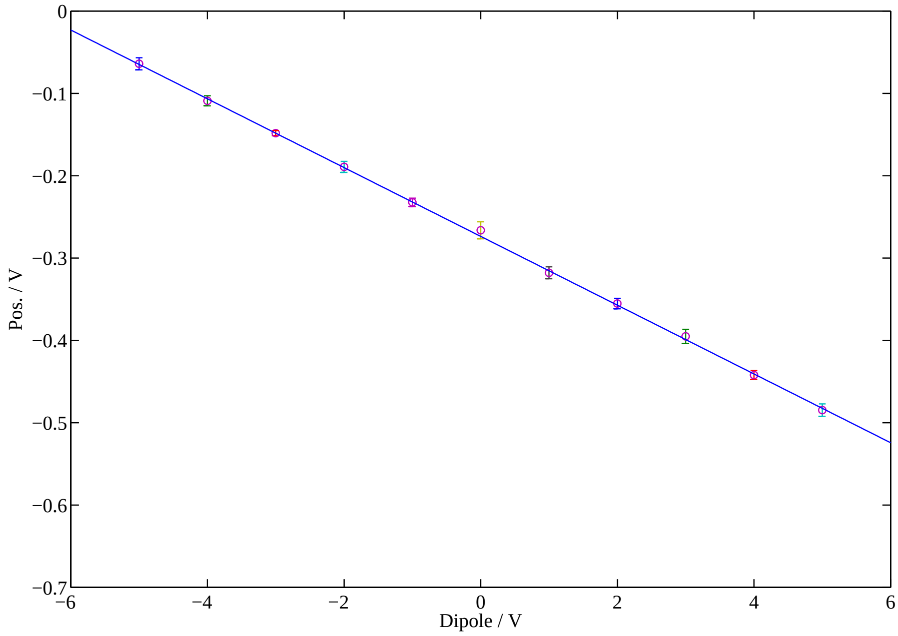

Figure 3.7: Voltage output vs. dipole voltage for BPM 2

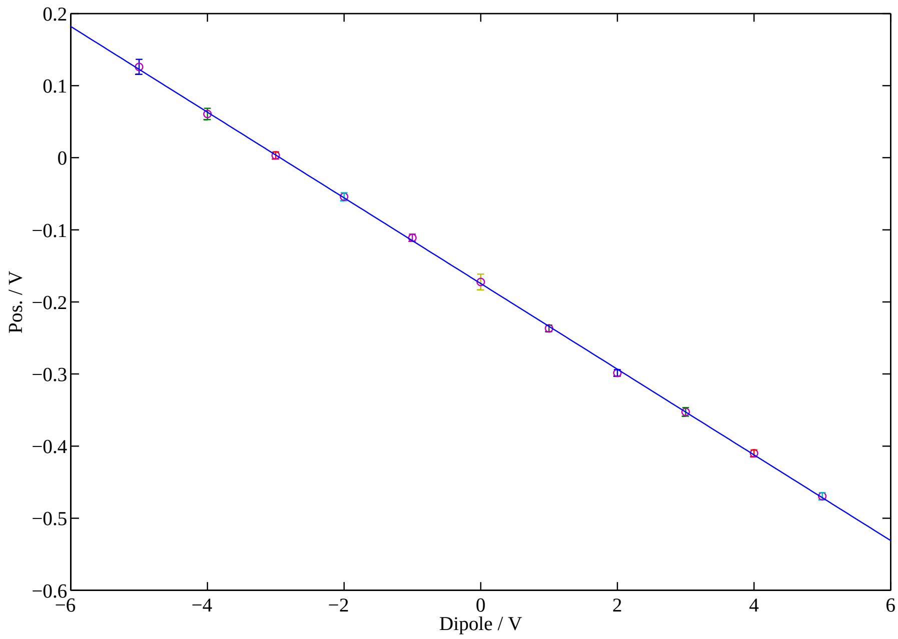

Figure 3.8: Voltage output vs. dipole voltage for BPM 3

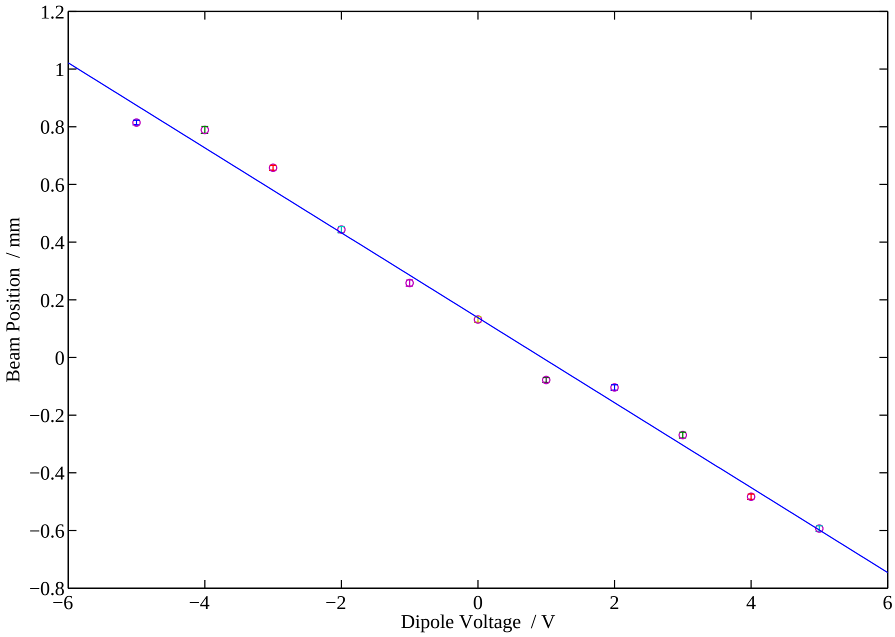

Figure 3.9: Position output vs. dipole voltage for the calibrated NLCTA BPM. The error bars are the statistical errors reported by the control system. The Chi-squared per degree of freedom, $X^2/N = 32.3$. It is known that, in the case of multibunch beam, these BPMs respond to the first bunch of the train, and the trajectory or charge of this bunch may not be typical of the remainder of the beam, and it is this that causes the scatter of the points around the linear fit.

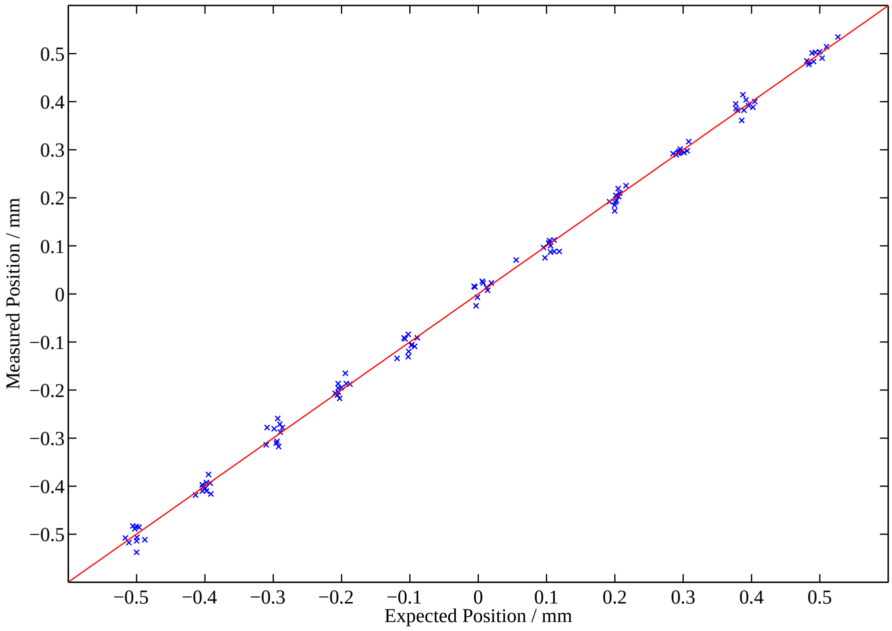

Figure 3.10: Actual position vs. expected position at BPM 2

BPM and the stripline position output were plotted against this (see figures 3.6 and 3.9). These show that, for a change in dipole voltage of +1 V, the output from BPM 1 will change by -0.0495 V and the position at the NLCTA BPM will change by -0.149 mm. Figures 3.7 and 3.8 show the equivalent data for BPMs 2 and 3. Thus K_f was calculated for each BPM with the following results,

- $K_1 = 1.89 \pm 0.03$ mm/V
- $K_2 = 2.38 \pm 0.03$ mm/V
- $K_3 = 2.10 \pm 0.03$ mm/V

3.3.5 Resolution

Using the calibration numbers it is possible to measure the BPM resolution in the following way.

If the position of the beam was measured at BPM 1 and 3 it was possible to calculate the expected position at BPM 2, as long as careful use was made of data acquisition triggers to make sure that it was the same beam pulse that was being measured by all BPMs each time.

The dipole was swept through a range of voltages while the beam position was recorded at each of the BPMs. Ten pulses were recorded at each dipole setting. Figure 3.10 is a plot of the actual position read out from BPM 2 against the position calculated from the measurements at BPMs 1 and 3, and it can be seen to be a straight line with a gradient of 1

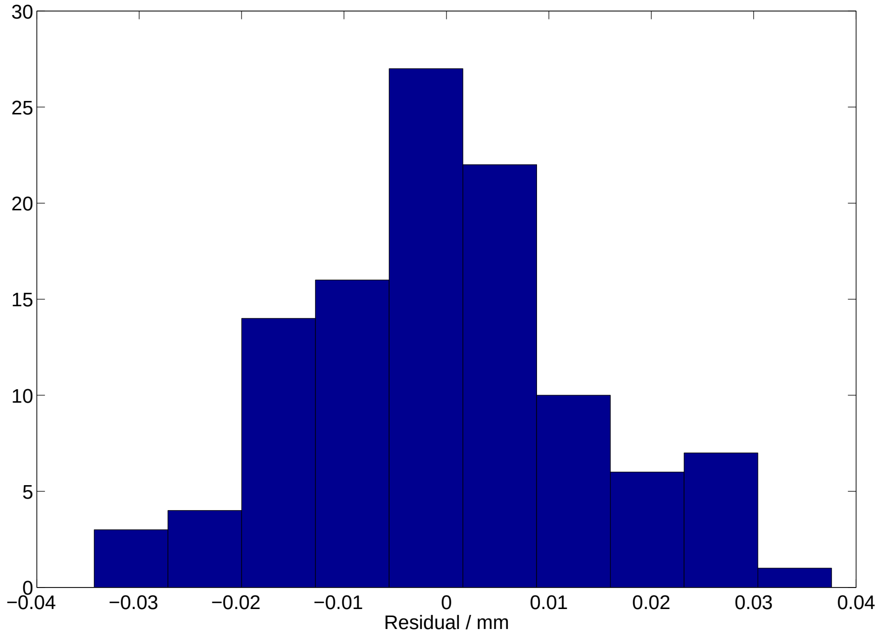

Figure 3.11: Distribution of the residuals, y_{res}.

and a zero intercept. Each dipole setting can be clearly seen as a cluster of points, and it is the spread of each of these clusters with respect to their mean that contains the resolution information. If the measured position is y_{meas} and the calculated position is y_{calc}, then the residual, y_{res}, is defined as,

$$y_{res} \equiv y_{meas} - y_{calc} \tag{3.4}$$

Figure 3.11 shows the distribution of the residuals. The mean of these, μ=0 μm, and the standard deviation, $\sigma_{y_{res}} = 14.2\ \mu m$. The standard deviation contains resolution information from all three BPMs, but the resolution of an individual BPM can be calculated by assuming that the noise in each BPM is uncorrelated, and that the resolution of each BPM is the same.

Using these two assumptions, it can be calculated[16] that,

$$\sigma_{BPM} = 77\% \times \sigma_{y_{res}} \tag{3.5}$$

Therefore the FONT2 BPMs have a resolution of ~10.9 μm, which is consistent with the expected value[16].

3.4 Adjusting System Settings

The FONT system had various settings that could be controlled externally in order to improve its performance. These settings, and the methods used to find their optimal values, will be described now.

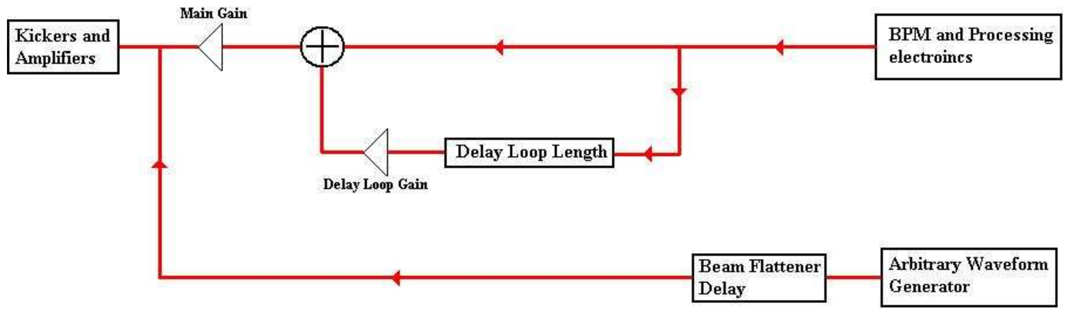

Figure 3.12: Block diagram of the basic signal path, including the parameters to be optimised.

3.4.1 Main Gain

See figure 3.12. The main gain controls the strength of the correction, and must be optimised to prevent the beam being over or under-corrected.

To optimise this, the beam was steered to nine different positions and ten pulses were recorded at each setting. The main loop was on in order to observe the effect of the corrections, and the delay loop was also on. Figure 3.13 shows beam position against time for the average of each of the ten pulses, and for a range of gains. In each case figure 3.13 shows the beam behaving as expected. For gains of 0.5 and 0.7 (figure 3.13a and 3.13b) the beam is being under-corrected, implying that the gain is too low. With gains of 1.1 and 1.3 (figure 3.13d and 3.13e) the correction is too large, with some evidence of the system being in a mode of damped oscillations.

A gain of 0.9 (figure 3.13c), however, seems to be reasonably optimised, as the beam is steered to nominal independent of the initial offset. There is no evidence of under- or over-damping.

In order to further investigate the system performance the effect of negative values of the main gain was studied in the same way. See figure 3.14.

For negative gain the beam would be expected to move further from zero with each latency period, as the negative value will cause a kick with the opposite polarity to that required to correct the beam. This behaviour can been seen in figure 3.14. For a setting of -0.8 (figure 3.14a) the curves only diverge slightly. As the gain decreases further, however, the divergence can be seen to increase. The final latency periods for gain settings -1.6 and -2.0 (figure 3.14c and 3.14d) look very similar due to the fact that the beam is so far from the zero position that the amplifiers, in attempting to apply the large negative correction, have saturated, and are unable to kick the beam any further.

3.4.2 Latency Measurement

The latency is measured as follows. First the dipole is moved through several voltage settings to introduce several offsets into the beam. The position profile of the bunch train is first measured with the main loop off to define a baseline, and is then measured with the main loop on (delay loop off). When the baseline pulses are subtracted from the pulses with the feedback on, the start and end of each period should be clearly visible. See figure 3.15.

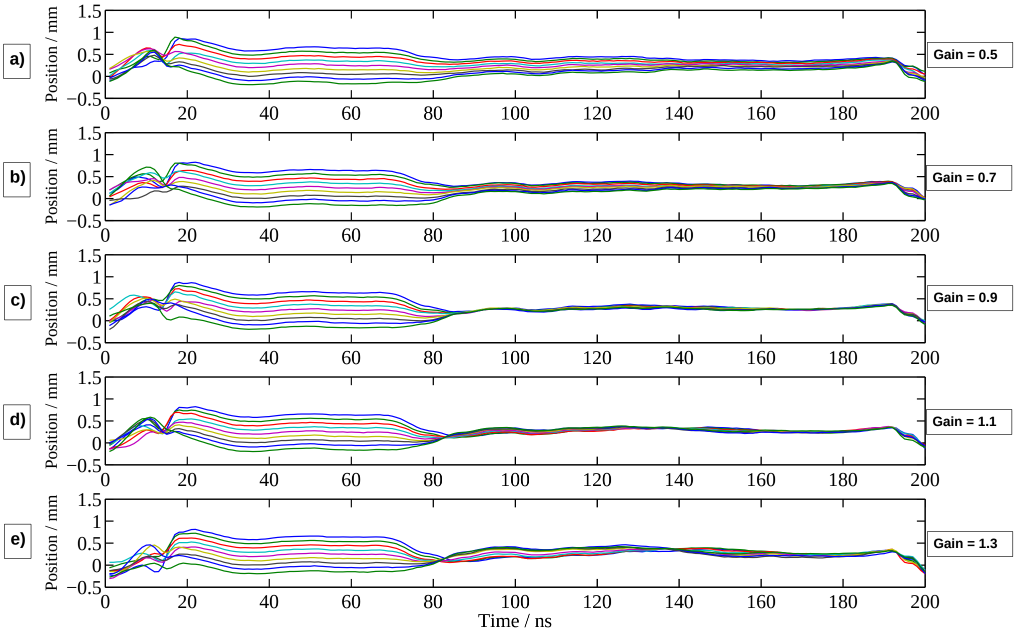

Figure 3.13: Beam position vs. time for BPM 1. The beam flattener, and delay loop were off, and the main gain was set to the following values: a) 0.5, b) 0.7, c) 0.9, d) 1.1, e) 1.3.

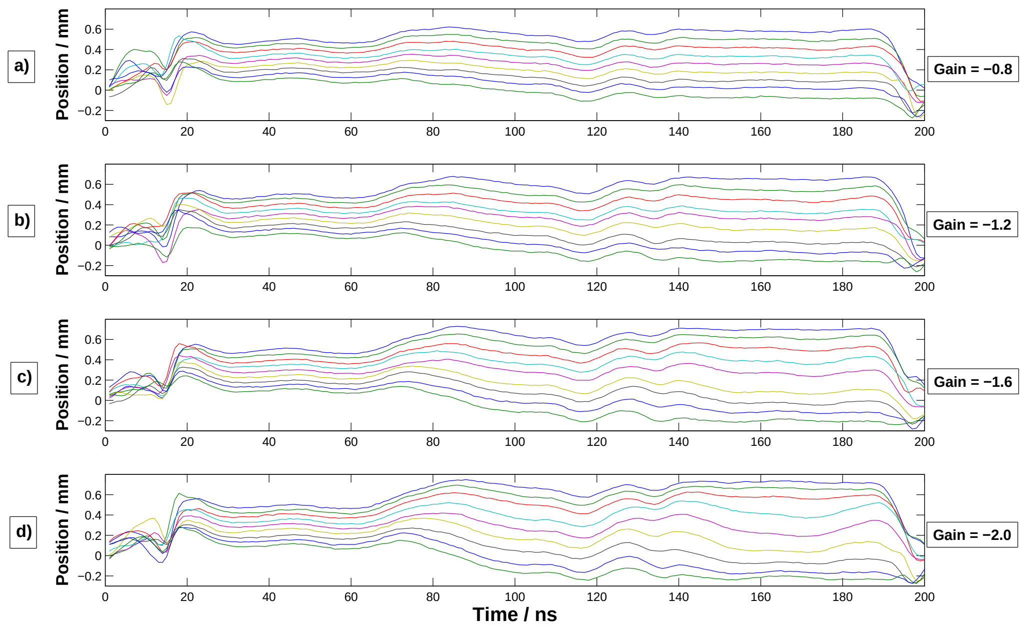

Figure 3.14: Beam position vs. time for BPM 1. The beam flattener, and delay loop were off, and the main gain was set to the following values: a) -0.8, b) -1.2, c) -1.6, d) -2.0.

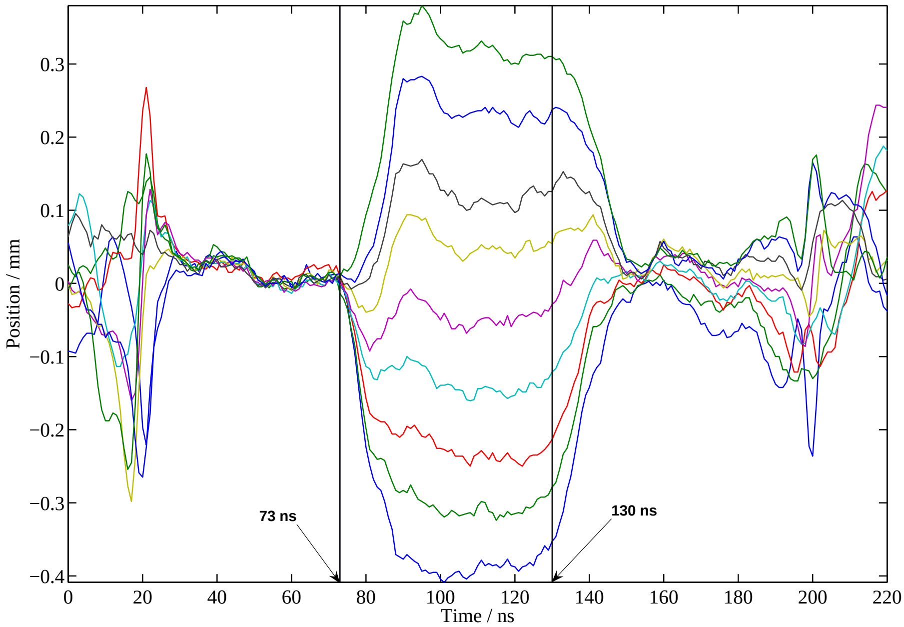

Figure 3.15: The difference between the beam's baseline position, and the position when the main FONT loop is on for nine dipole settings. The beam flattener and delay loop are off, and each curve is the average of ten pulses.

The onset of the kicker pulse is clearly visible at t=73 ns. The time at which the beam position begins to return to the original offset is more difficult to identify as some of the high frequency position variations are still visible despite the baseline subtraction. The vertical line marked at 130 ns serves as a reasonable estimation.

This gives a latency of ∼57 ns. This is 4 ns longer than predicted (see table 3.1), but this can be accounted for by additional cable not included in the prediction calculations. The original estimation assumed that the cable would travel in a straight line between the BPM outputs and the kicker, however this was not possible due to various beamline obstructions (e.g. quadrupole magnets, vacuum pumps, etc.). Also, the electronics had to be shielded from beam-induced radiation by several centimetres of lead, making it impossible for the cable to follow a straight line. 4 ns is equivalent to ∼0.72 m of additional cable (where the signal propagates at $\sim 0.6c$), and this amount of excess cable is consistent with what was used.

3.4.3 Delay Loop Length

See figure 3.12.

The system setup included a fixed amount of delay loop, as well as a control box that could select discrete lengths of cable to be added to this, where the additional length is controlled using four binary switches on the control box. It was estimated that the minimum delay possible with the loop was 36 ns, and that each switch adds ∼3.5 ns to the length of

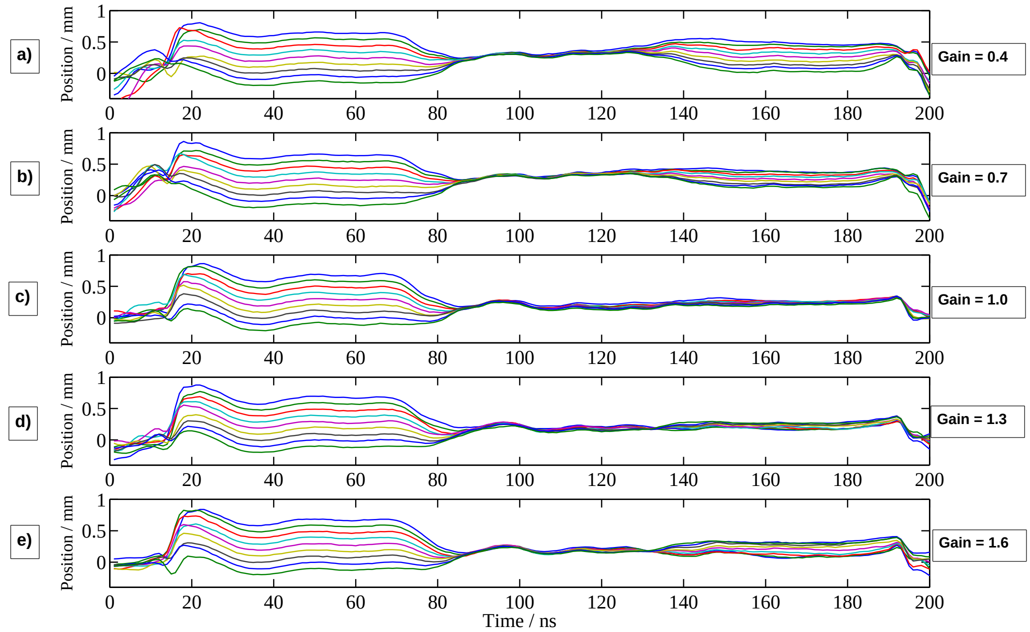

Figure 3.16: Beam position vs. time for BPM 1. The beam flattener was off, and the main gain and delay loop were on. The delay loop length was set to the value found in section 3.4.2, and its gain was set to the following values: a) 0.4, b) 0.7, c) 1.0, d) 1.3, e) 1.6.

the delay. The latency was measured and found to be ∼57 ns in section 3.4.2, and the delay loop was set to this value.

3.4.4 Delay Loop Gain

See figure 3.12.

The control box also allowed external control of the delay loop gain. Formally this should be unity, however it was decided to verify this with beam tests. The delay loop gain was optimised in a similar way to the main gain. A dipole scan was performed with the main loop and delay loop on, for various values of the delay loop gain. Ten pulses were averaged for each dipole setting, and the results are shown in figure 3.16.

It can be clearly seen (fig 3.16a and 3.16b) that 0.4 and 0.7 are too low, and do not apply the delayed signal to a sufficient degree to keep the beam at its nominal position. Gains of 1.3 and 1.6 (figure 3.16d and 3.16e) cause the magnitude of the delay loop signal to be too large, steering the beam through nominal. As expected the value of 1.0 (figure 3.16c) seems to be optimal as there is no apparent variation when the third latency period begins.

3.4.5 Timing of the Beam Flattener

The beam flattening signal (equation 3.1) was output by an AWG triggered to coincide with the beam. The AWG timing had to be synchronised with the beam, so a delay had to be

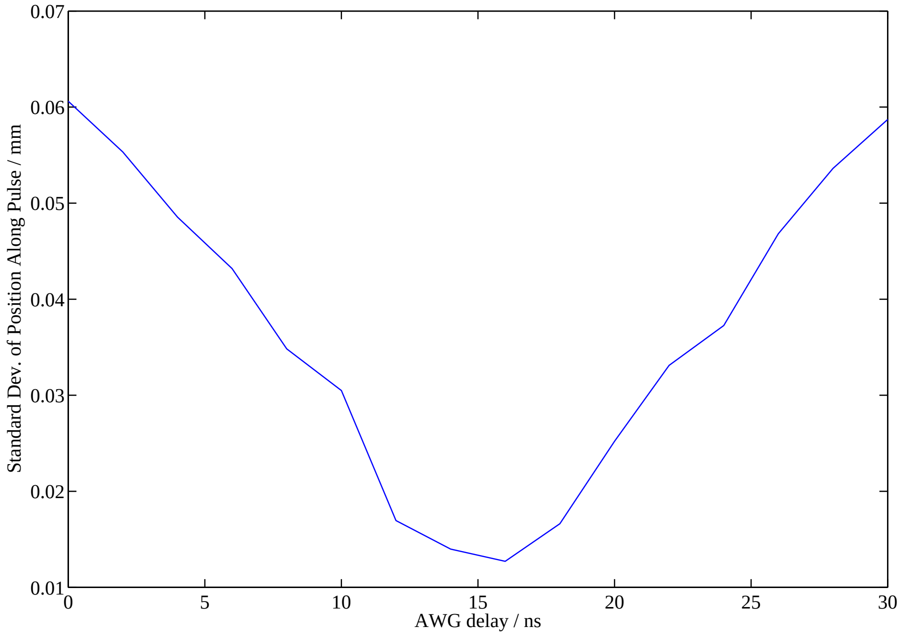

Figure 3.17: Standard deviation of the position along the pulse (measured at BPM 1) vs the delay between the AWG trigger and the output of the flattening signal.

set between the AWG receiving the trigger and outputting the pulse, and this delay had to be measured.

By varying the delay using the front panel of the AWG and observing the beam signal on a scope it was possible to narrow down the time to within a 15 ns window.

The beam signal was then carefully recorded as the delay time was swept finely through this window. For each delay time the standard deviation of the position along the beam pulse was calculated, and the results can be seen in figure 3.17. There is a minimum at a delay of 16ns, indicating that this delay is optimal in achieving a flat beam.

3.5 Results

3.5.1 Results from the Feedback BPM

Figure 3.18 demonstrates the best quality data taken with FONT2. These data are from BPM 1 using the external electronics (see figure 3.4) and were produced in the following way.

The dipole was swept through nine settings, taking ten pulses at each setting. The beam was first measured with the main loop and delay loop off in order to observe the baseline beam position. Next a dataset was taken to show the effect of the flattener. Then a third dataset was taken with both the flattener and the main FONT loop active, but without the delay loop. Finally the dipole sweep was performed with the flattener, main loop, and delay loop all on. Each of these four datasets are shown in figure 3.18, and each curve in this

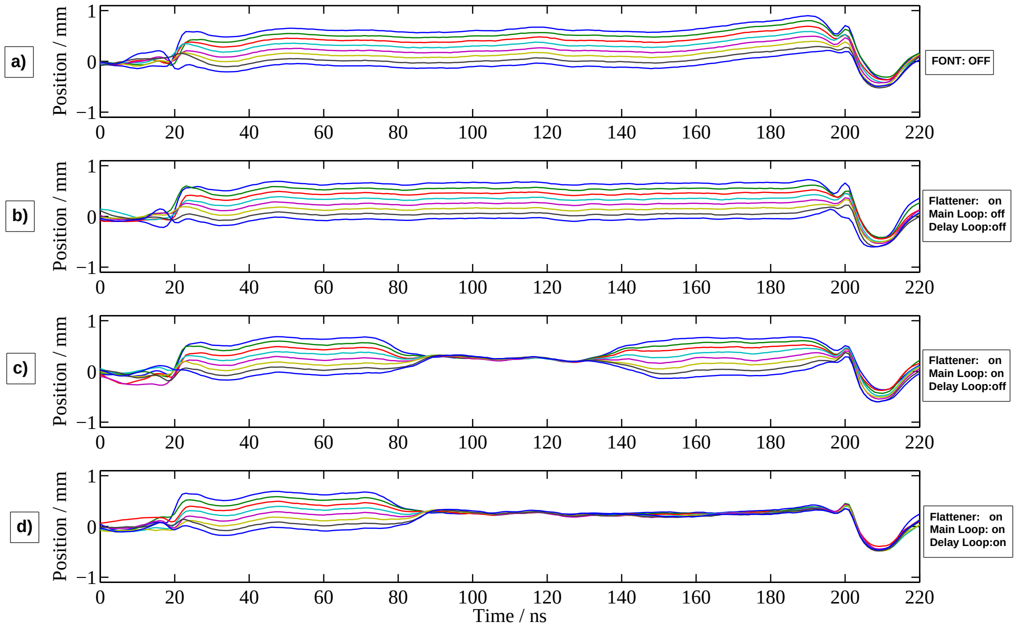

Figure 3.18: Beam position measured by BPM 1 vs. time for eight different dipole settings. a) Baseline beam position, b) Beam flattener on, c) Beam flattener, and main loop on, d) Beam flattener, main loop, and delay loop on.

figure corresponds to the average of the ten pulses.

In figure 3.18a a smooth slope can be seen over the last ∼40 ns of the bunch train, as well as faster position structures at the beginning and end. Figure 3.18b shows that the flattener has removed the slope, however it had difficulty with the fast structures at the start and end due to these being slightly faster than the 50 MHz bandwidth of the kicker amplifiers.

In figure 3.18c the various position offsets due to the dipole can be seen to converge to the same value the first time the kicker acts. The position the beam converges to is not zero on this plot as the position signal is being calculated by a set of electronics outside the beam-line tunnel whose offset will be different to that of the feedback processor due to slightly different inherent gains, attenuations, etc. At the end of the second latency period, the kicker switches off due to the absence of the delay loop, and the beam returns to the original offset.

Figure 3.18d shows all the FONT2 systems in operation. The beam corrects after the first latency period, and the delay loop begins to operate at the end of the second period, maintaining the beam at its nominal position.

Figure 3.19 shows a magnified version of figure 3.18d, and figure 3.20 shows the position spread of this plot at each point in time. It can be seen from this that a position spread of ∼0.7 mm is corrected to ∼0.05 mm, which is a 14:1 correction.

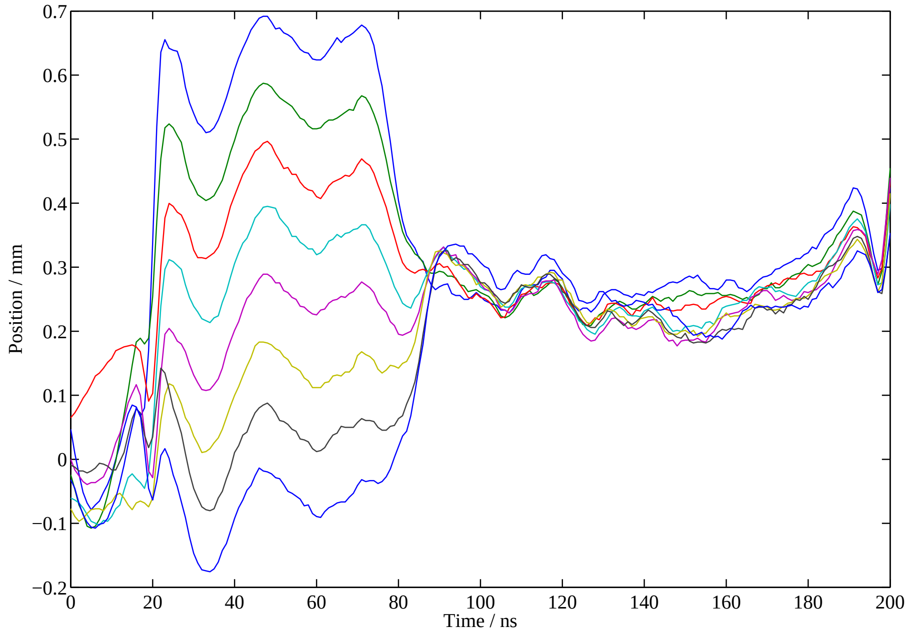

Figure 3.19: Beam position vs. time for eight different dipole settings. The flattener, main loop, and delay loop are all on.

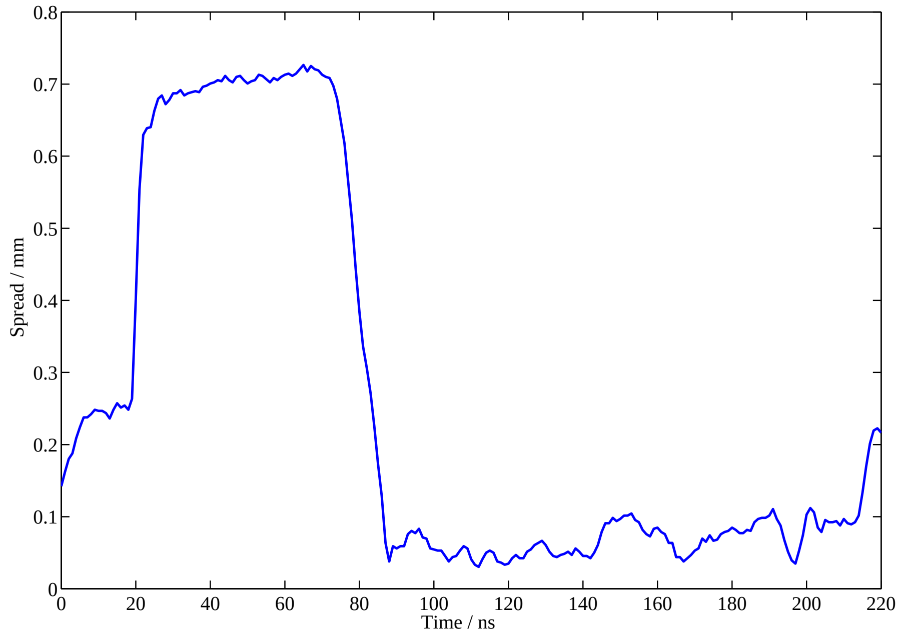

Figure 3.20: Position spread from figure 3.19.

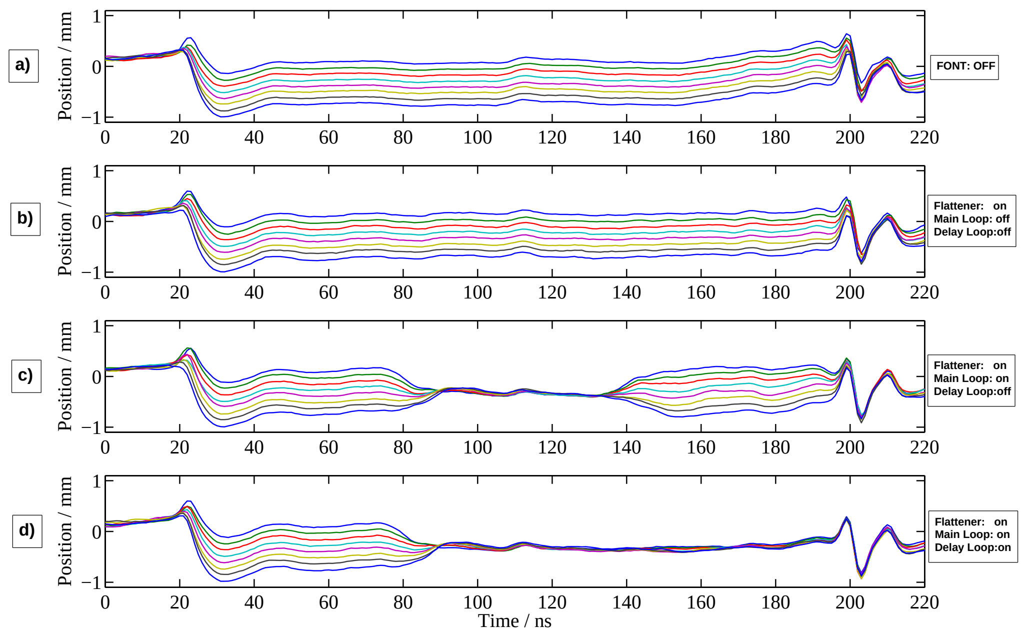

Figure 3.21: Beam position measured by BPM 2 vs. time for eight different dipole settings. a) Baseline beam position, b) Beam flattener on, c) Beam flattener, and main loop on, d) Beam flattener, main loop, and delay loop on.

3.5.2 Results from the Witness BPMs

Figure 3.21 shows the data from BPM 2 for the same dataset as shown in figure 3.18.

All the features described for the feedback BPM can be clearly seen in these plots, demonstrating that the system is operating as planned. The position correction is provided by the addition of an angle to the beam's trajectory at the kickers, resulting in a motion at the feedback BPM of the correct magnitude to bring the beam to zero. At other betatron phases, however, the position change will be of a different magnitude, and this is the case for BPM 2. There were no magnetic elements between the feedback BPM and witnesses, so the beam moved in a straight line, and the beam position will be diverging due to the extended lever-arm. This can be observed in figures 3.21c and 3.21d as a slightly greater amount of spread in the first correction period than in figures 3.18c and 3.18d.

Figure 3.22 shows the data from BPM 3, and the divergence is more pronounced as should be expected due to the larger distance to this BPM.

3.5.3 Coupling Between Planes

As the x-plane BPM pickoffs were also instrumented it was possible to measure how far the beam motion coupled between the planes. Figure 3.23 shows the main feedback dataset as measured in the x plane of BPM 3. In each plot all the curves overlap, implying that the dipole is correctly aligned within the resolution of the BPMs. Also, no difference is seen when the main FONT loop is on or off, or when the delay loop is on or off, which implies

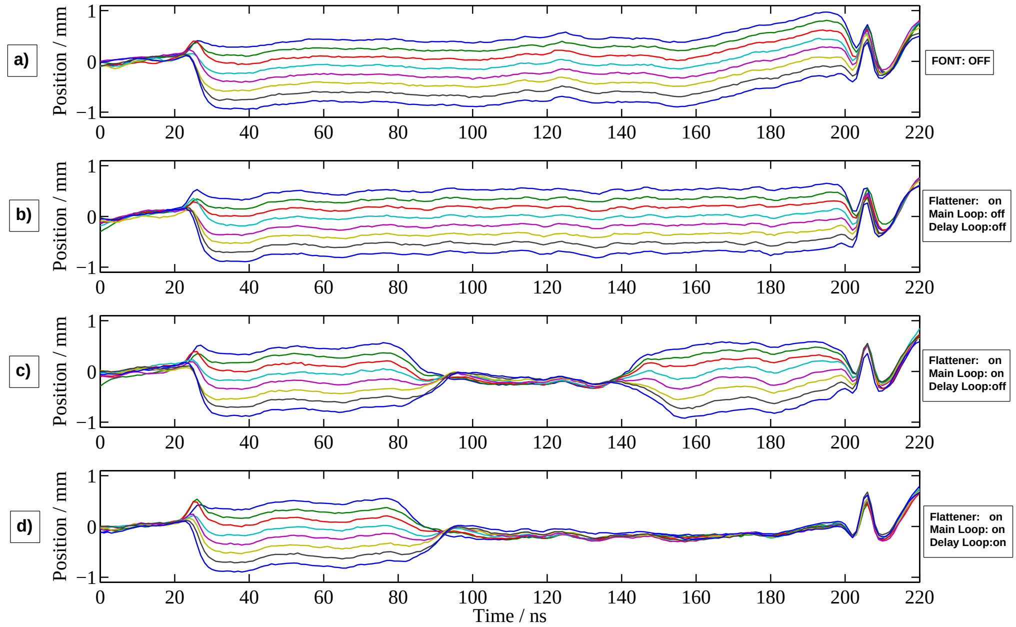

Figure 3.22: Beam position measured by BPM 3 vs. time for eight different dipole settings. a) Baseline beam position, b) Beam flattener on, c) Beam flattener, and main loop on, d) Beam flattener, main loop, and delay loop on.

that, within the resolution of the BPMs, each component is very well aligned.

If the components were misaligned by, say 0.5°, this would show up as a 1% coupling into x, or about 0.005 V difference between the most widely separated dipole settings. This is about the limit of the resolution of figure 3.23, and it can therefore be concluded that the alignment is equal to or better than this.

3.5.4 4th Latency Period

Using the measured value of 57 ns for the latency it can be seen that the fourth period should begin at ∼171 ns. As the beam pulse is only ∼177 ns long (see table 2.2) it is clear that only the ramp up of the kickers will be visible for the fourth latency, and in figures 3.18, 3.21, and 3.22 there is no strong evidence of a fourth latency period.

The following section details changes that were made in order to reduce the latency to a value that will allow observation of this final period.

3.6 Reduced Latency Setup

3.6.1 Changing the Beamline configuration

Figure 3.24 shows how the beamline setup was reconfigured in an attempt to reduce the latency further. Figure 3.24a shows the original setup, with the signal from the top and

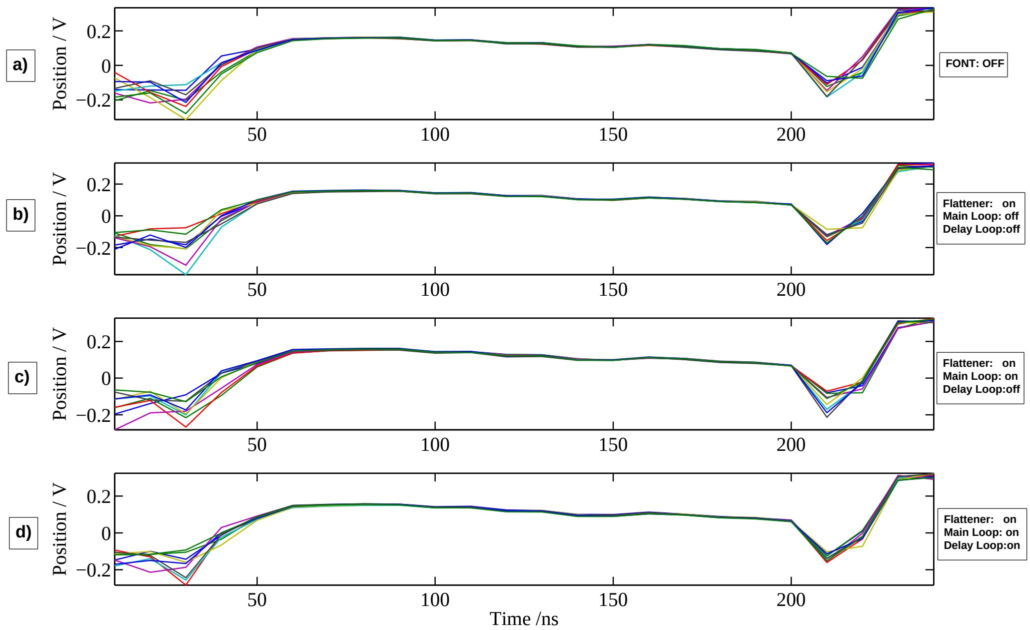

Figure 3.23: x-plane position data from BPM 3 vs. time for eight different dipole settings. a) Baseline beam position, b) Beam flattener on, c) Beam flattener, and main loop on, d) Beam flattener, main loop, and delay loop on.

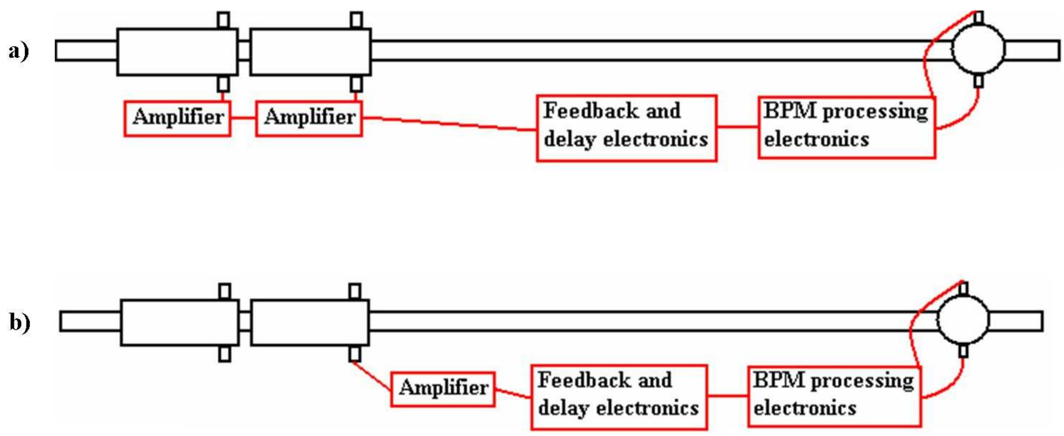

Figure 3.24: Original setup, (a), compared with the new setup, (b).

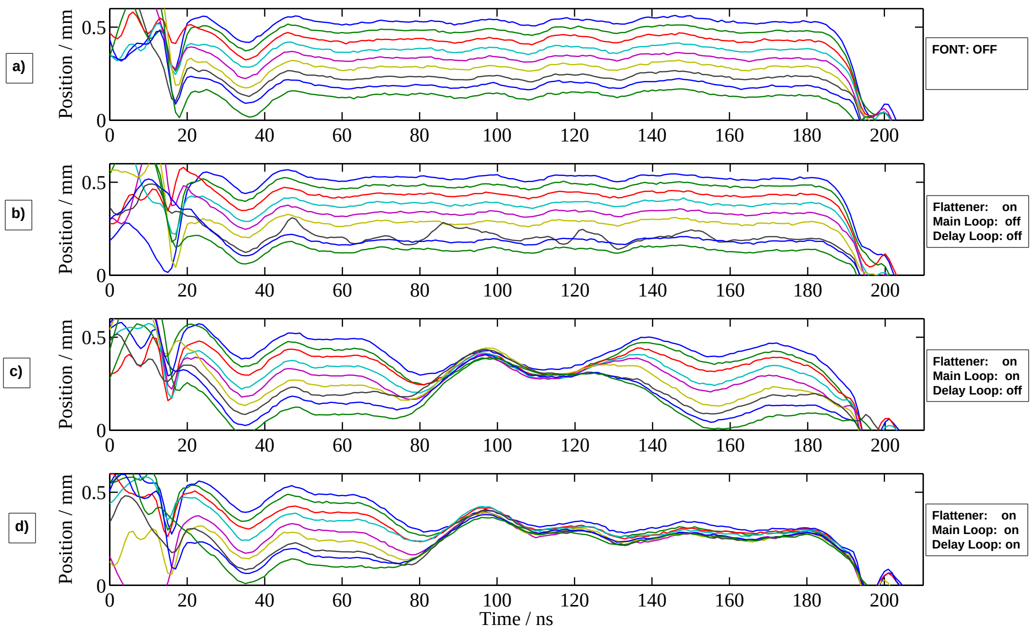

Figure 3.25: Feedback BPM data taken using the reduced latency setup. The behaviour of the black curve in plot (b) is due to the inclusion of a spurious pulse into the analysis.

bottom of the feedback BPM leading to the processing electronics, then into the feedback and delay box, and finally to the kicker amplifiers. It was seen that the time of flight of the signals could be reduced by disconnecting the upstream amplifier, and moving the downstream one to a position closer to the straight line between the feedback box and the kicker input. This can be seen in the figure 3.24b. In doing this a 1.5 ns length of cable was removed from the signal path. This also gave a beam time-of-flight time saving as the upstream kicker was now no longer part of the system. It was estimated that, in total, 3 ns could be saved on the latency – i.e the latency would become 54 ns.

The removal of a kicker caused half of the correction dynamic range to be lost. Also, it was necessary to increase the main gain by a factor of two in order to compensate for the smaller kick.

3.6.2 Results from the Reduced Latency Setup

Figure 3.25 shows the system operating with this reduced latency setup. Figure 3.25c shows the beam flattener and the main loop on, but with no delay loop, and the first three latency periods are clearly visible. When compared with the equivalent plot from the longer latency system (figure 3.18c) there is evidence of the start of a fourth latency period.

At the end of the third period the pulses are beginning to converge again. This is quite clear when compared to the baseline or beam flattener datasets, where the beam cutoff is very sharp.

In order to demonstrate this more clearly, figure 3.26 was produced. This shows ten pulses

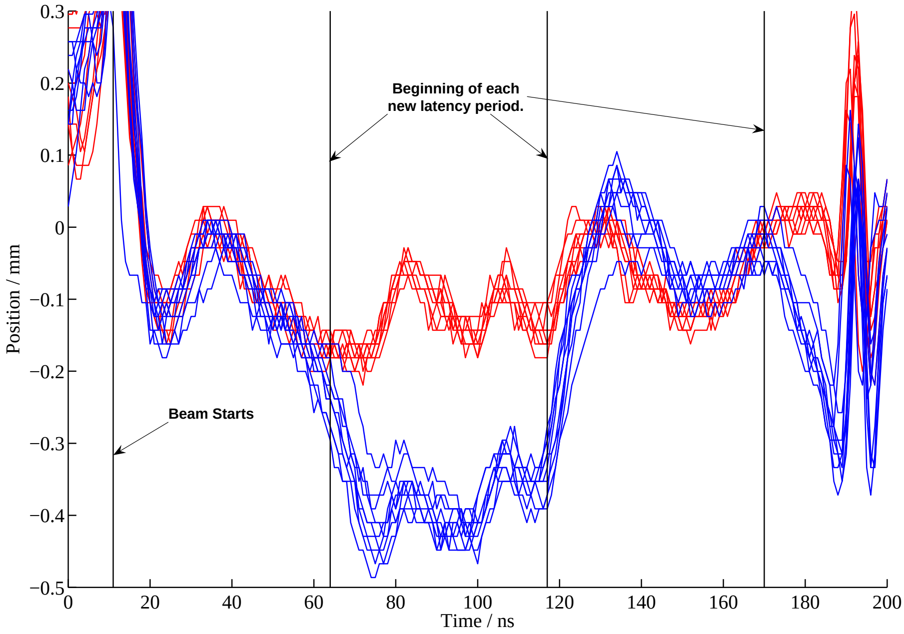

Figure 3.26: Ten un-averaged pulses showing the action of the main loop(blue) with ten baseline pulses overlaid (red). The fourth latency period is evident in the blue curves after t=170 ns as the correction kick begins.

with the main loop acting on the beam (blue), compared with ten baseline pulses (red). The vertical lines correspond to the start and end of the beam pulse, and the beginning of each latency period. The beginning of the fourth latency is apparent at t=170 ns where the beam begins to move back to the nominal position. From this plot, the latency was measured to be ∼53 ns.

3.7 Summary

To summarise, various changes were suggested to upgrade FONT1[15] to FONT2 with the aim of improving and further investigating the performance of the system.

1. Beamline components were brought closer together to lower the time taken by the time-of-flight of the beam.
2. Addition of second kicker to compensate for the loss of dynamic range.
3. Use of solid-state kicker amplifiers.
4. Addition of independent witness BPMs.
5. Real time charge normalisation in BPM processing electronics.
6. A system to flatten static position variations along the bunch train.

Due to all these improvements the following were observed,

1. Correction ratio of $\sim 14 : 1$ (figure 3.20).
2. Beam correction visible on independent BPMs (figure 3.21 and 3.22).
3. System latency of 53 ns (figure 3.26).
4. Additional latency period (figures 3.25, and 3.26).
5. Correct operation of the beam flattener (figures 3.17, 3.18 and 3.21).

Chapter 4

FONT3 – Experimental Goals and BPM Processor Design

In order to build a prototype that would more closely match the system designed for use in a warm linear collider, it was decided that the experiment should be moved to a machine with an energy of $\sim$1 GeV. If a correction with micron-level accuracy could be performed on a beam of this energy, then, given the same kicker-IP distance, this scales to a correction of nanometre-level accuracy of a $\sim$1 TeV beam as will exist at the ILC, thereby allowing the use of the same components in the final system.

4.1 Accelerator Test Facility (ATF), KEK

The ATF[17] was built at KEK[50] as a prototype damping ring for the Linear Collider in order to demonstrate the extremely small emittance required to achieve the planned luminosity. An ultra-low normalised vertical emittance of $\sim 6 \times 10^{-12}$ m·rad was achieved in 2003 for a bunch of 10^{10} electrons[51].

Figure 4.1 shows the beamline component map for ATF. The electron gun is located at the lower left of the image, and the long straight section following this is the linac in which the beam is accelerated to its final energy of $\sim$1.3 GeV. After leaving the linac, the beam then crosses underneath the ring into the beam transport line, which brings it to the injection point at the beginning (right hand side) of the top straight section of the damping ring. After many orbits of the ring, the beam is then extracted at the same point at which it was injected, and guided along the extraction line to the beam dump. It is in the extraction line that the many control and measurement devices used to maintain and measure the extremely low emittance are located.

4.1.1 Properties of the ATF beam

Table 4.1 compares the typical properties of the ATF beam on extraction from the damping ring during multi-bunch operation, with the typical properties of the beam used at NLCTA. It can be seen that the ATF beam has a much higher energy, and two orders of magnitude

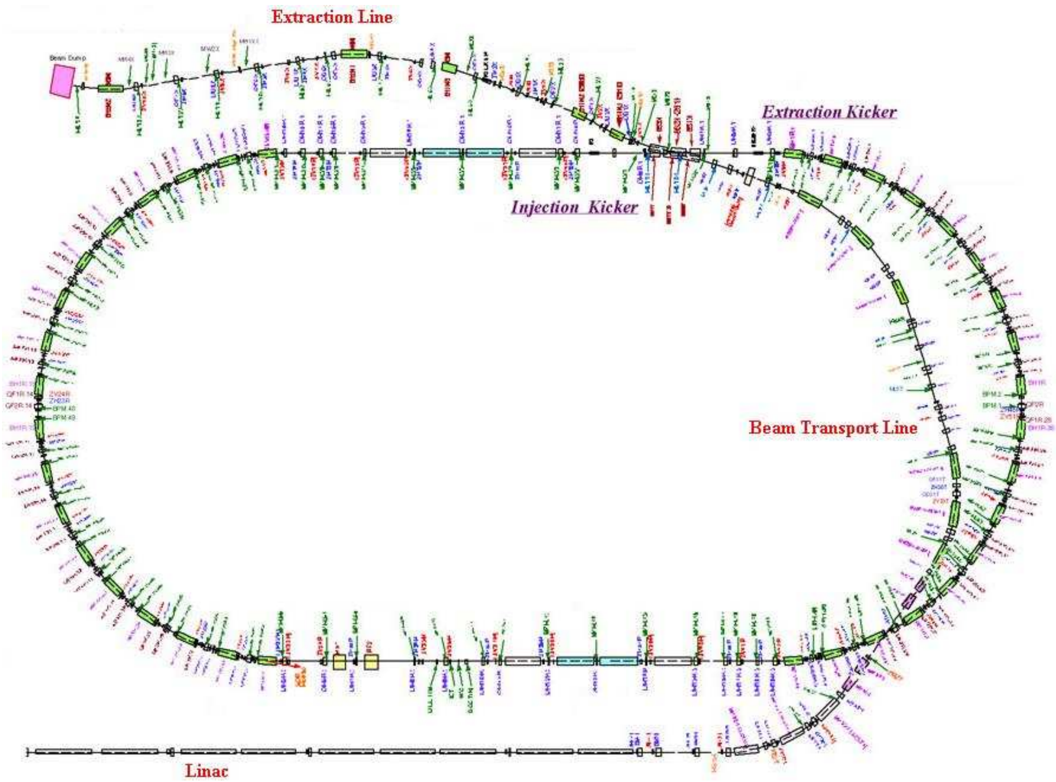

Figure 4.1: Plan of the ATF ramping ring showing the linac (bottom of the image), damping ring (centre), and the extraction line (top)[17]

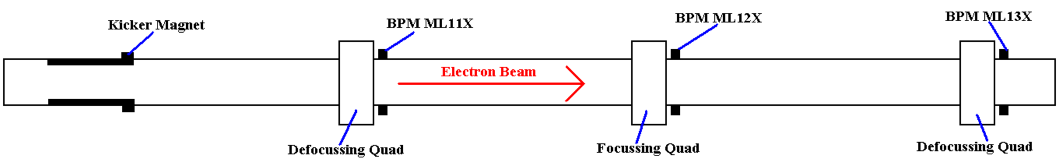

Figure 4.2: Layout of the extraction line[17]

Beam Property	**NLCTA**	**ATF**	**Units**
Beam energy	0.062	1.54	GeV
Typical bunch intensity	1.0×10^8	2.0×10^{10}	Electrons/bunch
Number of bunches/train	~ 1900	≤ 20	
Bunch – bunch separation	0.088	2.800	ns
Train length	177	56	ns
Bunch train repetition rate	$1 \sim 60$	$0.7 \sim 6.4$	Hz
Typical beam size ($x \times y$ rms)	500×1000	70×7	$\mu m \times \mu m$

Table 4.1: Comparison of the properties of the NLCTA and ATF beams[26].

more charge per bunch than the NLCTA beam, although the charge per train is very similar due to a larger number of bunches at the latter. Despite the much larger number of bunches per train, the train length at NLCTA is only a factor of three longer due to the fact that every single RF-bucket of the 11.4 GHz accelerating RF has been filled, resulting in a very short 88 ps between bunches. At ATF the accelerating RF has a frequency of 2.8 GHz, however bunches are only present in every eighth bucket, resulting in a bunch spacing of 2.8 ns.

Also of interest is the comparison between the beam dimensions. As the NLCTA was designed as a machine to process and test accelerating structures it was acceptable to have relatively large beam dimensions, however the ATF was designed as the proof-of-principle of a damping ring for a linear collider, and, as such, the emittance was reduced to an extremely low level, and much effort was made to control the betatron motion of the beam. Since beam jitter is proportional to the β-function it was expected that this high degree of control would lead to a very stable beam.

4.2 Differences with respect to FONT2, and experimental goals

While designing FONT3 it was important to take into account the large differences in beam parameters between ATF and NLCTA, and there were two parameters in particular that strongly influenced the goals of the experiment and its subsequent design.

4.2.1 Train length

At NLCTA a latency of 53 ns was achieved, allowing observation of the beginning of a fourth pass around the loop. In the case of the ATF beam, however, the entire train is only 56 ns in length, meaning that the system latency had to be decreased by a factor of $\sim$3 to achieve comparable results. A goal of a total latency of 20 ns was introduced, and divided amongst the components as follows,

1. 3 ns for the time of flight of the beam ($\sim$1 m at the speed of light).

2. 5 ns signal propagation time.

3. 5 ns to process the BPM signal.

4. 5 ns amplifier risetime.

5. 2 ns kicker fill time.

Items one and two are set by the kicker to BPM distance, and the speed of signal propagation in the BPM processor output cables, therefore these numbers are considered to be irreducible. Item five can also be considered irreducible as it depends only on the geometry of the kicker strips. Items three and four are, therefore, the only controllable contributions to the total latency, and they were set to these numbers so as to ensure the total latency would not exceed 20 ns, and as these seemed like challenging, yet achievable, goals.

4.2.2 Transverse beam dimensions

The much smaller transverse beam dimensions relative to the NLCTA beam results in much smaller beam jitter due to the smaller β-functions. This allows the possibility of observing a much more accurate correction, possibly of the level of one micron. It was made a goal of this experiment to make the correction accurate to roughly this level.

4.3 BPM Processor Design

In figure 4.2 the location of the kicker has been marked. It is approximately a metre upstream of the ATF stripline BPM ML11X (referred to as BPM11 from this point), and due to the availability of this stripline BPM at the required distance from the kicker it was decided that BPM11 would be used as the feedback BPM and an additional BPM would not have to be installed. In this case it would be necessary to design processing electronics that could use this stripline signal (see section 1.6).

Section 1.6 detailed several ways in which a stripline signal can be processed in order to retrieve beam position information:

1. Difference over sum

2. Log ratio

3. AM/PM

The log ratio method was rejected due to the excessive time lag through available logarithmic amplifiers, and the AM/PM method was rejected due to excessively tight constraints on the timing jitter of the components required to convert amplitude information to phase information. This leaves the difference/sum method.

Division of analog signals is costly in terms of time, so it would be beneficial if the charge profile of the beam was flat enough that charge normalisation was not necessary. A bunch charge error, $\frac{\Delta Q}{Q}$, (where Q is the nominal bunch charge) will appear as a position error, $\frac{\Delta y}{y}$, (where y is the measured bunch position) in the case of no charge normalisation. If the position signal is the difference between the signal from opposing striplines then the following is true,

$$y = f(t) \cdot Q \cdot R \cdot \left(1 + \frac{d}{r}\right) - f(t) \cdot Q \cdot R \cdot \left(1 - \frac{d}{r}\right) \tag{4.1}$$

$$y = 2f(t) \cdot Q \cdot R \cdot \frac{d}{r} \tag{4.2}$$

where $f(t)$ is the function that represents the output shape of the stripline, Q is the charge of the bunch, R is the impedance of the measurement device, r is the half-separation of the striplines (i.e. the radius of the BPM), and d is the displacement of the bunch from the electrical centre of the striplines in the plane of the striplines ($d \ll r$).

$$\frac{dy}{dQ} = 2f(t) \cdot R \cdot \frac{d}{r} \tag{4.3}$$

$$\frac{dy}{y} = \frac{2f(t) \cdot R \cdot \frac{d}{r} \cdot dQ}{2f(t) \cdot Q \cdot R \cdot \frac{d}{r}} = \frac{dQ}{Q} \tag{4.4}$$

Therefore, in the case where the un-normalised difference of two opposing striplines is used as the position signal, the fractional error in the position signal due to a charge deviation is equal to the fractional size of that deviation, as in equation 4.4.

A stated goal of the experiment is to correct the position of the beam to within 1 μm of zero in two correction periods. An initial position offset of 100 μm with a charge deviation of 10% below nominal will be measured as a position of 90 μm, and will, therefore, be undercorrected by 10 μm. On the second pass around the loop this will be measured as an offset of 9 μm and undercorrected by 1 μm. The amplifier is designed to provide a maximum kick of ~100 μm, so the maximum charge deviation allowed in order to achieve the goal of correcting to within 1 μm of nominal is 10%. Multibunch studies at ATF show that the charge profile can be tuned to this level (see figure 4.3), therefore it was decided that it was acceptable to neglect charge normalisation in the design of the processor.

4.3.1 Difference Between the Top and Bottom Stripline Outputs

The component used to produce the difference is a 180° hybrid – a device commonly used to subtract RF signals.

All hybrid designs rely on the manipulation of the phase of the input signal using various techniques (e.g. inductors, lengths of transmission line, etc.) and figure 4.4 shows the basic principle of operation. A signal input into B will be split between outputs C and D in such a way that the outputs have the same phase as the input. Input to A is also split between C and D, but the output from D will have a relative phase of 180°. If the inputs to A and B have the same spectra and are input simultaneously, the output at C will, therefore, be their sum, and the output at D will be the difference.

In the case of a hybrid with an infinite bandwidth, the spectrum of the difference signal will be the same shape as that of the raw signal. Figure 4.5 shows the theoretical output spectrum of the 12 cm ATF striplines. Due to the large power at relatively high frequencies (the lowest peak in this spectrum is at 625 MHz) this is not a satisfactory signal to input to the amplifier. It is preferable to shift this higher frequency power closer to DC (<100 MHz).

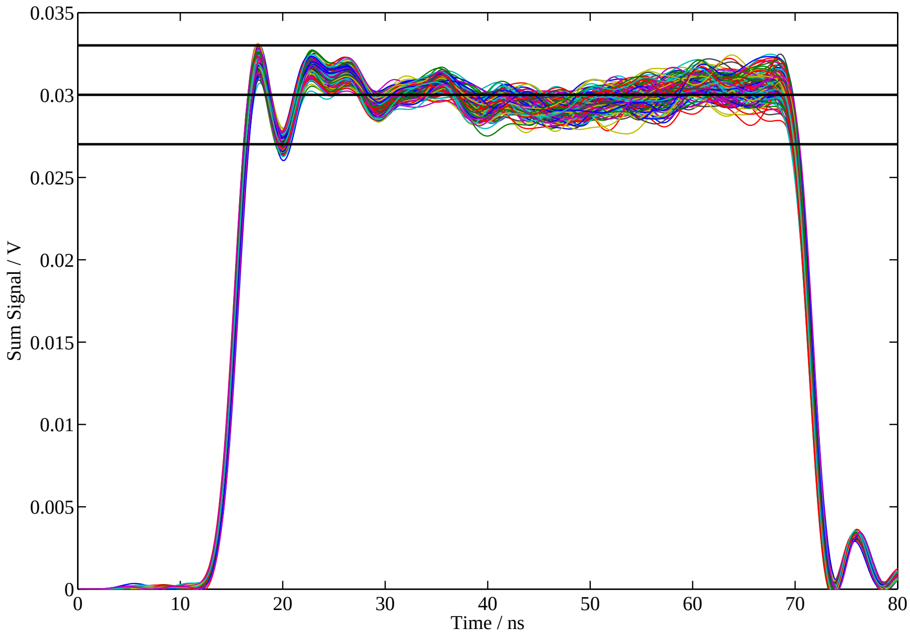

Figure 4.3: Measured sum profiles from the ATF extraction line. The lines mark the mean charge, and the mean ±10%. The positive and negative excursions at the beginning are due to the filters.

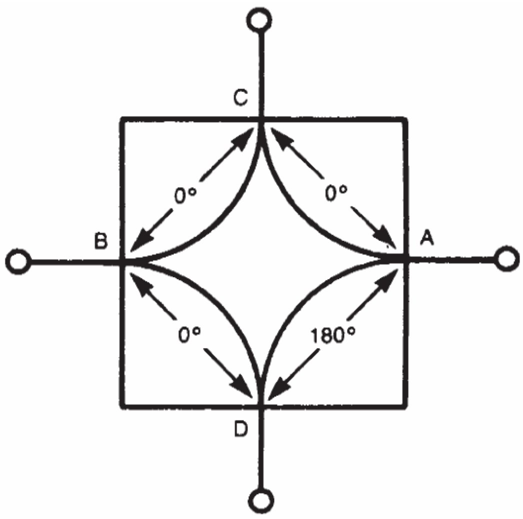

Figure 4.4: Diagram showing the operation of a 180° hybrid[18].

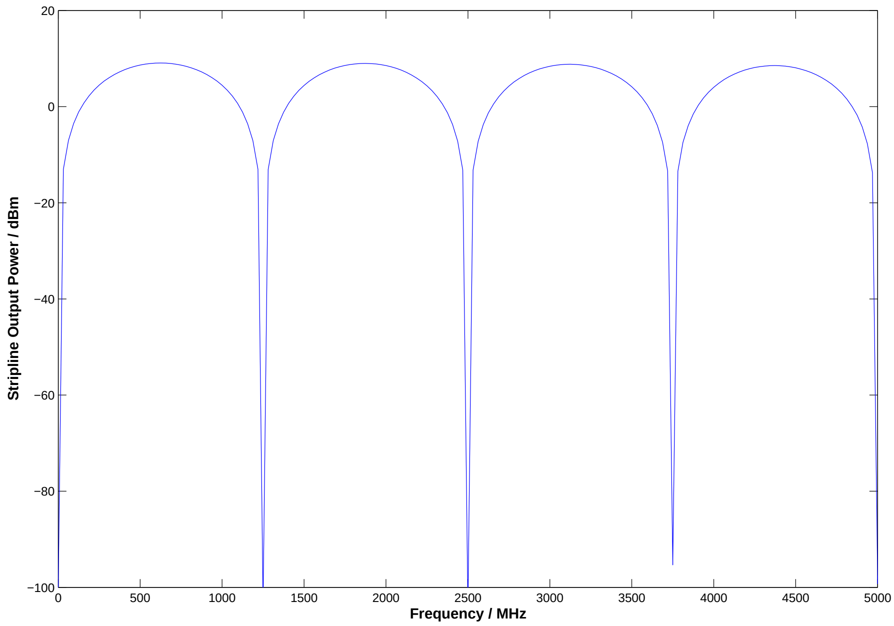

Figure 4.5: Theoretical power spectrum for a single bunch passing a 12cm stripline. This was calculated in Matlab[19] as the power spectrum of the time-varying function output from a stripline due to a delta-function bunch.

4.3.2 Mixers

A mixer is used to shift the frequency of the input power spectrum and is basically a circuit that takes an input signal, and outputs the product of this and a reference signal[1]. If, for example, the RF is a signal with a frequency of ω and an amplitude of A, and the LO is a signal with a frequency of ν and an amplitude of B, then the IF will be,

$$IF = A.\sin(2\pi\omega t) \times B.\sin(2\pi\nu t) = \frac{A.B}{2}\left[\cos(2\pi\omega - 2\pi\nu)\,t + \cos(2\pi\omega + 2\pi\nu)\,t\right] \qquad (4.5)$$

Equation 4.5 shows that the product of two sinusoidally varying signals contains a component with a frequency that is the sum of the input frequencies, and a component whose frequency is the difference of these. Therefore, to shift a signal to DC (zero frequency), it should be multiplied by a signal of the same frequency,

$$IF = A.\sin(2\pi\omega t + \phi) \times B.\sin(2\pi\omega t) = \frac{A.B}{2}\left[\cos(\phi) + \cos(4\pi\omega t + \phi)\right] \qquad (4.6)$$

where ϕ is the relative difference in phase between the two signals. Therefore, multiplying by a signal of the same frequency gives a DC component whose amplitude is proportional to the cosine of the phase difference between the inputs, and a component whose frequency

[1] The input signal is known as the RF (Radio Frequency) due to the mixer's original use as part of a radio receiver/transmitter, the reference is known as the LO (Local Oscillator), and the output is known as the IF (Intermediate Frequency).

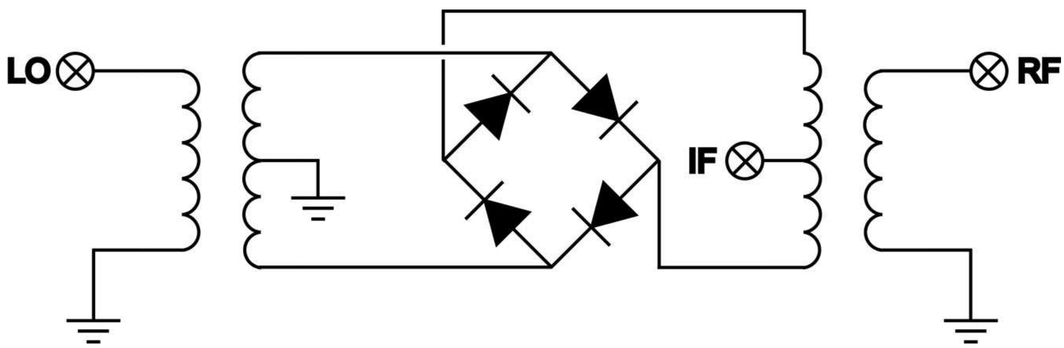

Figure 4.6: Circuit diagram for a mixer[20].

is twice that of the inputs. It is important to note that the amplitude of the DC component is maximum when the inputs have the same phase, i.e. $\phi = 0$.

Figure 4.6 shows the circuit diagram for a generic mixer. The diode ring in the mixer acts as a series of switches that are opened or closed depending on the direction of the potential difference across them. The LO amplitude will typically be much larger (by a factor of ten or more) than the RF input and so will govern the state of the diode switches, with the RF waveform providing only a small correction to this. When the LO is in its positive half-cycle, the two diodes to the right will be closed, allowing current to flow, and the RF input signal will only be induced on the lower half of the IF output coil. On the negative half cycle it is the opposite two diodes that are closed, forcing the RF induced current to flow through the upper half of the IF output coil, i.e. π out of phase with the previous case.

If the RF input is a signal of the same frequency and phase as the LO, the output during the positive half-cycle of the LO will be identical to the RF input, i.e. a positive half cycle. When the LO becomes negative, the output IF will be the same as the input, but with a phase shift of π, i.e. another positive half-cycle. The IF output will then be repeats of the positive part of the RF input, or, to a first approximation, a signal of half the amplitude of the input signal, but with twice the frequency and a DC offset of half the input amplitude, which is exactly as predicted by equation 4.6 (with $A = 1$).

Figure 4.6 shows that the LO signal should not be transmitted to the IF output, however, due to imperfections, a small amount of the LO will leak through. The amount of this leakage is normally specified by the manufacturer.

Equation 4.6 shows that any phase jitter in the LO will be transmitted to the mixer output as amplitude jitter. A 714 MHz reference signal exists at ATF and the broad peak at 625 MHz (see figure 4.5) will be shifted to $\leq$100 MHz as required when mixed with this LO. This signal is phase-locked to the accelerating RF in the linac, and should therefore be very well phase matched to the timing of the beam passing the extraction line BPMs. This signal was measured by triggering the scope with a beam signal (in this case the sum signal from the hybrid) and measuring the LO signal. The difference between a particular zero crossing of the measured LO and the peak of the sum signal was calculated, and figure 4.7 shows the spread of this difference for many beam pulses. The RMS phase jitter of the Gaussian fit is 1.52°. Equation 4.6 shows that the amplitude jitter is proportional to the

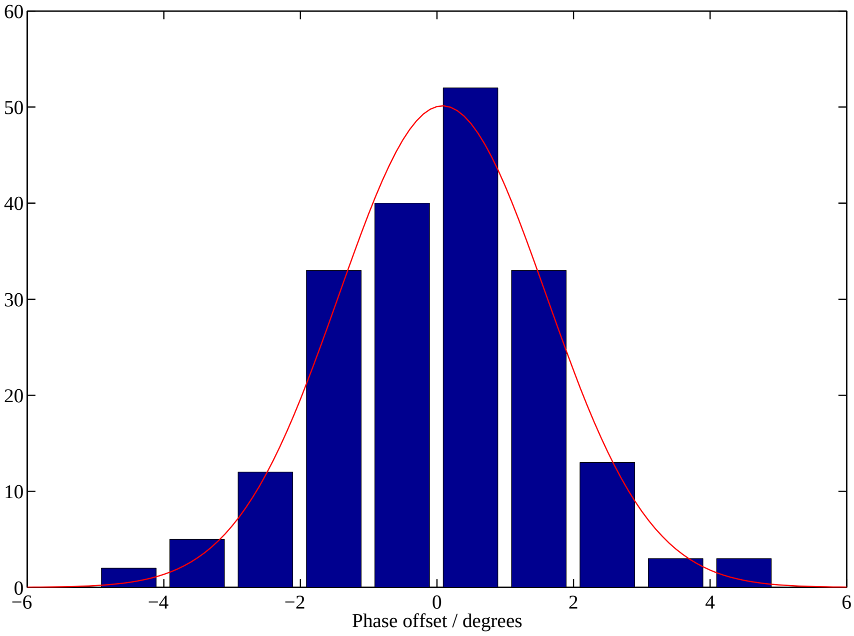

Figure 4.7: Stability of the LO phase with respect to the beam.

cosine of the phase jitter, therefore the scale of the amplitude jitter resulting from this will be $\leq 10^{-3}$, and it can be concluded that this signal is suitable for use as the mixer LO.

There will, therefore, only be a static phase mismatch between the two that can be trimmed to zero using suitable lengths of cable or dedicated phase shifters.

4.3.3 Filtering the signal

The output of the mixer will contain two significant components; the required DC-like signal, and a component with a frequency of $\geq$1 GHz. Other components output by the mixer will include a significant amount of LO leakage ($\sim$30 dB below the input LO power[2]) as well as mixer inter-modulation products, defined as,

$$n.RF_{freq} + m.LO_{freq} \tag{4.7}$$

where $n = 2, 3, 4, \ldots$ and $m = 2, 3, 4, \ldots$.

In a well constructed mixer these products will be considerably lower in power than the main output signal, they are, however, important to remember in this case due to the high accuracy required of the processor. An additional problem is that, as noted in table 4.1, the bunches repeat every 2.8 ns – equivalent to a repetition rate of 357 MHz. This will enter the

[2]RF power ratios are often referred to in units of dB, which is simply the logarithmic ratio between the two numbers, i.e. $LO_{leakage} = 10 \times \log_{10}\left(\frac{LO_{Input}}{LO_{Output}}\right)$. Absolute RF power, as its ratio to 1 mW, can also be measured in this way, and has units of dBm. This convention will be employed frequently in the following discussions.

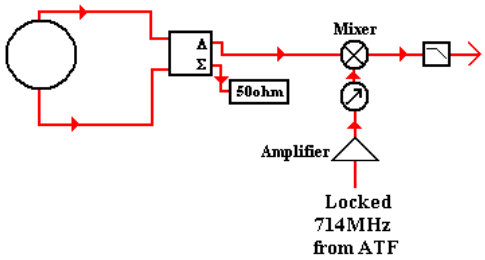

Figure 4.8: Basic BPM processor, showing the stripline pickups on the left, followed by the hybrid, with the difference (Δ) going into the mixer, and finally through the LPF. The sum signal (Σ) from the hybrid is terminated into 50Ω as it was decided not to normalise with the charge.

mixer with a large amount of power, and will be mixed with the 714 MHz to both 357 MHz and 1071 MHz.

These can be removed with the use of a low pass filter (LPF) on the mixer output. The filter must strongly reject the unwanted frequencies, and have a relatively broad bandwidth to maintain a low signal propagation time. The basic processor design developed thus far is shown in figure 4.8.

4.3.4 Filter Design

The filter must strongly suppress the beam bunching power mixed to 357 MHz, and 1071 MHz, as well as the LO mixer leakage at 714 MHz, however the bandwidth must be large enough to allow for a sufficiently low signal delay.

Systemview[52] is a package used for the simulation of RF devices, and was used to simulate various aspects of the processor. Figure 4.9 shows the simulated output of the mixer from a beam of 10^{-10} C/bunch with an offset of $\sim$100 μm, and the 714 MHz LO leakage can be clearly seen. The power spectrum of this signal is shown in figure 4.10, and the unwanted signals at 714 MHz and harmonics of 357 MHz are clearly observed. It can also be seen that this offset produces a low frequency power of $\sim$-40 dBm. A measurement of 1 μm is required implying that it is necessary to measure a voltage $\sim$100 times smaller, and, therefore, a power four orders of magnitude, or 40 dB, smaller. This is a power level of $\sim$-80 dBm and it was decided that the spurious mixer products should be reduced to this level.

It is, therefore, necessary to design the LPF to reduce the LO leakage by $\sim$50 dB, and the 357 MHz harmonics by $\sim$25 dB. The second design constraint is that the signal delay of the entire circuit should $\leq$5 ns. Figure 4.11 shows the simulated output of the stripline and the output after the hybrid and mixer, and the delay from these components can be seen to be $\leq$0.2 ns. To allow for a safety margin it was decided that the filtering should have a signal delay of $\leq$ 4ns.

Many styles of filter exist, and it was decided to design the LPF to have a Chebyshev response characteristic as this allows for strong suppression of out of band power with a wide passband. The strong suppression is achieved at the expense of the flatness of the frequency characteristic within the passband, however this passband ripple will be $\sim$0.5 dB and will

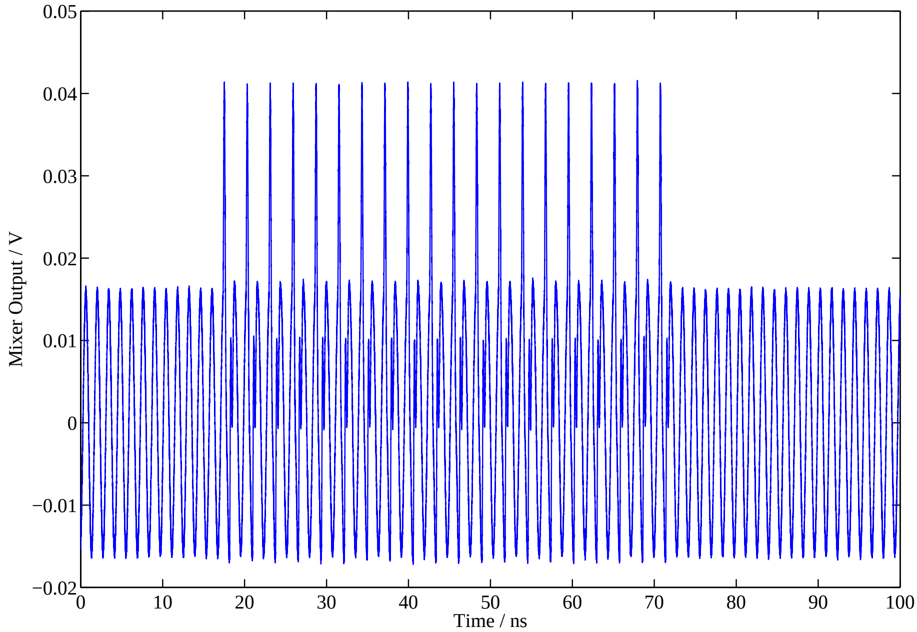

Figure 4.9: Simulated output of the mixer due to a 20-bunch beam with $\sim 10^{-10}$ C/bunch with an offset of $\sim$100 μm.

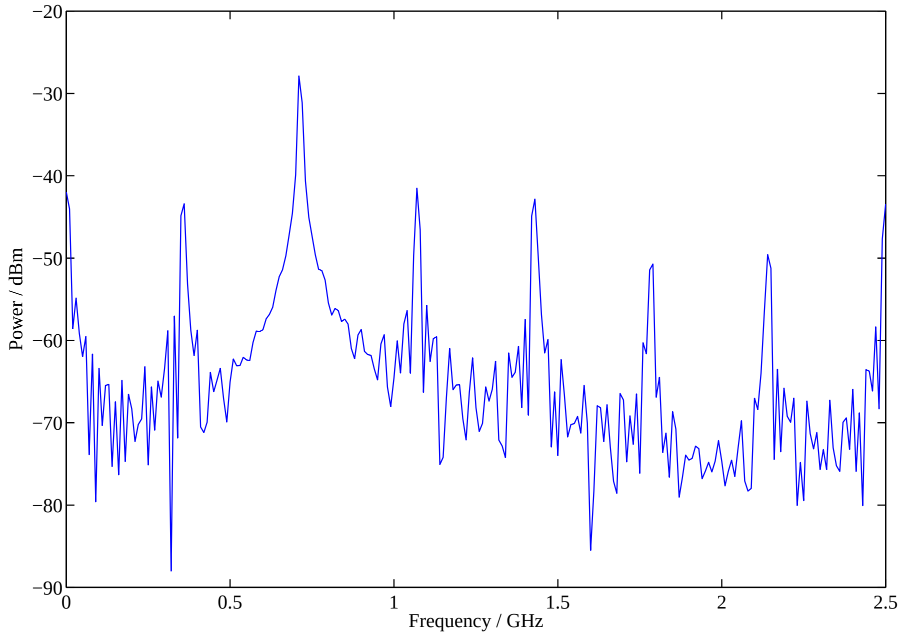

Figure 4.10: Output spectrum of the mixer calculated as the voltage from figure 4.9 into an impedance of 50 Ω.

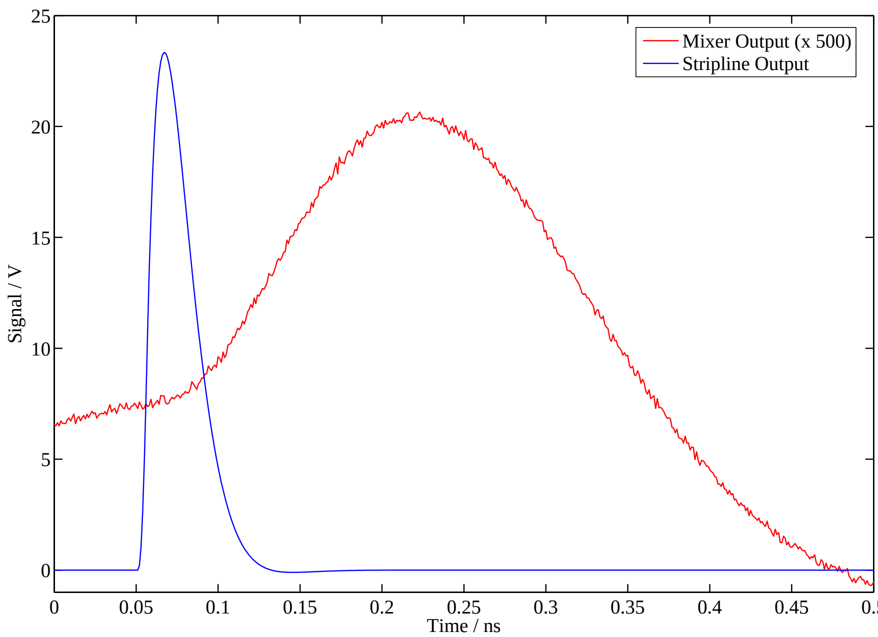

Figure 4.11: Simulated stripline output and processor output after a hybrid and mixer.

have little effect on the output.

A LPF circuit was designed to have a cutoff of 220 MHz, and 0.5 dB ripple in the passband, and is shown in figure 4.12. Figure 4.13 shows the attenuation curve, and figure 4.14 shows the signal delay of this circuit as simulated by SPICE[53].

Figure 4.13 shows that the attenuation of the filter at baseband is 6 dB, however this is a function of the way the SPICE handles the input and output impedances. In reality this plot will be uniformly shifted upwards by 6 dB. The attenuation at the frequencies of the unwanted mixer products can be seen to be stronger than required by $\geq$10 dB, which gives a substantial safety margin in building the devices. Figure 4.14 shows that the delay for signal components $\leq$100 MHz is $\leq$3.5 ns, also allowing for a margin of error in manufacture.

The component values in figure 4.12 are very precise, and are unlikely to be found in manufacturers catalogues, so the closest approximations to these values were found and inserted into the simulation, as shown in figure 4.15. The specified parasitic capacitance and inductance of each device has also been included in this figure. Figures 4.16 and 4.17 show the results of the SPICE simulation of this circuit, and it can be seen that there is very little change to the frequency response curve, and that the delay is slightly reduced.

4.3.5 Alternative Filtering Scheme

The LPF will limit the output spectra to $\leq$200 MHz, so any mixer input power with a frequency outside 714$\pm$200 MHz will be mixed to a frequency beyond the bandwidth of the LPF. A BPF with a bandwidth of $\sim$500 MHz–900 MHz placed before the mixer would remove these extraneous frequencies. This will have two benefits,

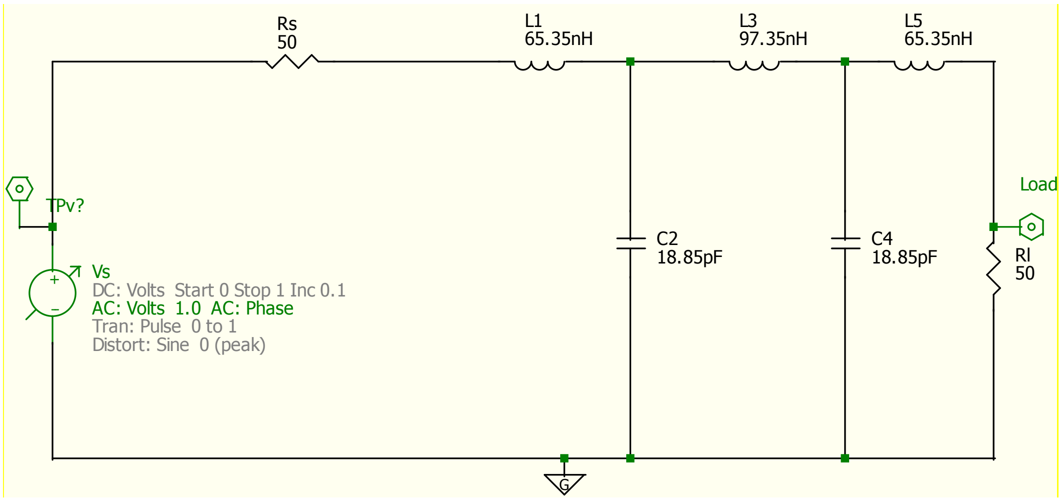

Figure 4.12: Calculated circuit for a Chebyshev LPF with a 220 MHz cutoff and 0.5 dB ripple in the passband.

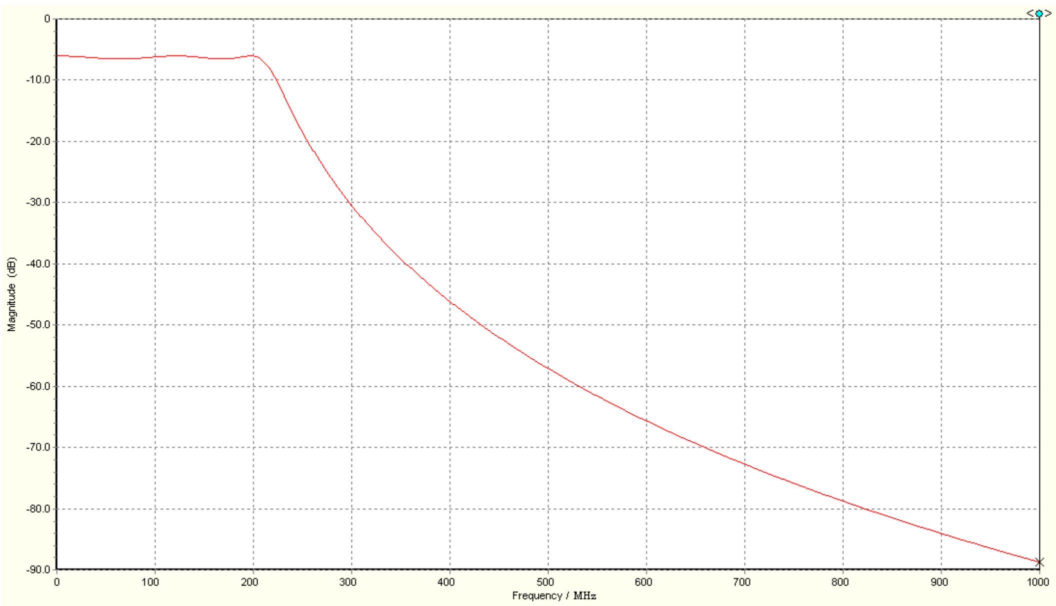

Figure 4.13: Simulated frequency dependence of the attenuation through the circuit shown in figure 4.12.

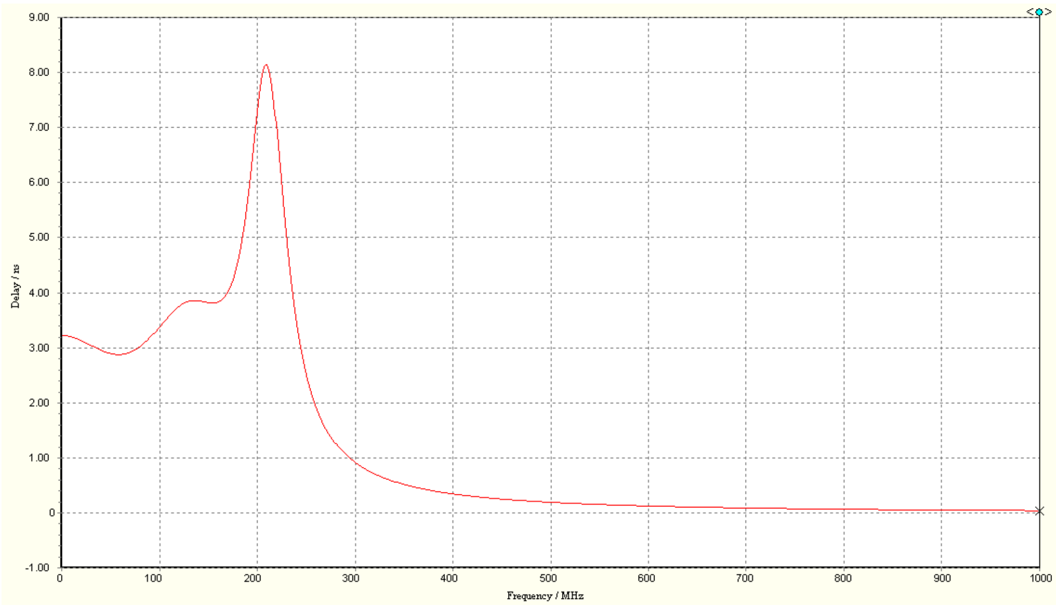

Figure 4.14: Simulated frequency dependence of the delay through the circuit shown in figure 4.12.

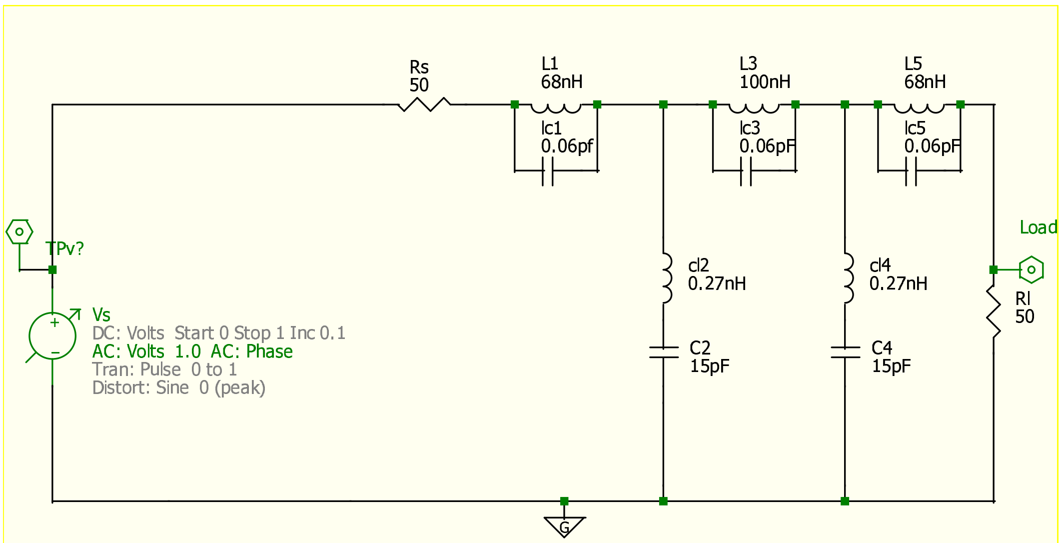

Figure 4.15: Figure 4.12 modified to use realistic component values. The parasitic capacitance and inductance of each device has also been included.

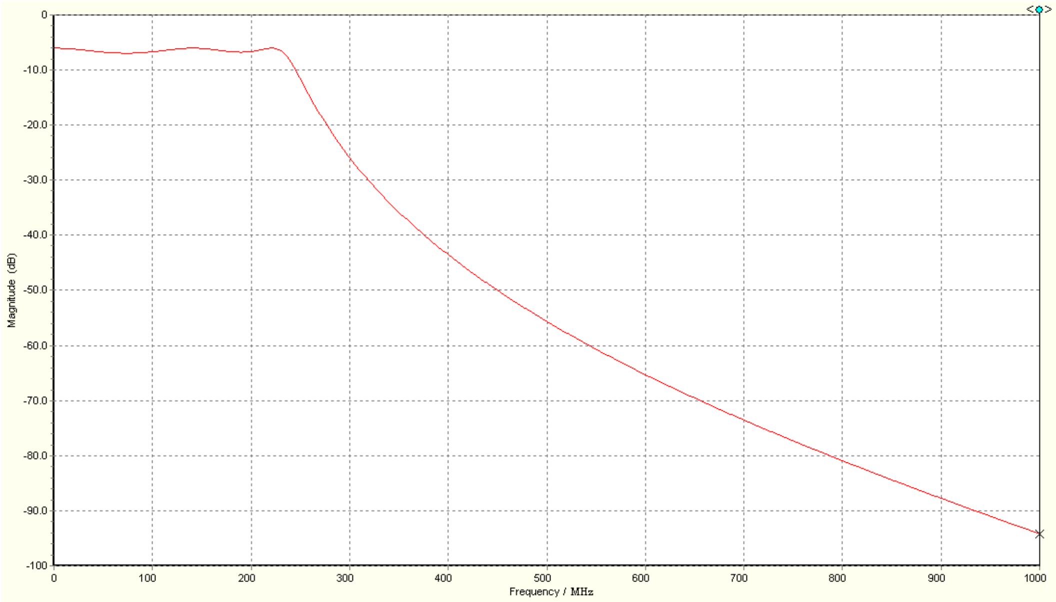

Figure 4.16: Simulated frequency dependence of the attenuation through the circuit shown in figure 4.15.

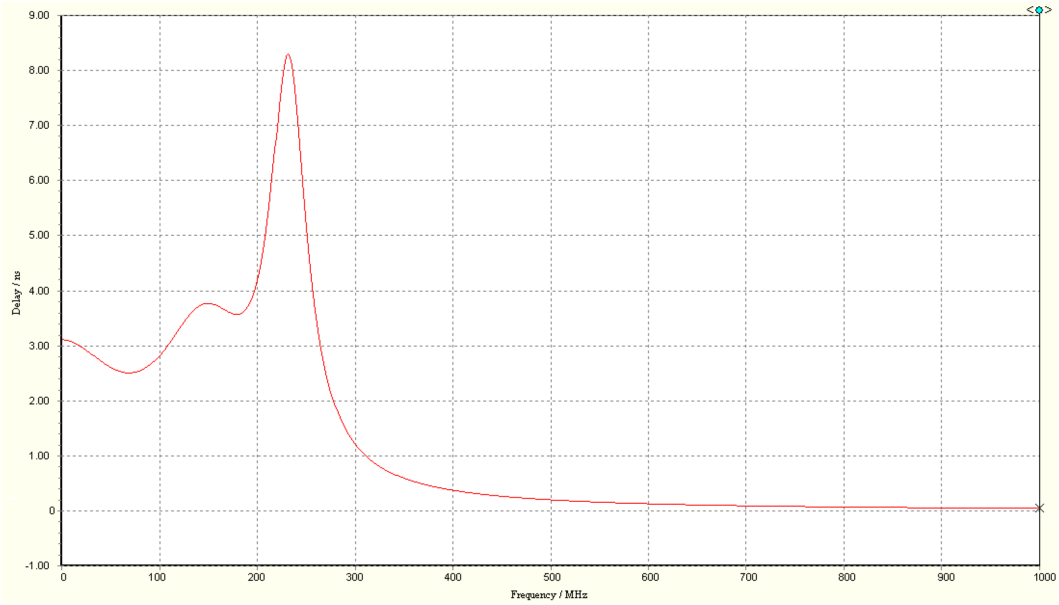

Figure 4.17: Simulated frequency dependence of the delay through the circuit shown in figure 4.15.

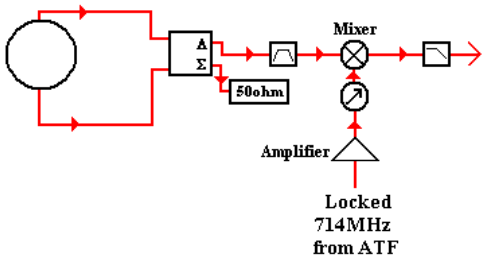

Figure 4.18: Similar processor to figure 4.8, but with the addition of a BPF before the mixer.

1. The spectrum entering the mixer, is potentially very broad, and could be beyond the frequency range of the mixer. Filtering protects the mixer from this out of band power.

2. The peak voltage coming from the subtraction device could be excessively large and damage the mixer. These large peak powers could also drive the mixer into a non-linear regime, thereby degrading the accuracy of the processor. Filtering will broaden the pulse, thus lowering the peak voltage and protecting the mixer from these effects.

The signal propagation time of the processor is not necessarily affected by the addition of an extra filter. Since the BPF attenuates the 357 MHz beam bunching power, this relaxes the strength with which the LPF must attenuate this frequency. The LPF can then be built to have a slower fall off and a faster signal propagation time. In this way the filtering load is shared between the two filters. Figure 4.18 shows this design.

Two new filters were designed for this scheme; the BPF, and a new, reduced signal delay, LPF.

The BPF should have a bandwidth of ~200 MHz, suppress the 357 MHz beam bunching component by ~15 dB, and have a signal delay of $\leq$2 ns at 714 MHz. Since a broad range of components centred on 714 MHz will be passed through this filter, it was designed to have the characteristics of a Bessel filter. These are designed to have equal signal delay for all components in the passband. Figure 4.19 shows the circuit, and figures 4.20 and 4.21 show the simulation results.

Figure 4.20 shows that this design suppresses 357 MHz power by ~15 dB (after accounting for the 6 dB baseline shift) as desired, and figure 4.21 predicts a delay of ~1.1 ns for the 714 MHz component.

Figure 4.22 shows figure 4.19 with the theoretical inductance and capacitance values replaced with realistic component values, including parasitic capacitances and inductances as specified by the manufacturer, and figures 4.23 and 4.24 show the simulated attenuation curve and signal delay for this circuit.

It can be seen from figure 4.23 that the centre of the bandpass has shifted downwards in frequency to ~650 MHz, however the attenuation at 714 MHz is only ~1 dB[3]. Note the large resonance at ~1.7 GHz caused by parasitic resonances in the inductors and capacitors. This strongly increases the attenuation at this frequency, but will have little effect at other frequencies.

[3]Remembering that the plot has been falsely shifted downwards by 6 dB due to the way in which the model handles the input and output impedances

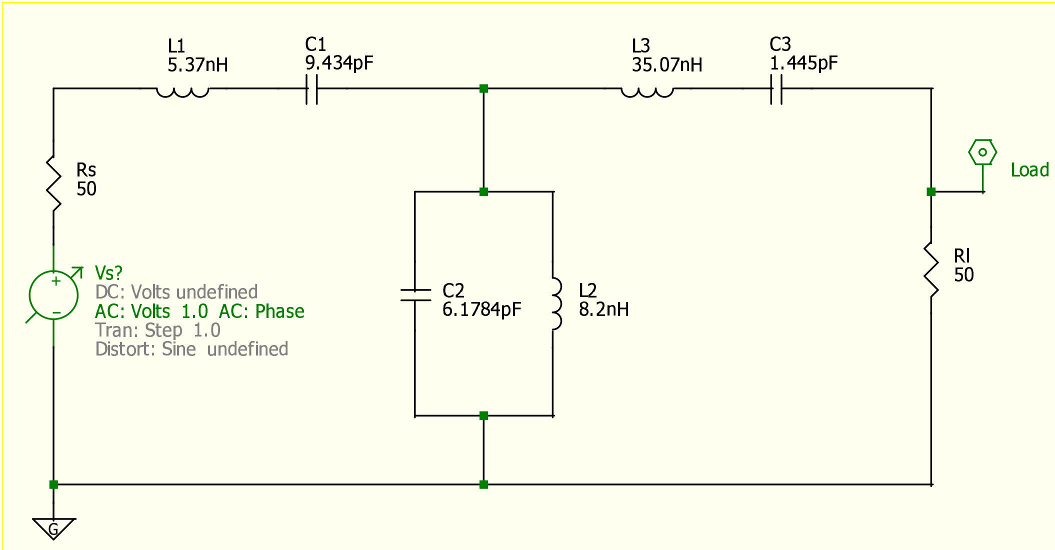

Figure 4.19: Circuit calculated for a Bessel BPF centred on 714 MHZ, 200 MHz bandwidth.

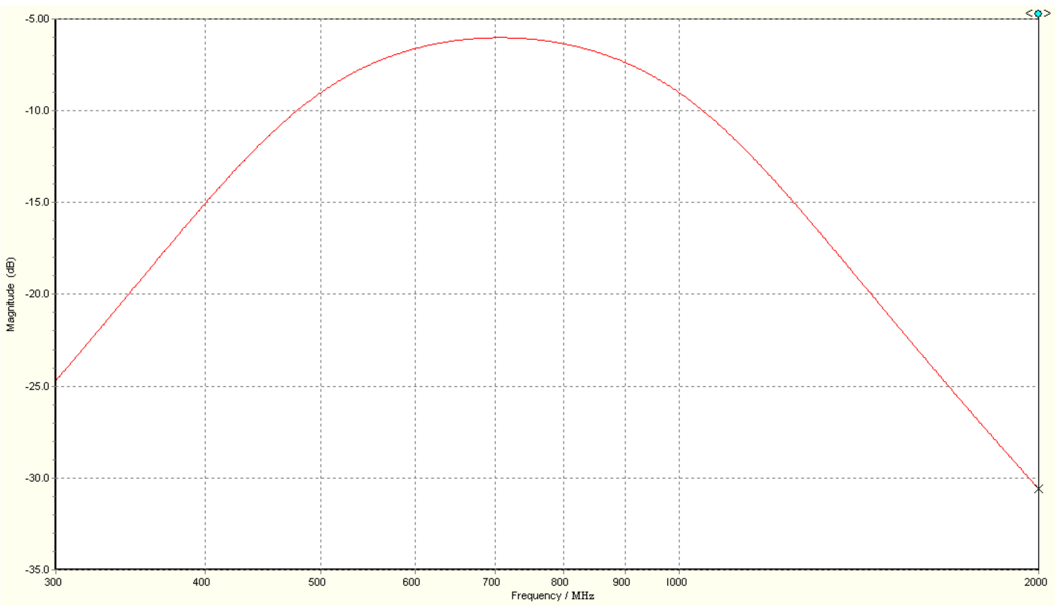

Figure 4.20: Simulated frequency dependence of the attenuation through the circuit shown in figure 4.19.

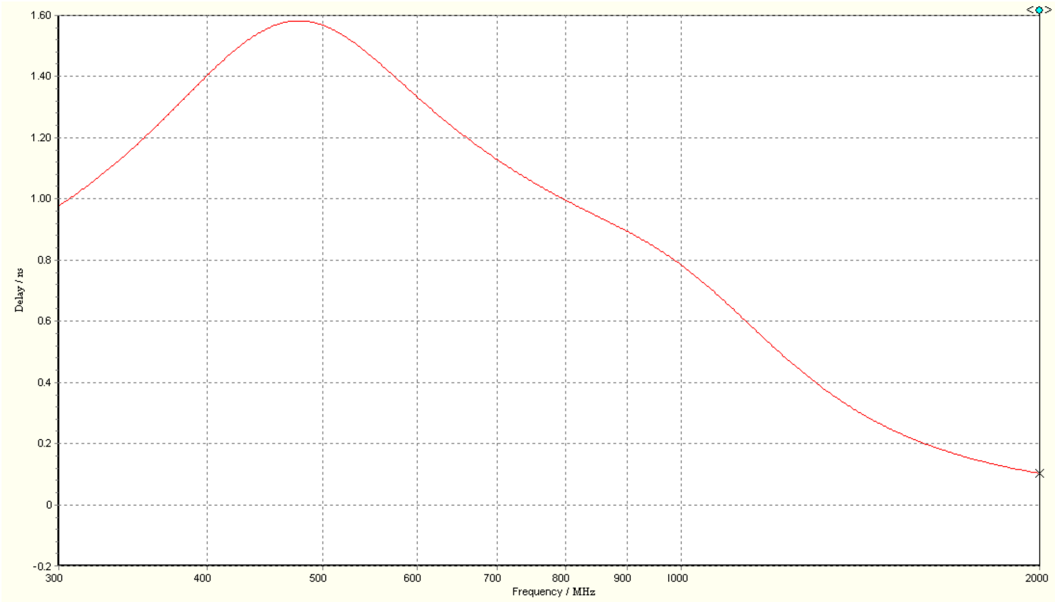

Figure 4.21: Simulated frequency dependence of the delay through the circuit shown in figure 4.19.

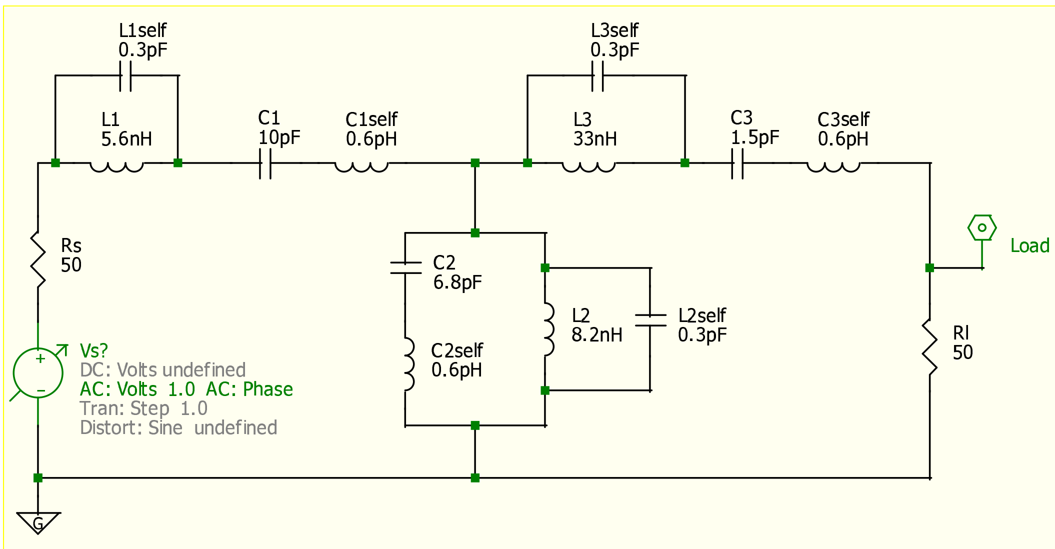

Figure 4.22: Circuit from figure 4.19 with realistic component values and the specified parasitic capacitances and inductances.

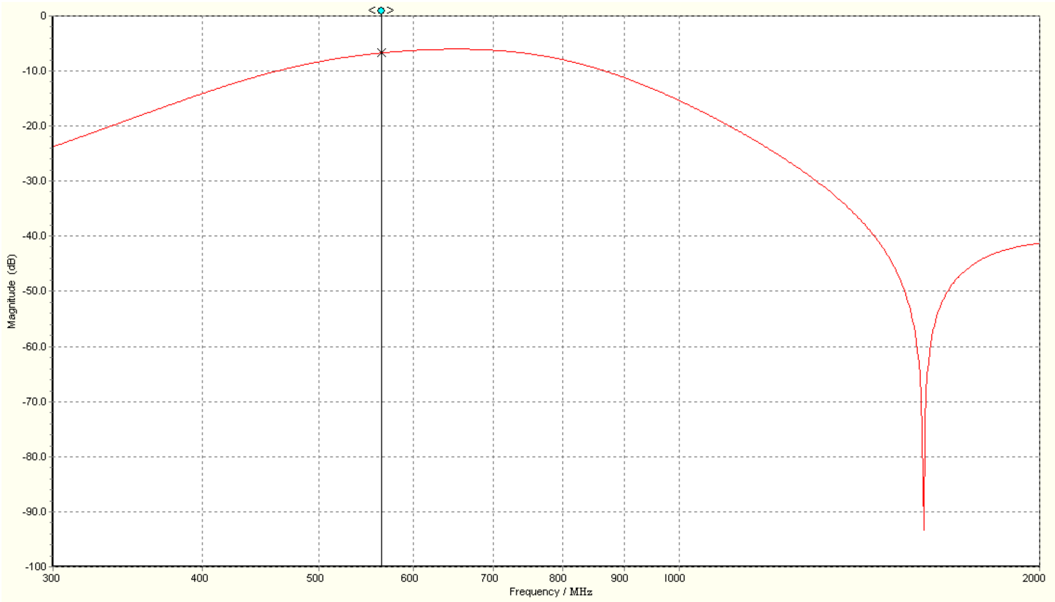

Figure 4.23: Simulated frequency dependence of the attenuation through the circuit shown in figure 4.22.

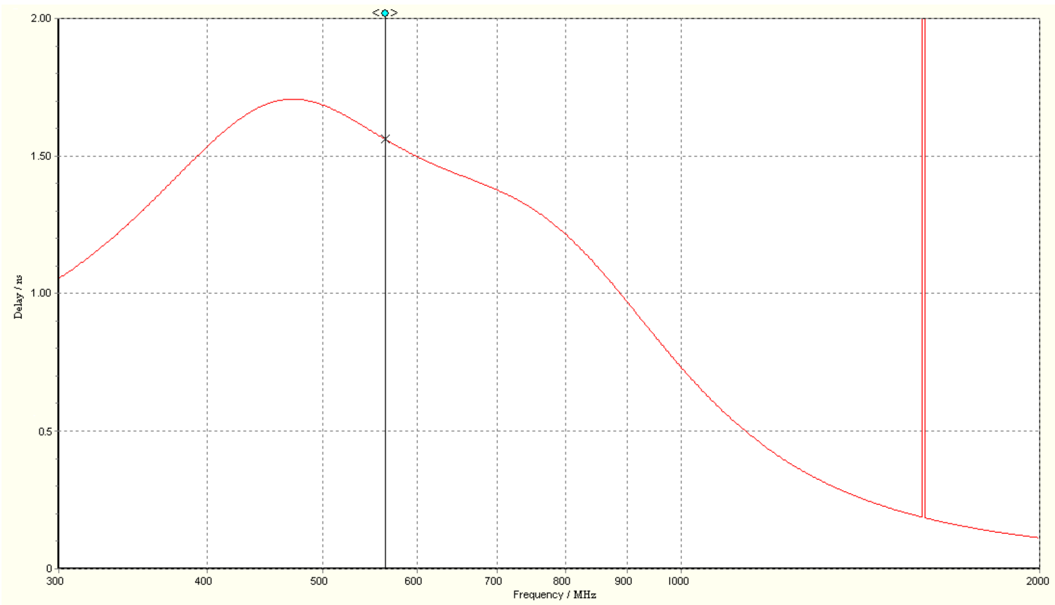

Figure 4.24: Simulated frequency dependence of the delay through the circuit shown in figure 4.22.

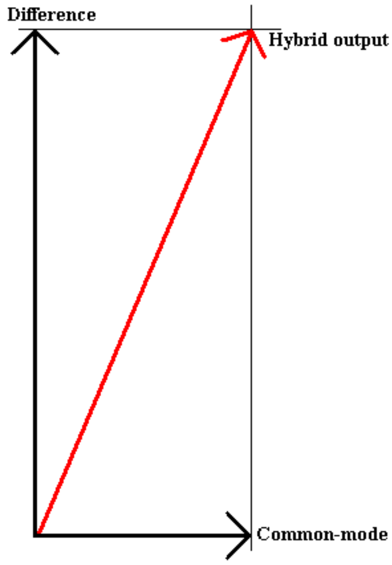

Figure 4.25: Phase space diagram of the hybrid output showing the difference and common mode signals.

Figure 4.24 shows the same resonance feature as figure 4.23. The signal delay at 714 MHz has increased slightly to ∼1.4 ns, however this is still within the acceptable range.

4.3.6 LO Phasing

The hybrid will input two signal components to the mixer,

1. The difference of the hybrid inputs.
2. A signal proportional to the differential of the input signal shape, referred to as the common mode signal.

The second component is due to mismatched cables lengths, and is therefore the result of the subtraction of two similar signals shifted in time by a small amount, dt, with respect to one another, i.e. the differential. Therefore, this signal will have a phase shift of $\pi/2$.

The hybrid output is shown in diagrammatic form in figure 4.25.

It can be seen that a judicious choice of LO phase will maximise the difference signal and minimise the common mode signal. Equation 4.6 showed that the maximum DC signal is obtained when $\phi = 0$, and zero output results from $\phi = \pi/2$, i.e. maximal selection of the difference signal (and maximal rejection of the common mode signal) is achieved by choosing a zero relative phase, $\phi = 0$, between the RF and LO mixer inputs.

One method to correctly phase the LO would be to move the beam to a large offset (where the difference signal amplitude is much larger than the common mode amplitude), and maximise the mixer output with respect to the LO phase. This method is problematic however, as it is difficult to know the location of the beam before the LO phase has been optimised, and it would be quite easy to move the beam to a position where the difference signal is large enough to saturate or damage one of the processing components. Also, it

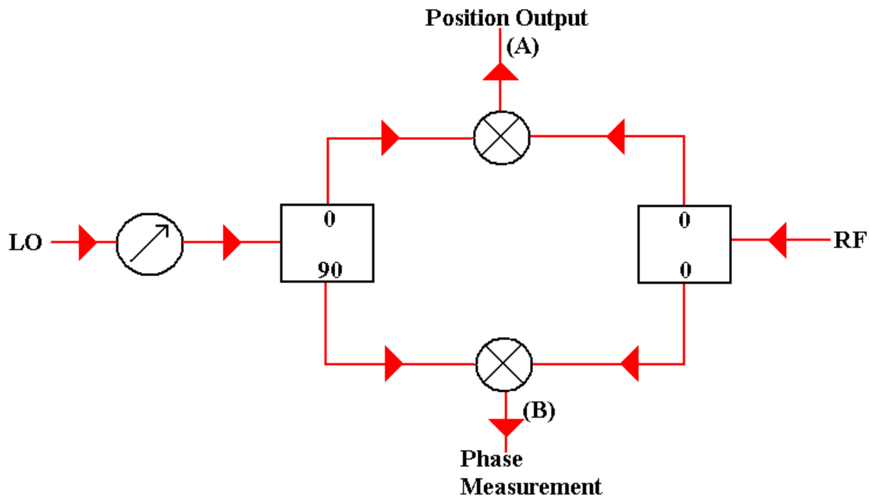

Figure 4.26: Block diagram of the mixing technique (IQ mixing) used to determine the 0 and $\pi/2$ components of the mixer output.

can be seen from figure 4.25 that if the beam is accidentally moved very close to zero, then maximising in this way will optimise for the amplitude of the unwanted common mode signal at the expense of the desired difference signal.

An alternative method can be developed by examining figure 4.25. The common mode signal is proportional to the charge of the beam, and not its position, therefore, if the LO is tuned to a phase that causes the mixer output to be independent of the beam position then the phase will be $\pi/2$ from the correct setting. If a length of cable equivalent to $\pi/2$ of 714 MHz was added to the cable carrying the LO to the mixer then the phase difference will be 0 or π. In this case $\phi = \pi$ is equivalent to a minus sign in the calibration coefficient of the processor (see equation 4.6).

Figure 4.26 shows the alternative mixing technique that was used to put this suggested method into practice. The output from the BPF is split with a 0° splitter, and each half is sent to the RF input of a mixer via cables of equal path length in order to preserve the 0° phase difference between them. The 714 MHz LO signal is also split by a 0° splitter, however the paths used to bring this signal to each of the mixers differ by $\pi/2$ of 714 MHz. This will result in the signal being mixed by orthogonal signals, and means that is output is tuned to have no dependence on the beam position, the output from the other mixer will be the required difference signal. This was especially useful for the feedback electronics as it allowed tuning of the phase without disconnecting the position output from the amplifier, and for the phase to be continually monitored throughout the experiment.

This method was implemented in the processor by using a voltage controlled phase shifter just upstream of the LO splitter (see figure 4.26). The mixed down common mode output was measured while the beam was stepped between $\sim +0.5$ mm and ~ -0.5 mm, and the LO phase was varied until the outputs were the same at each of these positions.

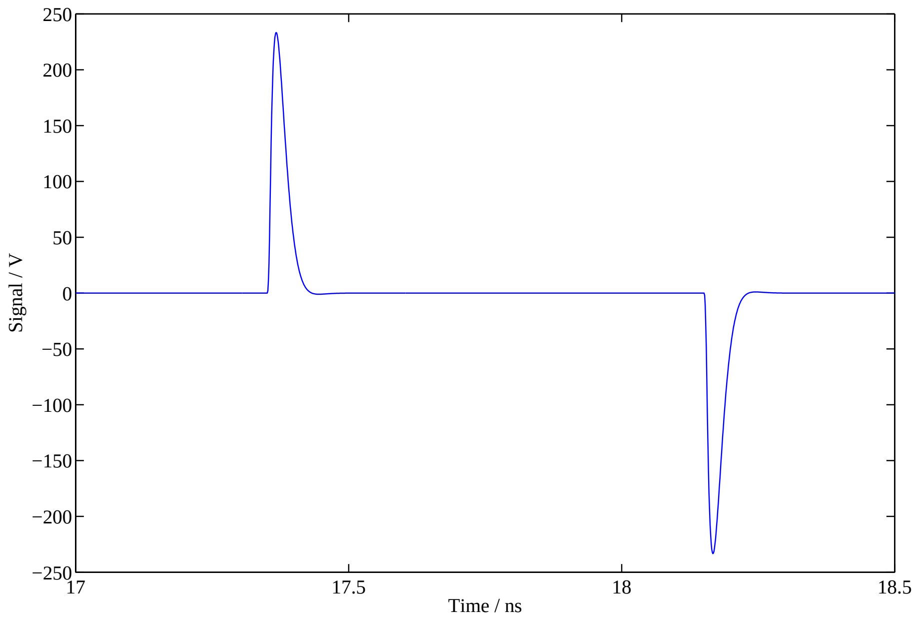

Figure 4.27: Simulated single bunch output of the stripline.

4.4 Processor Simulations

The SystemView RF simulation package was used to simulate the response of the processor.

To calculate the stripline signal, the beam was modelled as a twenty bunch train with a bunch charge of 1 nC ($\sim 6.25 \times 10^9$ e$^-$/bunch), and a bunch separation of 2.8 ns. The ATF BPMs have 12 cm long strip-lines with an opening angle of 52°. This opening angle means the strips will intercept ∼14.5% of the image charge. Using Ohm's Law it can be shown that,

$$V = C \cdot I \cdot R = C \cdot \frac{Q}{t} \cdot R$$

$$\therefore \int V \cdot dt = C \cdot Q \cdot R = 7.2 \times 10^{-9} \text{ Vs} \tag{4.8}$$

where the fraction of the total image charge intercepted by the strip, $C = 0.145$, the total bunch charge, $Q = 1$ nC, the input impedance of the measuring device, $R = 50\ \Omega$, and V is the measured instantaneous output voltage.

The stripline output was simulated by passing a 5 ps square pulse with a peak of 1440 V (i.e. $\int V \cdot dt = 7.2 \times 10^{-9}$ Vs) through a 10 GHz low pass filter. The negative of this square pulse was added 800 ps after the positive pulse, to simulate the reflection from the downstream end of the stripline. Thermal noise at 300 K was added to each signal, and the signal amplitude was adjusted by 1% (+1% for the top strip and −1% for the bottom strip) to simulate a beam offset of 1% of the distance between the opposing strips. In the ATF this distance is 2.4 cm, so this simulates a beam offset of 240 μm. Figure 4.27 shows the simulated signal, and figure 4.28 shows the spectrum of the twenty bunch beam.

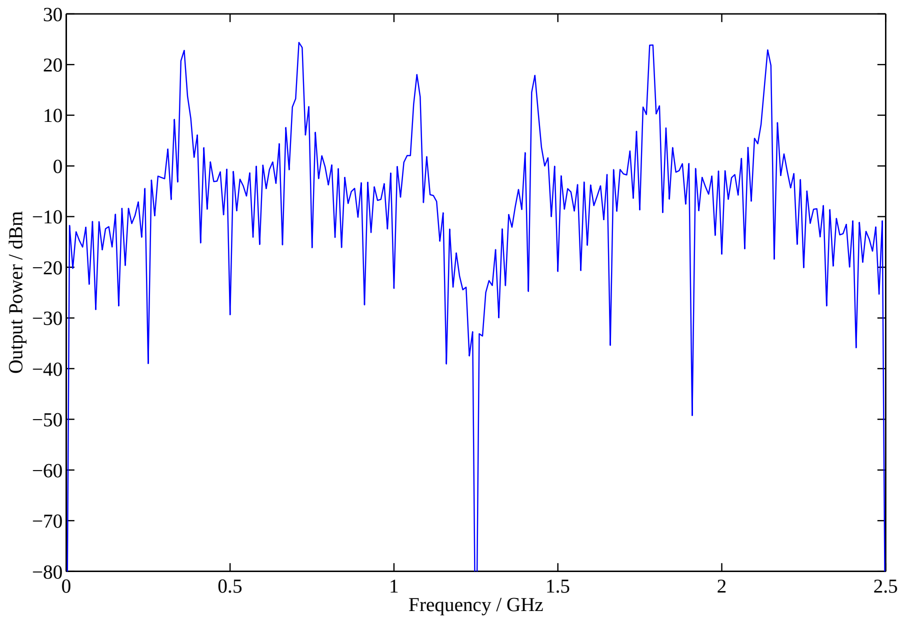

Figure 4.28: Spectrum of the simulated multi-bunch output of the stripline.

Figure 4.28 has the characteristic shape of the single bunch spectrum shown in figure 4.5 overlaid with thermal noise and harmonics of the 357 MHz bunching frequency.

The subtraction hybrid is simulated as a device that rotates the phase of the bottom stripline signal by 180° and adds this to the top stripline. The output power is reduced by 50% and low-pass filtered at 2 GHz to simulate the insertion loss and limited bandwidth of the device, however it is modelled as being perfect in all other ways. Figure 4.29 shows the simulated output of this device, and it can be seen to be very similar to the stripline output, but with wider pulses due to the reduced bandwidth. The reduced amplitude is a function of the small beam offset, the 3 dB insertion loss, and the narrower bandwidth.

Figure 4.30 shows the spectrum calculated from figure 4.29, and it can be seen to be very similar in shape to figure 4.28 with the power reduced by ∼40 dB. This is to be expected as voltage amplitude will have reduced by two orders of magnitude due to the 1% beam offset. Thus the power will be reduced by four orders of magnitude, or 40 dB, due to it scaling with the square of the voltage.

Figure 4.31 shows the simulated voltage output of the band-pass filter, and it can be observed that there is a reduction in amplitude due to the restricted bandwidth. The spectrum of this signal is shown in figure 4.32, and the attenuation of the out of band power is clearly visible. Also, note the power of the 714 MHz component has not been decreased.

Figure 4.33 shows the mixer LO and the output of the BPF, and it can be seen that they are in phase as required.

Figures 4.34 and 4.35 show the output from the mixer in the time domain and frequency domain respectively. In both of these plots the 714 MHz LO leakage signal is clearly visible. Also visible in figure 4.35 are the harmonics of 357 MHz. Importantly, the 714 MHz signal

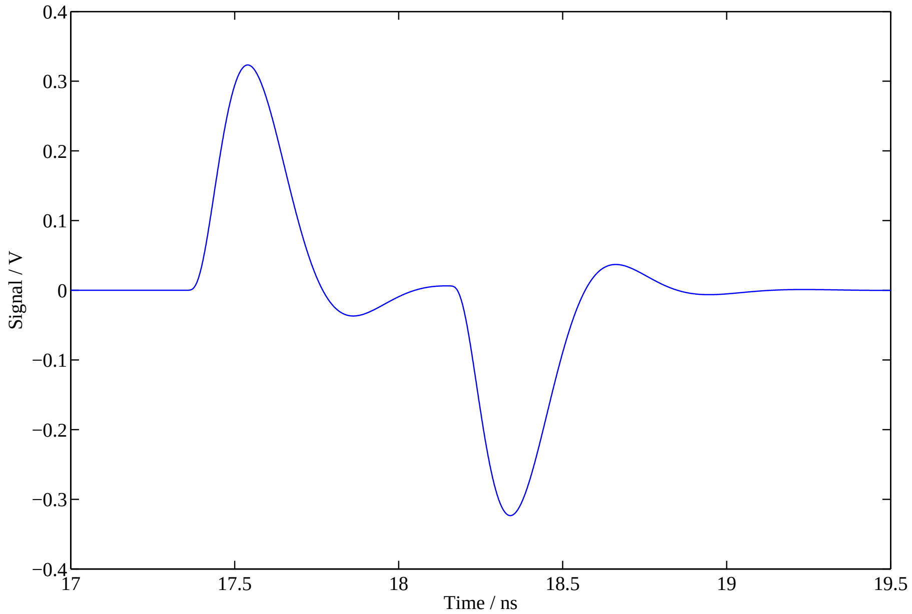

Figure 4.29: Simulated single bunch output of the hybrid subtraction device.

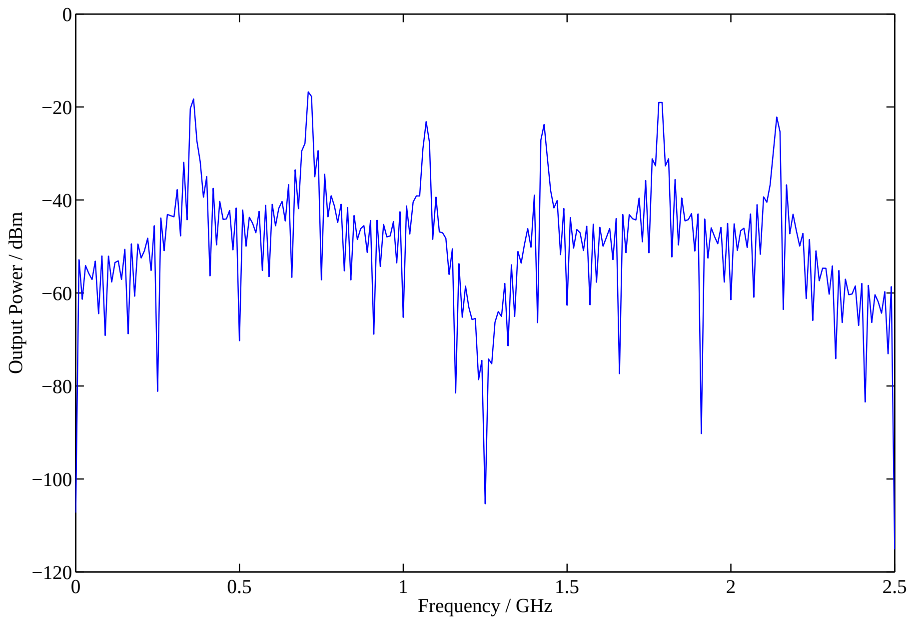

Figure 4.30: Spectrum of the simulated multi-bunch output of the hybrid.

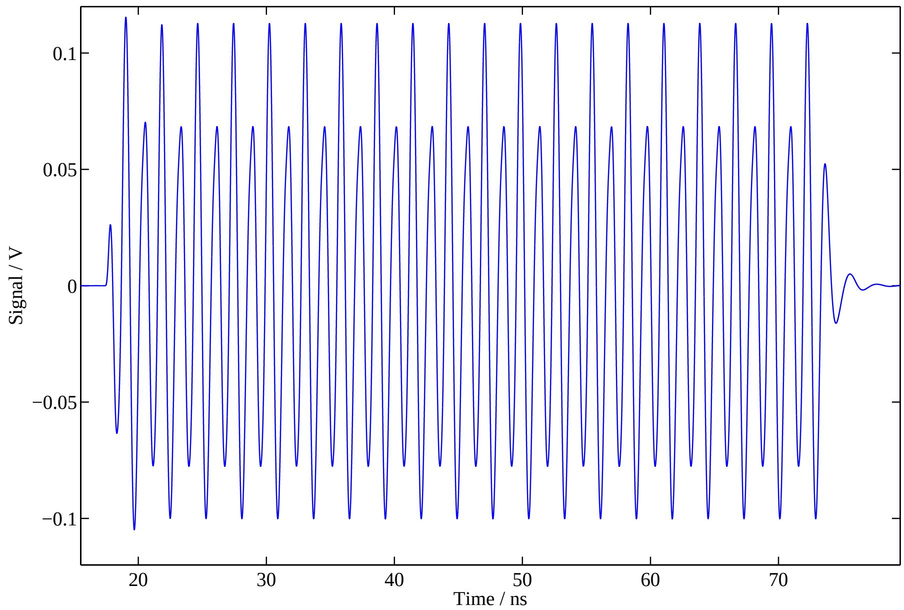

Figure 4.31: Simulated single bunch output of the BPF.

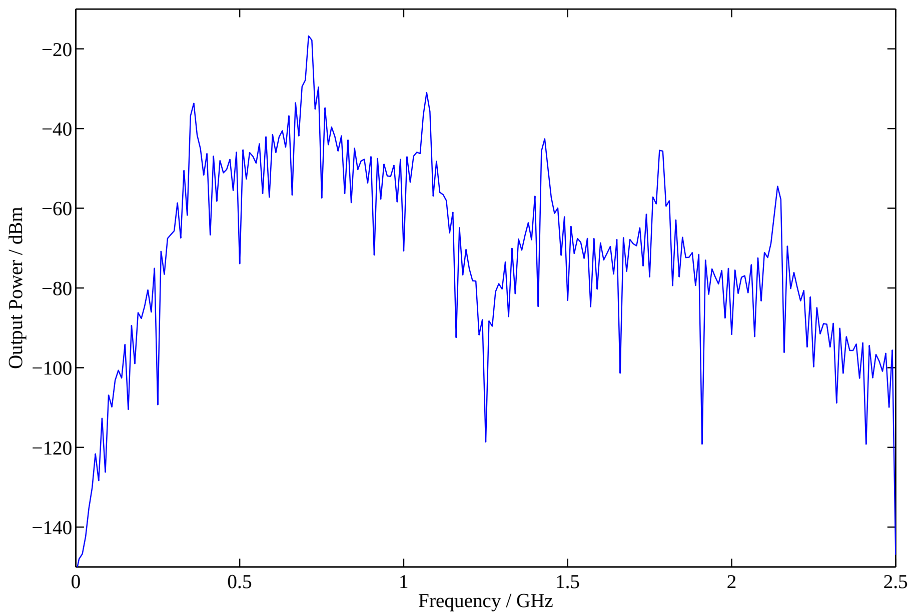

Figure 4.32: Spectrum of the simulated multi-bunch output of the BPF.

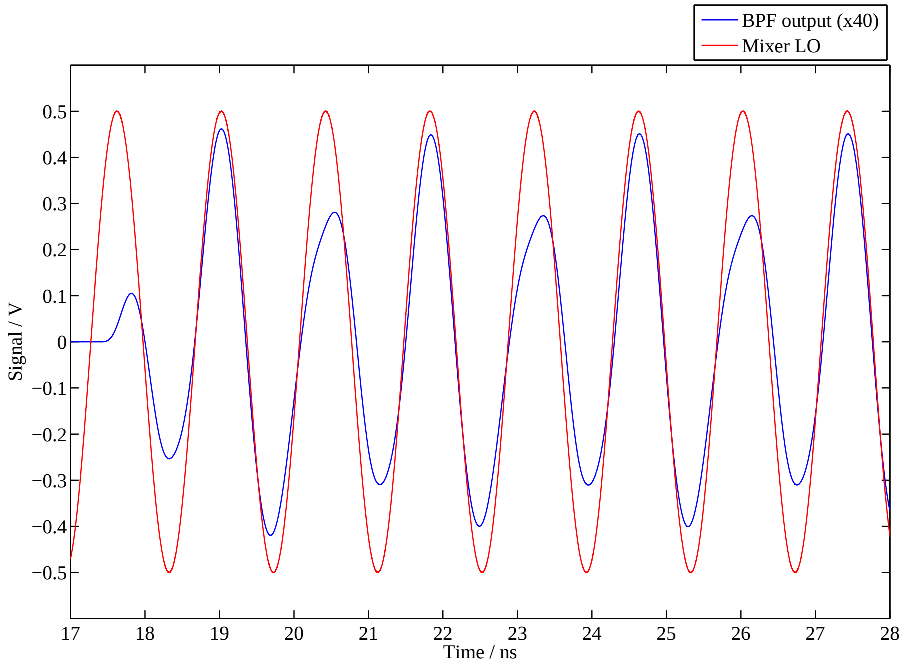

Figure 4.33: Comparison of the phase of the mixer LO and the input from the BPF for the first four bunches of the train.

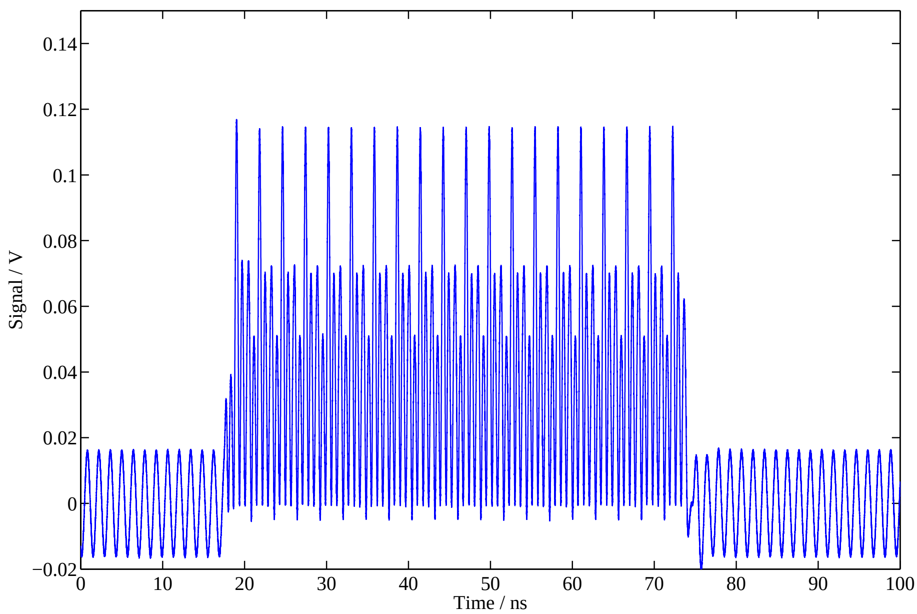

Figure 4.34: Simulated multi-bunch output of the mixer.

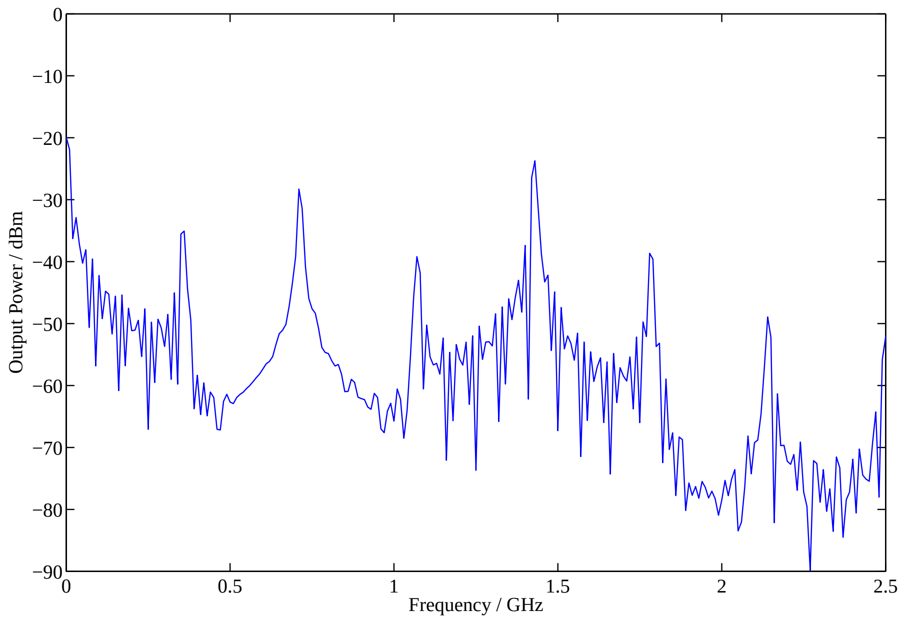

Figure 4.35: Spectrum of the simulated multi-bunch output of the mixer.

component from figure 4.32 has been mixed down to baseband with very little power loss[4].

Figures 4.36 and 4.37 show the output after the LPF, i.e. the final processor output. After some ringing at the front end of the pulse characteristic of a Chebyshev filter, the signal settles down to a value of ∼40 mV, giving a calibration constant of ∼0.167 V/mm. It can be seen from figure 4.37 that the baseband power has not been reduced, and that all components greater than ∼200 MHz have been strongly attenuated.

Reducing the beam offset to 1 μm – the design resolution – will cause a drop of slightly more than two orders of magnitude in the voltage output, or a ≥40 dB drop in the power. Therefore, the baseband power output due to a 1 μm beam offset will be ≤-60 dBm, which is within an order of magnitude of the power at 357 MHz and 714 MHz.

Figure 4.38 is a comparison of the noise floor of the processor with the output spectrum for a beam offset of 240 μm, and it can be seen that the baseband noise floor is ∼-80 dBm. The simulated resolution of the device is, therefore, ∼0.24 μm.

Figure 4.39 shows the start of the processor output overlaid with the raw stripline output in order to determine the signal delay. The vertical lines mark the time of arrival of the first stripline pulse and the point at which the processor output passes through half its maximum amplitude. A latency of 4.1 ns is predicted.

[4]The simulated insertion loss of the mixer is 3 dB.

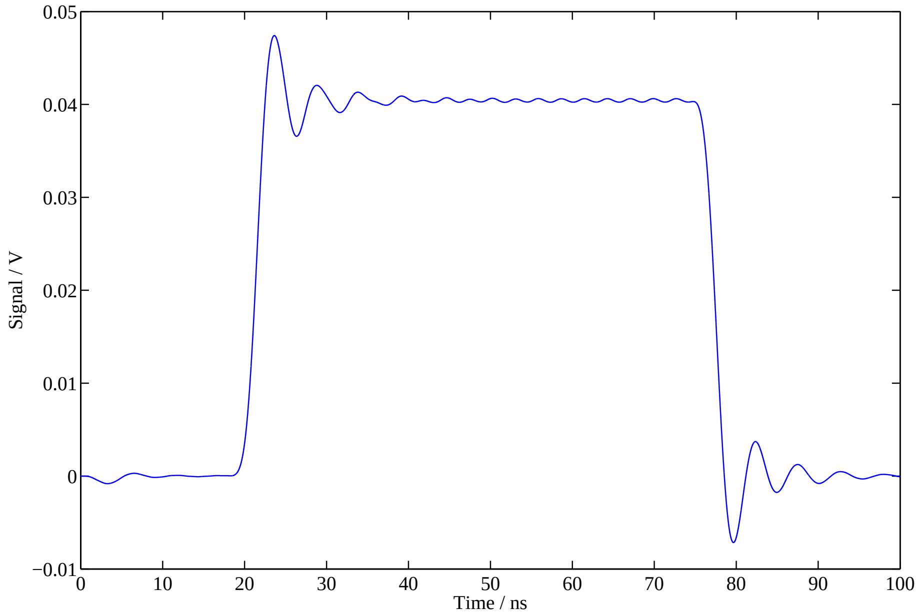

Figure 4.36: Simulated single bunch output of the LPF.

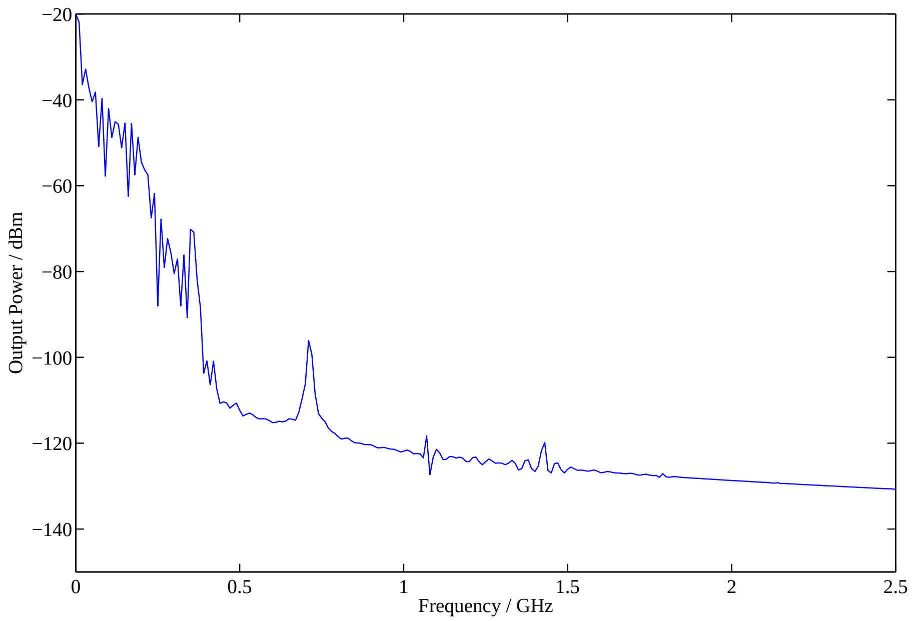

Figure 4.37: Spectrum of the simulated multi-bunch output of the LPF.

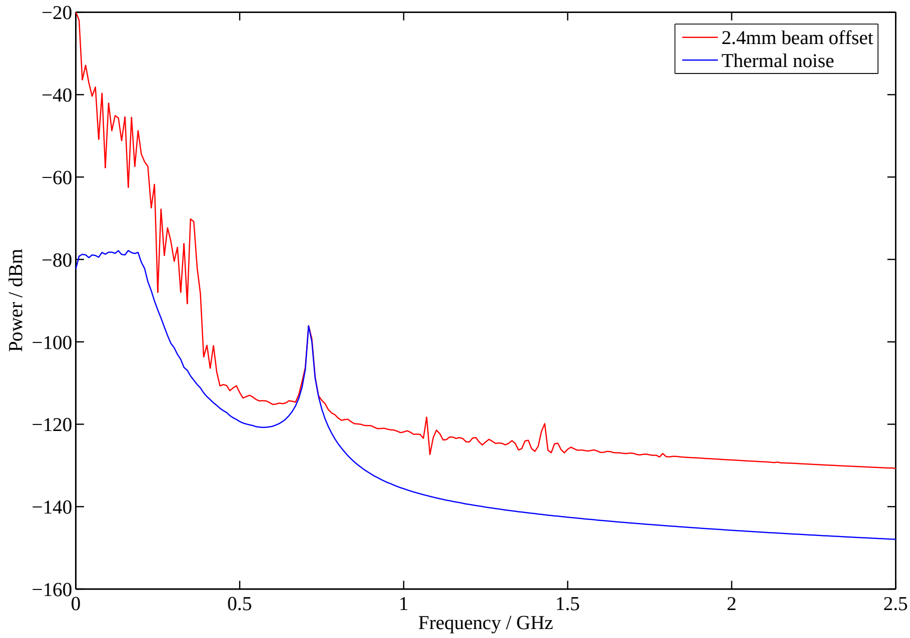

Figure 4.38: Comparison of the output spectra for a beam offset of ~240 μm and for pure noise.

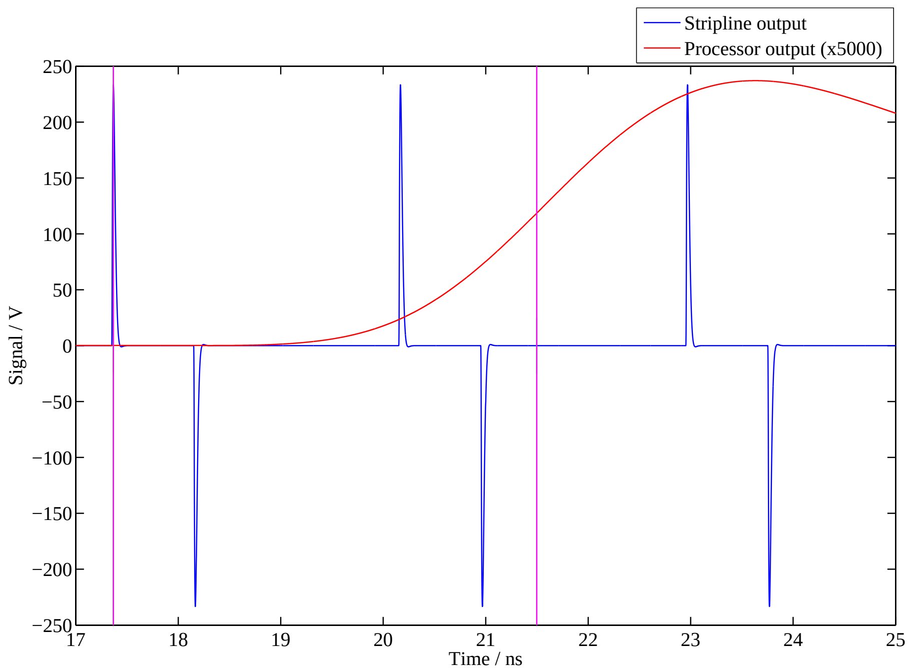

Figure 4.39: Comparison of the stripline output and processor output to determine the delay.

Chapter 5

FONT3 – Processor Performance

5.1 Single-Bunch Response

As the processor was designed to work near a maximum of the stripline output spectrum a valid position output can be obtained when either a single bunch or many bunches pass the striplines. In order to simplify the initial analysis, it was decided to begin the investigations of the processor's performance with single bunch beam, and move onto the multi-bunch performance at a later stage. Each component in the processor chain was measured in order to understand its response to a range of beam offsets. An initial calculation was performed to determine the dynamic range, i.e. the range of beam motion in which the processor will give an output linear with beam position, and there is no risk of damage to any components.

5.1.1 Calculation of the dynamic range

Table 5.1 gives the maximum power and voltage for each component in the processor, and it can be seen that the mixer sets the limit, so it can be said that the maximum deviation of the beam to maintain linearity is one that causes a BPF output of $\leq$0.25 V.

Section 4.4 showed the results of simulations of the processor's response to a 240 μm position offset, and figure 4.31 shows that this causes a maximum output of $\sim$0.12 V. Since the voltage output is linear with position, it can be concluded that the maximum power of the mixer is reached for an offset of $\sim$0.5 mm, and care should be taken to ensure the beam

Component	**Max. Power / dBm**	**Max. Voltage / V (into 50 Ω)**
Hybrid	33	10
Mixer	1	0.25
Filter components	24	2
Scope inputs	27	5

Table 5.1: Maximum power and voltage specifications for the processor components. The max power for the mixer is defined as the output power at which the mixer output is 1 dB below the output expected for perfect linearity.

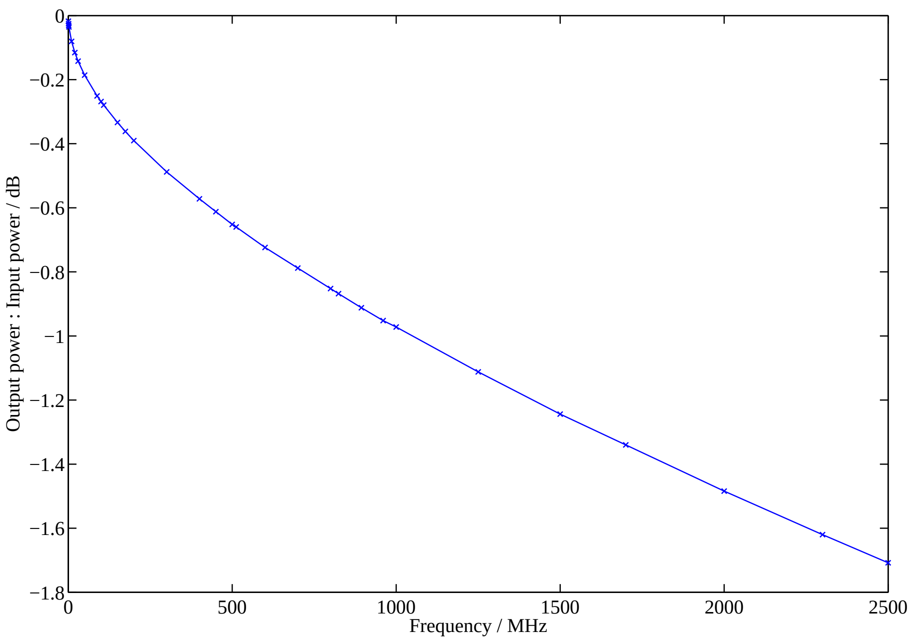

Figure 5.1: Attenuation as a function of frequency for the length of 5/8" heliax[21] used in the experiment (40 m).

position never exceeds this. The simulations shown in section 4.4 were for a bunch charge of 1 nC, which is $\sim 6\times 10^9$ e^-/bunch, and, since the processor output will also scale with the bunch charge, the safe beam offset will depend inversely on the bunch charge. For example, with a bunch charge of 1 $\times 10^9$ e^-/bunch, the safe offset will be ∼3 mm.

5.1.2 Induced signal on stripline

The BPM striplines were connected to an oscilloscope by 5/8" diameter heliax cable in order to measure their output as a function of the beam position. The signals measured will be deteriorated by both the frequency dependent attenuation of the heliax (see figure 5.1), and the sampling rate and limited bandwidth of the oscilloscope[1].

Figure 5.2 shows the average output of forty single pulses passing the BPM stripline (in this case the upper stripline of BPM13). For a 12 cm long strip, the predicted time between the positive and negative pulse, $t = 2l/c$, is 0.8 ns, and this agrees well with the measured delay seen in figure 5.2.

The positive and negative pulses in figure 5.2 have a FWHM of ∼0.7 ns, and this seems to imply that this is the bunch length, however, this measurement neglects the contribution of the cable and the bandwidth of the scope. The ATF Design and Study Report[26] gives a measured bunch length (FWHM) of $\sim 20\pm 5$ ps, which would require a system bandwidth

[1] In this case the sampling rate is 5×10^9 Samples/s, giving a Nyquist frequency of 2.5 GHz. The scope also has an input filter with a 3 dB point of ∼1.5 GHz to prevent signals with a frequency of greater than the Nyquist frequency from being measured, therefore the scope is said to have an input bandwidth of 1.5 GHz.

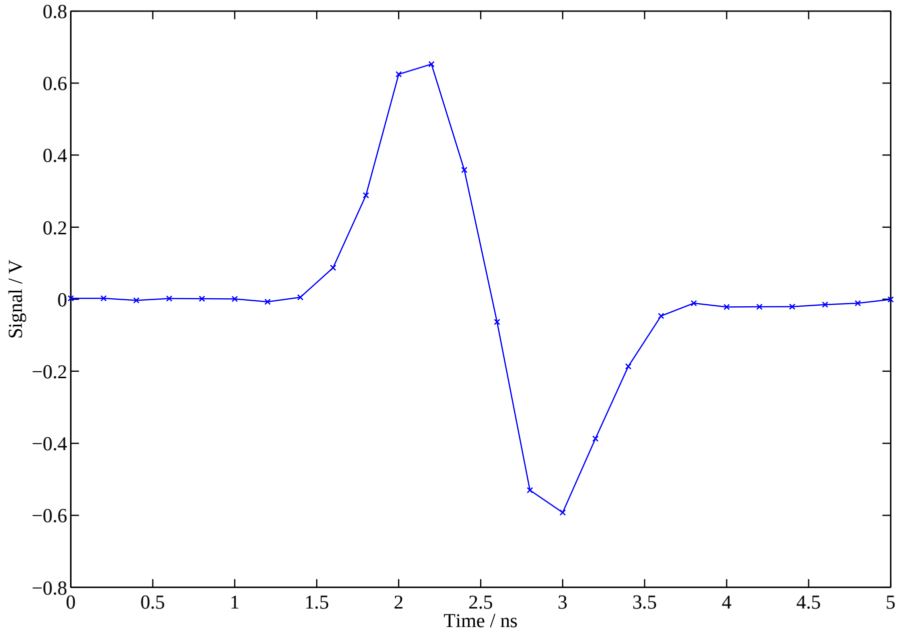

Figure 5.2: Example of the output of a single bunch passing a stripline. This is the average of forty such pulses.

of ∼50 GHz to measure effectively. The pulse shape seen in figure 5.2 is a rough indication of the true pulse shape, and is limited by the bandwidth of the scope. In the following discussion of the effectiveness of each component the signal spectrum will be demonstrated, however, this will not be done for this case as this would merely reflect the bandwidth of the scope and provide little information on the stripline signal itself.

5.1.3 Subtraction and summation by the hybrid

The beam position was measured at three points in the beamline, BPMs 11, 12, and 13 (see figure 4.2), and the charge was measured at BPM 13. M/A-COM[54] hybrid H-81-4-SMA was used to find the difference between the striplines, however, since the sum output of this device was internally terminated, the sum port of an Anzac H-1-4 hybrid[2] was used to measure the beam charge.

The electrical centre of the BPM was roughly determined by sweeping the beam to find a minimum in the signal output. The beam was then moved to several positions around this point, and the response of the hybrid to the beam was recorded.

Figure 5.3 shows the average power spectrum calculated for each beam position setting. The power was calculated using the measured voltage and the characteristic impedance of the system, Z_0 =50 Ω, and is given in units of dBm. At each position setting the theoretical shape of the stripline output shown in figure 4.5 is observable. The peak at ∼2 GHz is slightly reduced in power with respect to the ∼625 MHz lobe due to the bandwidth of the

[2]No specification document could be found for this device.

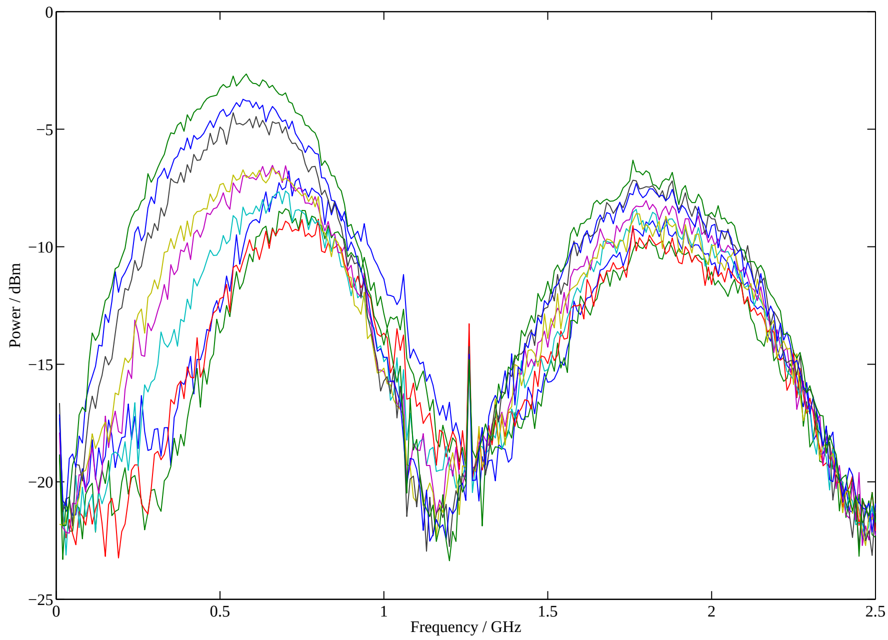

Figure 5.3: The average power spectrum output from the hybrid at various beam positions.

hybrid.

Also noticeable is the fact that the shape of the low frequency peak changes as the total output power reduces (i.e. as the beam moves closer to the electrical centre of the device), however, it is more accurate to consider this one lobe as the sum of two broad signals – one centred at ∼625 MHz containing the beam position information, and the common-mode component centred at ∼720 MHz. This common-mode component has no dependence on beam position and so should remain constant throughout the dipole scan. Its frequency is determined by a number of factors (including the scale of any impedance or path length mismatches in the hybrid, or in the signal paths from the BPM striplines to the hybrid) which make it extremely complicated to predict a priori.

The component at ∼625 MHz is strongly affected by beam position as expected.

5.1.4 Band-pass filter (BPF)

The beam was moved to several positions around the electrical centre and the output of the BPF was recorded. Figure 5.4 shows the average power spectrum for each of these positions. As expected it shares many of the same features of figure 5.3 – both the common-mode, and the position dependent signal are still present in the large lobe at ∼650 MHz – however, the spectrum is clearly limited below ∼300 MHz, and above ∼1.2 GHz due to the action of the BPF.

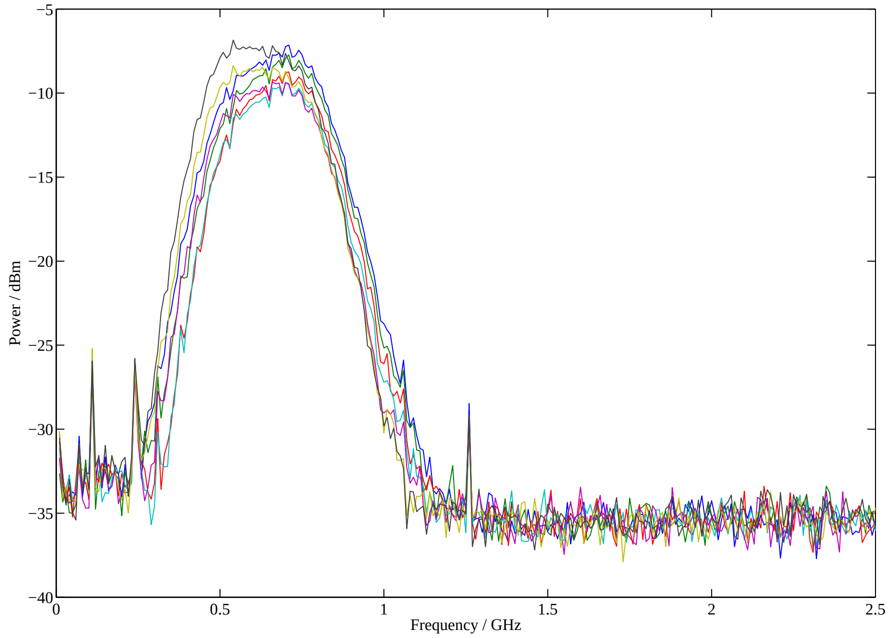

Figure 5.4: The average power spectrum output from the BPF at various beam positions.

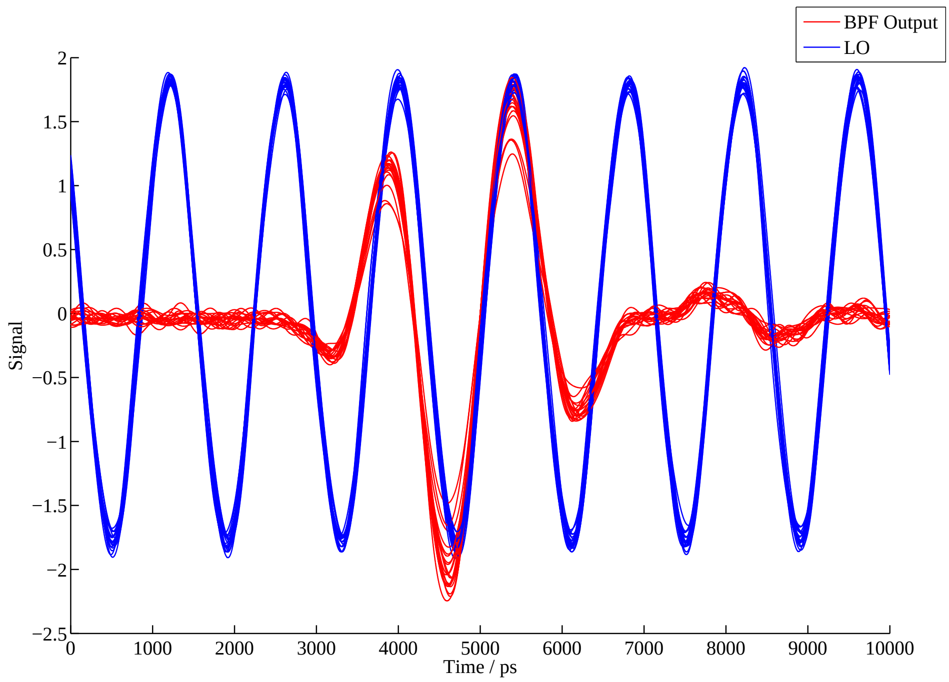

Figure 5.5: The output of the BPF for a relatively large, positive, beam offset (∼0.5 mm) and the 714 MHz LO as measured at the inputs to the mixer. The y-scale is in arbitrary units as the plots were scaled to make comparison simpler.

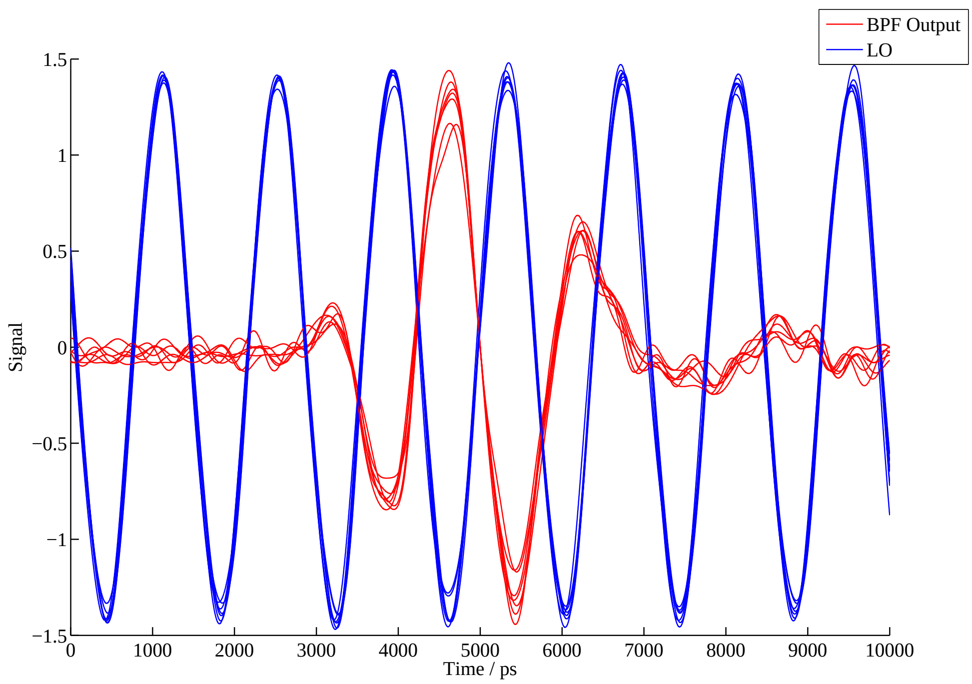

Figure 5.6: The output of the BPF for a relatively large, negative, beam offset (∼-0.5 mm) and the 714 MHz LO as measured at the inputs to the mixer. The y-scale is in arbitrary units as the plots were scaled to make comparison simpler.

5.1.5 Mixer and LO phasing

Figures 5.5, 5.6, and 5.7 show the output of the BPF compared to the LO input to the mixer for a positive (∼0.5 mm) beam offset, a negative (∼-0.5 mm) offset, and approximately zero position.

In figure 5.5 it can be seen that the peaks and troughs of both the LO and the BPF output are well aligned. Figure 5.6 shows the equivalent signals for the case where the beam is negatively offset, and it can be seen that, as expected, the phase of the BPF output has changed by π with respect to the case of a positive beam offset. A phase reversal of π is equivalent to a polarity change of the signal, and is due to the lower stripline having a larger amplitude than the upper one. This results in the same signal shape as in the case of a positive offset, but with the opposite sign.

The case where the beam is located very close to the electrical centre of the device is shown in figure 5.7. Using "perfect" electronic devices, a zero beam position should result in zero signal output from the hybrid, however this figure shows that there is indeed a component output at zero position.

Figure 5.7 also confirms the observation that the common-mode component is at a different frequency than the position dependent signal (see figure 5.4), as this can be seen to be affecting the smooth shape of the signal peaks and troughs. For example, the large, positive, peak in figure 5.7 can be seen to consist of two signals with a different phase. The same two signals are apparent in the final positive peak, and it can be seen that the time separation

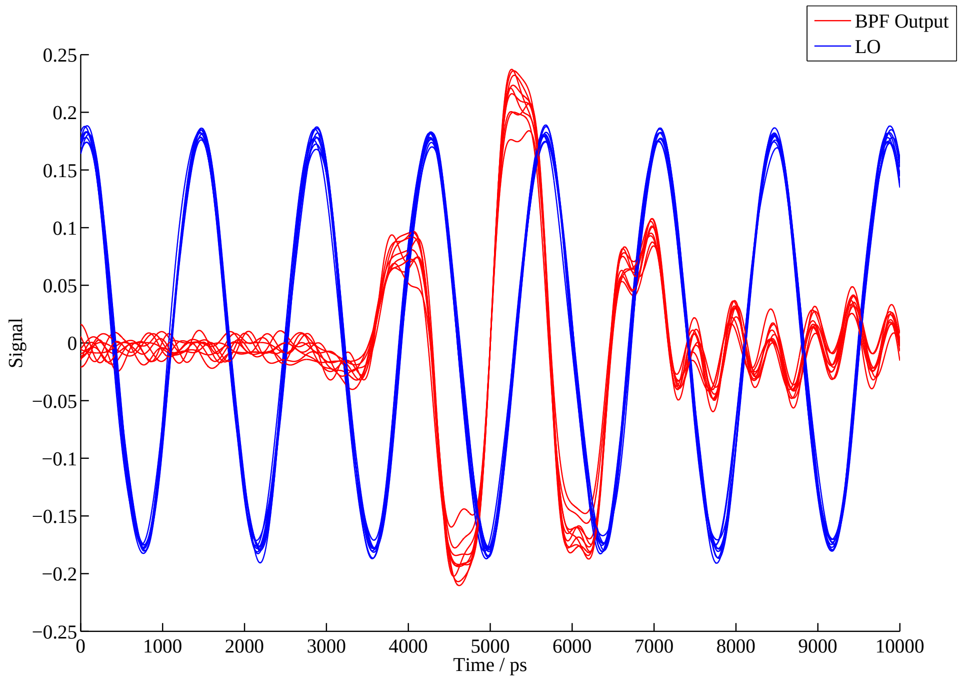

Figure 5.7: The output of the BPF for approximately zero beam offset and the 714 MHz LO as measured at the inputs to the mixer. The y-scale is in arbitrary units as the plots were scaled to make comparison simpler.

of the peaks of these signals is larger than in the large, central peak, i.e. these signals have different frequencies.

Figure 5.7 confirms that the phase of the common mode differs by $\pi/2$ from the phase of the position dependent signal.

The method described in section 4.3.6 was implemented in order to determine the correct LO phase, however, results to demonstrate the success of this method will not be shown for the case of single-bunch beam. See section 5.3.6.

Figure 5.8 shows the calculated average power spectrum output by the mixer at various beam positions, and several interesting features can be observed.

Three very narrow components can be seen at ~700 MHz, ~1400 MHz, and ~2200 MHz, which are harmonics of the 714 MHz LO signal leaking into the mixer output. The ~700 MHz component has a power of ~-1 dBm, which is ~8 dB below the input LO power of +7 dBm, i.e. the mixers can be said to have a LO leakage level of ~-8 dB.

The main point of interest in figure 5.8 is that the broad, beam-induced, response that was seen at ~650 MHz in figure 5.4 has been split into two broad peaks – one at ~1250 MHz, and one at $\leq$250 MHz. That is, a peak with a frequency of the sum of the frequency of the original and the LO, and a peak whose frequency is the difference of these[3]. Also note that the signal power in these components is significantly lower than the unmixed signal in figure 5.4. As this peak has been split into two components, conservation of energy requires that

[3]Since a negative frequency is not physical, it is actually the absolute value of this difference.

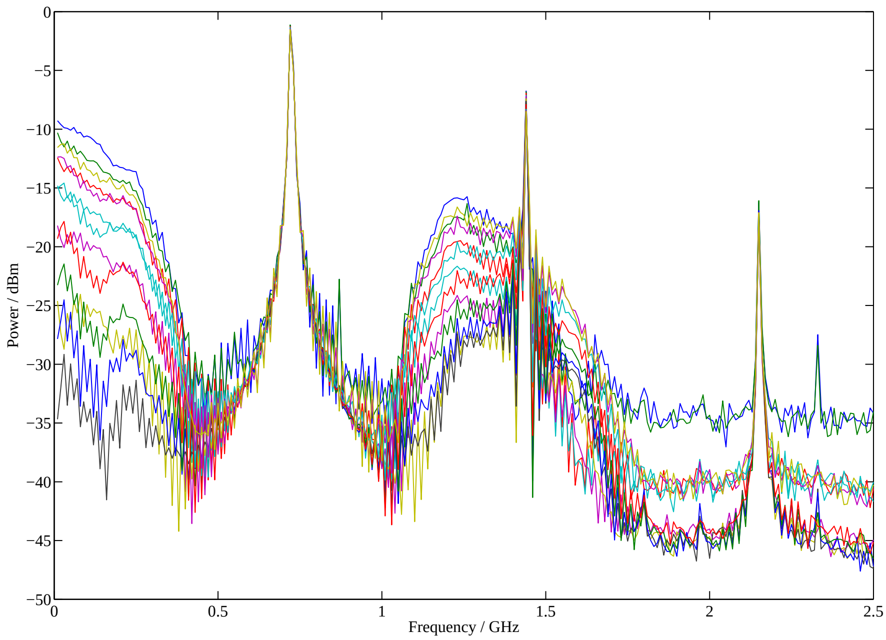

Figure 5.8: The average power spectrum output from the mixer at various beam positions.

they have at most half the power of the original, i.e. the mixed peaks should be at least 3 dB lower than the unmixed signal[4].

Note the high frequency noise in figure 5.8. The vertical scale of the scope was used to maximise the resolution of each waveform, so the resolution of the scope was increased for smaller amplitude signals, which leads to a lower measured noise level. This is why figure 5.8 shows several distinct power levels for the high frequency noise.

5.1.6 Low-pass filter (LPF)

The output of the LPF was measured by the same method as for the previous components.

Figure 5.9 shows the average spectra obtained at various beam positions. This demonstrates that the filter has reduced almost all the components above ∼300 MHz to a level where they are at the limits of the scope's resolution. The 714 MHz LO leakage component has been significantly attenuated to a level of ∼-35 dBm. Note that this is equivalent to the lowest power levels measured in the low frequency signal, i.e. the mixer leakage signal is similar in power to the noise level in the bandwidth of the position dependent signal. This is as predicted with the Systemview model (see section 4.4).

In figure 5.9 the noise is divided into the same distinct power levels as noted in figure 5.8.

The power output at 10 MHz for each of the beam positions was extracted from figure 5.9 and the RMS voltage associated with this power was calculated and normalised with the

[4]Due to imperfections in the mixer, this value is actually closer to 4 dB.

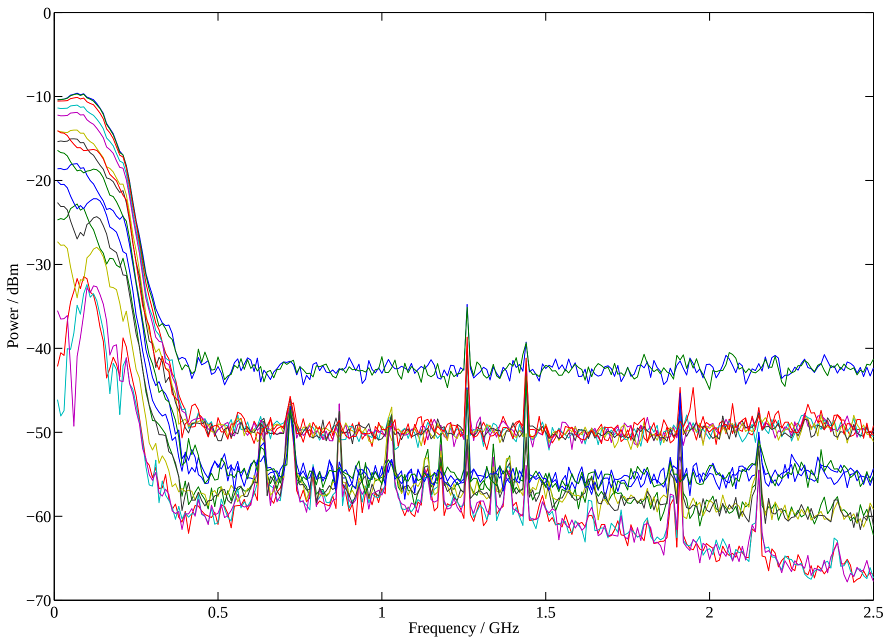

Figure 5.9: The average power spectrum output from the LPF at various beam positions.

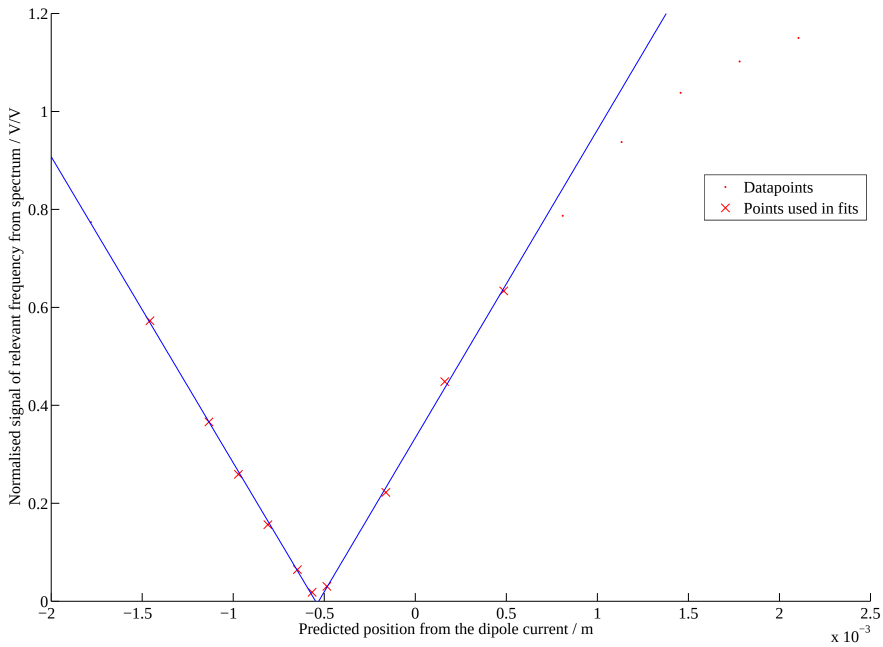

Figure 5.10: The charge normalised and averaged voltage of the 10 MHz component output of the LPF plotted against each beam position.

RMS voltage calculated from the sum spectrum[5]. Figure 5.10 shows how this normalised output varies with the beam position, where the beam position is calculated from the current change in the dipole used to steer the beam. It can be seen that the dynamic range is ∼2 mm. Straight lines where fit to the two linear sections and the gradients were found to be 629±17 m^{-1} and -625±8 m^{-1}. The absolute values of these numbers are consistent with one another as expected since the sensitivity of the processor should not depend on the polarity of the beam offset. Also, there is no evidence of the common-mode signal at any beam position. This is due to the fact that the common-mode signal had a higher frequency than the beam position dependent signal, therefore this signal will be more strongly suppressed by the LPF.

5.1.7 Processor delay

The processor delay was measured as follows. One stripline was connected to the scope, while the other was connected to an input of the hybrid and attenuated to a safe power level. The other hybrid input was terminated with 50 Ω to prevent signal reflections within the hybrid affecting the results. As there is a certain time delay across the protective attenuators, it was important to include the equivalent amount of attenuation on the cable bringing the raw stripline input to the scope – this has the additional benefit of protecting the scope from the excessive signal amplitude output from the stripline. The processor output was measured with only one stripline input attached, while the opposing stripline was attached directly to the scope. This allowed the time difference between the raw stripline signal and the processor output to be calculated. To maximise the accuracy of the measurement the input cables to the scope were swapped, i.e. attach the raw stripline signal back to the disconnected hybrid input (including the attenuators), and connect the other stripline signal to the scope. Once again, a number of pulses are recorded and the time difference between the pulses can be calculated.

For both of these cases the scope should be triggered with an external, beam related, signal. The processor delay is defined as the average of the measured delays.

Figure 5.11 shows an example of the input and output signals measured as described above. The delay was calculated as the time difference between the positive peak of the stripline output and the $1/e$ point of the rising edge of the processor output.

Figure 5.12 shows the spread of the results of the analysis of the processor signal delay. The 200 ps quantisation in the results is due to the 5 GHz sampling rate of the scope. The measured value for the latency of this device is 4.6±0.3 ns, which conforms to the 5 ns goal set in section 4.2. This value is higher than the value of 4.1 ns predicted by the simulations (see figure 4.39), however this simulation did not take the physical size of the components into account, so this would add a certain amount of time-of-flight delay to the latency. The difference between the measured and simulated values implies a signal path of ∼10 cm, and this is consistent with the actual processor.

[5]These data were taken at an early stage in the experiment before the beam had been tuned for low charge variation (see figure 4.3), so it was important to normalise by the sum signal.

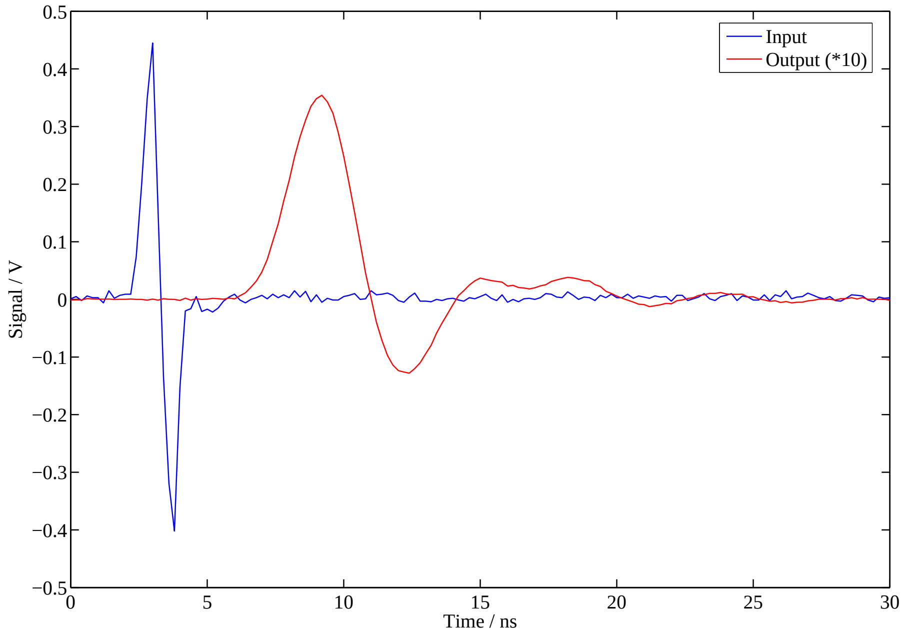

Figure 5.11: An example of the input and output signals of the processor.

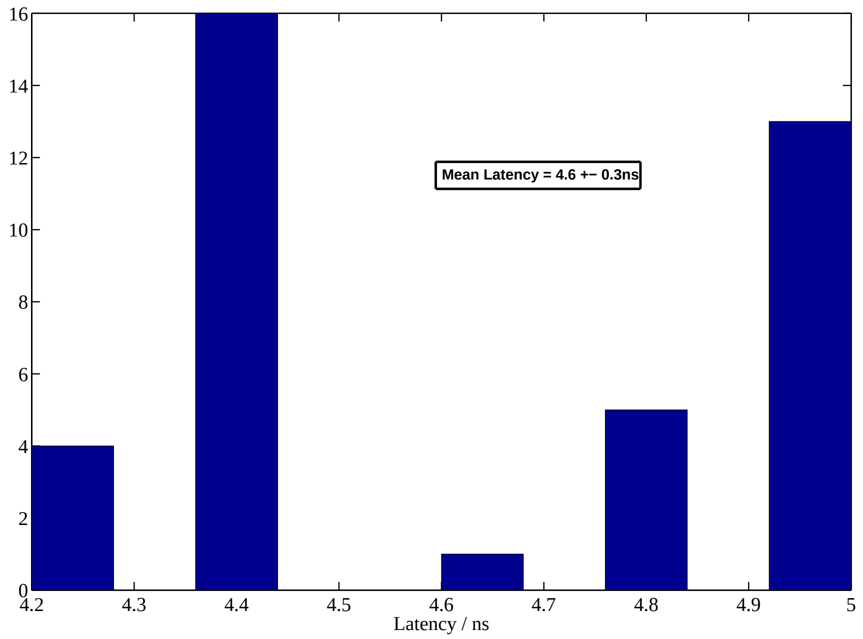

Figure 5.12: Spread of forty values of the processor signal delay.

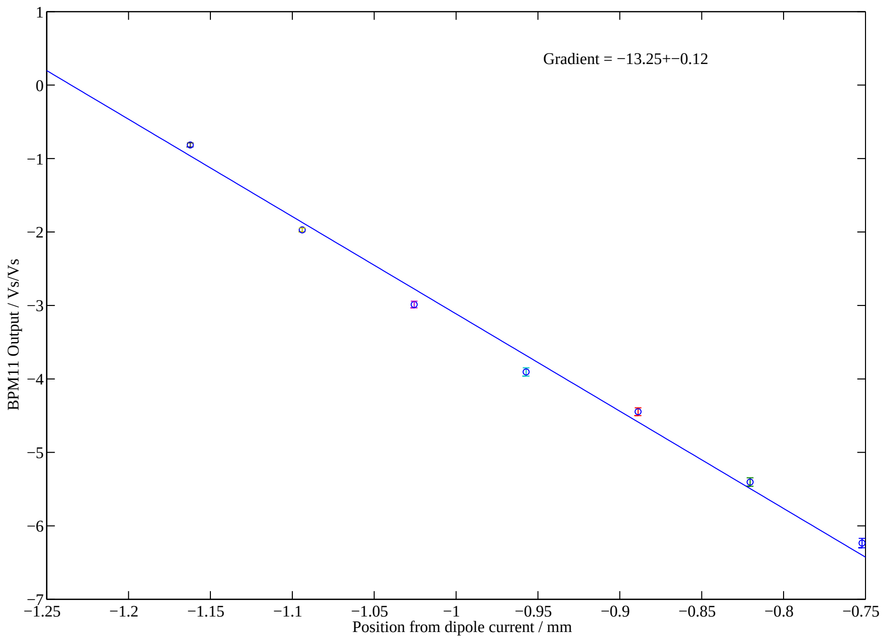

Figure 5.13: Output of BPM11, normalised by the charge as measured at BPM13, plotted against the position of the beam as calculated from the change of current in an upstream dipole.

5.1.8 Calibration and resolution

Three extraction line BPMs were instrumented with processors (BPMs 11, 12, and 13), and all magnetic elements between these BPMs were switched off in order to simplify the calibration. All the BPMs were instrumented to measure the y position, however an additional processor was installed on the x plane of BPM 13 with a summing hybrid to measure the total beam charge for each pulse. All 4 BPMs were setup as described in section 4.3.6[6].

The outline of the method used to calibrate the BPMs is the same as for the FONT2 processor (see section 3.3.4).

Due to timing jitter of the scope trigger and the presence of scope noise when observing single data points the BPMs were not calibrated with respect to the peak of the processor output. Instead the BPMs were calibrated with respect to the integral over 100 ns of the pulse. Figure 5.11 shows that the main peak of the output is ∼7 ns wide, and that the oscillations after the pulse last ∼20 ns, so the integral over 100ns will include the entire output of the BPMs, and will not be subject to the timing jitter of the data-acquisition.

Figures 5.13, 5.14, and 5.15 show the calibration curves obtained for BPMs 11, 12, and 13, respectively, during single bunch operation. The BPM output was calculated by dividing the integral of the difference signal from the BPM, by the integral of the sum signal from

[6]The LO phasing method applied to the difference BPMs could not be applied to the sum BPM since it is not possible to alter the polarity of the charge signal, so the phase was set so as to maximise its output.

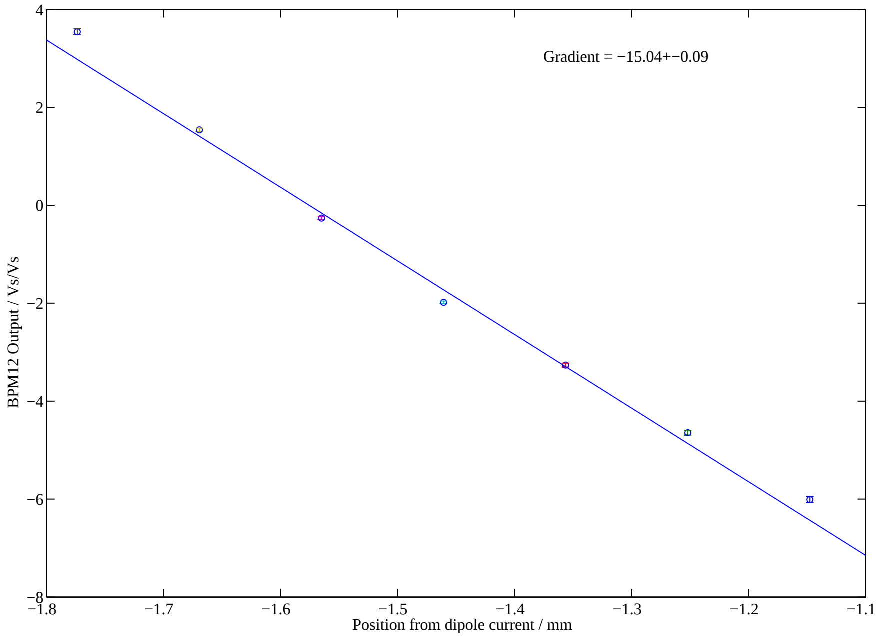

Figure 5.14: Output of BPM12, normalised by the charge as measured at BPM13, plotted against the position of the beam as calculated from the change of current in an upstream dipole.

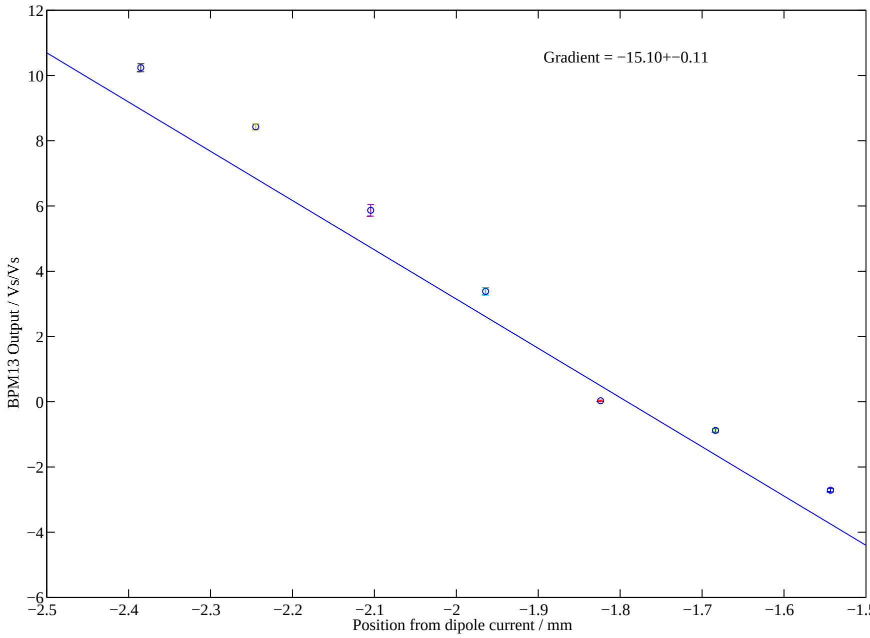

Figure 5.15: Output of BPM13, normalised by the charge as measured at BPM13, plotted against the position of the beam as calculated from the change of current in an upstream dipole. The data-point close to zero BPM output has a very small error, and it is this that causes the offset of the fit to be moved downwards.

the x plane of BPM 13. Each point in these figures is the average of forty beam pulses[7]. The errors, E on the points were defined as,

$$E \equiv \frac{\sigma}{\sqrt{N}} \tag{5.1}$$

where σ is the RMS spread of the data, and N is the number of data-points. These figures demonstrate the linearity of all three BPMs throughout a range of ~0.5 mm for BPM 11, ~0.7 mm for BPM 12, and ~0.9 mm for BPM 13. These differing values are due to the fact that all three devices were measured using the same beam pulses, and the spread is a function of their distance from the dipole causing the beam motion.

The plots shown in figures 5.13, 5.14, and 5.15 show that the processors are roughly linear, but it can also be seen that the calibration is far from perfect, with many of the points lying considerably off the straight line fit. The following values are the values calculated for the chi-squared per degree-of-freedom for each of the straight line fits, quantifying the poor quality of the fits.

- BPM 11 : 20.2
- BPM 12 : 54.2
- BPM 13 : 312.6

It is a known feature of the ATF dipoles that the resolution of the Digital to Analogue Converter (DAC) is relatively rough, and the magnet power supplies will not necessarily go to the exact current requested by the control system. The spread of the points around the straight line fit is, therefore, due to a lack of knowledge of the exact dipole current (thus the beam position is poorly known) and not due to processor noise. If the errors on the value of the dipole current are distributed randomly, then the straight line fit drawn through these points should average out the error and lead to a calibration constant very close to the true value.

Figures 5.16, 5.17, and 5.18 show the results of the resolution calculation. In each case a least squares fit was performed on the outputs of the three BPMs during a dedicated data run of two hundred beam pulses. The outputs were fit to a straight line as all magnetic elements between the BPMs had been switched off, and the position at each BPM was predicted from the outputs of the other two. This was subtracted from the measured position at the BPM. As this calculation uses the output from all three BPMs, the result will contain a noise contribution from all three devices. This is the same as for the FONT2 BPM (see section 3.3.5) and the geometrical factor is found to be $\sqrt{\frac{2}{3}}$. The values plotted in figures 5.16, 5.17, and 5.18 have been corrected with this factor.

The resolution of each device was defined as the standard deviation of the residuals, and the results are as follows,

- BPM11 : 20 μm

[7]Forty beam pulses were recorded at each beam setting, however, some of these pulses were rejected on the basis of low charge, or due to a machine fault causing the beam to temporarily make a large deviation from its regular orbit.

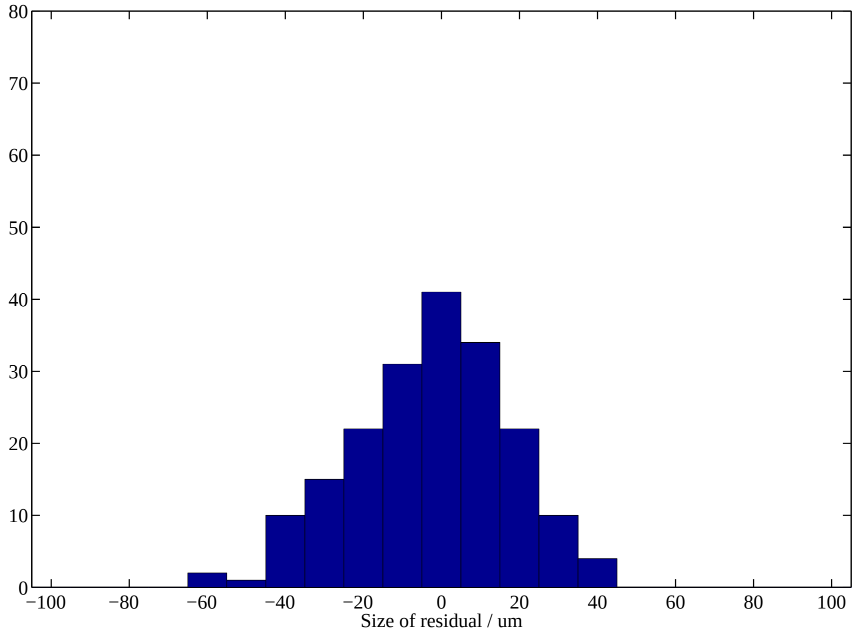

Figure 5.16: The spread of the residuals from the comparison of the position measured at BPM 11 to the expected position calculated from BPMs 12 and 13. The value for the residual has been multiplied by a geometrical factor of $\sqrt{2/3}$.

- BPM12 : 10 μm
- BPM13 : 28 μm

These resolution numbers are considerably higher than the 1 μm level stated as a goal of the project, however, these numbers should be considered an upper limit. This data was taken at an early stage in the experiment when the method used to correctly phase the LO had not been perfected. Also, the output of the processors was not amplified, allowing scope noise to contribute to the resolution. Later data taken with multi-bunch beam will show a much improved resolution (see section 5.3.7) due to the 714 MHz harmonic of the 357 MHz bunching frequency adding to the available power at the mixing frequency, thereby improving the signal to noise.

5.2 Summary of the Single Bunch Beam Tests

Single bunch beam tests were used to verify that each component operated as expected, and it was confirmed that the linearity of the processor is ∼2 mm (see figure 5.10). It was discovered that the hybrid creates a large common-mode component that is $\pi/2$ out of phase with the difference signal, and is removed by effective phasing of the mixer LO.

An electronic delay of 4.6 ns was measured by comparing the time of arrival of the input and output signals of the processor.

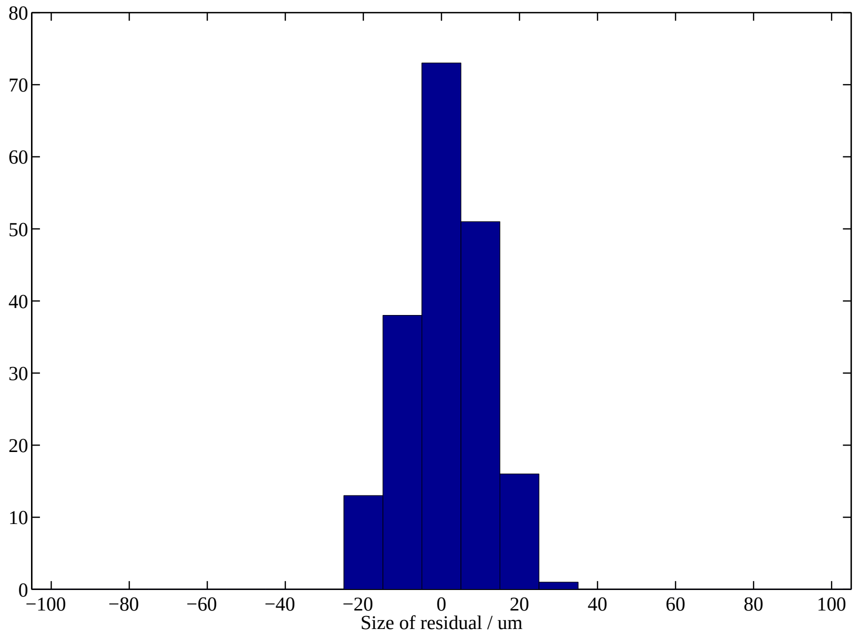

Figure 5.17: The spread of the residuals from the comparison of the position measured at BPM 12 to the expected position calculated from BPMs 11 and 13. The value for the residual has been multiplied by a geometrical factor of $\sqrt{2/3}$.

Using an integration method to analyse the processor output the processors were calibrated and a resolution of $\sim$20 μm was demonstrated. While this value is considerably larger than the goal of $\sim 1\mu$m, this represents an upper limit, which will be reduced using a better LO phasing scheme and with amplification of the signal before measurement (see sections 5.3.6 and 5.3.7).

5.3 Processor – Multi-Bunch Response

The processor response to multi-bunch beam (i.e. a train of twenty bunches separated by 2.8 ns) was analysed in the same way as for the single bunch beam tests, i.e. a step by step investigation of the response of each component, calibration of the devices, and the resolution.

5.3.1 Subtraction and summation by the hybrid

Figure 5.19 shows the output spectra from the hybrid during multibunch operation, and it can be seen to be quite different from the case of single bunch (see figure 5.3). In the case of multibunch, it is the bunching frequency of the beam, 357 MHz, and its harmonics, that dominate.

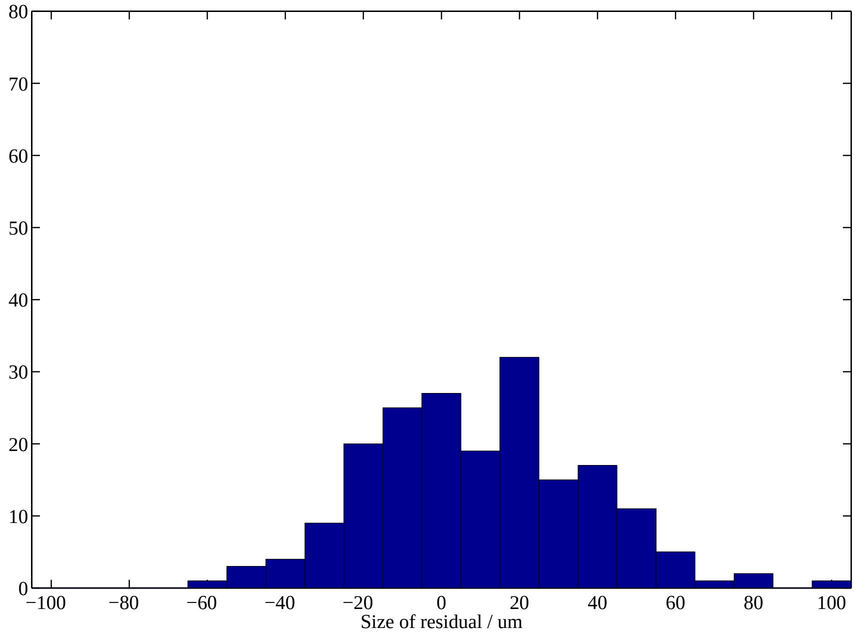

Figure 5.18: The spread of the residuals from the comparison of the position measured at BPM 13 to the expected position calculated from BPMs 11 and 12. The value for the residual has been multiplied by a geometrical factor of $\sqrt{2/3}$.

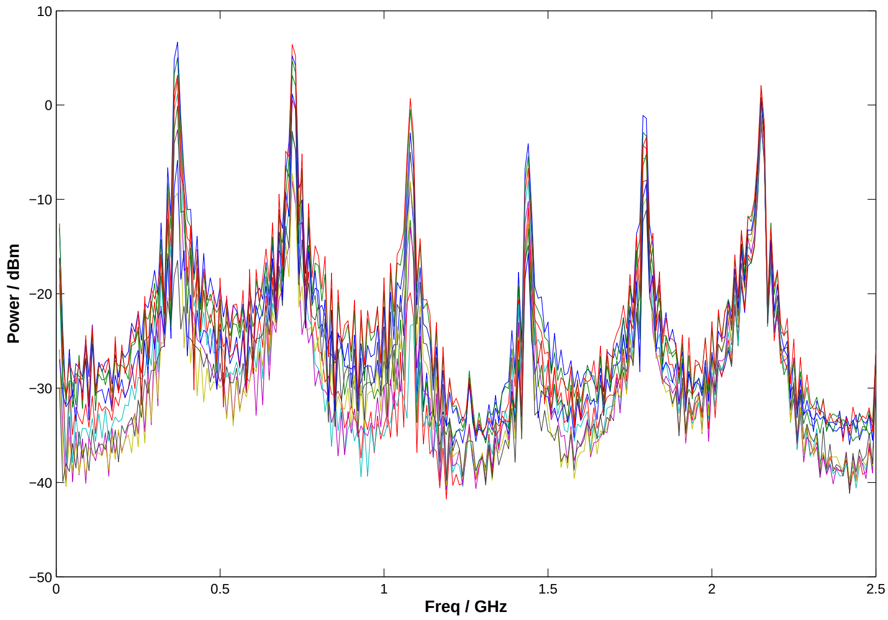

Figure 5.19: The spectra output from the hybrid for various beam positions during multi-bunch operation. Each curve is the average signal from forty beam pulses.

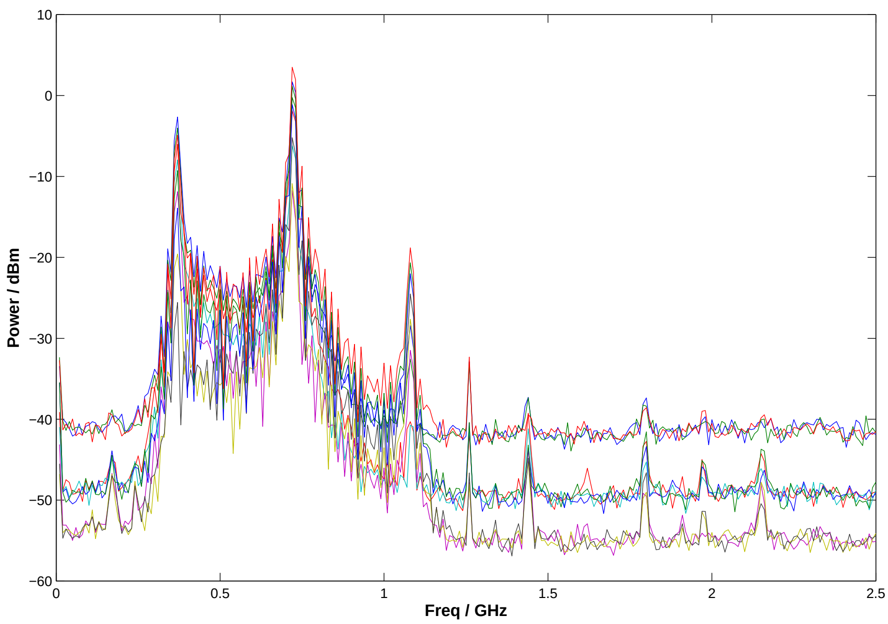

Figure 5.20: The spectrum output from the BPF for various beam positions during multi-bunch operation. Each curve is the average of forty beam pulses.

5.3.2 Band-pass filter (BPF)

Figure 5.20 shows the average spectra output from the BPF for various beam positions. Compare with figure 5.4.

In figure 5.20 the passband of the filter can be clearly seen, and the only major components to remain are the peaks at ∼357 MHz, ∼714 MHz, and ∼1071 MHz. Comparing with figure 5.19 it can be seen that the output at ∼357 MHz has been reduced by almost 10 dB.

Quantisation of the power of the high frequency noise due to different scope resolutions is also observed.

5.3.3 Mixer

Figure 5.21 shows the spectra output from the mixer as the beam is moved through various positions. The peak at 714 MHz is due to the mixer LO leakage. It can also be seen that the power observed at ∼714 MHz in figure 5.20 has been shifted to zero frequency and ∼1.4 GHz as expected. Also, the beam bunching component (357 MHz) has been slightly reduced in amplitude as some of its power has been mixed to ∼1.07 GHz.

5.3.4 Low-pass filter (LPF)

Figure 5.22 shows the spectra output from the LPF. In comparison with the mixer output (figure 5.21) the action of the filter is very clear. With the exception of numerous low power peaks (∼-40 dBm) there are only two components output from the processor – the position

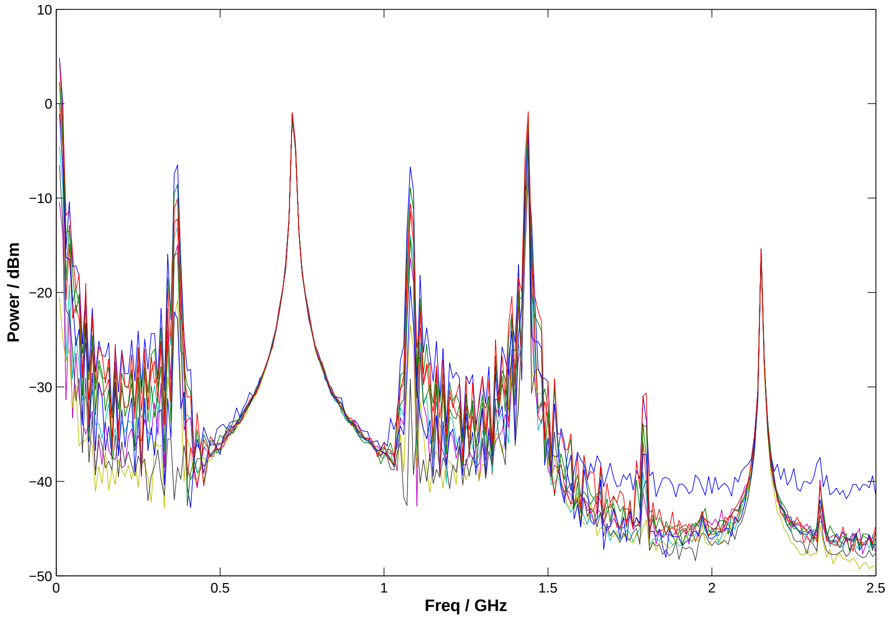

Figure 5.21: The spectrum output from the mixer for various beam positions during multi-bunch operation. Each curve is the average signal from forty beam pulses.

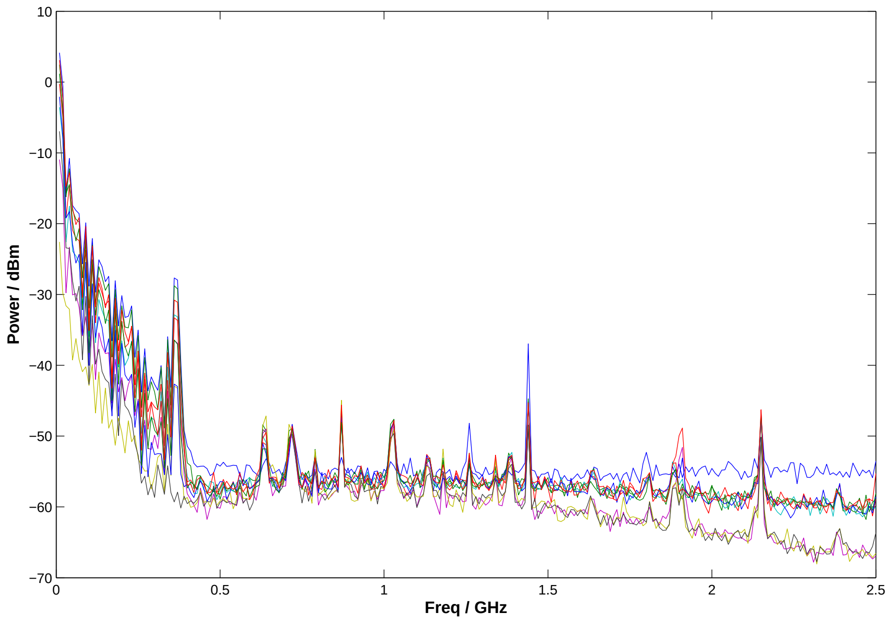

Figure 5.22: The spectrum output from the LPF for various beam positions during multi-bunch operation. Each curve is the average signal from forty beam pulses.

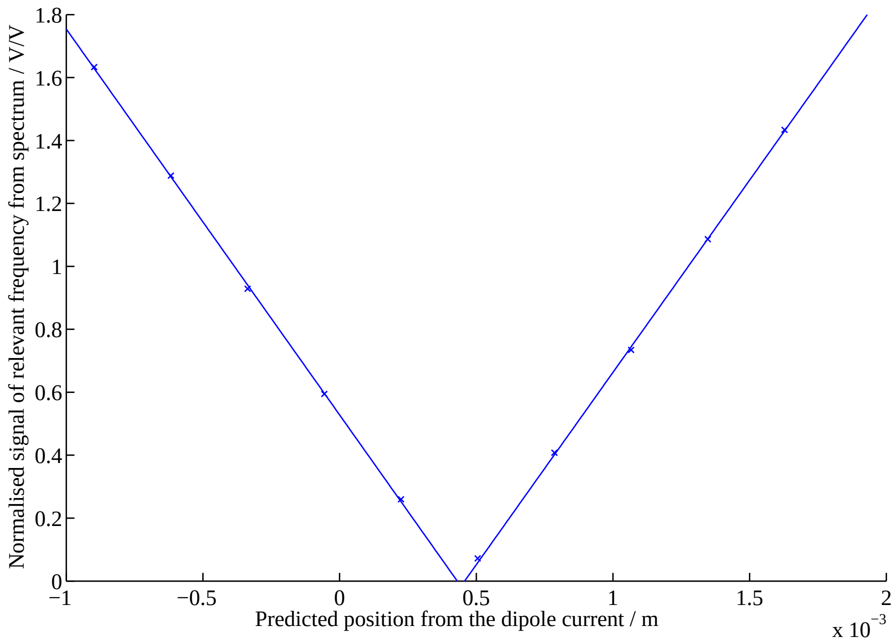

Figure 5.23: The charge normalised and averaged voltage of the 10 MHz component output by the LPF plotted against each beam position.

dependent component at ≤100 MHz, and the beam bunching component at 357 MHz. As predicted by the Systemview simulations (see subsection 4.4), the highest output power of the beam bunching peak is equivalent to the lowest low frequency position signal.

Figure 5.22 was analysed to find the power output of the processor at 10 MHz for each beam position, and figure 5.23 shows the RMS voltage corresponding to the power output plotted against the beam position. The linear range of these points is ≥2 mm, and there is no evidence of the common-mode component. The gradients of the straight lines (-1226±10 m^{-1} and 1222±14 m^{-1}) match well as expected.

5.3.5 Calibration

The method used for calibration is similar to sections 3.3.4 and 5.1.8. In this case it was not satisfactory to use the integral of the output as the pulse now lasts 56 ns, and this integral would include position information from every bunch in the train. Instead, a 4 ns[8] section beginning 8 ns after the start of the pulse[9] was averaged over to be used as the position output. This will be the convolution of the positions of the first several bunches. In order to normalise by the charge signal from BPM 13 the difference outputs from BPMs 11, 12, and 13, must be aligned well in time with the sum output from BPM 13.

[8]The input bandwidth of the amplifier was ∼250 MHz (i.e. a time constant of 4 ns), so beam position features whose characteristic time is shorter than this would not affect the output to the kicker.

[9]The start of the train was used in order to make the result less susceptible to multi-bunch instabilities (e.g. fast ion instability) that appear later in the bunch train. The start of the train contains transient effects, therefore the integration was started 8 ns after the start of the train in order to avoid these (figures 5.24, 5.25, and 5.26 show that these transients are not observed by this time).

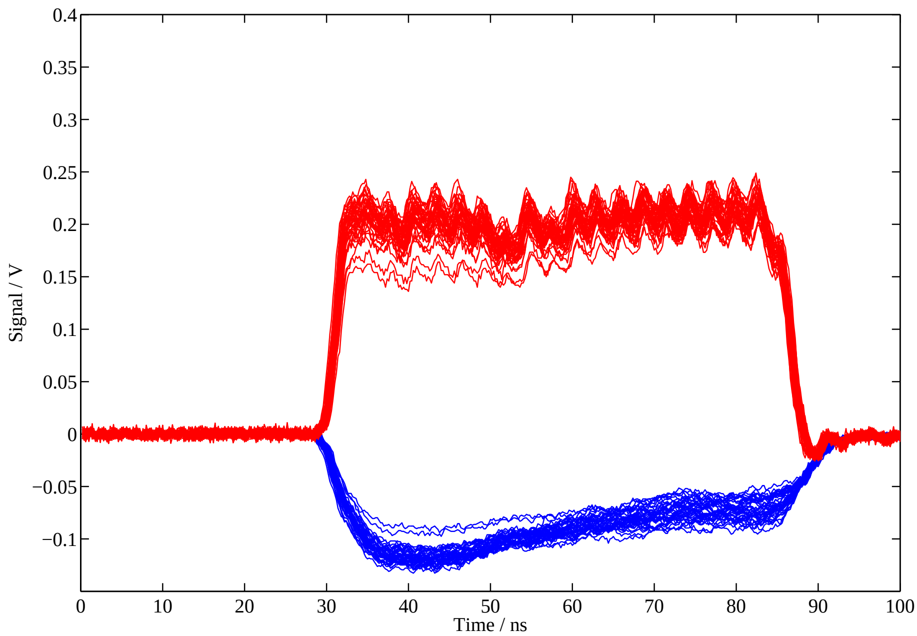

Figure 5.24: The BPM11 difference output (blue) plotted with the BPM13 sum output (red) to demonstrate the result of the time alignment.

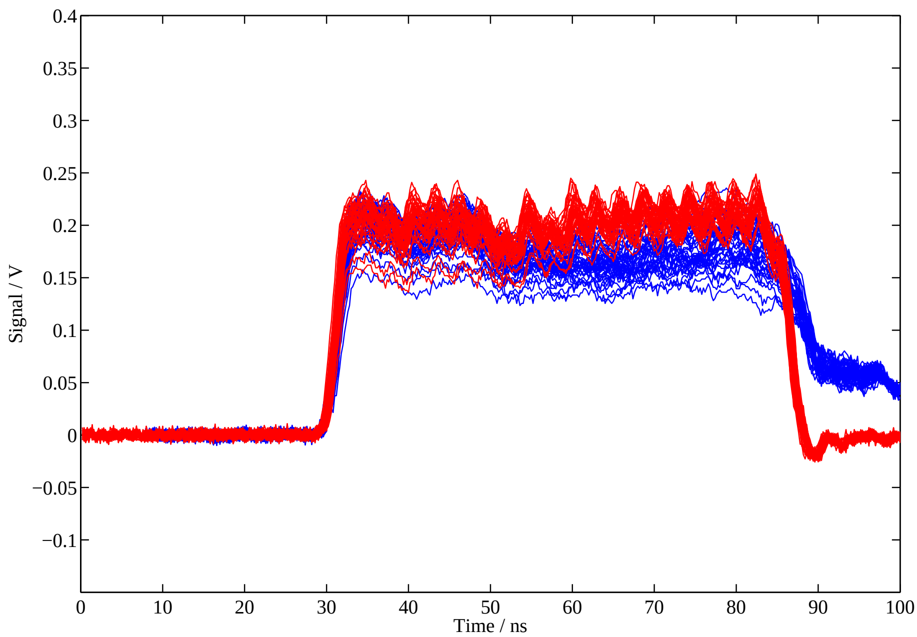

Figure 5.25: The BPM12 difference output (blue) plotted with the BPM13 sum output (red) to demonstrate the result of the time alignment.

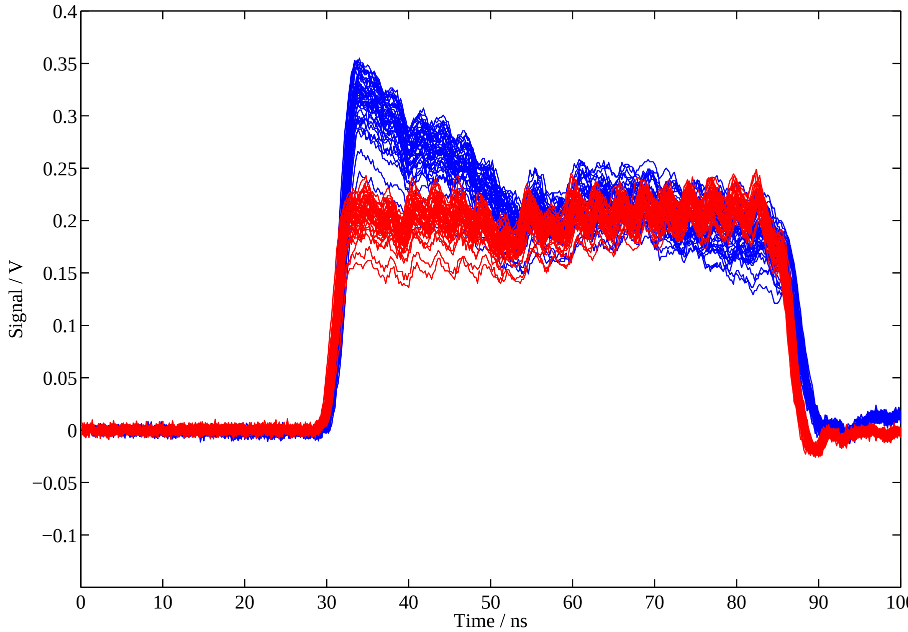

Figure 5.26: The BPM13 difference output (blue) plotted with the BPM13 sum output (red) to demonstrate the result of the time alignment.

Figures 5.24, 5.25, and 5.26 shows how the pulses were aligned by adding zeros to the beginning of the difference signal. This was done using a trial and error method, and the results can be seen to be good to the level of 1 or 2 scope samples, i.e. the alignment is good to $\leq$1 ns.

Figures 5.27, 5.28, and 5.29 show the results of the multi-bunch calibration, and the straight line fit through these points can be seen to be much better quality (as indicated by the $X^2/D.O.F$) than the single bunch beam calibration. The errors were calculated in the same way as for the single bunch calibrations (equation 5.1).

5.3.6 LO phase

The accuracy of the LO phasing during the previous calibration calculation was measured. Using the IQ mixer shown in figure 4.26 the effect of mixing with an LO whose phase is $\pi/2$ different from the chosen phase could be measured. This output, known as the Q signal, was analysed to determine its dependence on the beam position in the same way as the calibration constants, and the results are shown in figures 5.30, 5.31, and 5.32.

Table 5.2 summarises the results from figures 5.30, 5.31, and 5.32, and it can be seen that the calibration constants for the out of phase signal are $\sim$100 times less (the mean of the ratios) than those for the beam position signal, demonstrating that the LO phasing method is capable of finding the correct LO phase to within $\sim$ 1%. A phase error of 1% implies that the LO phase has been tuned to 1% of 90°, i.e. $\sim$0.9°of the correct value, and it can be seen by examining figure 4.25 that the amplitude of the mixer output will then contain 99.98%

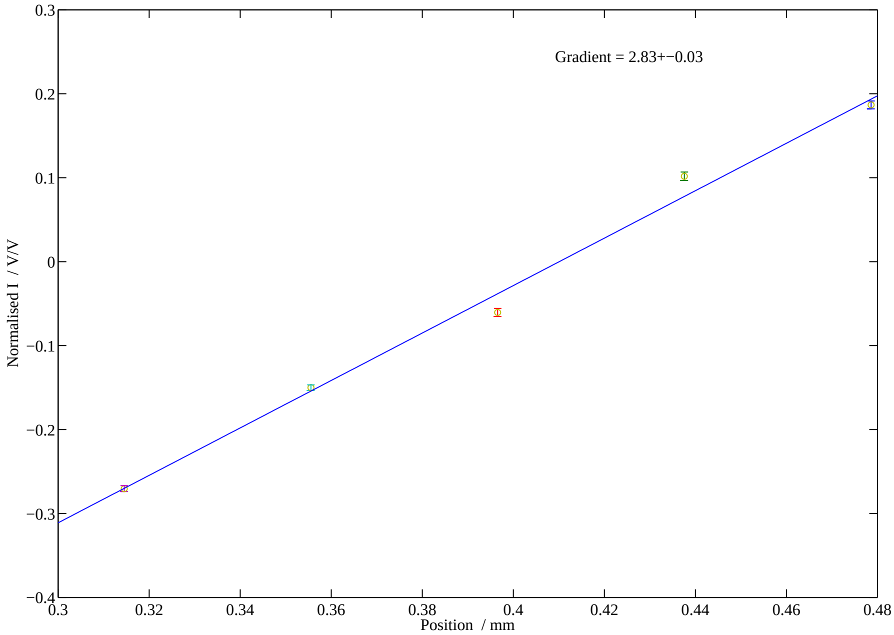

Figure 5.27: Average BPM11 position output of forty beam pulses plotted against the beam position. X^2/D.O.F.= 15.8

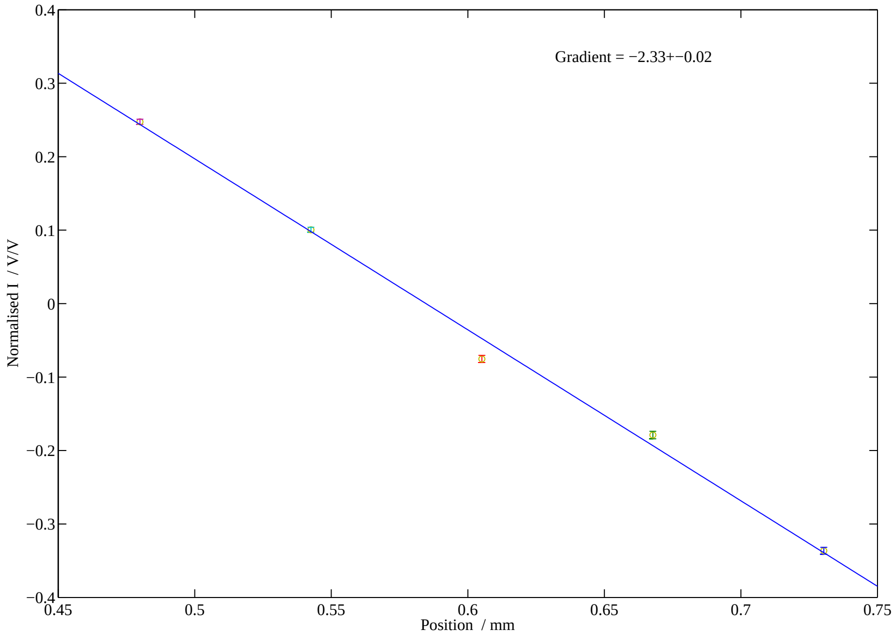

Figure 5.28: Average BPM12 position output of forty beam pulses plotted against the beam position. X^2/D.O.F.= 14.3

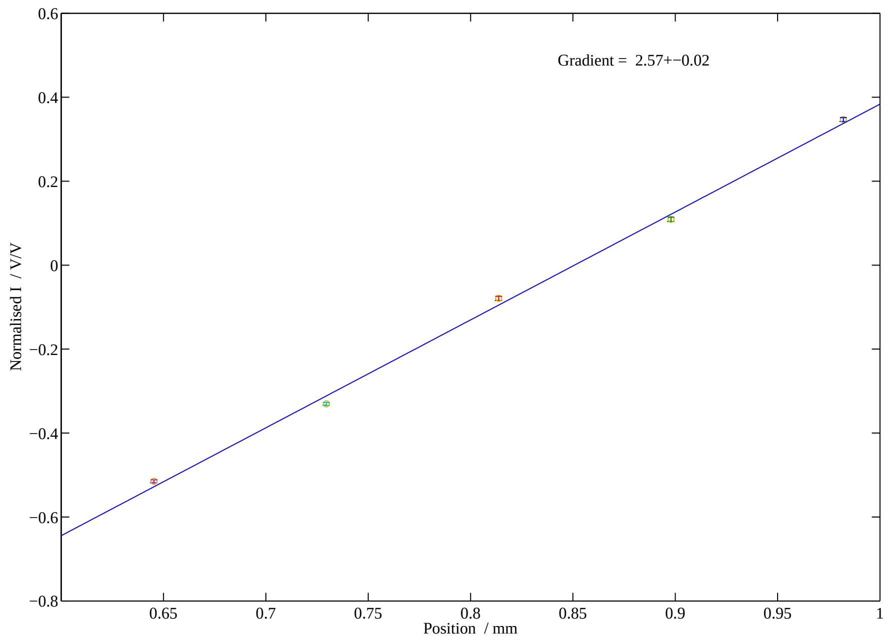

Figure 5.29: Average BPM13 position output of forty beam pulses plotted against the beam position. X^2/D.O.F.= 21.9.

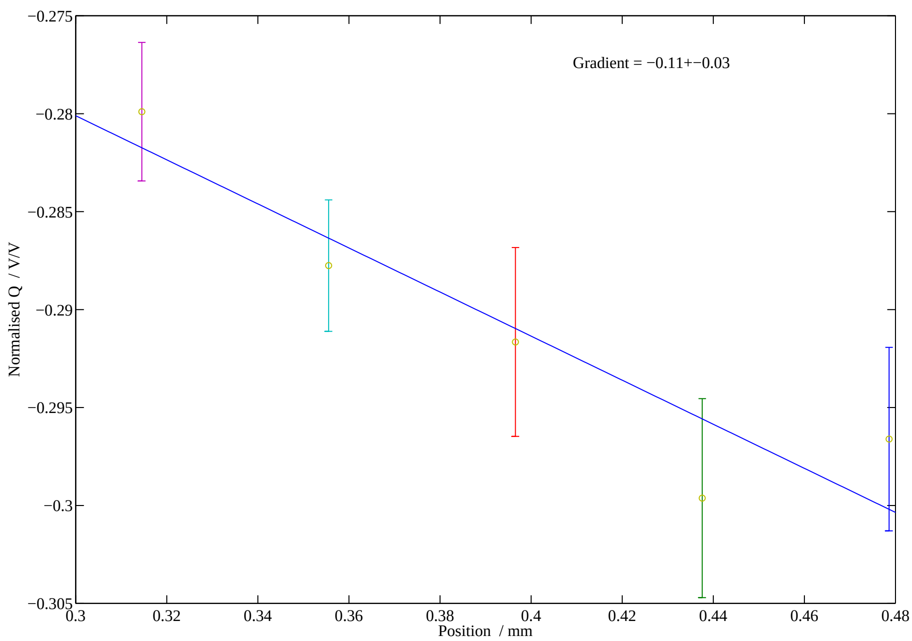

Figure 5.30: Response of the orthogonal component of the LO in BPM 11 to the beam position during a dipole scan. X^2/D.O.F.= 0.6

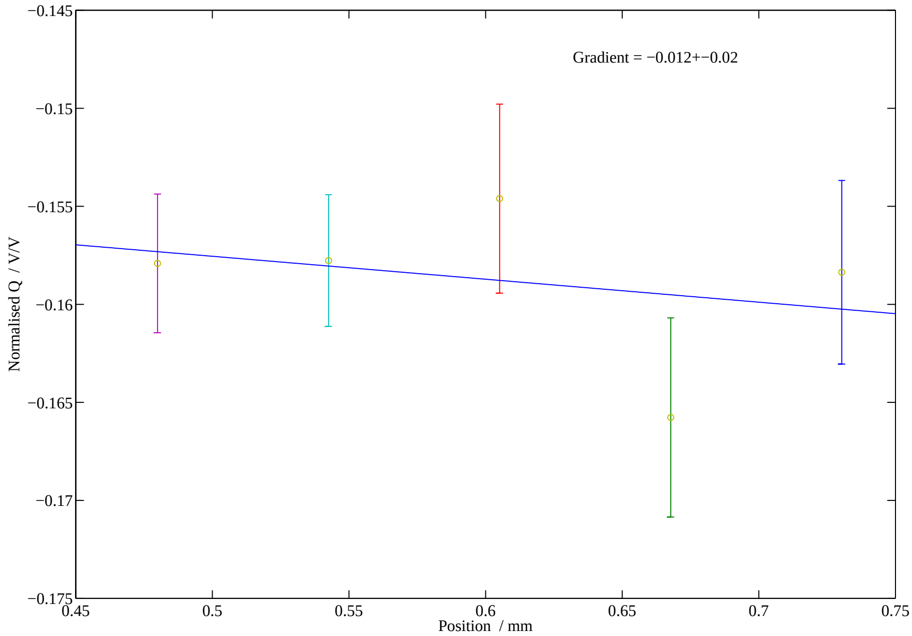

Figure 5.31: Response of the orthogonal component of the LO in BPM 12 to the beam position during a dipole scan. X^2/D.O.F.= 0.8

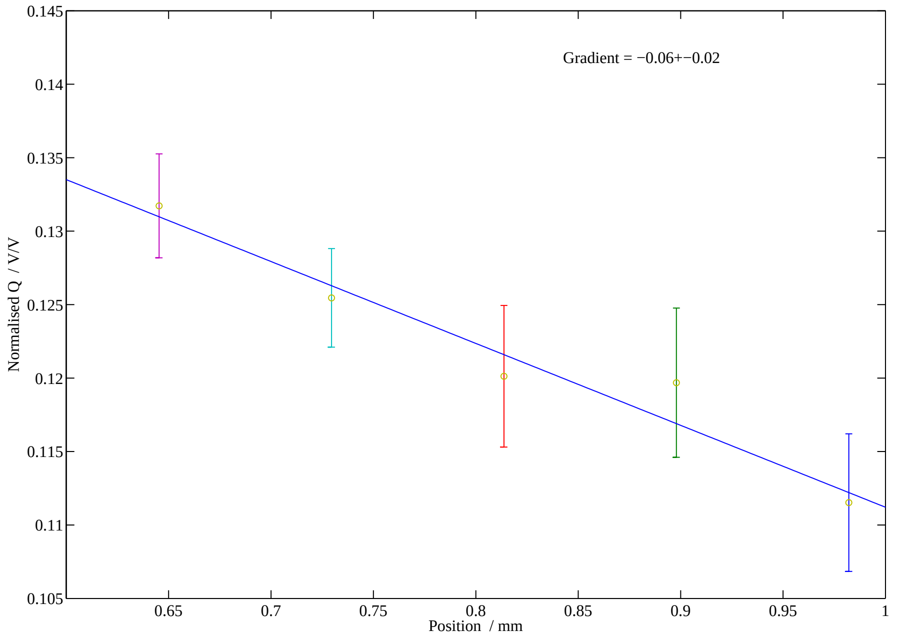

Figure 5.32: Response of the orthogonal component of the LO in BPM 13 to the beam position during a dipole scan. X^2/D.O.F.= 0.2

BPM	I Calibration / mm^{-1}	Q Calibration / mm^{-1}	Ratio
11	2.83 ± 0.03	-0.11 ± 0.03	-25.45 ± 7
12	-2.33 ± 0.02	-0.01 ± 0.02	230.00 ± 460
13	2.57 ± 0.02	-0.05 ± 0.02	51.40 ± 8

Table 5.2: Fit constants of the straight line fits to the plots of the I and Q output of each BPM against the beam position.

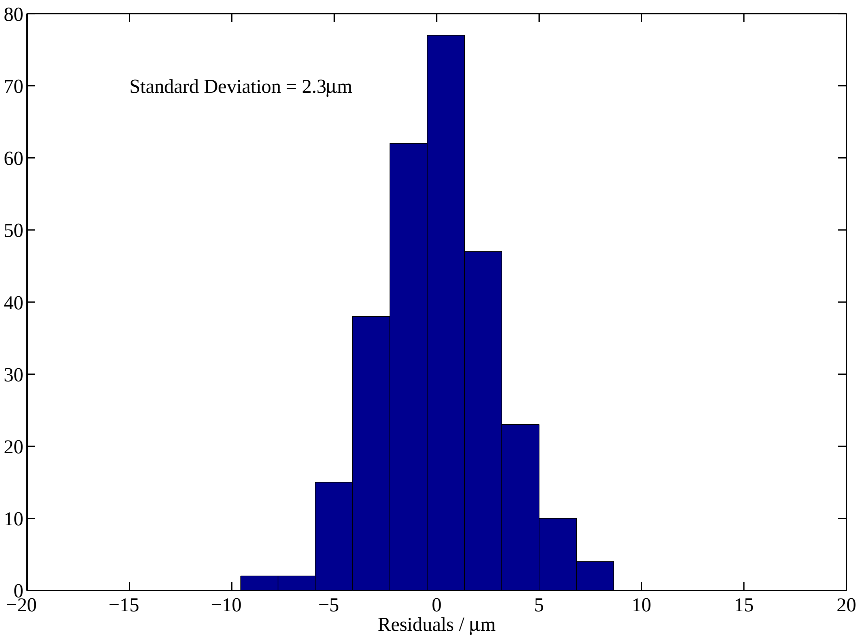

Figure 5.33: A histogram of the residuals after predicting the position of the beam at BPM11 from its position at BPMs 12 and 13.

(cos(1.0)) of the beam dependent signal, and 1.75% (sin(1.0)) of the common-mode signal. That is, the common-mode signal output from the hybrid has been reduced by ∼99% (or ∼20 dB) by the mixer.

5.3.7 Resolution

Using the calibration constants, the resolution of the processor during multibunch operation was calculated.

Figures 5.33, 5.34, and 5.35 show the spread of the residuals, and the results are a large improvement on those from the single-bunch operation. This improvement is due to the larger power available at 714 MHz due to a harmonic of the 357 MHz beam bunching, and due to additional experience in tuning the phase of the LO. Also, 20 dB Minicircuits[55] amplifiers[10] were used to increase the power of the processor output and minimise the contribution of scope noise.

[10] Model #ZHL-2010[56]

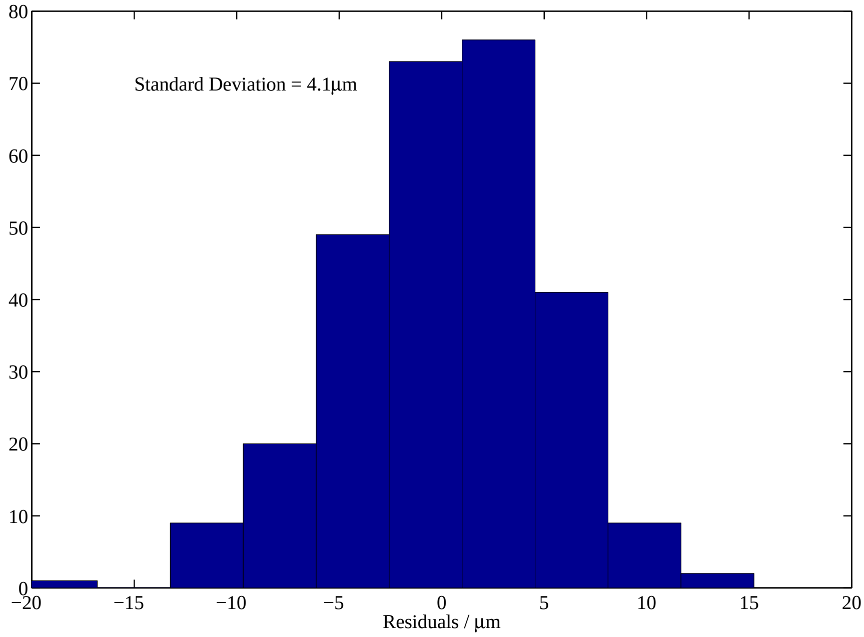

Figure 5.34: A histogram of the residuals after predicting the position of the beam at BPM12 from its position at BPMs 11 and 13.

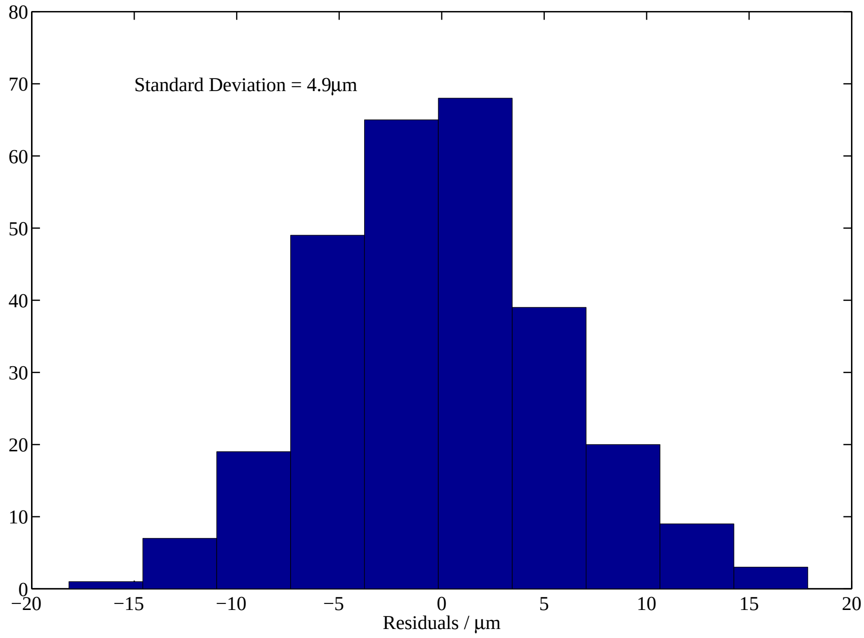

Figure 5.35: A histogram of the residuals after predicting the position of the beam at BPM13 from its position at BPMs 11 and 12.

Figure 5.33 shows that BPM 11 has a resolution that is approximately twice as good as that of BPMs 12 and 13. During these measurements, the processor on BPM 11 was installed directly onto the striplines, whereas BPMs 12 and 13 were installed outside of the accelerator enclosure and connected to the striplines by long runs (∼40 m) of 5/8” heliax cables. As shown in figure 5.1, this will cause a certain amount of attenuation of the stripline signal before it reaches the hybrid. Since no such attenuation will apply to the electronic noise in the processor, this will, therefore, reduce the signal to noise and give a worse resolution.

The calculated resolutions are,

- BPM 11: 2.3 μm
- BPM 12: 4.1 μm
- BPM 13: 4.9 μm

5.4 Summary of Processor Results and Conclusions

Single-bunch beam was used to measure the signal delay through the processor, and it was found to be 4.6±0.3 ns. This is within the goals defined in section 4.2.

The output of each component was investigated at a range of different beam positions, in both single-bunch and multi-bunch mode, and each device was found to behave as expected in terms of the magnitude of the common-mode signal, and the linearity of the output. It was found that the linear range was ±1 mm around the electrical centre (see figure 5.23).

A scheme was designed to determine the correct phase of the LO to minimise the common-mode output of the mixer. This was found to be successful, and to reduce the common-mode contribution to the output signal by better than, approximately, two orders of magnitude.

The resolution of the processor was measured, both when the device was attached directly to the stripline outputs, and when a length of heliax (∼40 m) was used to carry the stripline signal out of the accelerator tunnel. The resolution measured with multi-bunch beam was seen to be ∼5 μm for the externally connected processors (i.e. those separated from the striplines by the heliax), and ∼2 μm for the internal processor (i.e. connected directly to the stripline output). The difference between the external and internal processors was accounted for by the attenuation of the heliax causing a reduced signal power at 714 MHz to reach the subtraction hybrid.

Chapter 6

FONT3 – Feedback Results

During June 2005 the entire feedback system was assembled in the ATF extraction line, and the following details the experimental setup, optimisation of the feedback parameters, and the results of the full feedback test.

6.1 Experimental Setup

Figure 6.1 shows the basic layout of the experiment. The outputs of the three BPMs (known as BPM11, BPM12, and BPM13) were connected via $\sim$40 m lengths of heliax cable to processing electronics installed outside the accelerator enclosure, with the output of BPM11 also split to the feedback electronics installed locally ("Superfast BPM Processor"). This signal then travels to the feedback electronics, the "Superfast amplifier", and finally to the kicker. The "Adjustable-gap kicker" allows the distance between the strips to be varied so as to provide for a larger kick (narrow gap) or larger beam acceptance (wide gap), and thus allows optimisation of the kicker to the beam conditions. A dipole magnet upstream of the kicker was used to provide various beam offsets in order to test the ability of the feedback system to correct the position.

6.2 Trigger Jitter

During data recording the scopes were triggered from an external source locked to the extraction of the beam from the damping ring, so, in principle, the beginning of each recorded pulse will have a fixed time relationship to the beam. There will, however, be a certain amount of jitter in this trigger, as well as the scope trigger, that will cause each recorded pulse to begin at a slightly different time. This can be seen in the spread of the rising edges in the signals shown in figures 5.24, 5.25, and 5.26.

Since each pulse has the same charge to within $\sim$10%, the sum signal can be used to determine the timing of each pulse in order to remove the jitter. This was done by finding the time at which the waveform passes through a specific voltage[1].

[1] In this case the signal amplitude was $\sim$30 mV, and 3 mV was chosen as the alignment voltage.

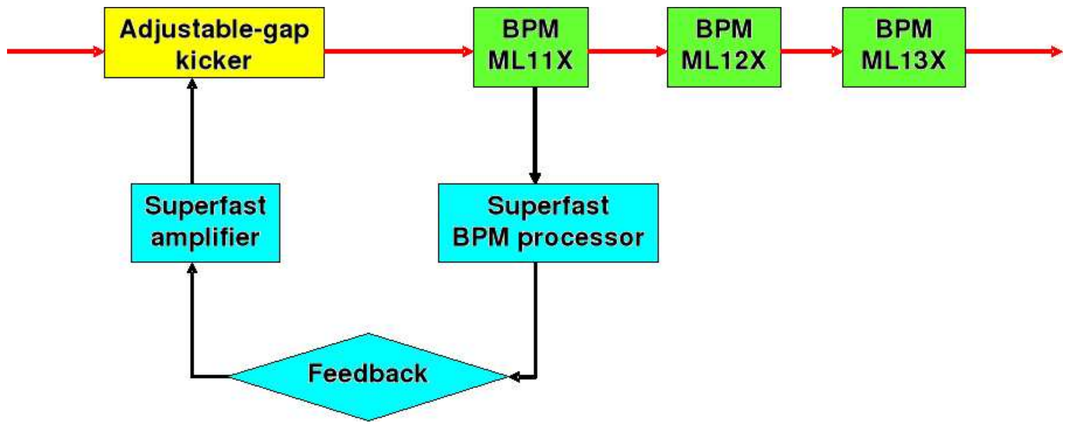

Figure 6.1: Block diagram of the experimental layout showing the beam moving through the kicker and BPMs, and the feedback signal travelling back from the feedback BPM (ML11X) to the kicker[22].

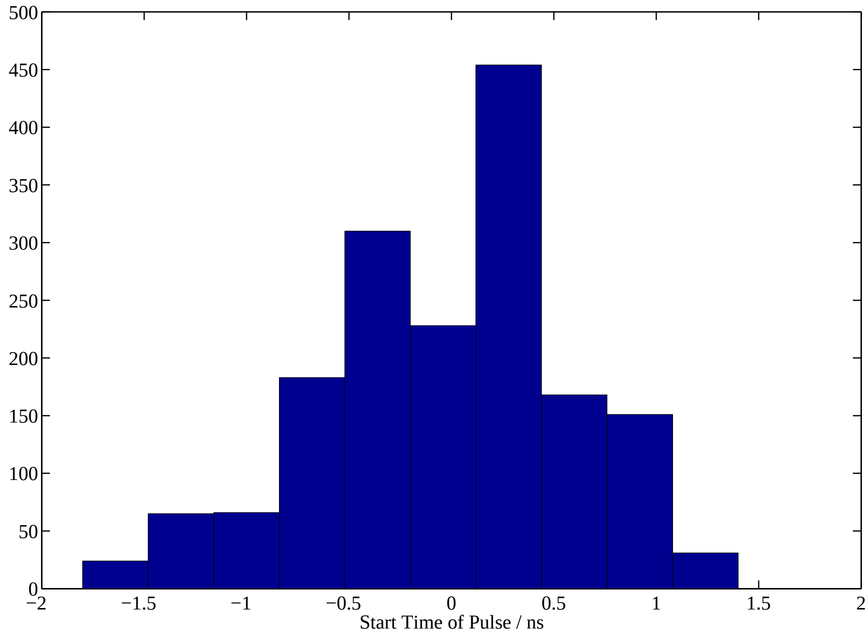

Figure 6.2: Spread of the rising edge of a large sample (∼1500) of charge profile pulses from BPM13 when measured by a scope triggered with the signal locked to the extraction of the beam from the damping ring. The mean time has been shifted to zero, and the spread is due to the jitter of the extraction trigger and the scope timing.

Figure 6.2 shows the spread of the beam times calculated with this method, and it is clear that there is over ±1 ns of jitter between the time the scope begins to acquire data and the passage of the beam through the BPM[2].

The pulses were then shifted in software to limit this jitter to the level of one scope sample period (≤0.2 ns).

6.3 Pulse Rejection

There are various ways in which a problem could develop in the ATF that will degrade the quality of the beam, e.g. a problem with acceleration resulting in the injection of an off-energy beam into the damping ring, or an electron gun issue resulting in a lower charge bunch train and inducing a BPM output error (see equation 4.3). Problematic pulses can be identified by examining the signal created by the sum processor installed on BPM13, and several criteria were developed with which to analyse this signal.

6.3.1 Missing Bunches

Occasionally the ATF electron gun undergoes a fault that causes several bunches at the end of the train to have very low charge or to be non-existent, and these pulses should be rejected. The penultimate bunch of each charge profile was examined to see if it contained any charge, and if not it was rejected. Figure 6.3 shows examples of charge profiles rejected using this technique.

6.3.2 Charge Uniformity

Some bunch trains were observed where the first few bunches had a significantly lower charge than the remainder of the train. It is especially important to reject these pulses, since this low charge will cause the kick to have a low initial magnitude, which will then be circulated around the delay loop. In this case, examination of the difference signals may lead to the erroneous conclusion that there is a problem with the feedback system.

The beginning of each charge profile was tested to ensure that its magnitude was ≥90% of the charge of the end of the train. Figure 6.4 shows examples of pulses rejected on the basis of this test.

6.3.3 Total Train Charge

Since there is no charge normalisation in the processor, it is important that the total charge of each train be similar to within ∼10% (see section 4.3).

Figure 6.5 shows the spread of the total train charge (calculated by summing along the output voltage of the BPM13 sum processor) after rejecting on the basis of missing bunches at the end of the train and a lower than expected initial charge. The median point has been indicated, along with the charge values ±10% of this median that were used as cutoffs.

[2]Note that this is not necessarily pure scope jitter, as there is likely to be a certain amount of jitter between the extraction trigger and the actual extraction of the beam.

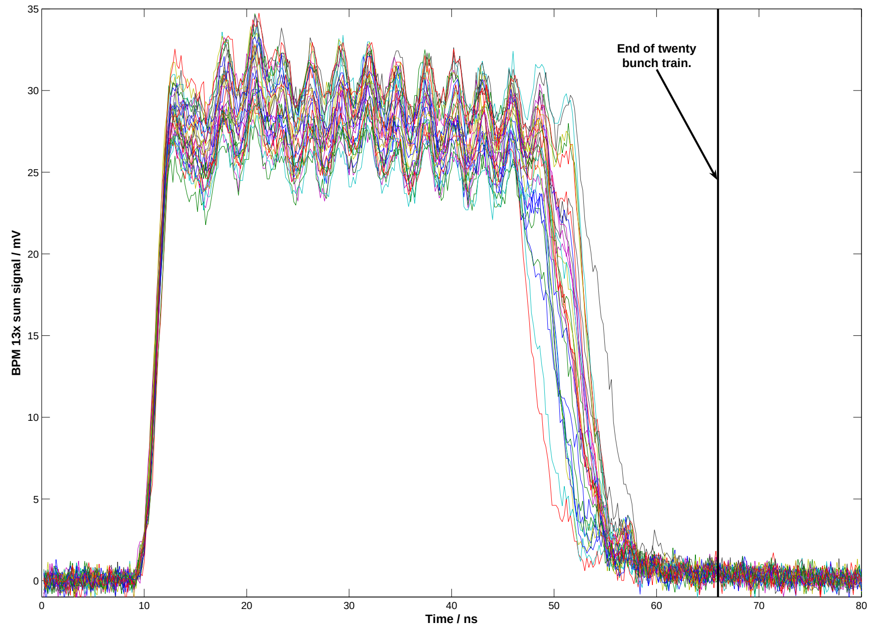

Figure 6.3: Examples of pulses rejected on the basis of the pulse train not containing all twenty bunches.

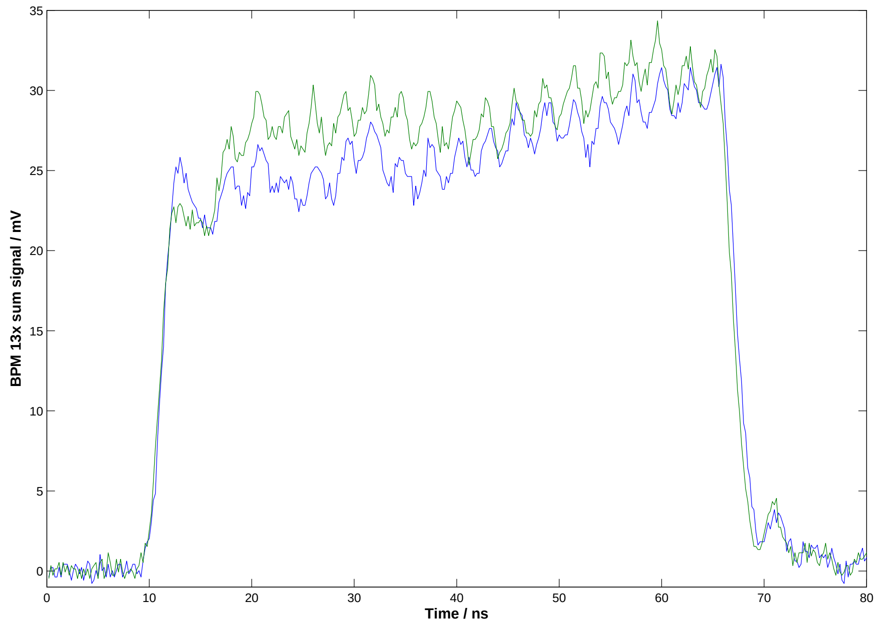

Figure 6.4: Examples of charge profiles rejected on the basis of low initial charge.

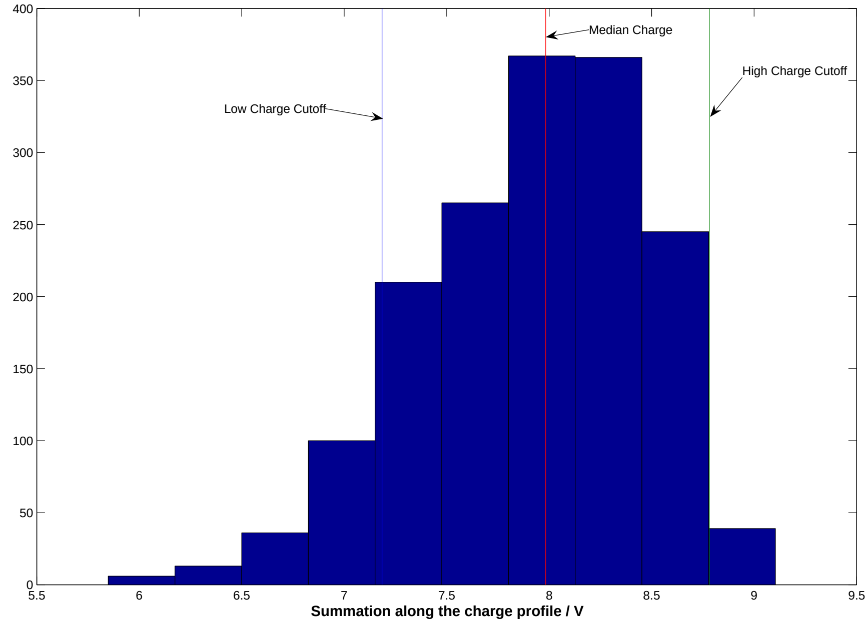

Figure 6.5: Spread of the total train charge after rejection on the basis of missing bunches and a low initial charge. The total charge was calculated by summing along the output from the sum BPM. The median charge, and the charge values ±10% of the median have been indicated.

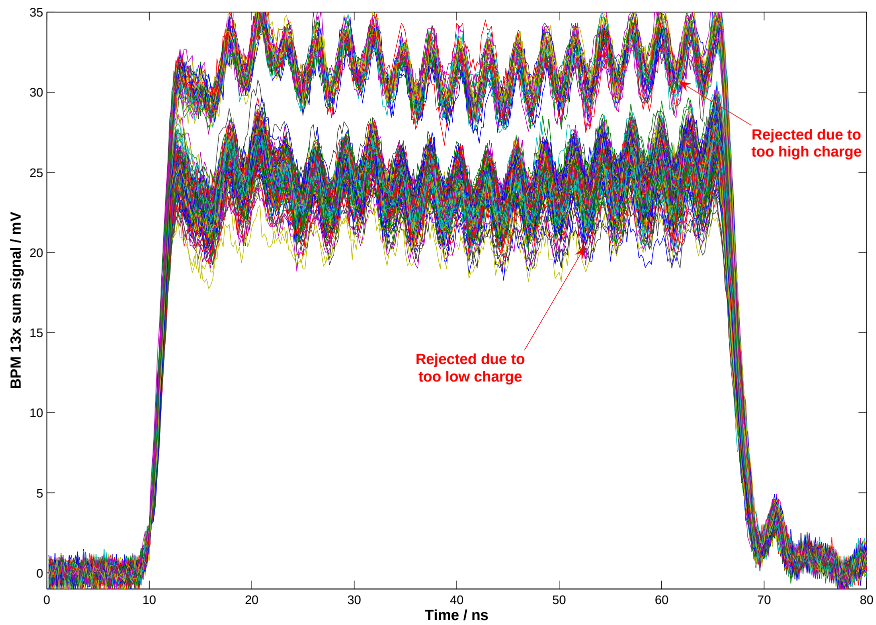

Figure 6.6: Bunch trains rejected on the basis of the total charge.

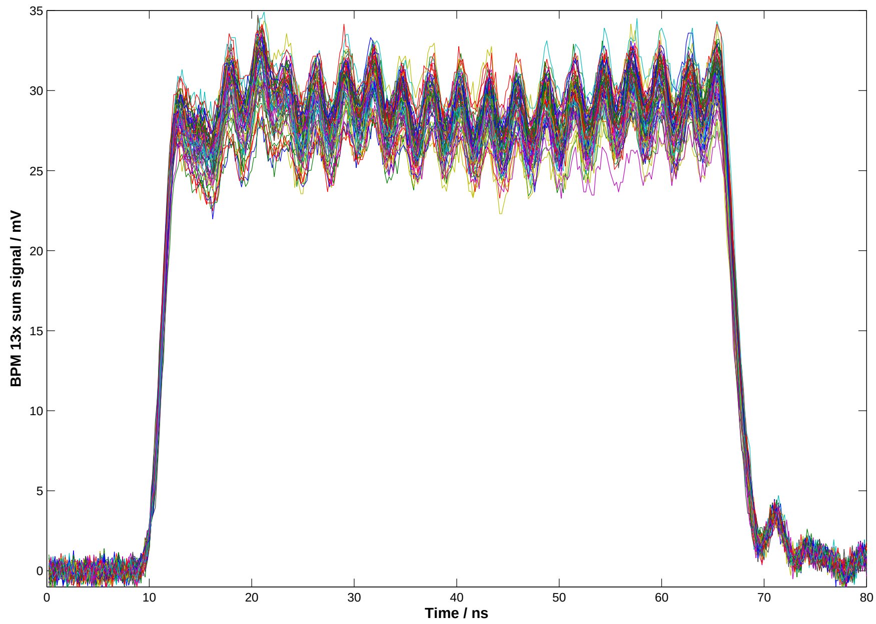

Figure 6.7: Examples of bunch trains not rejected from the analysis.

Figure 6.6 shows examples of pulses that were rejected on the basis of the total train charge.

6.3.4 Accepted Pulses

Figure 6.7 shows examples of pulses that have not been rejected.

6.4 System Optimisation

As with FONT2 there are various parameters to be optimised in order that the feedback system operates correctly (see figure 3.12). In the case of FONT3 it is necessary to optimise the same parameters as with FONT2, with the exception of the beam flattener, which was not applied to the ATF beam,

- Main Gain
- Delay Loop Length
- Delay Loop Gain

The method used to set each parameter correctly is identical to that for FONT2; i.e. scan that parameter through various settings in order to find its optimal value.

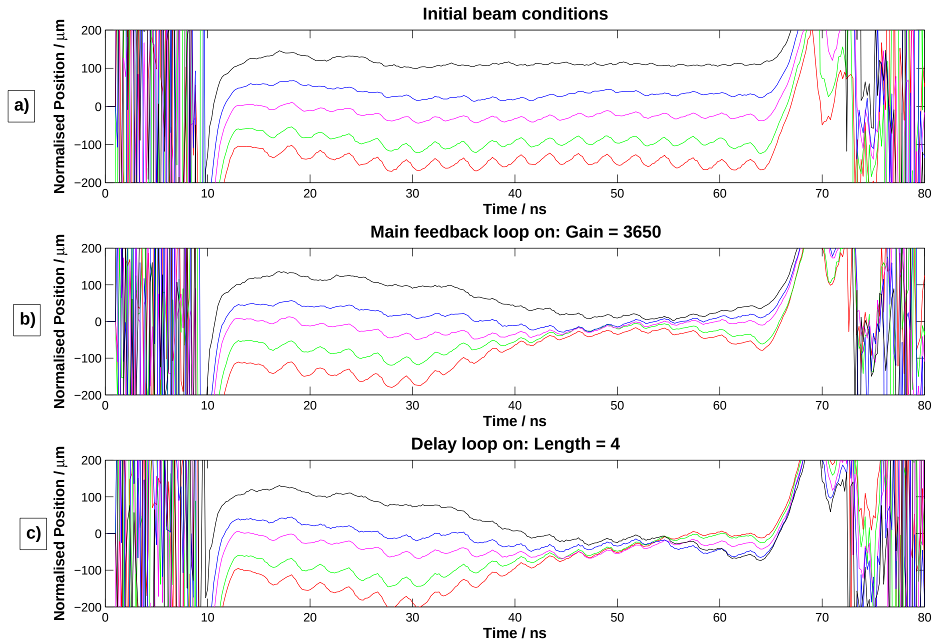

Figure 6.8: Average beam position measured by BPM11 at each of a range of beam positions. The plots show the beam position in three feedback conditions: a)no feedback, b)feedback but no delay, and c)feedback and delay. For this data set the main gain had been set to a value thought to be optimal.

6.4.1 Main Gain

In order that the delay loop be optimised, it is important that the main gain be well understood, and set to the correct value. It was, therefore, the first parameter to be optimised.

An initial estimation of the performance of the amplifier suggested that the optimal gain was 3650[3], so this value was used along with a value significantly lower than this, and one significantly higher (1890, 3650, and 5480).

In each case, the the dipole was used to scan the beam through five positions (the total range of motion induced by the dipole was $\sim$250 μm.), first with the main loop off, then with the main loop on and delay loop off, and finally with the main and delay loops on. Forty pulses were recorded at each setting, with several pulses being rejected due to the criteria detailed in sections 6.3.1, 6.3.2, and 6.3.3. The results are shown in figures 6.8, 6.10, and 6.11.

Figure 6.8 shows the effect of setting the main loop gain to 3650. The five curves in each plot represent the position of the bunch train at five different settings of an upstream dipole, and each is an average of up to forty pulses[4]. The curves are the result of normalising the difference signal by the sum signal in order to show how the position of the beam is affected

[3]The gain values refer to the multiplicative factor the amplifier applies to the BPM output, i.e. a processor output of 1mV and a gain of 1000 would result in a potential difference of 1V between the upper and lower kicker strips.

[4]Forty pulses were recorded at each beam setting, however, as explained previously, some errant bunch trains have been rejected from the analysis.

by the feedback. Although it is the un-normalised difference that provides the feedback signal, it is the aim of this experiment to optimise the actual beam position, and not the difference signal, therefore it will be normalised plots that are shown. Since the sum signal contains only noise before and after the bunch train, division by this signal results in a wildly oscillating output before and after the beam. The large peak at the end is due to normalisation by the rapidly falling charge.

The dipole was used to steer the beam to five different vertical positions while the feedback electronics are off, and figure 6.8a shows the output of BPM11. This shows the input condition of the beam. It can be seen that the position profile is flat to the level of $\sim$50 μm. The normalised output of negative beam positions contains a larger 357MHz component than the positive positions, and this is due to the π phase change that occurs as the beam position changes sign. Dividing this with the sum signal, whose phase does not depend on position, removes the 357 MHz when the beam is above zero and reinforces it when the beam is below zero.

The feedback loop was then turned on with a gain of 3650, and the dipole was set to the same five settings. The delay loop remained off. Figure 6.8b shows the output of BPM11, and the action of the feedback is quite clear. It can be seen by comparison that the position of the beginning of the input beam has shifted a small amount since the data from figure 6.8a was taken, and analysis of the amplifier output has shown that this is due to a small amount of 'droop' in the baseline when the amplifier is powered on.

If the middle three curves are examined (i.e. ignoring the black and red curves), it can be seen that the kicker begins to act sometime between 30 ns and 35 ns, and moves the beam toward zero position. After a $\sim$10 ns delay for the processor and amplifier output to maximise and the kicker to fill, it can be seen that these three curves now overlap – i.e. the feedback has moved all three initial positions to the same place, implying that this main gain value is very close to optimal. After a further latency period the beam begins to return to its initial location.

The position change of the black and red curves caused by the kicker is never large enough to coincide with the other curves due to their large initial offset. This large absolute position value caused the processor to input a signal that is beyond the linear range of the amplifier. The non-linearity of the amplifier at these large offsets means the kick is insufficient to correct the beam position. This effect was investigated in simulations of the feedback system[57]. The feedback loop and amplifier were simulated and the initial conditions shown in figure 6.8a were input. The results of this are shown in figure 6.9.

Figure 6.9b shows that a perfectly linear amplifier will perform a good position correction of all the initial offsets, while figure 6.9c shows that an amplifier that becomes increasingly non-linear at large offsets will not provide sufficient kick to fully correct the large offsets. The similarity between figure 6.9c (the simulation) and 6.9d (the measured output) strongly suggests that the failure of the feedback system to fully correct the large beam offsets is indeed due to the amplifier having a limited linear range.

The delay loop was then turned on with a gain of $\sim$0.9 (see section 6.4.4) and this is shown in figure 6.8c. The action of the first pass of the loop is clear and is similar to the previous case. At $\sim$55 ns the signal from the delay loop reaches the amplifier and acts on

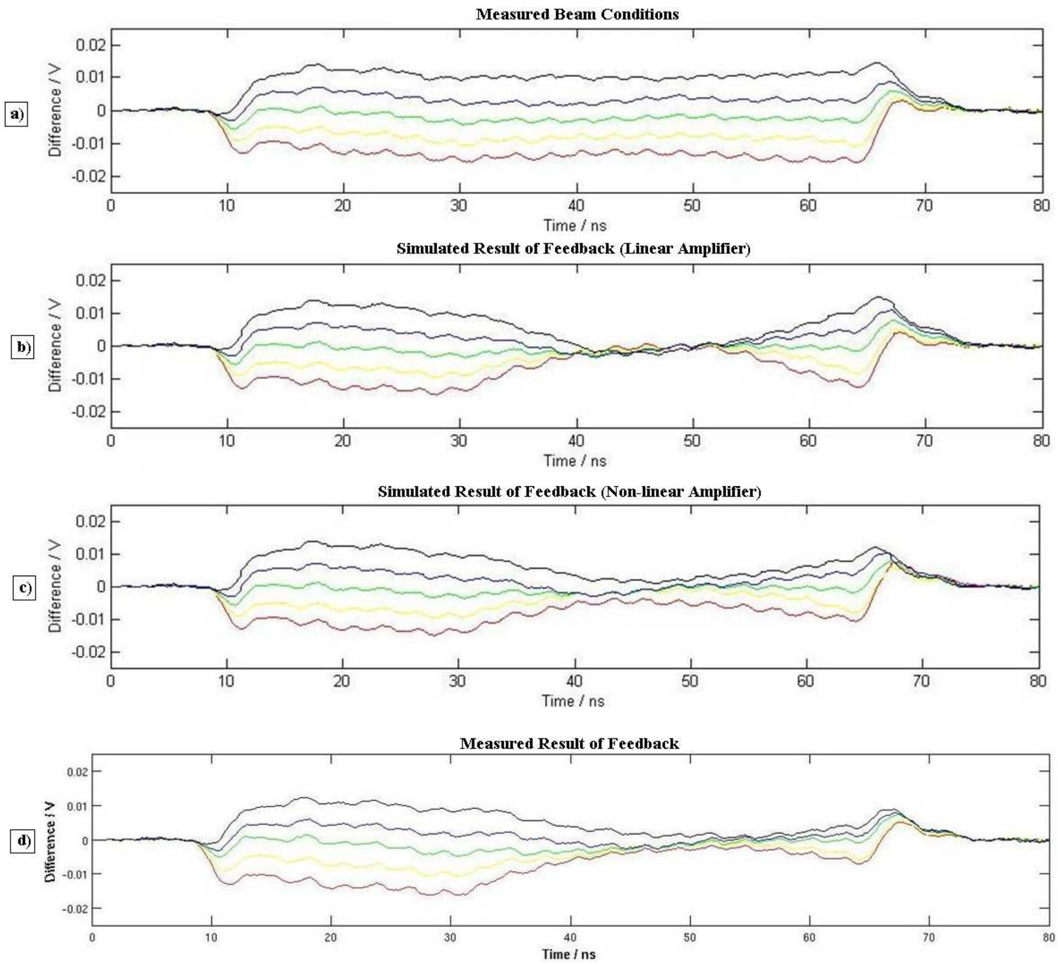

Figure 6.9: Simulations of the action of the feedback system on measured initial conditions. a)The initial conditions, b)Simulated output with a perfectly linear amplifier, c)Simulated output with an amplifier that becomes non-linear at large amplitudes, and d)The measured output of the feedback system. The main gain was set to a value thought to be optimal in the simulation and the measured data sets.

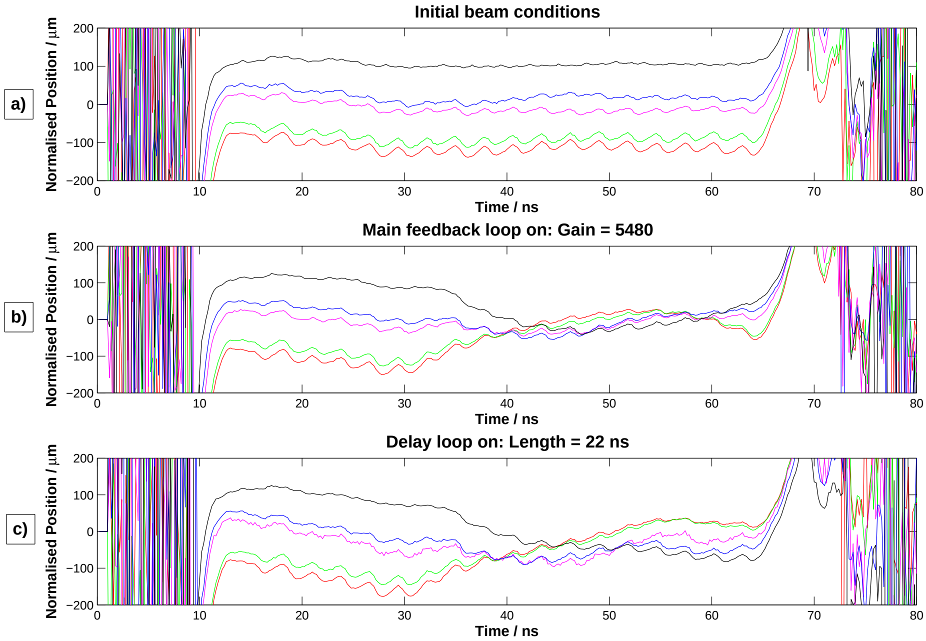

Figure 6.10: Average output of BPM11 at a range of beam positions. The plots show the beam position in three feedback conditions: a)no feedback, b)feedback but no delay, and c)feedback and delay. This shows the effect of too high gain.

the beam. The bunch trains that were initially lower (i.e. the red and green curves) are pushed above zero by the introduction of this signal, and the positive position trains (black and blue) are pushed negative, which suggests that the delay loop settings are incorrect. This will be investigated in sections 6.4.3, and 6.4.4.

Figure 6.10 shows the effect of increasing the main gain beyond the optimal value. Figure 6.10a shows the beam position with no feedback, figure 6.10b shows the effect of the main loop without the delay loop, and figure 6.10c shows the full system in operation.

On examining figure 6.10b, it can be seen that the large gain value causes the system to over-correct, thereby kicking the beam beyond zero position. Since the delay loop is not operating for this data set the system will begin oscillating, as the beam position will be over-corrected every latency period.

Figure 6.10c shows that the delay loop acts to maintain the larger than necessary correction signal. With a correctly optimised delay loop, it would be expected that the beam would move back through zero but with a smaller absolute offset during the final latency period – i.e. the system would undergo a damped oscillation. As this is not observed in these data, it can be concluded that the delay loop parameters are set incorrectly, which is consistent with figure 6.8.

Figure 6.11 shows the case of lower than optimal gain. It is clear that the gain is excessively low, and that the correction is not sufficient to fully correct the position. With the delay loop in operation (figure 6.11c) it can be seen that the system does converge toward

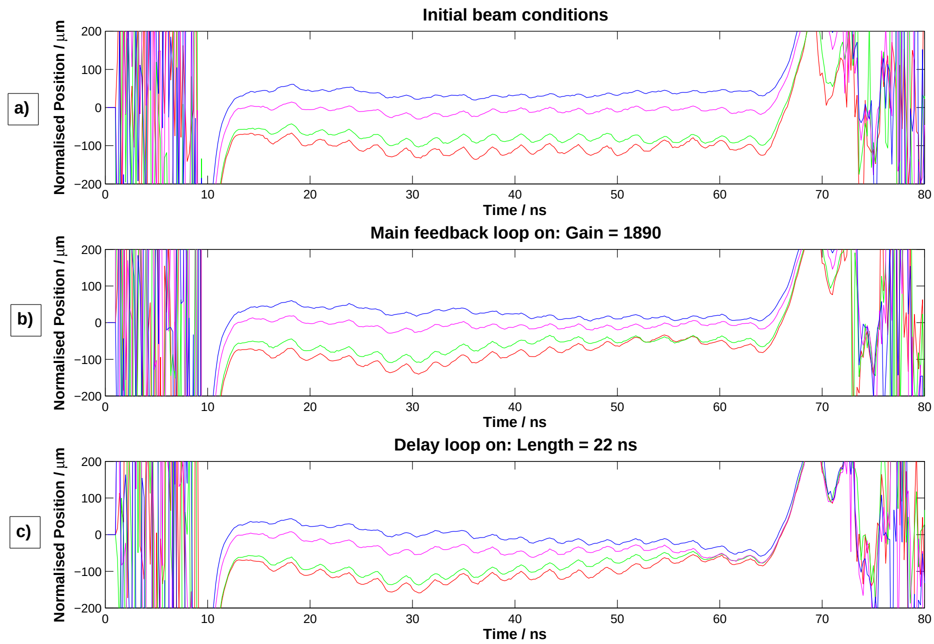

Figure 6.11: Average output of BPM11 at a range of beam positions. The plots show the beam position in three feedback conditions: a)no feedback, b)feedback but no delay, and c)feedback and delay. For this data set the main gain had been set to a value known to be lower than optimal.

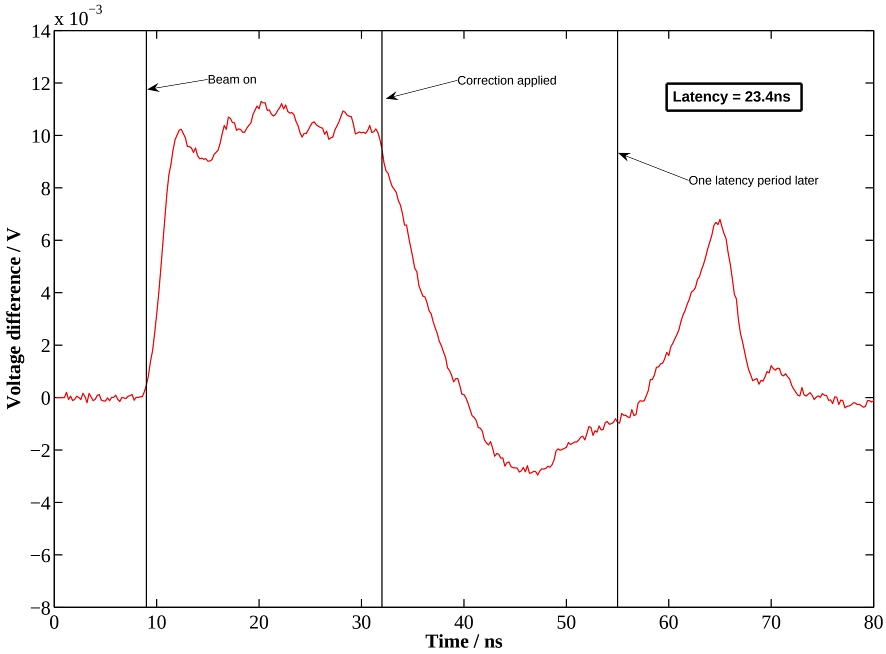

Figure 6.12: Difference between the lowest and highest beam position. The main gain was set at a high value and the delay loop was off. The calculated positions of the start of the beam pulse and the beginning of the correction have been marked with vertical lines.

a zero position offset.

6.4.2 System Latency

The case of a high main gain and no delay loop (see figure 6.10b) was analysed to calculate the total latency of the system. As explained in section 6.4.1, the uppermost and lowermost curves should be discounted as the amplifier output has become non-linear for these data. The highest and lowest position settings that are still within the dynamic range of the amplifier can then be examined, and the position difference between them has been plotted in figures 6.12 and 6.13[5].

In figure 6.12 the time at which the signal emerges significantly from the initial noise, and the time at which the position difference falls below 90% of the plateau that lies between approximately 17 ns and 30 ns have been indicated. If these points mark the beam start time and the beginning of the first correction, respectively, then the time separation between these points is the system latency. The separation measured from this plot is 23.4 ns.

A third vertical line has been drawn on figure 6.12 to mark the beginning of the last latency as predicted by the calculated value of 23.4 ns.

Figure 6.13 shows the latency calculation with different definitions of the beginning and end of the first latency period. They are defined as the points at which the position difference

[5]In this case the un-normalised difference is used, as subtracting the large amount of initial noise that exists in the case of the normalised difference would disguise the rising edge of the signal.

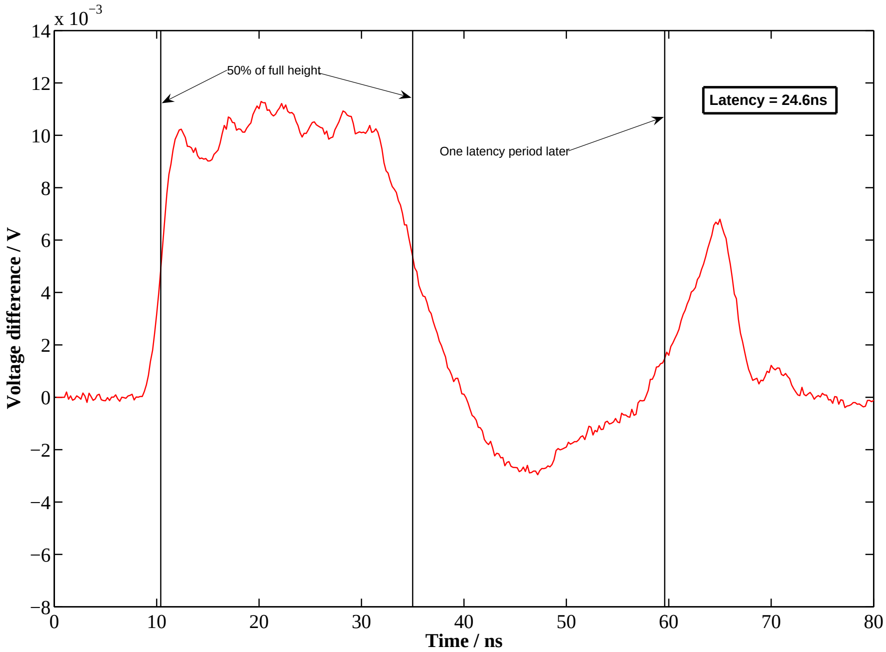

Figure 6.13: Difference between the lowest and highest beam position. The main gain set at a high value and the delay loop was off. The vertical lines show where the beam passes through 50% of the maximum value.

crosses through the half-max value of the plateau. This results in a calculated latency of 24.6 ns. A third line has been drawn in figure 6.13 to show the beginning of the final latency period as calculated with the value of 24.6 ns.

The latency was defined as the average of these two values, i.e. 24±0.5 ns. This is ∼4 ns greater than the design latency (see section 4.2.1) due to a less than optimal beamline layout being used. Time constraints on the experiment prevented full optimisation of the cabling, etc. that would have led to a lower latency.

6.4.3 Delay Loop Length

The main gain was set to the value determined in section 6.4.1 and the beam was moved to two widely separated positions. Figure 6.14a shows the output of BPM11. The action of the feedback loop can be seen at ∼30 ns as before, and it can be seen that the positions begin to diverge again at ∼55 ns.

The delay loop was switched on and set to three different lengths – 22 ns, 24 ns, and 26 ns. At each delay loop length the beam was steered to two positions. The results are shown in figures 6.14b, 6.14c, and 6.14d, and it can be seen how different values for the delay loop length affect the final ∼10 ns of the bunch train.

For a delay of 22 ns it appears that the end of the train is kicked too strongly, implying that this value is smaller than optimal. This causes the beginning of the delayed correction signal to be added to the end of the initial feedback signal, and is observed as an excessively

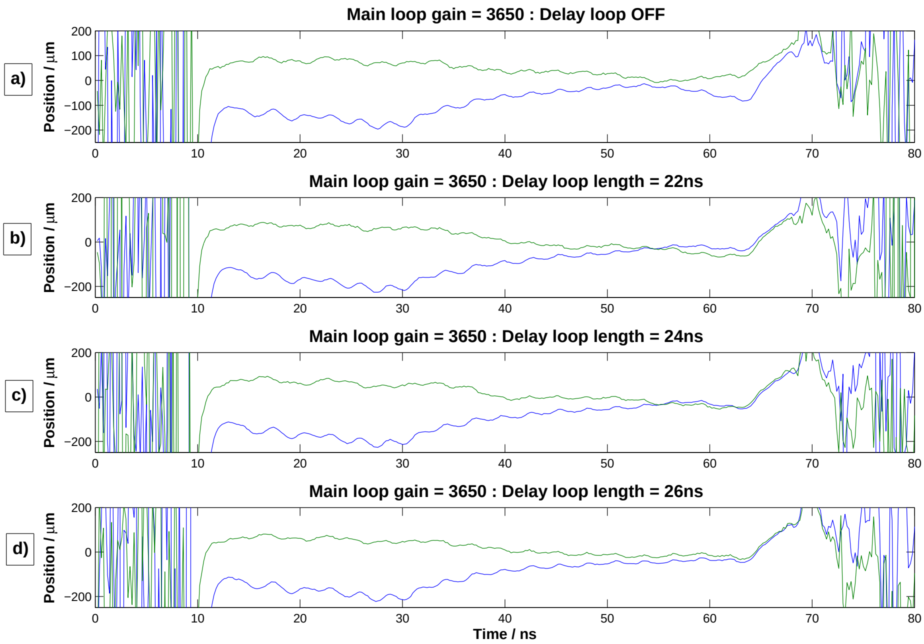

Figure 6.14: Average output of BPM11 at two beam positions for a range of delay loop length settings. The main loop was on with a gain of 3650 for all data sets. a)No delay, b) Delay length = 22 ns, c) Delay length = 24 ns, d) Delay length = 26 ns. The delay numbers are accurate to ~1ns

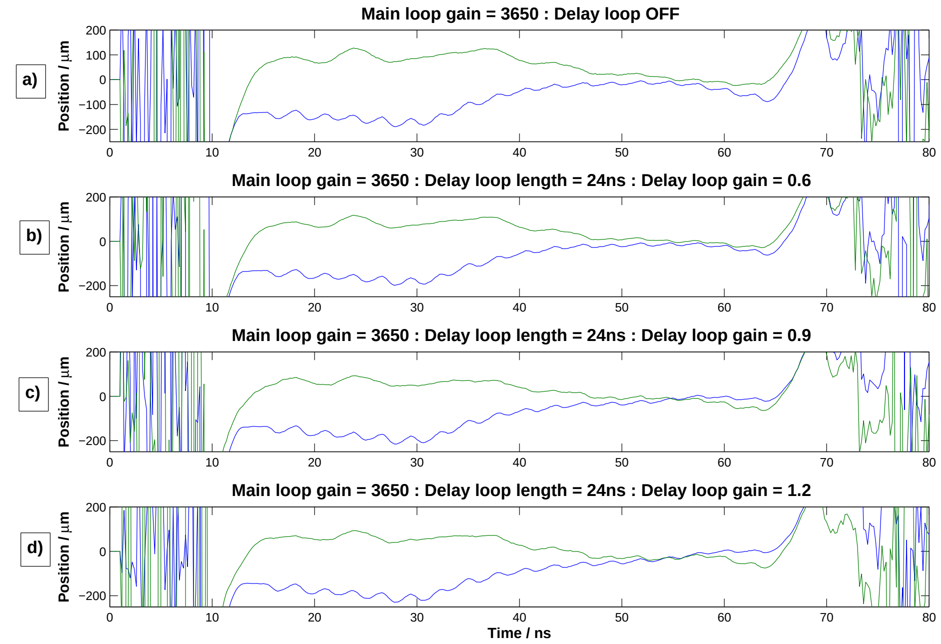

Figure 6.15: Average output of BPM11 at two beam positions for a range of delay loop gain settings; a)No delay, b)Delay gain = 0.6, c)Delay gain = 0.9, d)Delay gain = 1.2.

large kick to the end of the train.

A delay length of 26 ns appears to be too large. When compared with the case of no delay loop, it can be seen to be insufficient to bring the beams to zero position. This is due to the fact that an excessively long delay results in the final latency correction being applied too late, and not having sufficient time to ramp up to the necessary value.

A delay loop length of 24 ns is apparently the optimal value, as it is clear that the bunch trains are brought to zero position and remain there for the remainder of the pulse. This agrees well with the result from section 6.4.2.

6.4.4 Delay Loop Gain

Formally the gain of the delay loop should be unity in order that the delayed correction be identical to the correction of the initial latency period. This was checked with beam tests.

The main loop gain was set to the optimal value from section 6.4.1, and the delay loop was set to the length found in section 6.4.3. The beam was measured at two widely separated positions, and the gain of the delay loop was varied to observe its effect. Figure 6.15 shows the output of BPM11. Figure 6.15a shows the output when the delay loop is off, and figures 6.15b, 6.15c, and 6.15d show the output with the delay loop switched on at the various gain settings. On examination of the final ∼10 ns of the bunch train, it can be seen that increasing the value of the delay loop gain results in a larger kick being applied. It appears that a gain of ∼0.9 is too large and that a value of ∼0.6 is closer to optimal, since it is this value that results in the positions being brought together until the end of the train.

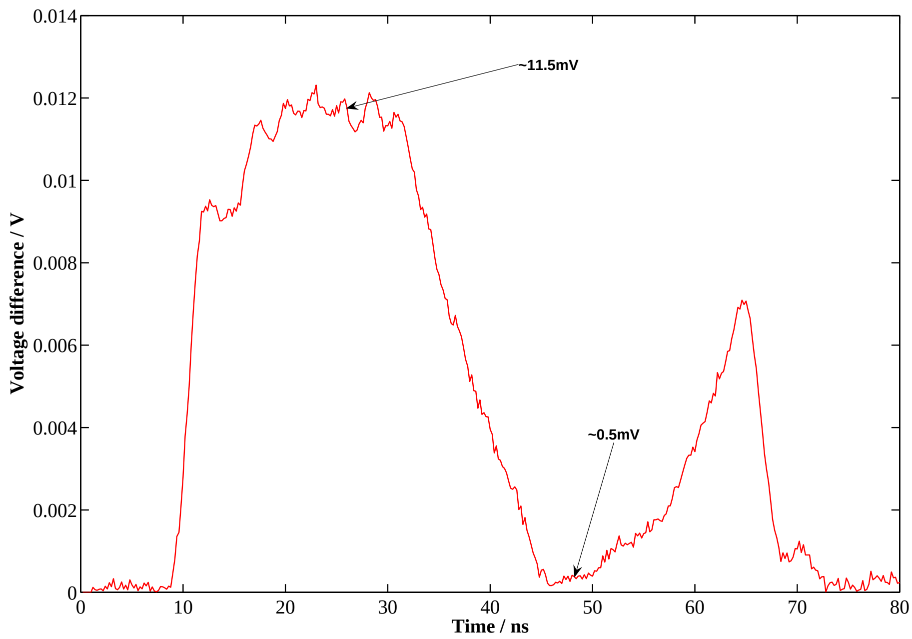

Figure 6.16: Difference in position of the two most widely separated bunch trains in figure 6.8.

That the apparently correct value of ∼0.6 is in conflict with the expected value of unity is not indicative of errors in the feedback electronics or in the data analysis. Instead it is the result of the small amount of time during which the final latency is visible – ∼7 ns. As this is only a fraction (∼0.5) of the time taken for the kicker output due to the delay loop signal to ramp up (observe the speed of the first correction), transient effects relating to various differences between the delay loop and the main loop may play a considerable role. Such differences could include different values for their bandwidths or different frequency characteristics. It is postulated that if more of the final correction period could be observed (for example, in the case of a longer bunch train, or shorter feedback latency), then a delay loop gain value of ∼1.0 would prove to give the cleanest correction. In the case of figure 6.15, the value of ∼0.6 gives the cleanest appearance of the beam, however it may not actually be the formally correct value.

6.5 Correction Ratio

The time-varying, absolute position difference between the two most widely separated position settings in figure 6.8b (ignoring the saturated curves) has been plotted in figure 6.16. In order to more clearly illustrate the rising edge of the beam it was decided that the unnormalised difference would be used instead of the normalised position as the large amount of noise on the normalised curve would disguise the beginning of the beam. It can be seen in figure 6.8b that the feedback causes the position curves to overlap at ∼45 ns, and then

Item	Value
Main Gain	3650
Latency	24±0.5 ns
Delay Gain	0.6
Correction Ratio	23:1

Table 6.1: Feedback system results.

separate again at ∼50 ns due to the lack of the delay loop, so a point in this time range has been chosen to illustrate the effectiveness of the position correction. Figure 6.16 shows that the separation of the signals is initially ∼11.5 mV, and is reduced to ∼0.5 mV at ∼48 ns by the action of the feedback. This implies a correction ratio of ∼23:1 at best.

6.6 Summary of Feedback Results

A feedback system was set up on the extraction line of the ATF. A very fast (≤5 ns) beam measurement system with a resolution of ∼3 μm was developed for this system, along with a fast, high-power amplifier.

Table 6.1 shows the results from the investigations to determine the optimal values for the feedback experiment.

The system was designed to have a latency of ∼20 ns, and table 6.1 shows that this value was exceeded by ∼4 ns (see section 6.4.2). This excess is due to tight time constraints on the experiment preventing the optimisation of the signal path between the BPM output and the amplifier input.

The optimised delay gain of 0.6 is considerably lower than the expected value of 1.0 (see section 6.4.4). This is due to the fact that only ∼7 ns of the final latency period is visible, and transient effects related to, for example, the differing frequency characteristics of the main and delay loops will have a strong effect. It is suggested that given a shorter system latency, or a longer bunch train (allowing a longer portion of the final latency period to be visible) a delay loop gain of ∼1.0 would be seen to be optimal.

The result of 23:1 for the correction ratio is for the case for two bunch trains separated by ∼130 μm (see figure 6.8b) implies a correction accuracy of ∼6 μm, which compares well with the aim of performing a micron level correction of the ATF beam.

Chapter 7

Summary and Conclusions

7.1 Summary

Since the initial FONT experiment in 2002 (see section 2.2), successive FONT prototypes have expanded on the initial idea with great success. Some of the most important specifications and results from the three FONT prototype tests are shown in table 7.1.

The first FONT prototype was installed on the NLCTA beamline, and was a very successful demonstration of the principle of a fast beam position correction. A 10:1 position correction was performed on the 65 MeV beam with a system latency of 67 ns. Figure 2.10c shows the results of the full system. Despite this success it was acknowledged that there were several shortfalls of this scheme that could be addressed in a second feedback test on the same beamline. Several details of the experiment that were targetted for improvement included,

1. The relatively long (with respect to the length of the bunch train) latency.
2. The lack of independent witness BPMs.
3. The lack of real-time, pulse-by-pulse, charge normalisation of the position signal.
4. The large size of the tube amplifier would not permit its installation at the IP of a linear collider.
5. A static position profile existed on the NLCTA beam, which disguised the full action of the correction.

	FONT1	**FONT2**	**FONT3**
Beamline	NLCTA	NLCTA	ATF
Beam Energy	65 MeV	65 MeV	1.3 GeV
Bunchtrain Length	170 ns	170 ns	56 ns
BPM Technology	Button	Button	Stripline
Processor Algorithm	Difference over sum	Log ratio	Difference
Amplifier Technology	Tube	Solid state	Solid state
System Latency (ns)	67	53	24
Correction Ratio	10:1	14:1	23:1

Table 7.1: A comparison of the specifications and results of the three FONT prototype tests.

Improvements were put in place to address all of these issues, and the resulting system was known as FONT2. Table 7.1 shows that an improved correction ratio of 14:1 was achieved. The shorter latency measured in FONT2 is dominated by the factor of two reduction in the distance between the kicker and the BPM, and is not a result of faster BPM processing electronics or amplifier (see table 3.1). The BPM processor was actually substantially slower than in FONT1 (18 ns compared to 5 ns), due to the use of logarithmic amplifiers to perform the log-ratio processing used to make sure the charge normalisation occurred in real time.

A considerable difference was made to the results by the use of a so-called 'beam flattener'. An algorithm was developed (see equation 3.1) that measured the average position profile of a number of bunch trains, and caused a signal to be applied to the kicker that would remove this position profile. As long as the position variations are slower than the bandwidth of the kicker amplifier then this system should cause them to be flattened. Results of this are shown in figure 3.18b.

Figure 3.18d shows the results of the full system – beam flattener, main loop, and delay loop.

FONT1 and FONT2 provided proof that a position feedback system could correct the position of the beam within the bunch train expected at the NLC ($\sim$ 266 ns), however it was important to perform a test that would provide a direct insight into how successful such a system would be at the NLC. For this reason it was decided to move the experiment to the ATF in KEK, which provides a $\sim$1.3 GeV beam with a bunch spacing of 2.8 ns. A scaling law implies that a micron-level correction of a $\sim$1 GeV beam scales to a nanometre-level correction of a $\sim$1 TeV beam, therefore a prototype that is shown to be capable of providing a micron-level correction of the ATF beam would be capable of performing the necessary nanometre-level correction of the linear collider beam. It would, in other words, be almost identical to the system installed at the linear collider.

Table 7.1 shows that the ATF beam is 56 ns long, which means that it would not be possible to install the FONT2 electronics on this beamline due to their long latency (see table 3.1). Also, the much higher energy of the beam means that an amplifier of much higher power would be needed to provide sufficient kick to the beam.

The FONT3 prototype was designed to overcome these obstacles and provide a micron-level correction to the ATF beam. Due to the very short train length available, it was decided that the total system latency should be $\leq$20 ns. Due to the geometry of the system $\sim$10 ns of this time was considered irreducible (see section 4.2.1), so the remaining time (10 ns) was divided evenly between the kicker amplifier and the BPM processor.

A BPM processor was designed that processed stripline signals to measure the position of the beam, and it was shown to have a resolution of $\sim$4 μm, and a signal delay of 4.6$\pm$0.3 ns.

Figure 6.15b shows the cleanest operation of the FONT3 system. For this data-set the main loop gain, and delay loop length were set to values measured to be optimal, and the delay loop gain was set to a value that gave the cleanest looking results. This is not necessarily the formally correct value, since transient effects may be coming into play.

FONT3 demonstrated a $\sim$6 μm correction of a $\sim$1.3 GeV beam with a total correction delay of 24 ns, which is $\leq$10% of the full train length of the NLC. Also of note is that all of the equipment – BPM processor, amplifier, feedback electronics – is of a physical size that could be installed in the rather crowded space of the interaction region of a linear collider.

7.2 The Future of FONT

International Linear Collider

As explained in section 1.5.1, an international panel in August 2004 recommended that the international high-energy physics community move forward to designing and building a TeV energy linear collider based on superconducting accelerating cavities. Since the problem of fast position feedback was much more challenging in the alternative, room temperature cavity machine, FONT had concentrated on providing a proof-of-principle for this design. The ITRP decision meant that the beam parameters FONT had been designing for were no longer relevant in the case of the linear collider, which meant a change in direction for the fast feedback experiment.

The new machine design (known as the ILC) has a much longer bunch separation than planned for by FONT ($\geq$100 ns instead of 2.8 ns), so it is now possible to use digital electronics instead of analogue RF circuitry. Digital devices allow for the possibility of more complicated algorithms to calculate the kick to be applied to the beam, and it could be imagined using the measured position and angle of the beam, as well as information fed-forward from the damping rings, to determine the optimal correction to apply to the beam at the IP.

Due to the ITRP decision, the ATF has upgraded to provide a bunch train that more closely resembles the ILC beam parameters. The FONT group intends to use this beam to begin tests of a digital feedback system, and work is currently underway to develop an improved BPM processor and amplifier, and a set of digital electronics based on Field Programmable Gate Arrays (FPGAs).

Since the processor designed for FONT3 has been shown to operate in single bunch mode (see section 5.1) it is possible to use the same processor for the FONT4 beam. One problem, however, is that the use of digital feedback electronics will require the position signal to be sampled at its peak, and the narrow peak provided by the processor as designed for FONT3 will be difficult to sample with the required accuracy. For this reason the BPF and LPF in the processor were re-designed to have a smaller bandwidth in order to stretch the output pulse. This was simulated, and the results can be seen in figure 7.1.

It can be seen that, as desired, a reduction in the output bandwidth from 170 MHz to 100 MHz causes a stretching of the pulse. Due to the conservation of energy, the signal amplitude is reduced by a corresponding amount, therefore it is expected that this bandwidth reduction will slightly impact the resolution of the device.

Beam tests of this processor are currently underway, and a full test of the digital feedback system is planned by the end of 2006.

Compact Linear Collider (CLIC)

An alternative design for a TeV-scale linear collider, known as the Compact Linear Collider (CLIC)[58], has been put forward by CERN, and this machine is rather novel, in that the

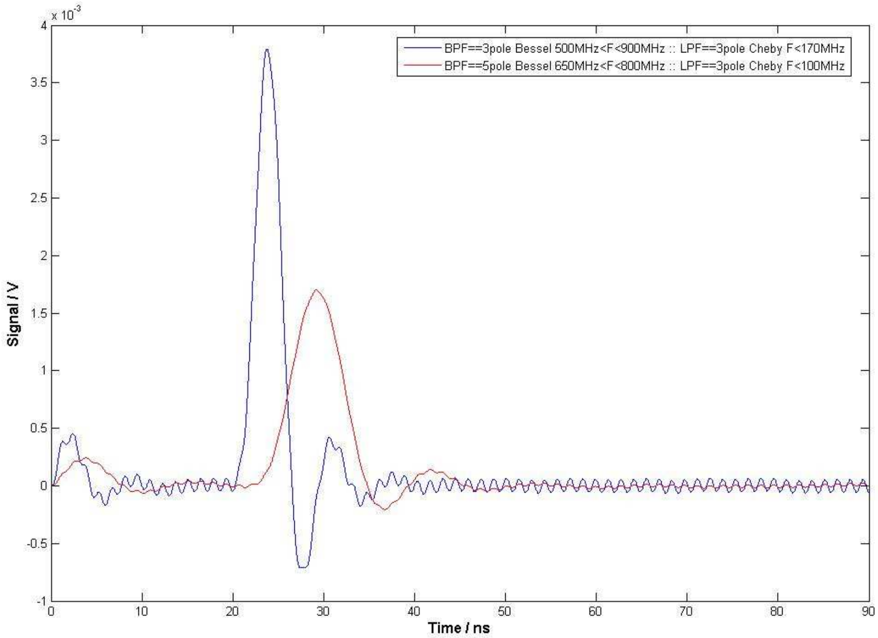

Figure 7.1: Simulated output of the FONT3 (blue) and FONT4 (red) processors.

accelerating power for the colliding beams is extracted from another beam with lower energy but higher charge. RF power at a frequency of 30 GHz is taken from this 'drive' beam in decelerating cavities and transfered to accelerating structures in the main linac, where the high RF frequency allows very high accelerating gradients (≥150 MV/m), and, therefore, much higher energies to be reached[1] (∼3 TeV).

Table 7.2 shows some parameters of the CLIC design. Note that the extremely small horizontal and vertical IP sizes imply that ground motion induced collision offsets will have an even greater impact on the luminosity than in the case of the ILC, so beam-based feedback

[1]Obviously arbitrarily high energies can be reached by a linear collider if there are no restrictions on its length, but the very high gradients predicted for CLIC allow it to reach a much higher energy than ILC for the same length.

	CLIC		
Property	**500GeV**	**3TeV**	**Units**
Electrons/bunch	0.4	0.4	$\times 10^{10}$
Bunches/train	154	154	
Train Repetition Rate	200	100	Hz
Bunch Separation	0.67	0.67	ns
Train Length	102	102	ns
Horizontal IP Beam Size (σ_x)	202	60	nm
Vertical IP Beam Size (σ_y)	1.2	0.7	nm
Longitudinal IP Beam Size	35	35	μm
Luminosity	21.0	80.0	10^{33}cm^{-2}s^{-1}

Table 7.2: Some parameters of the CLIC design for a high energy linear collider.

is extremely important. The very short train length planned for CLIC means that such a feedback system would have to be based on very low latency analogue electronics, just as tested for FONT3 at ATF. Due to the shorter bunch separation, changes would have to be made to the BPM processor, however, the basic system as demonstrated at ATF would be applicable to the CLIC beam delivery system.

Bibliography

[1] F. Richard J.R. Schneider D. Trines A. Wagner. "TESLA Technical Design Report Part I: Executive Summary" (2001).
http://tesla.desy.de/new_pages/TDR_CD/PartI/exec.html

[2] "Zeroth order design report for the Next Linear Collider" (1996). SLAC-R-0474-VOL-1.

[3] S. Baird. "Accelerators for Pedestrians". Notes from a lecture course given to the CERN PS division.

[4] K. Jobe. "NLCTA Photo Scrapbook".
http://www.slac.stanford.edu/accel/nlc/local/Projects/NLCTA/pictures/

[5] H. Weise. "TESLA Technical Design Report: Technical Layout" (2000). Page I-28.
http://tesla.desy.de/new_pages/TDR_CD/PartI/partI.pdf

[6] G. Vismara. "Signal Processing for Beam Position Monitors". Presented at the 9th Beam Instrumentation Workshop, MIT (2000).

[7] M. Wendt. "BPM Instrumentation: A Short Introduction for Non-Experts" (2000).

[8] M. Ross *et al.* "RF cavity BPM's as beam angle and beam correlation monitors". Presented at Particle Accelerator Conference (PAC 03) (2003).
http://www.slac.stanford.edu/cgi-wrap/getdoc/slac-pub-9886.pdf

[9] L. Hendrickson *et al.* "Simulations of IP feedback and stabilization in the NLC". Presented at the 9th European Particle Accelerator Conference (EPAC 2004), Lucerne, Switzerland, 5-9 July 2004.
http://www.slac.stanford.edu/cgi-wrap/getdoc/slac-pub-10492.pdf

[10] R. Dubois. "The Stanford Linear Collider Home Page".
http://www-sld.slac.stanford.edu/sldwww/slc/slc.html

[11] T. Johnson. "Stanford Large Detector Information Page".
http://www-sld.slac.stanford.edu/sldwww/sld-working.html

[12] A. Seryi. "A shaky road to subnanometer beams: NLC ground motion, vibration and stabilization studies" (2001).
http://www.slac.stanford.edu/cgi-wrap/getdoc/slac-pub-8825.pdf

[13] G. White. Private Communication.

[14] S. Smith. "Design of an NLC intra-pulse feedback system". SLAC-TN-03-067.
http://www-project.slac.stanford.edu/lc/temp/NLC%20Tech%20Notes/LCC-0056%20final.pdf

[15] S. Jolly. "An Intra-Pulse Fast Feedback System for a Future Linear Collider". Oxford University (2003). D.Phil thesis.
http://hepwww.ph.qmul.ac.uk/lcdata/FONT/thesis/contents.html

[16] A. Kalinin. Private Communication.

[17] "ATF Collaboration Home Page" (2005).
http://www-atf.kek.jp/atf/index.html

[18] M. Gellis. "180° Hybrid Tutorial" (2000).
http://members.tripod.com/michaelgellis/180hybrid.html/

[19] "Matlab".
http://www.mathworks.com/

[20] Pulsar Microwave Corporation. "Mixers" (2002).
http://www.pulsarmicrowave.com/products/mixers/mixers.htm

[21] "Andrew Corporation Homepage".
http://www.andrew.com/catalog38/cat38viewer.aspx?PageNum=520

[22] P. Burrows. "Beam Based Feedback Systems". Presented at Snowmass (2005).
http://alcpg2005.colorado.edu:8080/alcpg2005/program/accelerator/GG2/phil_burrows20050818043330.ppt

[23] S. Eidelman *et al.* "Review of Particle Physics" (2004). http://pdg.lbl.gov

[24] "NLC BDIR Configuration" (2003).
http://www-sldnt.slac.stanford.edu/nlc/configs/2003/index.html

[25] T. Raubenheimer. "Suggested ILC Beam Parameter Range" (2005).
http://www-project.slac.stanford.edu/ilc/acceldev/beampar/Suggested%20ILC%20Beam%20Parameter%20Space.pdf

[26] J. Urakawa *et al.* "Accelerator Test Facility, Design and Study Report" (1995).
http://lcdev.kek.jp/ATF/Pub/KEK-I-95-4.pdf

[27] K. Wille. "The Physics of Particle Accelerators". Oxford University Press (2000).

[28] "International Linear Collider Technical Review Committee, Second Report" (2003). Prepared for the International Committee for Future Accelerators (ICFA).
http://www.slac.stanford.edu/pubs/slacreports/slac-r-606.html

[29] "Large Hadron Collider".
http://lhc.web.cern.ch/lhc/

[30] "CERN homepage".
http://public.web.cern.ch/Public/Welcome.html

[31] "International Linear Collider Homepage".
http://www-project.slac.stanford.edu/ilc/

[32] "Stanford Linear Accelerator Center homepage".
http://home.slac.stanford.edu/aboutslac.html

[33] A.D. Yeremian R.H. Miller. "Handbook of Accelerator Physics and Engineering; Electron Guns and Preinjectors" (2002). Pages 423-428.

[34] H. Wiedemann. "Synchrotron Radiation". Springer-Verlag (2003).

[35] "Deutsches Elektronen-Synchrotron homepage".
http://www.desy.de/html/home/fastnavigator.html

[36] "International Technology Review Panel Recommendation" (2004).
http://www.interactions.org/cms/?pid=1014290

[37] M. Minty and F. Zimmerman. "Measurement and Control of Charged Particle Beams". Springer-Verlag (2003).

[38] R. Shafer. "Beam Position Monitoring" (1990).

[39] M. Wendt. "BPM Read-out Electronics Based on the Broadband AM/PM Normalisation Scheme". Presented at DIPAC – ESRF Grenoble (2001).
http://accelconf.web.cern.ch/accelconf/d01/papers/CT01.pdf

[40] G. Vismara. "The Comparison of Signal Processing Systems for Beam Position Monitors".
http://hepwww.ph.qmul.ac.uk/lcdata/FONT/reference/vismara_bpmproc_dipac99.pdf

[41] S. Smith *et al.* "ILC Cavity BPMs". Presented at Snowmass (2005).
http://alcpg2005.colorado.edu:8080/alcpg2005/program/accelerator/GG2/steve_smith20050818042632.ppt

[42] D. Shulte. "Study of Electromagnetic and Hadronic Background in the Interaction Region of the TESLA collider". University of Hamburg (1996). PhD Thesis.

[43] G. White. "FONT Homepage".
http://hepwww.ph.qmul.ac.uk/∼white/FONT/default.htm

[44] R. Tatchyn. "Selected parameter studies of the NLCTA as a FEL driver over the 130-40 Angstrom wavelength range". Given at the Workshop on 4th Generation Light Sources (1992).

[45] J. Frisch. Private Communication.

[46] C. Perry (2002).

[47] C. Perry (2003).

[48] Hewlett-Packard. "Tutorial Description of the Hewlett-Packard Interface Bus". Hewlett-Packard (1987).

[49] N. Holt. "The Architecture of OpenVME". ICL (1995).

[50] "KEK: High Energy Research Organisation" (2004).
`http://www.kek.jp/intra-e/index.html`

[51] J. Urakawa *et al.* *Physical Review Letters*, **92**, 5.

[52] Eagleware-Elanix Corporation. "Systemview Professional Edition V6.0" (2005).
`http://www.eagleware.com/products/systemview/index.html`

[53] T. Quarles *et al.* "SPICE General Purpose Circuit Simulation Program".
`http://bwrc.eecs.berkeley.edu/Classes/IcBook/SPICE/`

[54] "MACOM Homepage".
`http://www.macom.com`

[55] "Mini–Circuits homepage" (2006).
`http://www.minicircuits.com`

[56] "ZHL-2010 specifications".
`http://www.minicircuits.com/zhl-2010.pdf`

[57] C. Clarke. "FONT3 Simulations". Internal FONT note.

[58] "Compact LInear Collider (CLIC) Study".
`http://clic-study.web.cern.ch/CLIC-Study/`