The Development of a Fast Intra-train Beam-based Feedback System Capable of Operating on the Bunch Trains of the International Linear Collider

Douglas Robert Bett
Lady Margaret Hall, Oxford

Thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy at the University of Oxford

Trinity Term, 2013

Abstract

This thesis will describe the latest work from the Feedback On Nanosecond Timescales project, commonly known as FONT. The goal of the FONT project is the development of a beamline feedback system to be installed at the interaction point (IP) of a future linear collider in order to maximize the luminosity that can be achieved. The prototype FONT feedback system is beam-based, meaning that the correction is determined from direct measurement of the position of the beam, and intra-train, meaning that the correction is applied within the duration of the current bunch train.

The FONT system, consisting of three stripline beam position monitors, a digital processor unit built around a Field Programmable Gate Array (FPGA) and a pair of electromagnetic kickers, is described. Recent improvements to the position measurement process are detailed and the performance of the feedback system is presented.

The modification of the firmware to operate on a machine with a large number of bunch trains, such as the International Linear Collider, is described and the design is verified through the use of a laboratory test bench developed to simulate such a machine. The FONT5 digital board is proved capable of operating on a train resembling the specification for the International Linear Collider: 2820 bunches separated in time by 308 ns.

Acknowledgements

It almost goes without saying that a great deal of the credit for the work presented here must go to my supervisor, Phil Burrows. I am truly indebted to him for giving me the opportunity to work on such an interesting project with such a great group of people. This includes Colin Perry - whose vast depth of knowledge was only exceeded in usefulness by his willingness to share it - as well as Glenn Christian and Javier Resta López.

Thanks also to my fellow students, Ben Constance, Mikey Davis and Neven Blaskovic Kraljevic, whose dedication and enthusiasm were truly inspirational. A special thank you goes to Robert Apsimon who showed me the ropes when I first started as a graduate student and continued to provide me with advice and assistance right through to the end. I would also like to take this opportunity to thank our Japanese hosts at KEK and our international colleagues from the ATF2 collaboration.

A final thank you goes to the people to whom I owe everything - my parents, Tom Bett and Denise Bett.

Contents

1 Introduction **1**
- 1.1 Particle accelerators . . . 1
  - 1.1.1 Linear and circular accelerators . . . 1
  - 1.1.2 Colliders . . . 2
- 1.2 The International Linear Collider . . . 3
  - 1.2.1 Feedback systems for linear colliders . . . 5
- 1.3 The Accelerator Test Facility . . . 6
  - 1.3.1 ATF . . . 6
  - 1.3.2 ATF2 . . . 7
- 1.4 Summary . . . 7

2 Description of the FONT5 feedback system **9**
- 2.1 BPMs and analogue processor electronics . . . 9
  - 2.1.1 Stripline BPMs . . . 10
  - 2.1.2 Analogue BPM processing electronics . . . 14
  - 2.1.3 Latency of the processor modules . . . 17
- 2.2 FONT5 digital board . . . 17
  - 2.2.1 Hardware . . . 17
  - 2.2.2 Firmware . . . 20
- 2.3 Kickers and kicker amplifiers . . . 28
  - 2.3.1 Kickers . . . 28
  - 2.3.2 Kicker amplifiers . . . 28
- 2.4 Operational concerns for the FONT5 system . . . 30
  - 2.4.1 Eliminating the static ADC voltage offsets . . . 30
  - 2.4.2 Optimization of the LO phase . . . 31
  - 2.4.3 Optimization of the ADC sample timing . . . 32
  - 2.4.4 Centring of the beam in the BPMs . . . 35
- 2.5 Summary . . . 35

3 BPM Calibration and Resolution **37**
- 3.1 The FONT installation at the ATF . . . 37
  - 3.1.1 Layout of beamline components . . . 37
  - 3.1.2 FONT processing hardware at the ATF . . . 38
  - 3.1.3 Analysis of FONT5 data . . . 40
- 3.2 BPM calibration . . . 41
  - 3.2.1 Theoretical calibration . . . 41

3.2.2 Calibration methods . . . 43
3.2.3 Calibration results . . . 48
3.3 Resolution of the FONT BPM system . . . 48
3.3.1 Model method . . . 52
3.3.2 Fit method . . . 53
3.4 Effect of LO phase offset on measured position . . . 58
3.4.1 Effect of LO phase on position for an ideal processor . . . 58
3.4.2 Effect of a path length difference at the processor inputs . . . 60
3.4.3 LO phase scans . . . 60
3.4.4 Phase compensation . . . 64
3.5 Summary . . . 70

4 Performance of the FONT5 Feedback System 71
4.1 The FONT5 feedback algorithm . . . 71
4.1.1 Derivation . . . 71
4.1.2 Implementation . . . 72
4.2 Latency of the feedback system . . . 73
4.3 Performance of the feedback system . . . 77
4.3.1 Kicker calibration . . . 77
4.3.2 Observed effect of the feedback system . . . 80
4.4 Factors influencing feedback performance . . . 87
4.4.1 Expected jitter given observed bunch to bunch correlations . . . 88
4.4.2 Phase sensitivity . . . 89
4.4.3 Alternative feedback algorithm . . . 89
4.5 Summary . . . 91

5 Design and Simulation of a Long Train Feedback System 92
5.1 Simulation of a feedback system using FONT hardware . . . 92
5.1.1 FONT4 digital board . . . 92
5.1.2 Generation of FONT5 timing signals . . . 94
5.1.3 Simulation of BPM signals . . . 96
5.2 FONT5 firmware for long bunch trains . . . 104
5.3 Results of simulated feedback system . . . 107
5.3.1 Simulated bunch trains . . . 107
5.3.2 Simulated kicker calibration . . . 109
5.3.3 Simulated feedback . . . 110
5.3.4 Simulated gain scan . . . 112
5.4 Summary . . . 113

6 Conclusions 116
6.1 Summary . . . 116
6.2 Outlook . . . 117

Bibliography 118

List of Figures

1.1 Feynman diagrams for the principal Higgs-production mechanisms at the ILC 4
1.2 Schematic representation of the International Linear Collider 4
1.3 Ground motion measured at a range of locations 5
1.4 Schematic of the FONT feedback system for the ILC 6
1.5 A schematic representation of the ATF. 7

2.1 Simplified schematic of the operation of the FONT5 system 10
2.2 Diagram of the BPM co-ordinate system 11
2.3 Photographs of BPMs at the ATF 11
2.4 Technical drawing of the ATF stripline BPMs. 12
2.5 Stripline output signal plotted in both time and frequency domains 13
2.6 Schematic diagram of the FONT analogue processor modules 15
2.7 Photograph of FONT BPM processor module 3 16
2.8 Measurement of the latency of the BPM processor modules 18
2.9 Schematic of the FONT5 digital board 19
2.10 Photograph of a FONT5 digital board 21
2.11 Photograph of the primary PCB of the FONT5 digital board 22
2.12 Photograph of the front panel of the FONT5 digital board 23
2.13 Schematic diagram of the FONT feedback firmware 25
2.14 Schematic of the FONT5 timing signals on the time scale of the 2.16 MHz clock 26
2.15 Schematic of the FONT5 timing signals on the time scale of the fast clock . 27
2.16 Technical drawing of the kickers used in the FONT feedback system 29
2.17 Photograph of FONT kicker K1 29
2.18 Photographs of a FONT kicker amplifier 30
2.19 Schematic of the LO distribution for the FONT processor modules 32
2.20 Calibration plots for the LO phase shifters 33
2.21 Processor output signals from a single BPM as digitized by a FONT5 board 34
2.22 Diagram illustrating the function of the scan delay parameter 35

3.1 Schematic representation of the beamline in the FONT region 38
3.2 Schematic representation of the beamline in the MFB1FF region 38
3.3 Schematic indicating how the FONT hardware is connected to ATF beamline components 39
3.4 A single bunch extraction data set. 42
3.5 Uncalibrated vertical position of the bunch over the duration of a mover calibration run 44

3.6 Mean value of the uncalibrated vertical position as a function of BPM mover vertical setting . . . 45
3.7 Measured magnetic flux density as a function of corrector current . . . 47
3.8 Mean value of the uncalibrated vertical position as a function of the expected displacement due to the ZV8X current . . . 49
3.9 Histogram of measured BPM calibration constants . . . 50
3.10 Comparison of theoretical calibration constant with observations . . . 51
3.11 Effect of the drift removal procedure . . . 52
3.12 Comparison of the measured position distributions with the prediction obtained using the model method . . . 54
3.13 Comparison of the measured position distributions with the prediction obtained using the fit method . . . 56
3.14 Distribution of the scaled residuals for both model and fit methods . . . 57
3.15 Value of the bunch samples for P1 as a function of master LO phase . . . 61
3.16 Behaviour of the ratios $\frac{\Delta}{\Sigma_I}$ and $\frac{\Sigma_Q}{\Sigma_I}$ as a function of master LO phase . . . 63
3.17 Effect of the phase compensation procedure in the linear range of the LO phase scan . . . 66
3.18 Correlation of the position y or the phase compensated position ξ for each pair of BPMs . . . 67
3.19 Effect of the phase compensation procedure for a calibration using the P1 mover 68
3.20 Distribution of the scaled residuals using the phase compensated position ξ . 69

4.1 Schematic of the simplified FONT5 feedback module . . . 73
4.2 Schematic illustrating the principle behind the latency measurement . . . 75
4.3 Plot of the latency measurement for the P2-K1 loop . . . 76
4.4 Plot of the latency measurement for the P3-K1 loop . . . 77
4.5 Kicker calibration plots . . . 79
4.6 Plot of the vertical beta function in the FONT region. . . . 81
4.7 Histograms of the positions during the feedback run for P1 . . . 82
4.8 Histograms of the positions during the feedback run for P2 . . . 83
4.9 Histograms of the positions during the feedback run for P3 . . . 84
4.10 Histograms of the positions during the feedback run for MQD14X . . . 85
4.11 Histograms of the positions during the feedback run for MQF15X . . . 86
4.12 Comparison of the performance of feedback algorithms . . . 91

5.1 Schematic illustrating the principle of simulation of a feedback system using FONT hardware . . . 93
5.2 Photograph of the FONT4 digital board . . . 94
5.3 Schematic of the simulated dual BPM, dual kicker system . . . 95
5.4 Diagram illustrating the significance of the timing parameters . . . 97
5.5 Diagram indicating the role of the bunch strobe and associated signals . . . 102
5.6 Actual beam and simulated beam signals digitized by the FONT5 board . . . 103
5.7 Schematic illustrating the behaviour of the FONT4 DACs . . . 104
5.8 Schematic of the setup used to investigate performance of the long train FONT5 firmware . . . 108

5.9 Σ_I output of the FONT4 test-bench digitized using the long train FONT5 firmware 109
5.10 Mean recorded bunch position for the first ten bunches of the ILC-like train during a scan of the FONT5 constant DAC out value 111
5.11 Comparison of Δ values reported by FONT4 board and FONT5 board . . . 112
5.12 Bunch position as a function of bunch number while nominal gain feedback was active 113
5.13 Standard deviation of bunch position as a function of bunch number for various values of the train jitter 114
5.14 Mean bunch position as a function of bunch number for various values of the feedback gain parameter 114
5.15 Mean and standard deviation of bunch 2 and bunch 2820 as a function of the feedback gain parameter 115

List of Tables

1.1 List of design parameters for the ATF accelerator 7

2.1 Trim DAC settings for the FONT5 board ADCs 31

3.1 Value of the gain ratio parameter for all BPM processor modules 43
3.2 Measured values of the BPM calibration constants for P1, P2 and P3 44
3.3 Measured values of the BPM calibration constants for MQD14X, MQF15X and MFB1FF 46
3.4 Corrector calibration constants 47
3.5 Comparison of prediction coefficients obtained for both model and fit methods 55
3.6 Resolution predicted using both model and fit methods 55
3.7 Phase offset between Σ_I and Σ_Q outputs 62
3.8 Phase sensitivity of each combination of processor and BPM 64
3.9 Jitter due to phase predicted from the phase dependence of each BPM . . . 64
3.10 Correlation coefficients for each pair of y or ξ measurements 65
3.11 Resolution predicted from the phase compensated position using both model and fit methods 68

4.1 Estimates of the latency of the FONT5 feedback system 77
4.2 Contributions to the latency of the FONT5 system 78
4.3 The betatron phase advance between selected pairs of FONT beamline components 78
4.4 H values and derived G values 80
4.5 Calculated values of the mean position 87
4.6 Calculated values of the position jitter 87
4.7 Effect of the feedback system on the mean position and the position jitter . . 88
4.8 Comparison of non-phase compensated jitter achieved with that predicted . . 88
4.9 Comparison of phase compensated jitter achieved with that predicted 89

5.1 Values of the k parameters for simulation of the FONT region at ATF 98
5.2 List of the control registers of the FONT4 test bench firmware 105
5.3 Data returned by the FONT4 test-bench firmware 106
5.4 Beamline model parameter values 110

Chapter 1

Introduction

Particle physics is the study of the building blocks of the universe. Research in the field of particle physics concerns the very nature of matter and energy, along with the interactions between them. Modern particle physics is summarized in the theory known as the Standard Model which posits that the universe consists of quarks, leptons and the force-carrying particles, the gauge bosons [1]. Recent successes of the Standard Model include the discovery of the top quark in 1995 [2] and the tau neutrino in 2000 [3]; in March 2013, the final particle predicted to exist, the Higgs boson, was tentatively confirmed to have been discovered [4].

First theorised in 1964, the Higgs boson is the elementary particle whose existence was predicted in order to explain the masses of the fundamental particles of the Standard Model [5]. The Large Hadron Collider (LHC), the highest-energy particle accelerator in the world, was constructed with the particular aim of finding the Higgs boson. Having verified that it exists at a mass of ~ 125 GeV, there is a compelling case for a new machine dedicated to its study: a so-called "Higgs factory".

1.1 Particle accelerators

Particle accelerators are the principal tool with which particle physics research is conducted. The precursors to modern accelerators first appeared in the 1930s and ever since the technology of particle accelerators has been improving in order to deliver higher energy and higher intensity.

1.1.1 Linear and circular accelerators

Modern accelerators used in the study of high-energy physics come in two main types. A linear accelerator, or linac for short, consists of a straight beamline containing many accelerating elements. The particle has low energy at the start of the linear accelerator and attains maximum energy at the end after traversing all the accelerating elements. Circular accelerators consist of a beamline that forms a closed loop. In contrast to the single-pass nature of the linear accelerator, the beam in a circular accelerator makes many circuits

around the machine; thus, the beamline need contain only a few accelerating elements but must incorporate many components whose function is to deflect the beam so that it follows a circular path. These components are called bending magnets.

As a consequence of circulating a beam, circular accelerators must deal with several unique challenges. First, a particle whose energy increased each turn would experience a larger radius of curvature each time it traversed a magnetic field of fixed strength. In order for the beam to follow the same path around the machine as it is accelerated, the strength of these magnetic fields must therefore be increased in tandem with the beam energy. Circular accelerators where the strength of the magnetic fields is synchronized with the beam energy in this fashion are called synchrotrons. Second, the radial acceleration of the beams by the bending magnets results in the emission of radiation tangential to the beam trajectory known as synchrotron radiation [6]. In order to accelerate the beam, the energy provided each turn must exceed these losses. Synchrotrons where the energy provided to the beam each turn is only enough to replace the loss due to synchrotron radiation are known as storage rings.

1.1.2 Colliders

Particle physics experiments almost invariably involve colliding pairs of particles and observing what is created as a result. The point where the two particles collide is called the interaction point (IP). The energies of the beams of particles being collided determines the interactions that can be studied as there must be enough energy in the laboratory frame to provide both the rest mass energy of the particles created as well as any kinetic energy they must have in order to conserve momentum. The most efficient arrangement for a collider is thus that of a pair of oppositely directed beams of equal energy, in which case the potential rest mass energy is twice the beam energy.

For this reason a linear collider essentially consists of a pair of linacs pointed at each other. Circular colliders are able to make more efficient use of hardware by having two oppositely-directed beams within a pair of beam pipes that are adjacent to each other and so are able to share beamline components such as bending magnets. At certain points the two beam pipes are combined to enable collisions; thus, a linear collider has a single interaction point while circular colliders may have several.

An important parameter for any collider is the luminosity, L. The likelihood of an interaction X between a pair of particles is quantified by the reaction cross section, $\sigma(X)$. A rare interaction is one that has a small cross section. The rate at which a particular interaction takes place at the IP of a collider is given by $E(X) = L\sigma(X)$, where the luminosity [7] is given by:

$$L = H\frac{fN^2}{4\pi\sigma_x\sigma_y} \tag{1.1}$$

In the above expression, σ_x and σ_y are respectively the horizontal and vertical size of the bunches at the IP, f is the frequency at which pairs of bunches arrive at the IP, N is the number of particles in a bunch and H is an enhancement factor that accounts for electro-

magnetic forces present in the collision ($H = 1$ for low intensity collisions [7]). For a circular collider, f is given by the product of frequency of revolution and the number of bunches in the circulating beam; once the beam is accelerated up to full energy, it can be stored and used for collisions for many hours [8]. This is not the case for a linear collider where the two beams are disposed of post-IP; in this case, the frequency of collisions is limited by the number of bunches in a train and the rate at which the linacs are able to generate trains. For this reason, in order to achieve the luminosity required to gather reasonable amounts of data from the interactions of interest, it is necessary to use a smaller beam size at a linear collider relative to the circular case. The size of the beam is generally quantified by the emittance parameter, ϵ, which describes the variation in position and momentum that exists amongst the particles of the beam [7]. For example, the vertical emittance ϵ_y may be defined as the product of the vertical beam size σ_y and the vertical divergence σ'_y.

Given that the LHC exists and has already produced Higgs bosons, it is reasonable to ask why a second machine with this ability is necessary. The answer lies in the fact that the LHC is a proton-proton collider [9]. Protons, like all baryons, have internal structure: the proton is usually described as a bound state of three quarks, but proton interactions can also happen by means of the gluons that bind those quarks together. In short, proton-proton interactions are inherently complicated. One may be interested in the specific interaction that includes a Higgs in the final state but these interactions are invariably accompanied by a large background of less interesting events. This background can be reduced by colliding instead truly point-like leptons. The heavier μ and τ generations of charged leptons are unstable with lifetimes too short for it to be feasible to construct a machine for colliding them using current technology [10]. A future lepton machine intended for study of the Higgs would thus collide electrons and positrons.

Having established that an electron-positron collider is the preferred design allowing precision measurements to be made of the property of the Higgs, the case for a linear collider is easily made: for a circular collider, the amount of energy lost due to synchrotron radiation for a particle of charge q and mass m is proportional to $(q/m)^4$ [6]. As the electron is approximately 2,000 times lighter than the proton, an electron beam must radiate away 1×10^{13} times more energy per turn than a proton beam of the same energy following the same trajectory. Such a machine would have a very high power consumption as well as a very large beamline due to the great many accelerating and beam controlling elements that would have to be accommodated within to ensure a net increase in energy of the beam per revolution. A circular electron-positron collider is thus prohibitively expensive.

1.2 The International Linear Collider

The design for a linear lepton collider closest to fruition is that of the International Linear Collider (ILC) [11]. The initial design calls for two oppositely directed 250 GeV linacs, one for electrons and one for positrons, while leaving open the possibility of a later upgrade to 500 GeV each. The machine will thus at first primarily produce Higgs particles via the dominant mechanisms of "Higgs-Strahlung" ($ee \to ZH$) and W-W fusion ($ee \to \nu\nu H$),

diagrams Figure 1.1(a) and Figure 1.1(b) respectively [12].

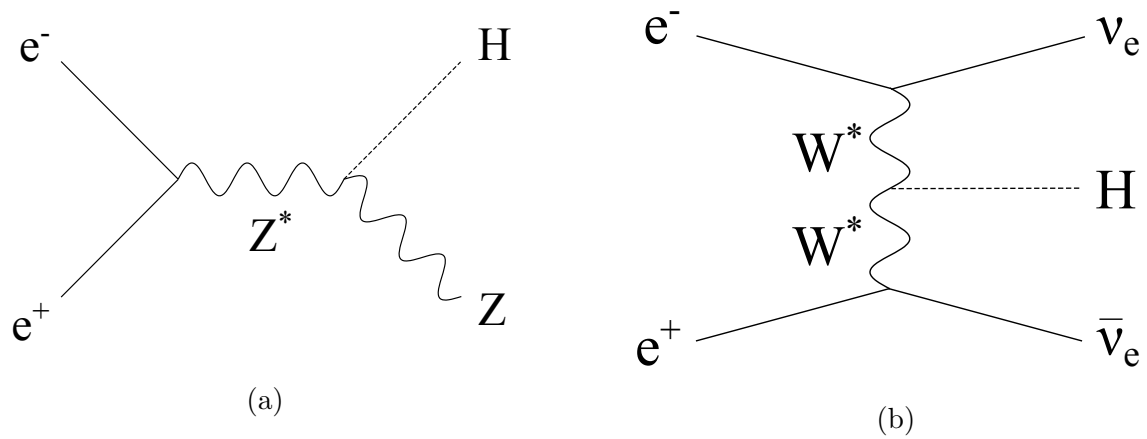

Figure 1.1: Feynman diagrams for the principal Higgs-production mechanisms at the ILC.

The design for the ILC is illustrated schematically in Figure 1.2. The blue line indicates the electron beamline while the green line corresponds to the positron beamline. Electrons are ejected from a photocathode by a laser pulse at the location identified by the pyramid. They are then accelerated up to 5 GeV prior to injection into the electron damping ring. The damping rings are a pair of storage rings whose function is to accept an injected beam with large emittance and reduce it to the ultra-low level necessary for high luminosity collisions at the IP. This occurs through a process known as radiation damping. As synchrotron radiation photons are emitted in a narrow cone around the direction of motion of the beam but acceleration is provided only along the axis defining the ideal orbit there will be a gradual loss of momentum in the transverse directions [13]. The damping process takes less than 25 ms [14] after which the beam is extracted from the damping ring and transported to the beginning of the 11 km long main electron linac.

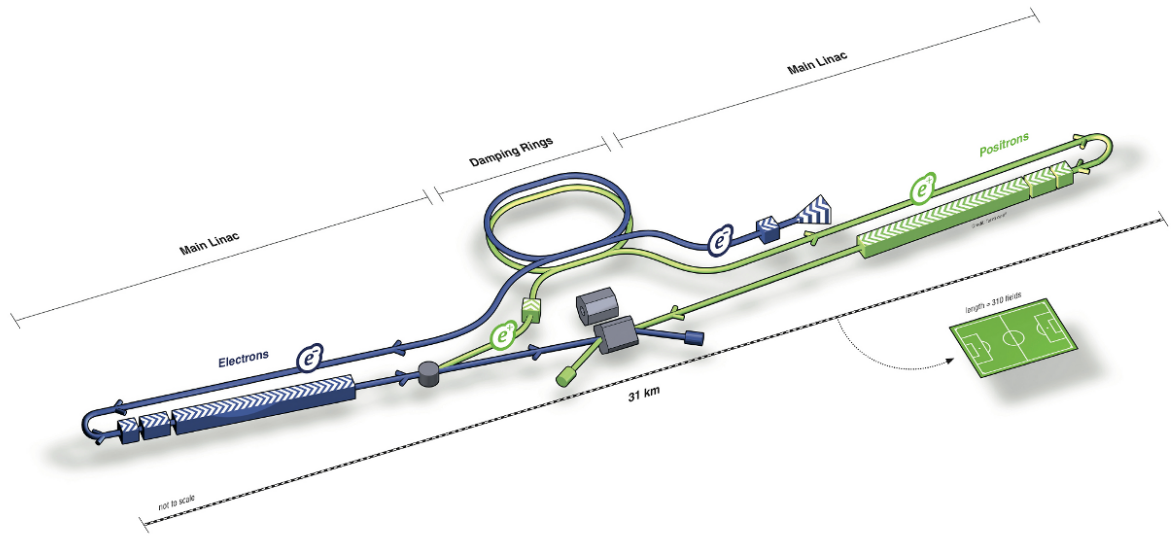

Figure 1.2: Schematic representation of the International Linear Collider [15].

The linac consists primarily of radiofrequency (RF) cavities. In each cavity a longitudinal electric field is induced by an RF source and imparts energy to the passing beam. The current design for the ILC calls for superconducting cavities. These offer several advantages over normal conducting cavities due to their low electrical loss. They can support an electromagnetic field for a longer period of time without damage to the cavity and thus permit operation with long bunch trains; at the same time, they can provide a high accelerating field even with a large aperture which leads to lower wakefields where the passage of the previous bunch has an adverse effect on subsequent bunches [16]. After the electron beam has been accelerated to high energy it passes through a long helical undulator, a device consisting of a periodic structure of dipole magnets of alternating polarity used to stimulate the emission of synchrotron radiation. The resultant photon beam is then impinged on a titanium alloy target to generate positrons via the mechanism of pair production; the procedure for accelerating the collected positrons is then identical to that outlined so far for the electrons. After the main linacs, the two beams will be focused such that the beam size is reduced to $\sigma_x = 640$ nm, $\sigma_y = 5.7$ nm [11] at the IP, where one of two detectors will be in place to observe the result of the interaction.

1.2.1 Feedback systems for linear colliders

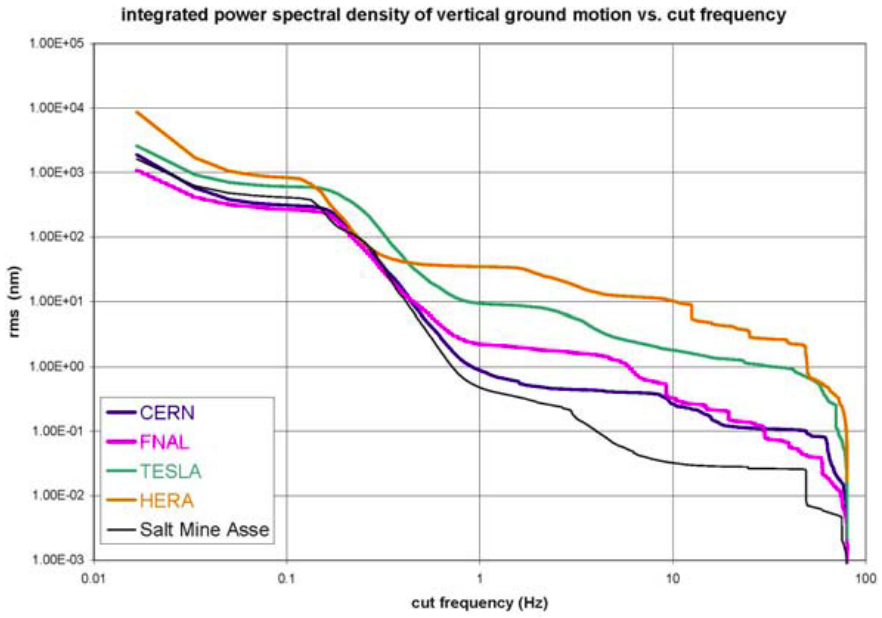

Figure 1.3: Ground motion as a function of frequency as measured at a range of locations [17].

The extremely small beam size at the IP of the ILC makes the beam stability an area of great concern as a small position offset between the two beams results in a dramatic loss of luminosity [18]. Figure 1.3 shows the ground motion spectra measured at five different sites integrated up to a cut-off frequency in order to give the RMS vertical ground motion as a function of frequency. The data shows that motion on the order of tens of nanometres could

be experienced at the ILC train repetition rate of 5 Hz. An intra-train feedback system is necessary to correct for such motion.

The FONT fast intra-train beam-based feedback system is intended to ensure high luminosity at the IP and is shown schematically in Figure 1.4. The blue and green arrows represent respectively the electron and positron beams and the IP is located at the point where the arrows cross. In the FONT scheme, the position of the electron bunch is measured several metres downstream of the IP; at this point, the deflection due to the Coulomb interaction that results from a nanometre-level offset between the two bunches becomes a change in position on the order of tens of microns, greatly relaxing the resolution requirement for the position measurement.

The processor of the feedback system then determines from this position the signal that, once amplified and applied to the kicker located in the positron beamline just upstream of the IP, will deflect the next bunch of positrons to compensate for the offset of the previous bunch relative to the last electron bunch. The delay loop ensures that the required deflection is applied to subsequent bunches as well; in the absence of a delay loop, the system would interpret the improvement it causes as evidence that no kick is required. The delay loop thus allows the correction to be maintained throughout the bunch train while being continually refined based on the latest measurements.

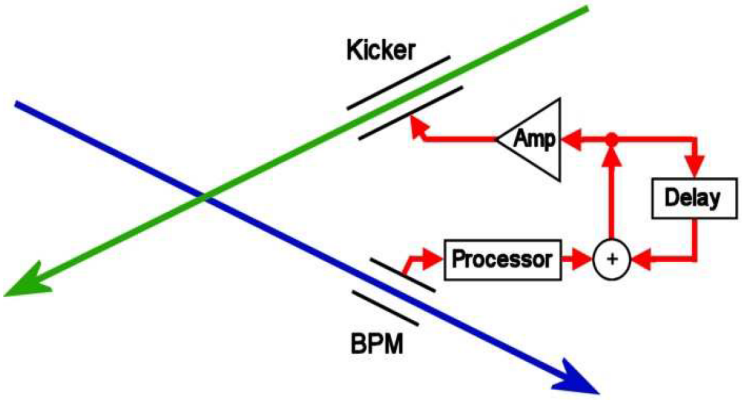

Figure 1.4: Schematic of the FONT feedback system for the ILC.

As the ILC is currently still at the design stage, an implementation of the FONT system was deployed at the Accelerator Test Facility for testing and development.

1.3 The Accelerator Test Facility

1.3.1 ATF

The Accelerator Test Facility (ATF) is a test accelerator located at the High Energy Accelerator Research Organization (KEK) [19] in the city of Tsukuba, Japan and depicted schematically in Figure 1.5. The ATF is an electron machine and consists of a linac, a

damping ring and an extraction line. Commencing operation in 1997, the original goal of the ATF was the super low emittance beam required in order to realize a future electron-positron collider [20]. The design parameters for the ATF are given in Table 1.1; the design emittance was achieved in 2001.

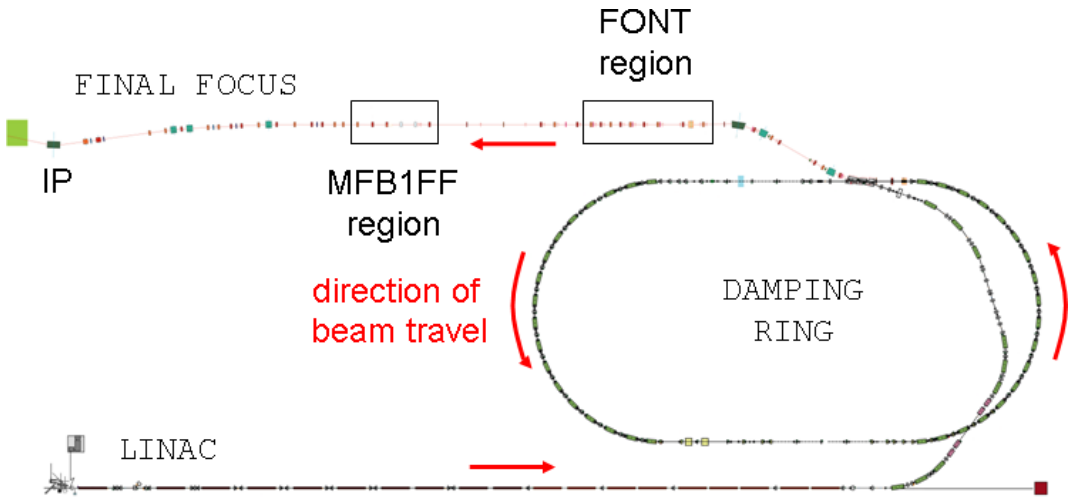

Figure 1.5: A schematic representation of the ATF.

Parameter	Design value
Energy	1.28 GeV
Intensity	1×10^{10} electrons/bunch
Repetition rate	$0.7 - 6.4$ Hz
ϵ_x	1×10^{-9} m rad
ϵ_y	1×10^{-11} m rad
Beam size ($\sigma_x \times \sigma_y$)	70 μm × 7 μm (rms)

Table 1.1: List of design parameters for the ATF accelerator.

1.3.2 ATF2

In 2008 as part of the ATF2 project the existing ATF extraction line was extended to include a much longer section of beamline of the type that would be used at the final focus of the ILC. The FONT feedback system was installed in the location of the original extraction line in order to gather the data contained in this thesis.

1.4 Summary

This chapter provides a very brief overview of the motivation for a particle accelerator such as the proposed International Linear Collider. The Accelerator Test Facility, the prototype

accelerator designed to demonstrate how some of the key challenges involved in the development of such a future linear collider may be overcome, was introduced and the basic principle of the FONT intra-train feedback system was outlined.

Chapter 2

Description of the FONT5 feedback system

The FONT5 feedback system was developed in order to demonstrate the concept of the FONT IP feedback system. It differs from the system envisaged at the IP of a linear collider (Figure 1.4) in that a second BPM and kicker are used to allow for correction of the beam angle as well as the beam position; also, the BPMs and kickers are installed on the same beamline. The problem of ensuring two bunches collide is thus substituted for stablisation of the position and angle of a single beam at some arbitrary location. The FONT5 system is illustrated schematically in Figure 2.1. The system consists of three stages: firstly, stripline beam position monitors (BPMs) and analogue processing electronics deliver the signals required to determine the position of the beam. Secondly, digital logic operating on a hardware module built around a field-programmable gate array (FPGA) processes the position signals and uses this information to calculate the kick required to bring the beam on axis. Thirdly, the kick signal is transmitted via a custom-made high voltage, low latency amplifier to an upstream kicker in order to modify the position of the next bunch in the current train. A brief description of each of these three stages along with a discussion of the system as a whole forms the material of this chapter.

2.1 BPMs and analogue processor electronics

A beam position monitor (BPM) is a device installed in the beam pipe of a particle accelerator and used to determine the transverse location of the beam as it traversed the BPM [21]. Figure 2.2 defines the co-ordinate system that will be used throughout this thesis. The beam pipe is coaxial with the z-axis and the beam propagates in the direction of increasing z. The trajectory of the beam is such that at any given instant it will be displaced from the z-axis in both the x and y directions and will travel with a direction vector that is not parallel to the z-axis. In this thesis, the transverse vertical displacement of the beam will be referred to as the position of the beam and represented by the symbol y. The angle of the direction vector projected in the $y-z$ plane relative to the z-axis will be referred to as the angle of the beam and represented by the symbol y'.

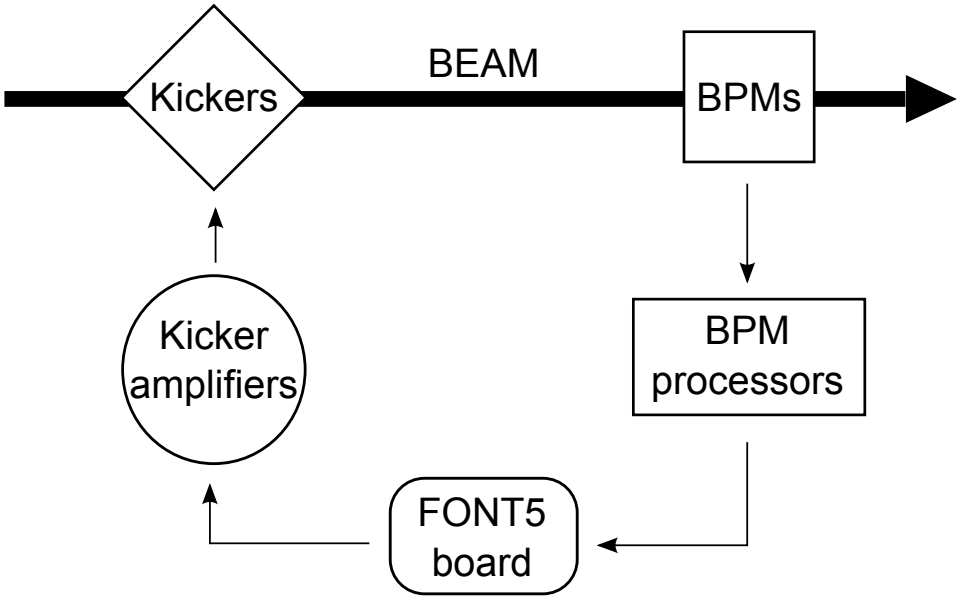

Figure 2.1: Simplified schematic of the operation of the FONT5 system.

The beam position monitoring system in the extraction line of the ATF makes use of two varieties of BPM (Figure 2.3):

- **Stripline BPMs**
  The bunch of charged particles induces voltages on strips of conducting material located on opposite sides of the beampipe [22].
- **Cavity BPMs**
  The bunch of charged particles sets up a standing wave within an RF cavity [22].

2.1.1 Stripline BPMs

The stripline BPMs used at the ATF consist of four strips in a symmetrical arrangement around the inside of a cylinder that forms the local beam pipe (Figure 2.4). The striplines form two pairs; the designations "top" and "bottom" are used for the striplines that define the y-axis of the BPM and "left" and "right" are used for those on the x-axis, where these directions are relative to a bunch travelling in the direction of increasing z. Each strip lies along the length of the cylinder (i.e. in the z-direction) and conforms to the circular cross-section of the beam pipe. The ends of each strip serve as breaks in the electrical continuity of the beam pipe and as a result the passage of a bunch of charged particles induces a voltage pulse. The form of the voltage pulse is determined by the induced image current (dQ/dt), the geometry of the BPM (radius R, width w), the impedance of the measurement electronics (ρ) and the location of the bunch along the axis defined by the pair of diametrically opposed striplines (x, y). An oscilloscope was used to measure the voltage signal induced on a single

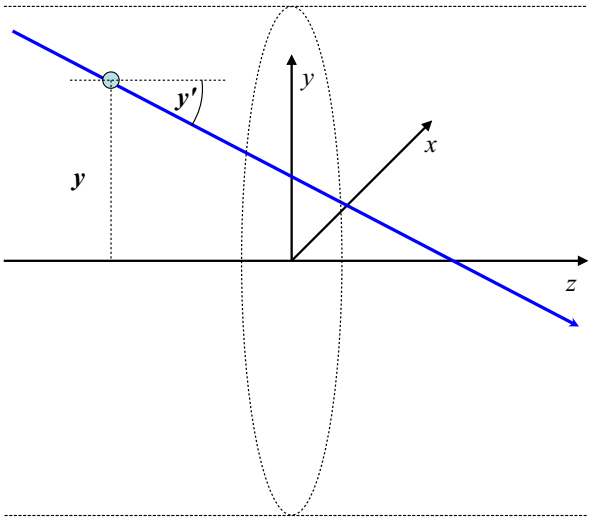

Figure 2.2: Diagram of the co-ordinate system used in this thesis. The cylindrical beam pipe contains a bunch (blue circle) the direction of travel of which is indicated by the blue arrow.

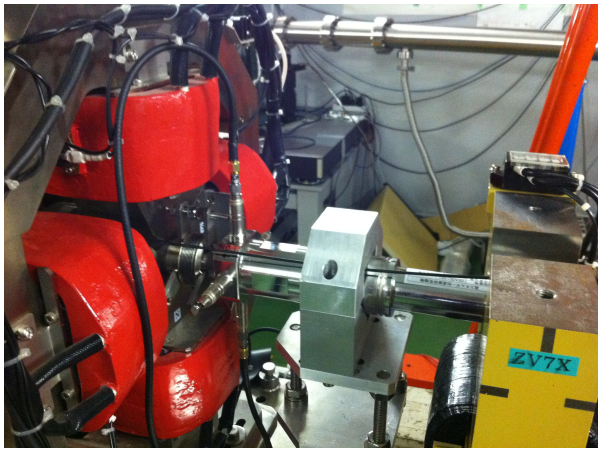

(a)

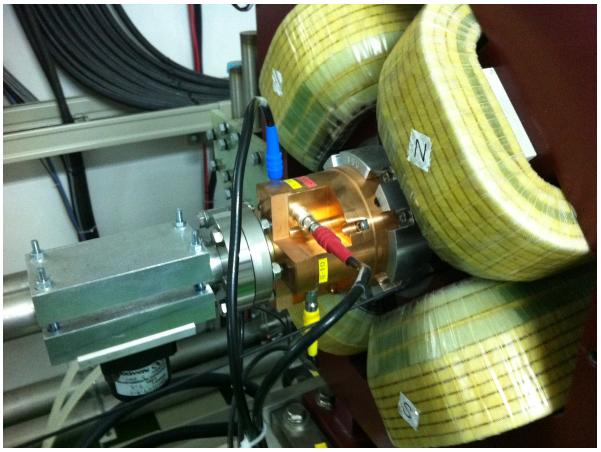

(b)

Figure 2.3: Photographs of BPMs at the ATF. (a) stripline FONTP1; (b) cavity MQF16FF.

stripline (i.e. the voltage with respect to ground) by the passage of a bunch; the result is shown in Figure 2.5(a). The pulse is bipolar and may be explained as follows: the voltage is zero prior to the arrival of the bunch at the BPM at which point a positive peak is generated as the bunch crosses the electrical discontinuity between beam pipe and stripline. This is followed by a brief period of zero voltage as the bunch travels the length of the BPM and then a negative peak of equal magnitude but opposite sign to the first as the bunch crosses the boundary between stripline and beam pipe. The width of the peaks is artificially broad due to the limited bandwidth of the oscilloscope; the typical bunch length at the ATF is ~ 8 mm [23].

A bunch passing through the BPM with a vertical displacement of y from the centre thus gives rise to a pair of voltage pulses at the top (Equation 2.1) and bottom (Equation 2.2) outputs that may be designated V_A and V_B respectively [24]:

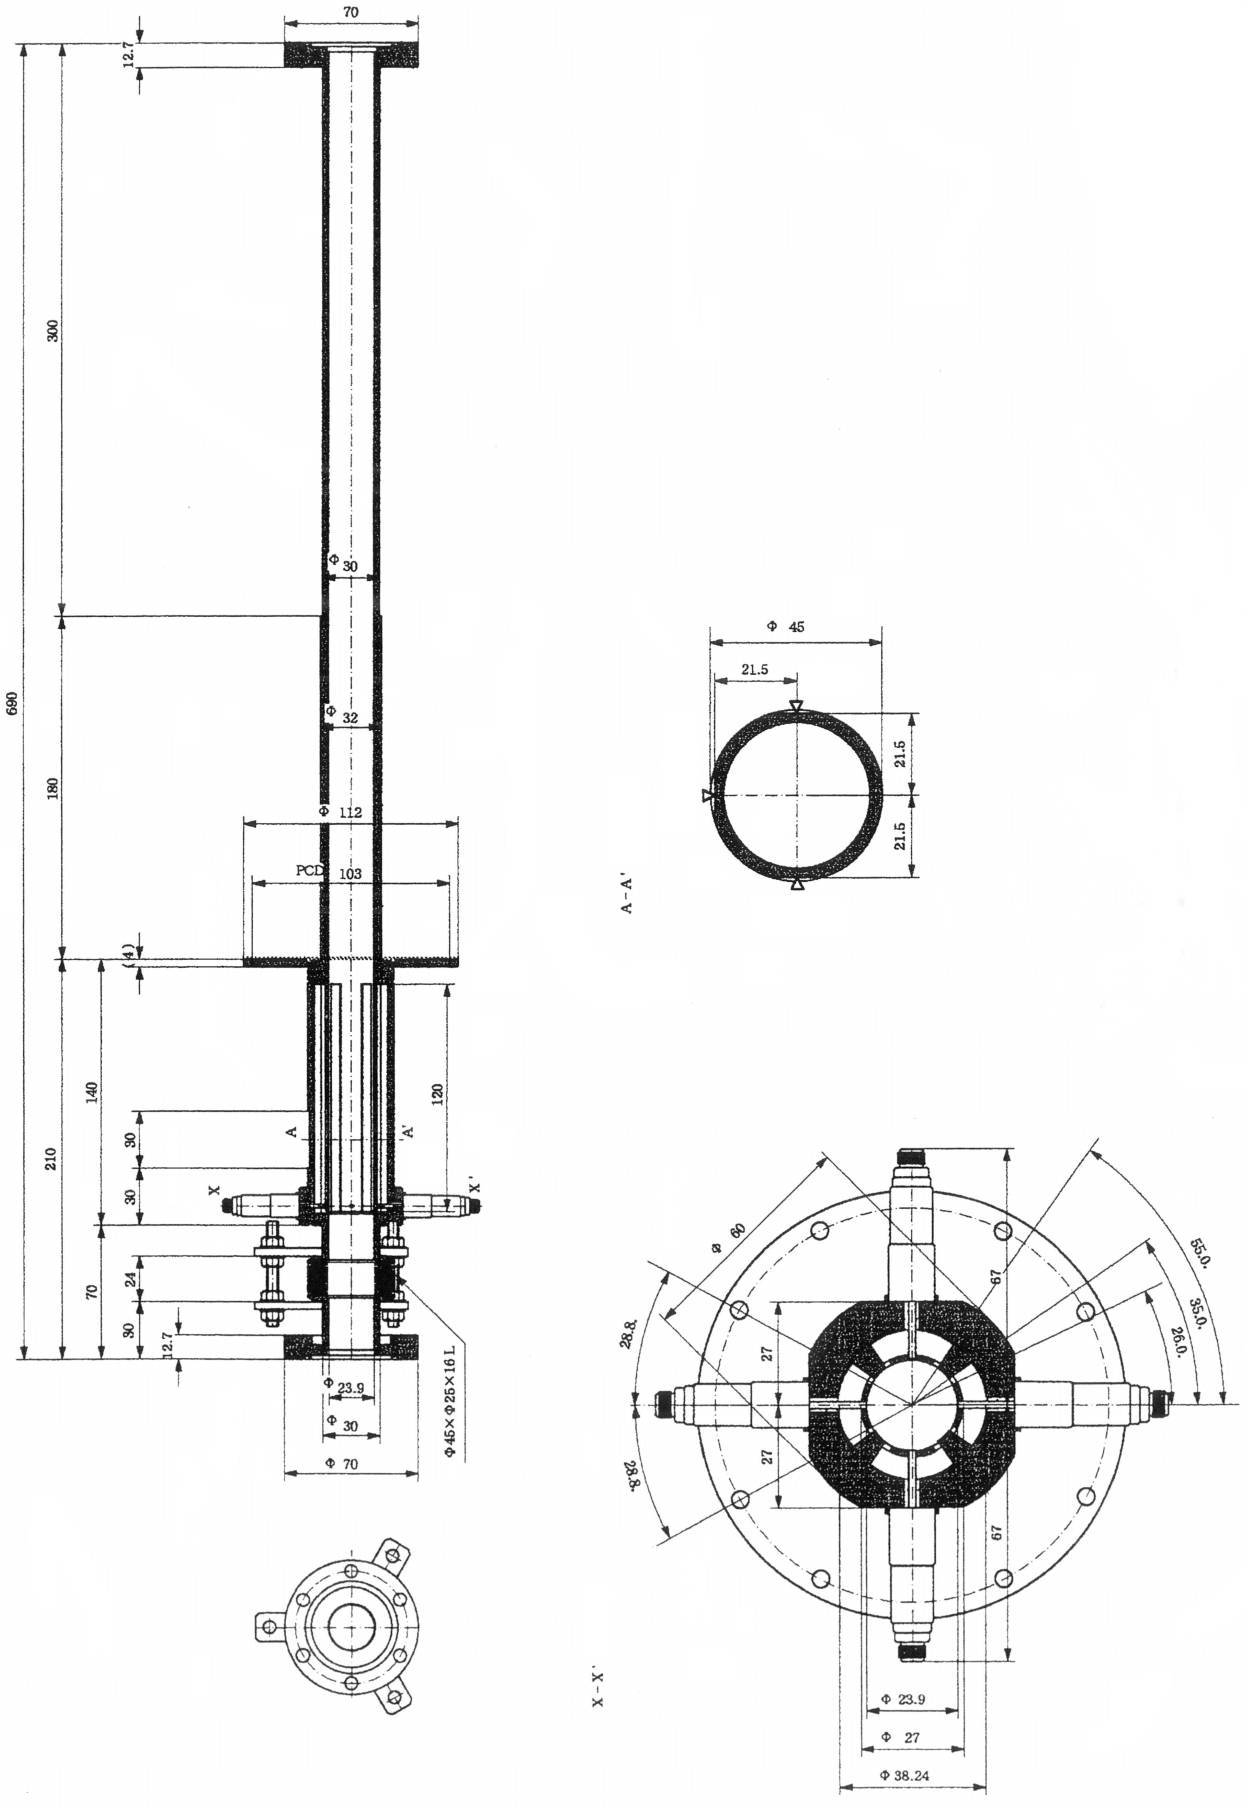

Figure 2.4: Technical drawing of the stripline BPMs used as part of the FONT system at the ATF.

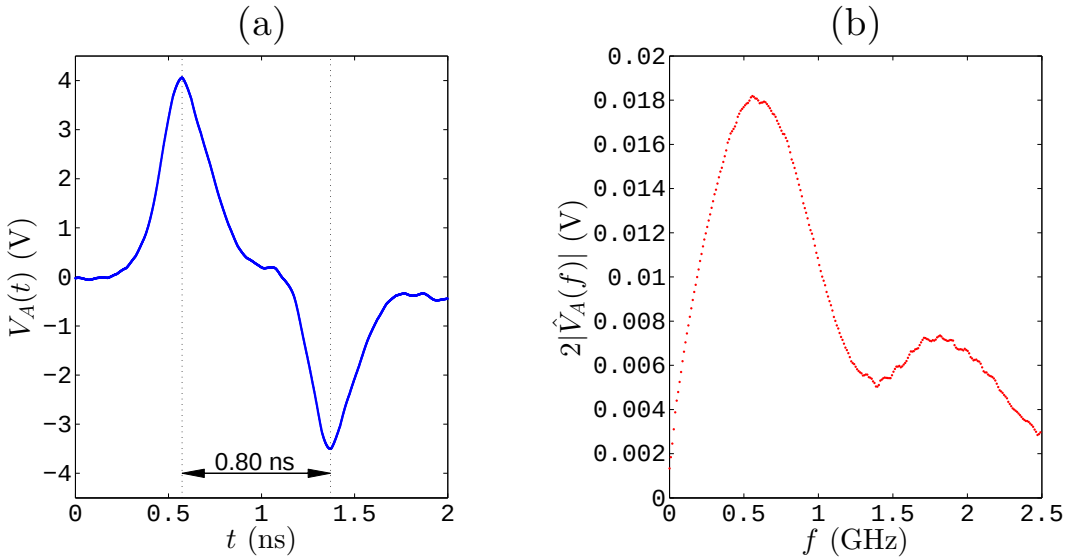

Figure 2.5: Stripline output signal. (a) the voltage at a stripline output measured using an oscilloscope (the interpolated sample interval of 1 ps gives the appearance of a thick continuous line); (b) the single-sided amplitude spectrum of the Fourier Transform of the stripline signal.

$$V_A = \rho \frac{w^2}{(R-y)^2} \left(\frac{dQ}{dt} \right) \tag{2.1}$$

$$V_B = \rho \frac{w^2}{(R+y)^2} \left(\frac{dQ}{dt} \right) \tag{2.2}$$

A similar pair of expressions exists for the left and right outputs where the vertical displacement y is replaced by the horizontal displacement x; however, throughout this thesis only the y position will be considered. The top and bottom outputs vary only as a function of y and (dQ/dt); the two may thus be combined so as to cancel out the effect of the charge and ultimately give a single signal proportional only to position. This is accomplished by taking the sum (Equation 2.3) and difference (Equation 2.4) of the two.

$$V_A + V_B = 2\rho w^2 \frac{R^2 + y^2}{(R^2 - y^2)^2} \left(\frac{dQ}{dt} \right) \tag{2.3}$$

$$V_A - V_B = 4\rho w^2 \frac{yR}{(R^2 - y^2)^2} \left(\frac{dQ}{dt} \right) \tag{2.4}$$

Taking the ratio of these two and making the simplification $R \gg y$ (acceptable since the radius of the FONT BPMs is 12 mm and the typical measured beam position is much less than 100 μm) yields a quantity that is linearly dependent on the position y (Equation 2.5).

$$\frac{V_A - V_B}{V_A + V_B} = \frac{2y}{R} \tag{2.5}$$

The FONT system uses analogue electronics to construct signals proportional to the sum and the difference of the stripline signals. These signals go on to be digitized by the ADCs of the FONT5 board. As the bunch position is required in order to perform feedback, it is calculated in real time by the FONT5 board by taking the ratio of a pair of samples, one each from the digitized sum and difference waveforms (the board is informed manually which samples correspond to the passage of the bunch).

2.1.2 Analogue BPM processing electronics

The FONT BPM processing electronics are constructed from analogue electronic components according to the design [25] depicted schematically in Figure 2.6. Each BPM processor module takes as inputs the raw BPM stripline signals V_A and V_B along with an RF signal phase-locked to the beam known as the local oscillator (LO) and delivers signals corresponding to the difference (Δ) and sum (Σ_I) of the two inputs along with a third signal, the sum quadrature (Σ_Q), whose description is postponed for now. These signals are recorded using the FONT5 board. A total of eight analogue processor modules were constructed [1] according to the design laid out in Figure 2.6 and have been verified operationally; one processor module is required per axis of each BPM to be monitored. The original design for the processing electronics included amplifiers on the processor outputs but due to performance issues with the selected components they were abandonded. Their functionality is now implemented by the final stage of the processor consisting of amplifiers, attenuators and filters on each of the three outputs; these components are located on an auxiliary board as can be seen in Figure 2.7. The term processor module will be used throughout this thesis to refer to the unit formed by the combination of the original design of the processor and the auxiliary board.

The operation of the processor itself proceeds as follows: both input stripline signals (V_A and V_B) are immediately passed through Mini-Circuits LFCN-1000 [26] 1 GHz low-pass filters (LPFs) and then split into two. One half is input to a Mini-Circuits SCPJ-2-9 [27] 180° hybrid, a passive device that delivers the difference $\Delta_0 = V_A - V_B$ that will ultimately be output as the Δ signal. The other half enters a resistive coupler, a simple circuit consisting only of resistors which produces the sum $\Sigma_0 = V_A + V_B$.

Next, both Σ_0 and Δ_0 signals are filtered using band-pass filters, the central frequency of which is intended to match the maximum of the power spectrum of the FONT BPMs. This frequency is determined by the geometry of the BPM using the relation $f = c/4L$ where L is the length of the striplines (their extent in the z-axis). For the stripline BPMs at ATF, $L = 120$ mm and so $f = 625$ MHz (Fig 2.5). The measured central frequency of the actual band-pass filter implemented in the processor is 685 MHz. The Σ_0 signal is then split into two to form the Σ_I and Σ_Q signals. When the system is correctly configured, the

[1] Initial production work was performed for ten modules; the ones designated 6 and 9 were found to be defective and not completed as a result.

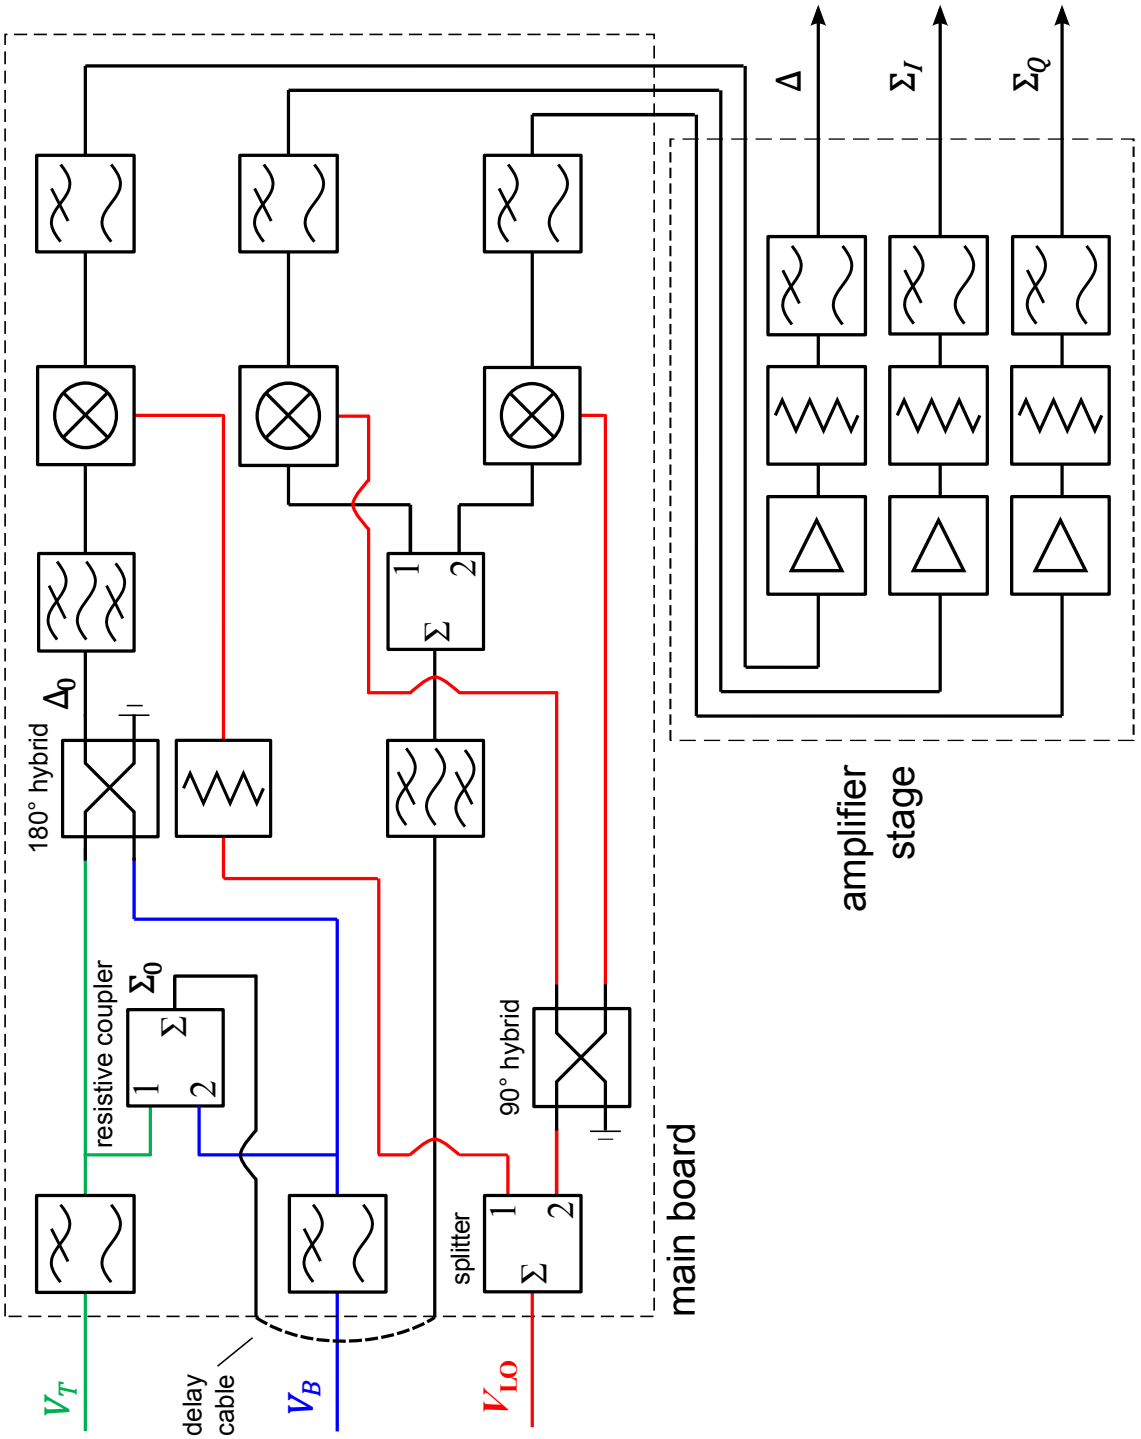

Figure 2.6: Schematic diagram of the FONT analogue processor modules.

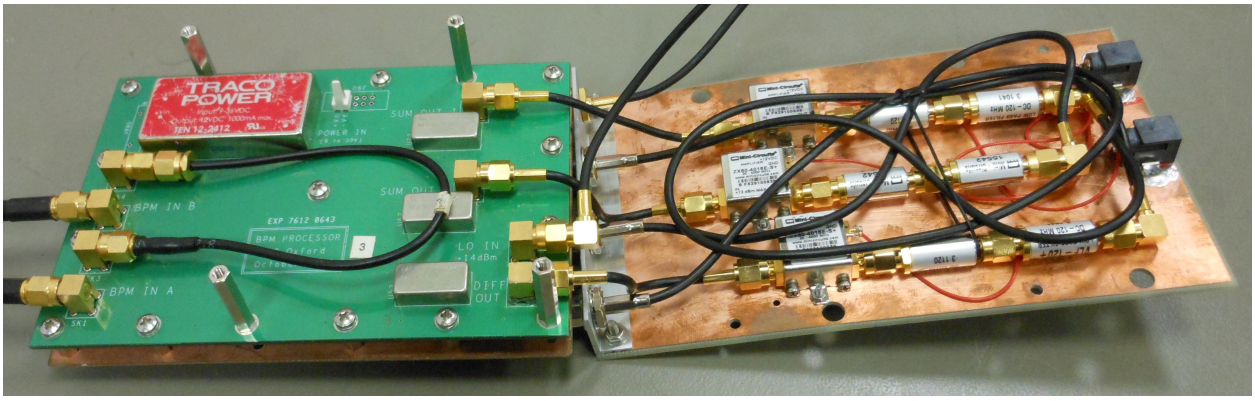

Figure 2.7: Photograph of processor module 3. The unit is partially disassembled for viewing purposes; the main board is on the left and the auxiliary board is on the right.

Σ_I signal corresponds directly to the $V_A + V_B$ term of Equation 2.5; the Σ_Q output is used as an indicator of when this configuration is achieved.

Each of the three signals (Δ, Σ_I and Σ_Q) is then input to a Mini-Circuits SYM-2 [28] radio-frequency (RF) mixer. These devices generate an output signal from a pair of input signals by shifting the frequency spectrum. In this case, each input is mixed with a copy of the local oscillator (LO) signal to generate a signal whose frequency spectrum consists of two terms - an up-mixed term whose frequency is the sum of the two input frequencies and a down-mixed term whose frequency is the difference of the two. The local oscillator signal is an additional input to the processor module; it is immediately split into two using a Mini-Circuits JPS-2-4 splitter [29], with one half feeding the mixer for the Δ signal and the other half entering a Mini-Circuits HPQ-08 [30] 90° hybrid coupler, the output of which is a pair of signals 90° out of phase which are used to feed the Σ_I and Σ_Q mixers. The local oscillator signal must be phase-locked with respect to the beam to ensure that variation in the processor output signals is solely due to variation in the stripline signals. A 714 MHz signal from the ATF master timing system is used as in principle it is stable in time with respect to the beam and is also close to the peak of the stripline signal frequency spectrum (Figure 2.5(b)). As a result, the output of the mixer varies relatively slowly; the Σ_I signal is well-approximated by a Gaussian with a full width at half maximum of 4.5 ns and can be easily sampled at the frequency of the ADCs on the FONT5 board (357 MHz).

The three mixer output signals are then passed through custom ∼100 MHz low-pass filters to remove the up-mixed term. At this stage, the signals pass from the main body of the processor to the auxiliary board that consists of Mini-Circuits ZX60-4016E-S [31] low noise amplifiers in order to maximize the signal-to-noise ratio, Mini-Circuits VAT-3+ [32] 3 dB attenuators in order to match the final amplitude of the signals to the range of the ADCs on the FONT5 board and a final pass through Mini-Circuits VLF-120+ [33] 120 MHz low-pass filters to remove any high frequency noise introduced by the amplifiers. The output signals of the processor modules are then sent to the FONT5 board, the operation of which is the subject of the next section.

2.1.3 Latency of the processor modules

The latency of the processor module is the time interval between arrival of the V_A and V_B signals at the main board inputs and the appearance of the Δ, Σ_I and Σ_Q signals at the auxiliary board outputs. One of the principal design goals of each component of the FONT system was that the latency should be as low as possible as the total latency of the system must be less than the time interval between consecutive bunches (a bunch spacing of 187.6 ns applies for the multi-bunch data within this thesis) in order for correction on the bunch to bunch timescale to be possible. As the latency of the analogue processors was increased by the addition of the auxiliary board, a procedure for measuring the latency of the processor module is described here.

The configurations used are depicted schematically in Figure 2.8 along with the oscilloscope traces from which the results were obtained. Schematic Figure 2.8(a) shows how an Avtech AVM-2 [34] ultra high speed pulse generator, set to provide BPM-like signals, was used as the source for the V_A and V_B inputs of the processor module as well as the trigger input for the oscilloscope. Plot Figure 2.8(b) shows the trigger (blue) and the Σ_I signal (red) monitored on the scope (which was set to infinite persistence in this case to compensate for the fact that the signal generator providing the LO is not locked with respect to the pulse generator). From this trace the interval t_1 between the trigger and the centre of the output pulse is measured. Schematic Figure 2.8(c) shows the configuration adjusted to bypass the processor by connecting the cable that was formerly connected to the V_A input to the one which was previously used for the Σ_I output while making no other changes; t_2, the interval between the trigger and the erstwhile processor input signal is then measured from plot Figure 2.8(d). The latency of the processor module is given by $t_{\mathrm{lat}} = t_2 - t_1$ and is found to be 15.2 ± 0.1 ns. This represents an increase of $\sim$5 ns compared to previous measurements of the latency of the main processor board alone [35] and is in line with the additional delay expected from the components in the auxiliary board.

2.2 FONT5 digital board

The FONT5 digital board was designed and constructed at the University of Oxford for use as the control unit of the FONT5 position and angle correction system [36]. The description of the FONT5 board is divided into two subsections. The first will briefly deal with the physical hardware that makes up the digital board including the core component, the FPGA. The second will outline the specific configuration of the FPGA that is used during operation of the FONT system at ATF. This configuration is held in non-volatile memory and loaded into the FPGA when the board is powered on; it is thus referred to as the FONT5 firmware.

2.2.1 Hardware

A functional schematic of the FONT5 board illustrating the interaction of the various hardware components is given in Figure 2.9. A description of each component follows below:

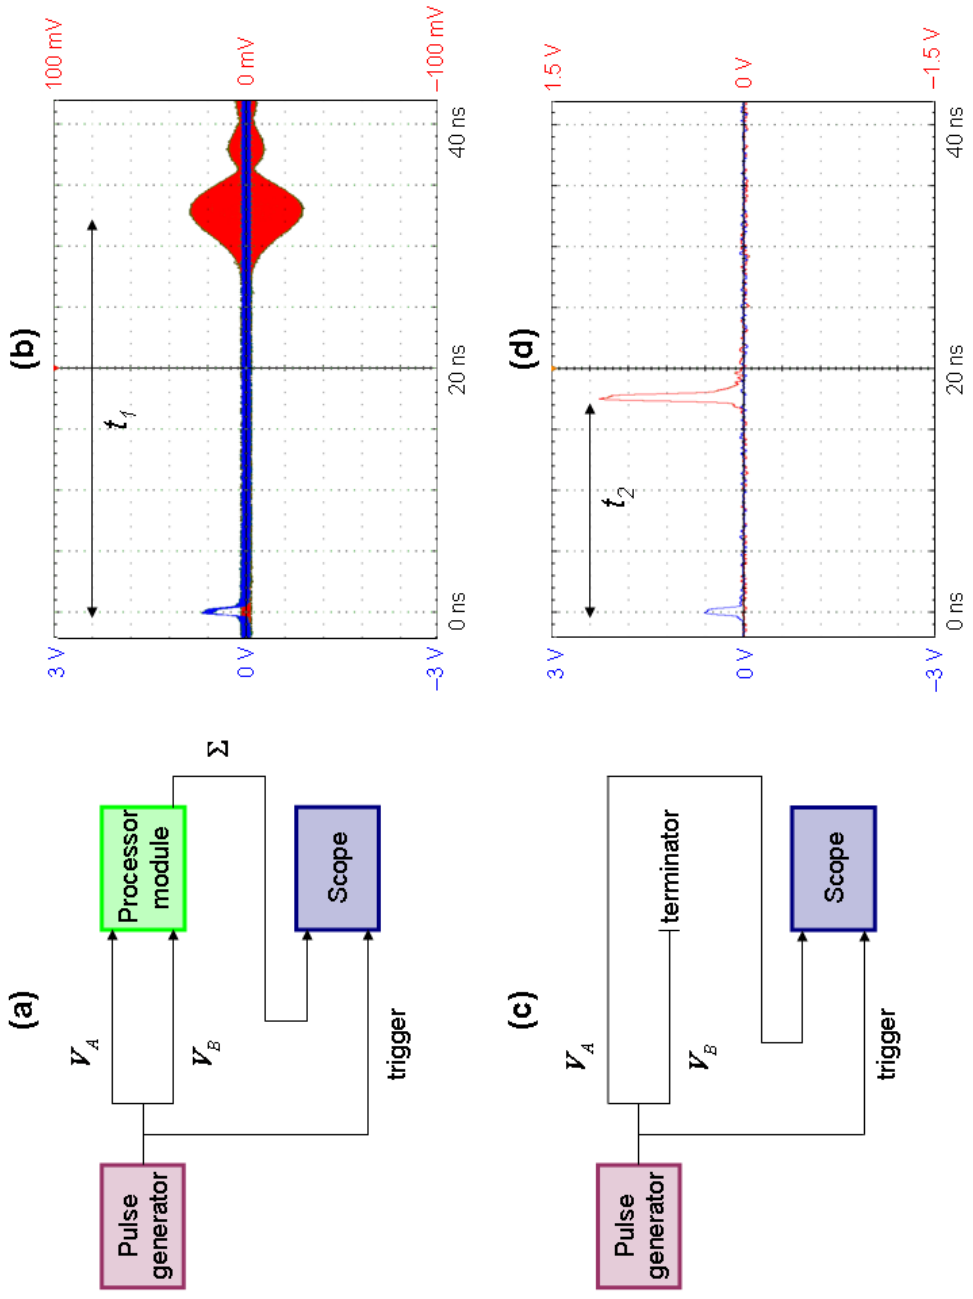

Figure 2.8: Schematics of the configurations used for the latency measurement of the processor module and the oscilloscope traces obtained in each case. (a) schematic when including the processor module; (b) oscilloscope trace when including the processor module; (c) schematic when bypassing the processor module; (d) oscilloscope trace when bypassing the processor module.

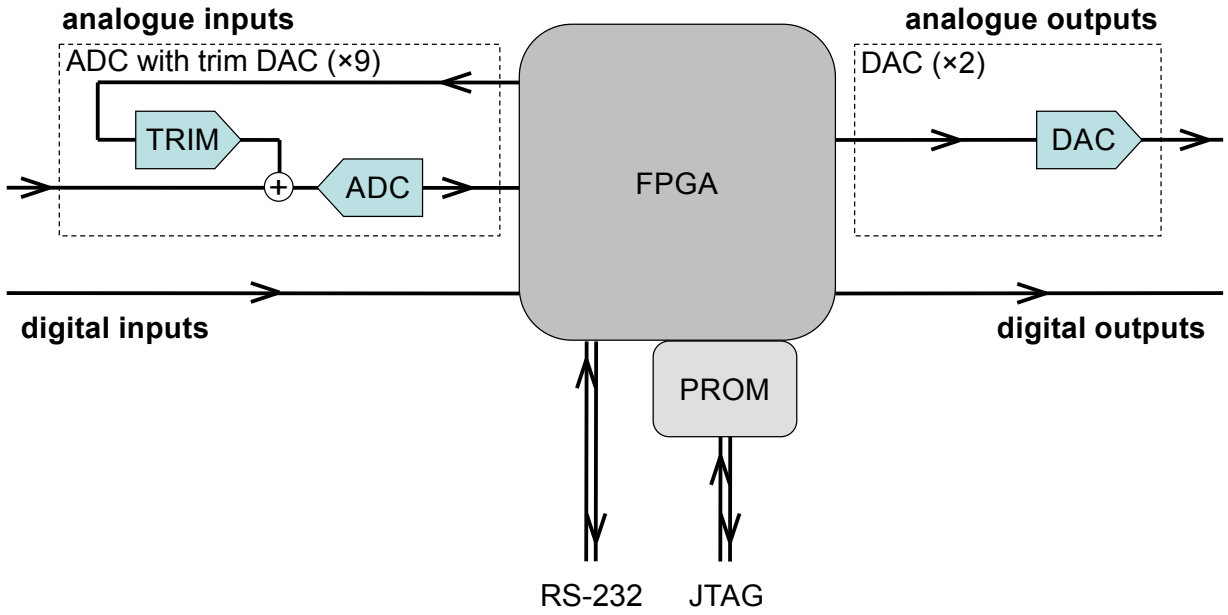

Figure 2.9: Schematic of the FONT5 digital board.

- **FPGA**
  The FONT5 digital board is built around a Xilinx Virtex-5 XC5VLX50T [37] field-programmable gate array (FPGA). An FPGA is an integrated circuit designed to be reconfigurable after manufacture and consists of programmable logic components that may be connected in many different ways. The configuration is specified using a design written in hardware description language. The Virtex-5 can operate up to a maximum internal clock speed of 550 MHz [38].

- **PROM**
  The FPGA is accompanied by a Xilinx XCF32P [39] programmable read-only memory chip (PROM). This device is used for non-volatile storage of a single firmware design which is then loaded automatically when the board is powered on.

- **Analogue-to-digital converters (ADCs)**
  The board incorporates nine Texas Instruments ADS5474 [40] 14-bit ADCs with a maximum clock speed of 400 MHz. They are arranged in three groups of three with each group sharing a common clock. The voltage range that can be measured by the ADCs is ± 1.1 V.

- **Digital-to-analogue converters (DACs)**
  The board includes four Analog Devices AD9744 [41] 14-bit DACs with a maximum clock speed of 210 MHz and an output range of ± 0.5 V. Only two of the four DACs are connected for use at any one time.

- **Trim DACs**
  Also included are three LTC2624 [42] quad 16-bit DACs with a maximum clock speed of 50 MHz (a quad DAC is essentially four DACs packaged together). Three of the

twelve DACs are unused and the remaining nine are each connected to a different one of the nine ADC channels. By setting the value of each trim DAC appropriately, the average value of the samples from each ADC channel when no signal is present can be made equal to zero.

- **RS-232 serial interface**
  An RS-232 port exists to provide a means of communication with the deployed board for monitoring and control purposes. An RS-232 cable connects the FONT5 board to a serial device server and hence to a computer running the DAQ software.

- **JTAG interface**
  Joint Test Action Group (JTAG) is the common name for the IEEE 1149.1 Standard Test Access Port and Boundary-Scan Architecture [43]. The FONT5 board includes a JTAG port to allow a computer to communicate with the FPGA/PROM; this is the means by which the FONT5 firmware was loaded in to the PROM.

- **Digital inputs and outputs**
  The digital inputs and outputs are used to carry binary signals, specifically clock and trigger signals.

A photograph of a FONT5 board is provided in Figure 2.10. On the left of the photograph is the protective case. In the centre of the photograph sits the secondary printed circuit board (PCB) which houses the power supply. This PCB is placed with the inner surface face up; when the FONT5 board is assembled, it is folded over and held parallel to the primary PCB, which is on the right of the photograph. At the extreme right of the photograph the back of the front panel of the case is visible. The primary PCB is partly obscured by the jumper cables that carry the signals from the external connectors on the front panel.

A more detailed photograph of the primary PCB indicating the locations of some key components is given in Figure 2.11. The Xilinx Virtex-5 FPGA is the silver square located roughly in the centre of the PCB surrounded on three sides by the nine ADC chips arranged in three groups of three. Along the remaining side lie the four DAC chips. As only two of these are required for operation of the feedback system, the front panel has connectors for two of the DACs and the other two are spares. The PROM resides on the bottom left of the PCB and connects via the twisted blue and green wires to the JTAG input on the front panel. A photograph of the front panel of the assembled FONT5 board is given in Figure 2.12. The labels indicate the various input and output connectors and their function during operation of the FONT feedback system.

2.2.2 Firmware

The standard FONT5 feedback firmware is a highly modular design written in the hardware design language Verilog [44] and compiled using the ISE Design Suite [45], a software package created by Xilinx. The FONT5 firmware was written by B. Constance but includes modules developed by G. Christian for earlier versions of the FONT digital hardware. A more detailed description of the features of the FONT5 firmware can be found in [46]. A simplified view

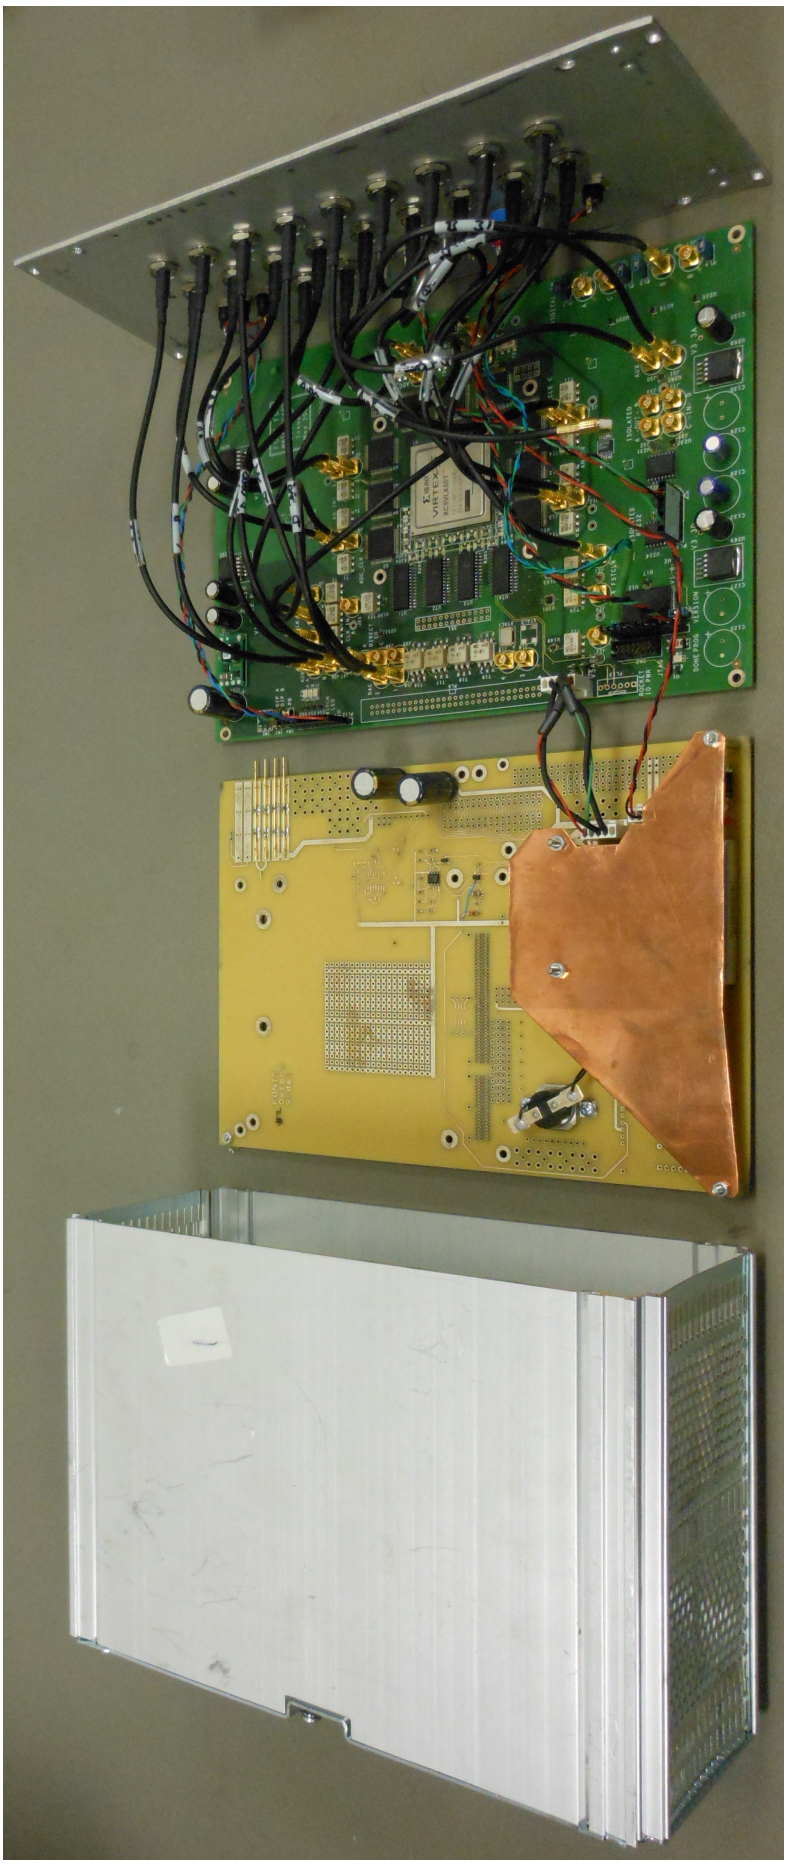

Figure 2.10: Photograph of a FONT5 digital board disassembled for viewing purposes.

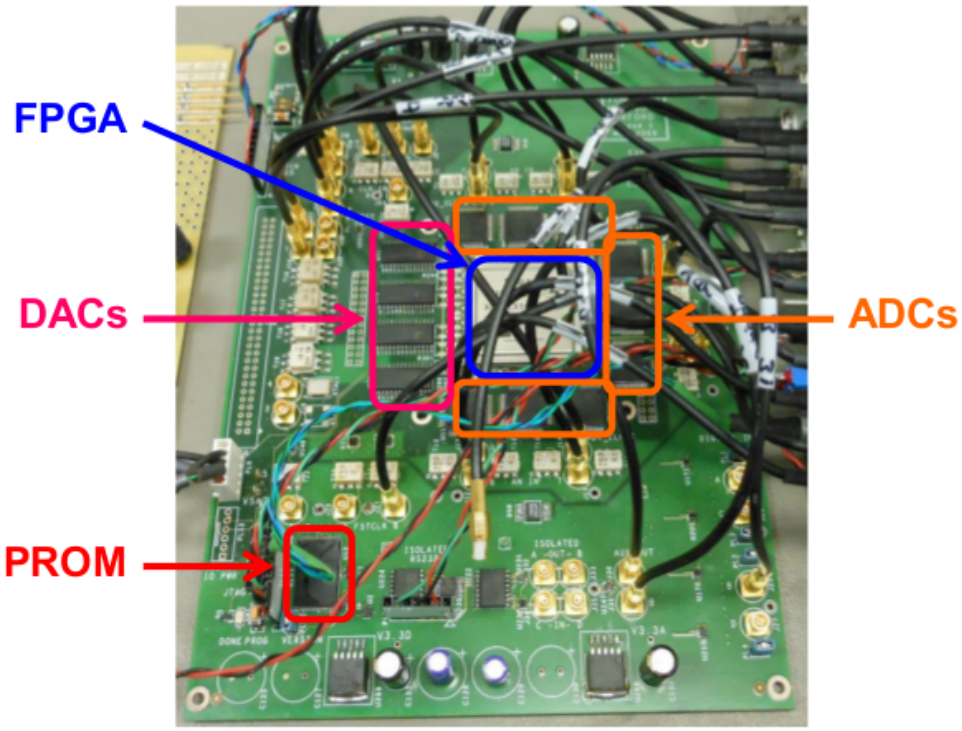

Figure 2.11: Photograph of the primary PCB of the FONT5 digital board.

of the interaction between the modules is illustrated in Figure 2.13; the modules and signals have been colour coded to roughly indicate their function: control (red), timing (blue), data (black) and communication (green). The primary inputs of the top-level module of the firmware design are listed on the left hand side of the diagram and the primary outputs on the right hand side. Each named square represents either a single module or several instances of the same module contained within the top-level module; only the most important constituent modules are shown in the diagram. These may be described as follows:

- **ADC block module** (`adc block`)
  The ADC block module contains the logic relating to control of the ADCs. There are three ADC block modules, each of which deals with one of the three sets of three ADCs: channels 1-3, 4-6 and 7-9.

- **DAQ RAM module** (`daq ram`)
  The DAQ RAM module contains the memory used to store both the sample values measured by the ADCs and the outputs of the feedback calculation that are applied to the DACs. There are a total of 11 DAQ RAM modules; one for each of the 9 ADCs and 2 DACs.

- **DAQ sequencer module** (`daq seq`)
  The DAQ sequencer module formats the data stored within the firmware for transmission over the serial connection.

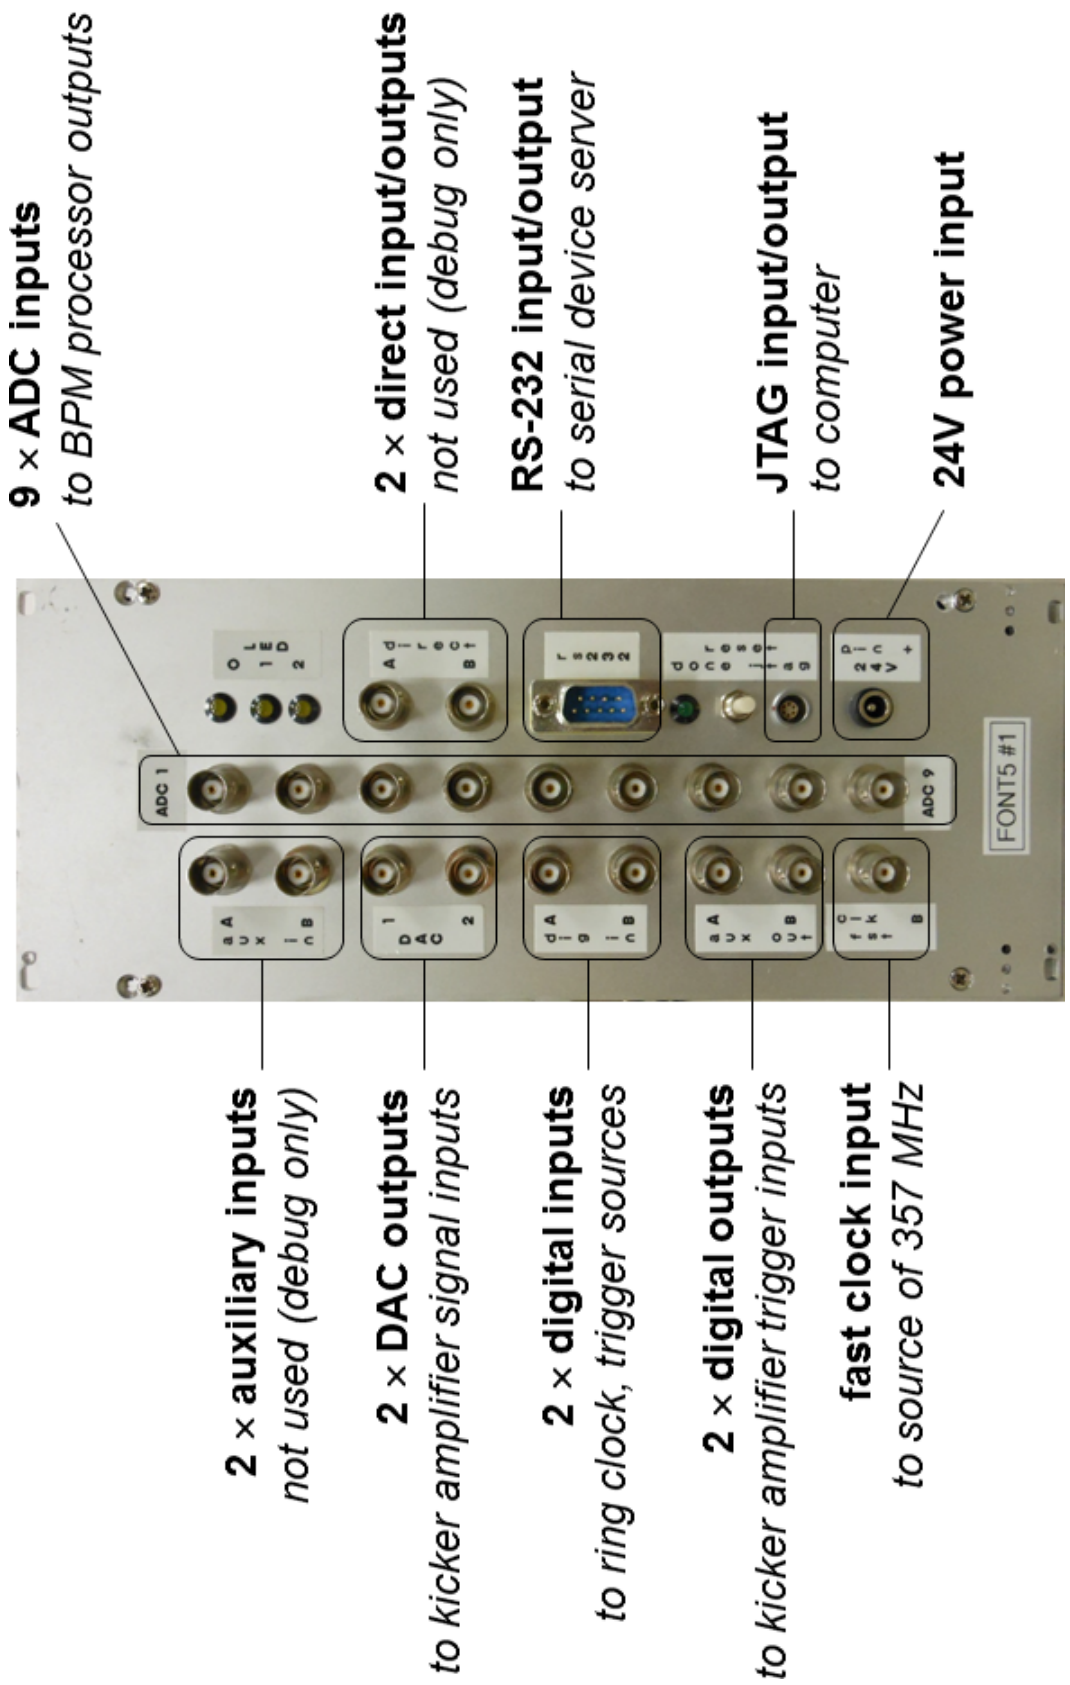

Figure 2.12: Photograph of the front panel of the FONT5 digital board.

- **Timing and synchronization module** (`timing sync`)
  The timing and synchronization module contains the logic responsible for the generation of important control signals. This includes the trigger signals for the kicker amplifiers and the store strobe which defines when the ADCs are active.

- **Feedback calculation module** (`data proc`)
  The feedback calculation module contains the feedback algorithm which generates the values to be applied to the DACs through processing of the input ADC data.

- **UART module** (`uart`)
  Universal Asynchronous Receiver/Transmitter (UART) normally refers to a piece of computer hardware that translates data between parallel and serial forms; on the FONT5 board, this function is performed on the FPGA by the UART module. The UART module contains the logic handling the byte-by-byte receipt and transmission of data over the serial connection.

- **UART decoder module** (`uart decode`)
  The UART decoder module is where the list of commands recognized by the FONT5 board and the appropriate action to take in each case are defined.

- **Control register module** (`ctrl regs`) The control register module is where the dynamic control parameters are assigned their values. The parameters are dynamic in the sense that they are not hard coded in the firmware; an example of a control register is the trigger in delay that sets the start of the board sampling window relative to the external board trigger.

In order for the logical design to be mapped to the physical hardware in such a way that all signals are able to assume their correct values within a single clock cycle, it is desirable for as much of the logic as possible to be operating on a time domain with as slow a clock speed as is practical. The exception to this is the logic related to the feedback calculation which must run as fast as possible in order to minimize the latency of the system. As there is an additional requirement for temporal stability relative to the beam and an upper limit on the clock speed deriving from the maximum clock rate of the ADCs [2], the time-critical logic is clocked using a 357 MHz signal external to the board referred to as the fast clock. This signal is a frequency divided version of the 714 MHz LO used to downmix the BPM signals. The fast clock distribution includes a phase-locked loop configured as a timing jitter filter. As well as clocking the feedback logic, the fast clock is replicated three times for use as the ADC clocks. Each of these three copies is passed through a variable delay element and then itself split three ways resulting in three ADC clocks with an adjustable phase that each drive three ADCs. The remaining logic is clocked using a 40 MHz oscillator built in to the FPGA (referred to as the slow clock).

[2]The maximum clock speeds of the DACs (210 MHz) and the trim DACs (50 MHz) are not limiting factors as their required refresh rate is very much less than that of the ADCs. The DACs need only operate at a frequency corresponding to the bunch spacing ($\sim$ 5 MHz) and the trim DACs are clocked only once per board power cycle.

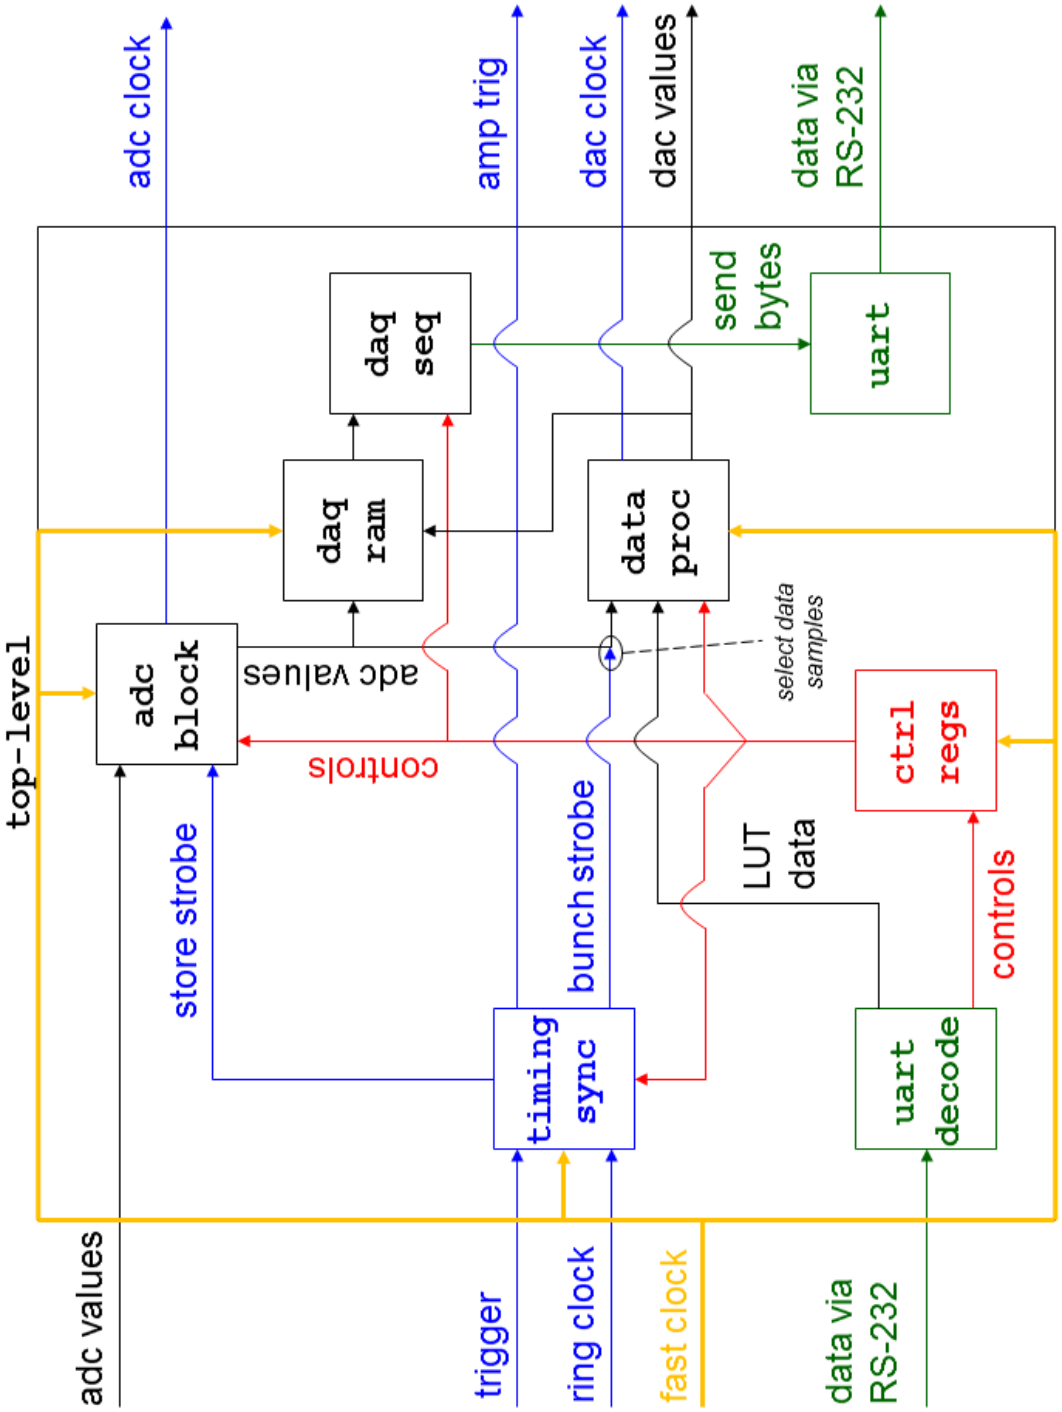

Figure 2.13: Schematic diagram of the FONT feedback firmware. Squares represent Verilog modules and arrows represent signals.

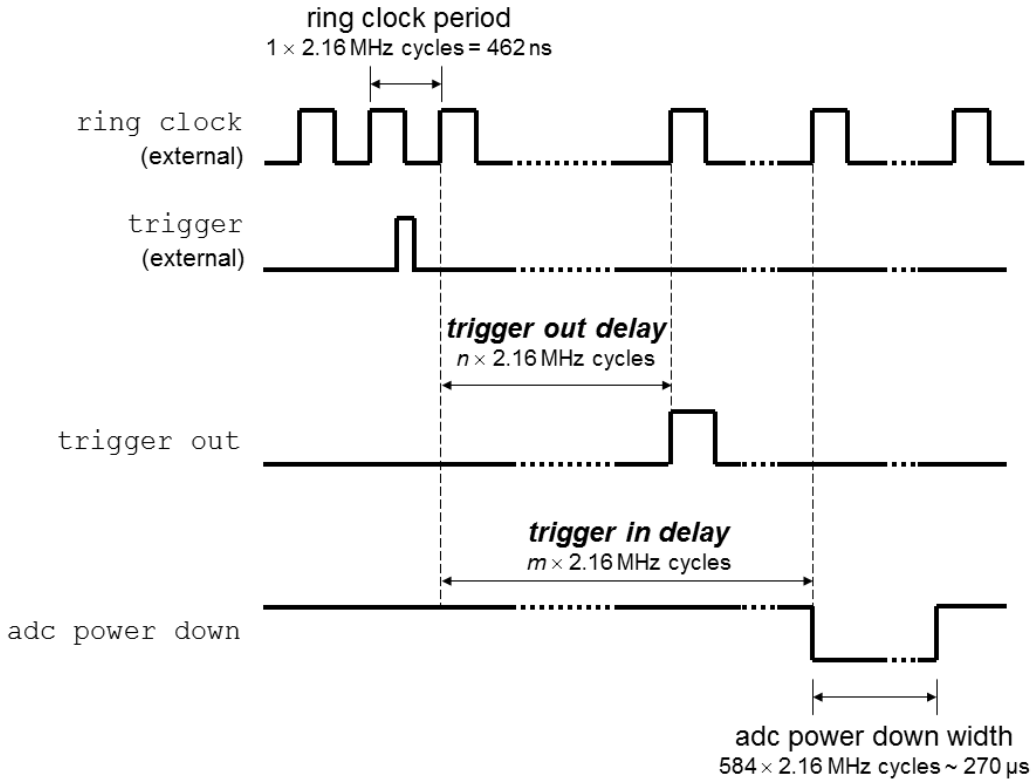

Figure 2.14: Schematic of the FONT5 timing signals on the time scale of the 2.16 MHz clock. Each thick black line represents the binary state as a function of time of the signal specified on the left; the dotted lines indicate that a number of clock cycles have been omitted from the diagram.

The bulk of the firmware design may be considered as an algorithm that runs when an external signal referred to as the trigger is received. Once the trigger signal is detected, the timing and synchronization module starts counting cycles of an external clock with a frequency of 2.16 MHz. As this is the frequency of revolution of the ATF damping ring the 2.16 MHz clock is known as the ring clock. The ring clock is sourced from the master ATF timing system; its primary purpose is to make the system insensitive to timing jitter of the external trigger signal. This timing jitter is known to be much less than one period of the ring clock and so by defining delays relative to the first ring clock edge after the trigger is detected, rather than relative to the trigger edge itself, the arrival time of the bunches relative to the start of the FONT sampling window remains constant from trigger to trigger (the bunch arrival time is locked relative to the ring clock).

As the ring clock cycle counter reaches certain values the module sends signals to other modules that deal with control of the hardware. The most important values are the trigger in delay, which changes the state of the ADC power down signal to zero and hence switches the ADCs on, and the trigger out delay, which results in the assertion of the trigger signal that is sent to prepare the kicker amplifiers for operation. The behaviour of the signals mentioned in this paragraph is illustrated in Figure 2.14.

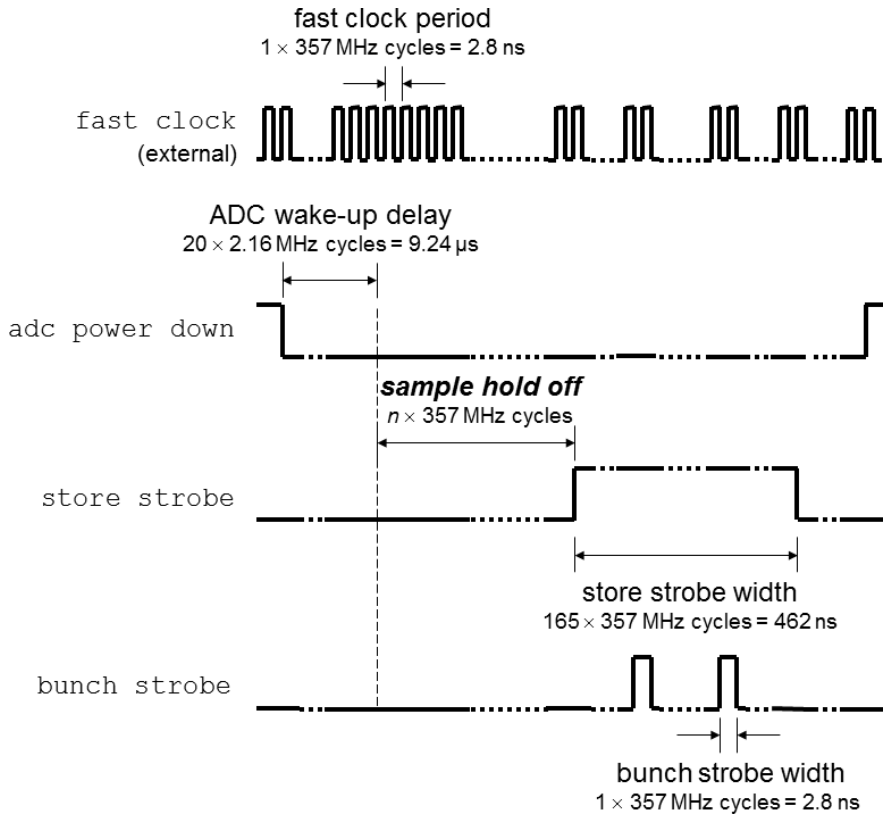

Figure 2.15: Schematic of the FONT5 timing signals on the time scale of the fast clock. Each thick black line represents the binary state as a function of time of the signal specified on the left; the dotted lines indicate that a number of clock cycles have been omitted from the diagram.

The ADCs are powered for a period of $\sim$ 270 µs but the sampled data is only stored during the single ring clock cycle when a signal called the store strobe is high. The delay from the ADC power down signal going low to the store strobe going high is equal to twenty ring clock cycles (to account for the ADC wake-up time of 5 μs [40]) plus the number of fast clock cycles specified by the sample hold off parameter. The sample hold off parameter thus allows the timing of the sampling window for all ADC channels to be shifted together by one cycle of the 357 MHz fast clock (2.8 ns). The scan delay parameters control the variable delay elements on the ADC clocks and thus allow the timing of the sample window for each set of three ADCs to be shifted with an increment of 70 ps.

Once asserted, the store strobe remains high for 165 cycles of the fast clock, the full duration of a ring clock cycle. In addition to the store strobe, an additional signal is required for each BPM used to perform the feedback calculation in order to specify which of the 164 samples stored corresponds to the bunch. These signals are the bunch strobes and they effectively trigger operation of the feedback modules. The behaviour of the signals explicitly mentioned in this paragraph is illustrated in Figure 2.15.

The kicks to be provided in the two kickers are calculated in two instances of the feedback module and then stored in memory as well as being sent to the DAC. A more detailed description of the operation of the feedback modules will be given in Chapter 4. After data

collecting has ceased the UART module supervises the transmission of the stored ADC and DAC data over the RS-232 link.

Outside of the triggered logic, certain modules operate continuously. This includes the two control register modules (one each for signals operating on the fast and slow clock domains) which store the values of all modifiable parameters, such as the trigger in delay, and make them continuously available to other modules that require them. The UART decoder module is also active all of the time; this module interprets signals received by the board via the serial link and performs particular actions in response, typically the updating of the value of a control register.

2.3 Kickers and kicker amplifiers

2.3.1 Kickers

FONT uses a pair of stripline kickers provided by the SLAC National Accelerator Laboratory in order to modify the direction of the beam at two locations. A technical drawing of the kicker is given in Figure 2.16. Each kicker consists of a pair of electrodes on opposite sides of a vacuum chamber. The electrodes are shorted at one end and the other end has multiple connections to support external coaxial cables. A deflection of the beam along the axis defined by the electrodes can be provided by the application of a pair of pulses to the electrodes. By convention a positive voltage pulse is applied to the top of the kicker and a negative voltage pulse to the bottom.

A photograph of the kicker K1 installed in the extraction line of the ATF is given in Figure 2.17. In the centre of the image some of the connectors for the coaxial cables are visible; for added flexibility as to the location of the kicker amplifiers, there are a total of sixteen such connectors (four for each of the connector blocks that form a cross at the downstream end of the kicker) but only eight are required to carry the kicker drive signals from a FONT kicker amplifier [3].

2.3.2 Kicker amplifiers

In order to achieve the required voltage across the kicker, amplification of the FONT5 board DAC output signal is necessary. This is performed in two stages. First a Mini-Circuits ZPUL-21 [47] pulse amplifier provides a gain of 21 dB and then the resulting signal is input to the kicker amplifier (the amplifier whose outputs directly drive the kickers). To achieve the necessary level of drive signal on the timescale needed for bunch-to-bunch feedback, the kicker amplifier has to be capable of generating a high current output within a timescale of the order of tens of nanoseconds. The design and production of these units was performed by TMD Technologies Ltd [48] to a specification provided by Colin Perry and the amplifiers

[3] A single cable transports only a quarter of the total drive voltage in both positive and negative cases in order to reduce the risk of cable failure.

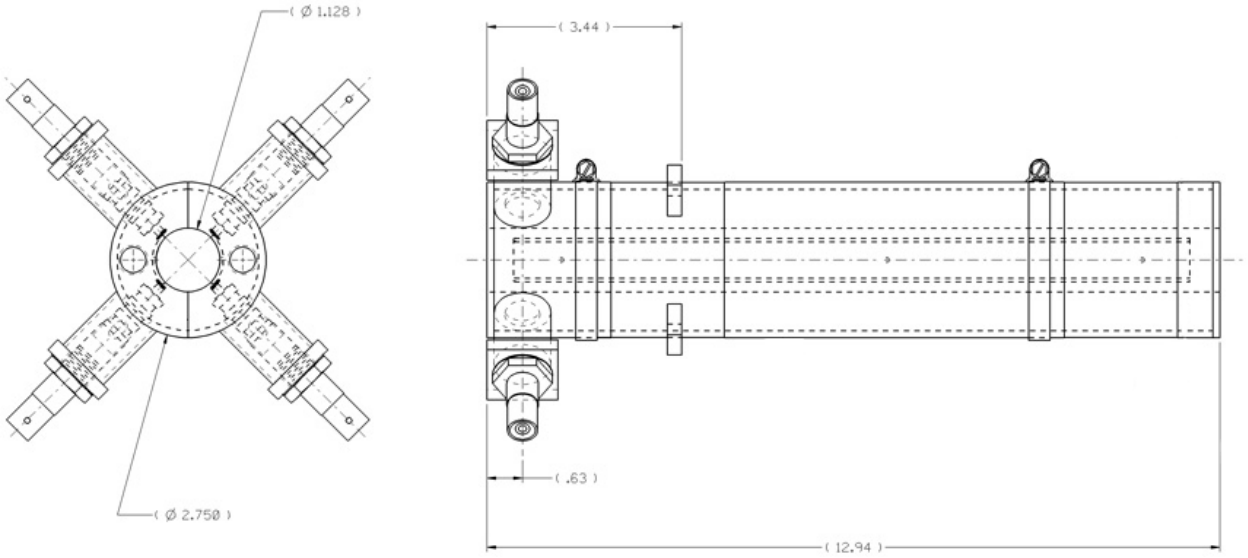

Figure 2.16: Technical drawing of the kickers used in the FONT feedback system.

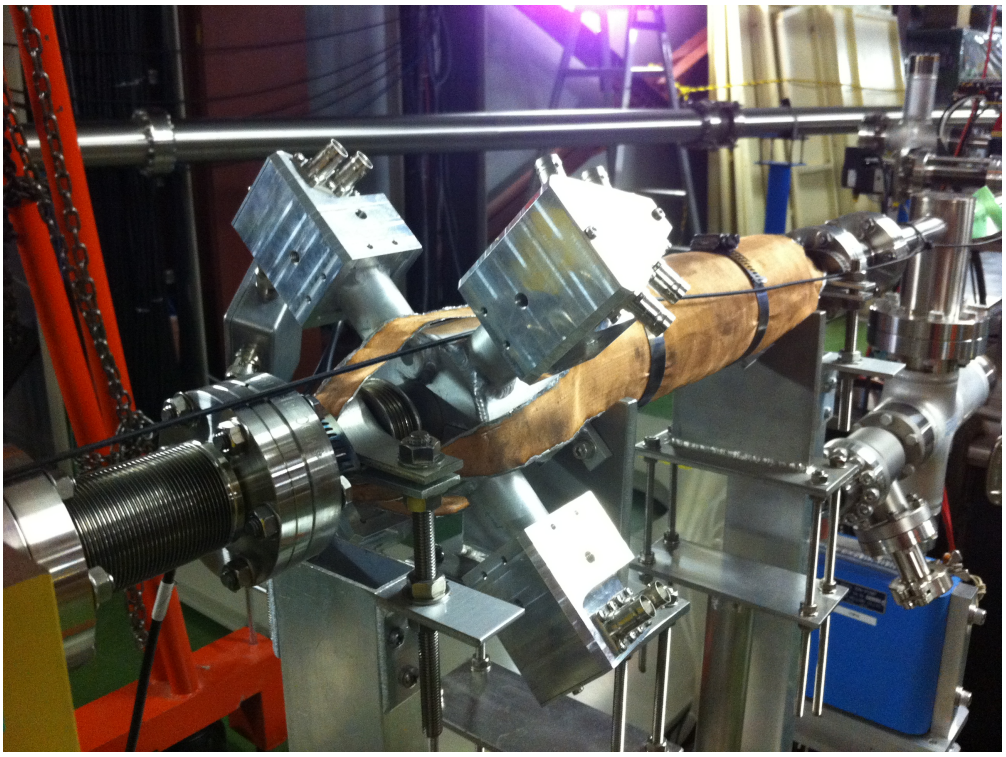

Figure 2.17: Photograph of FONT kicker K1 installed in the extraction line of the ATF.

are capable of providing ±30 A of drive current with a rise time of ∼35 ns. A total of three were produced; Figure 2.18 contains a pair of photographs of one of the units showing it from both the back and the front. The signal and trigger inputs sit on the back of the unit, along with the 24 V dual lead power supply input. The positive and negative signal outputs (four of each) along with an output for current monitoring purposes are located on the front panel.

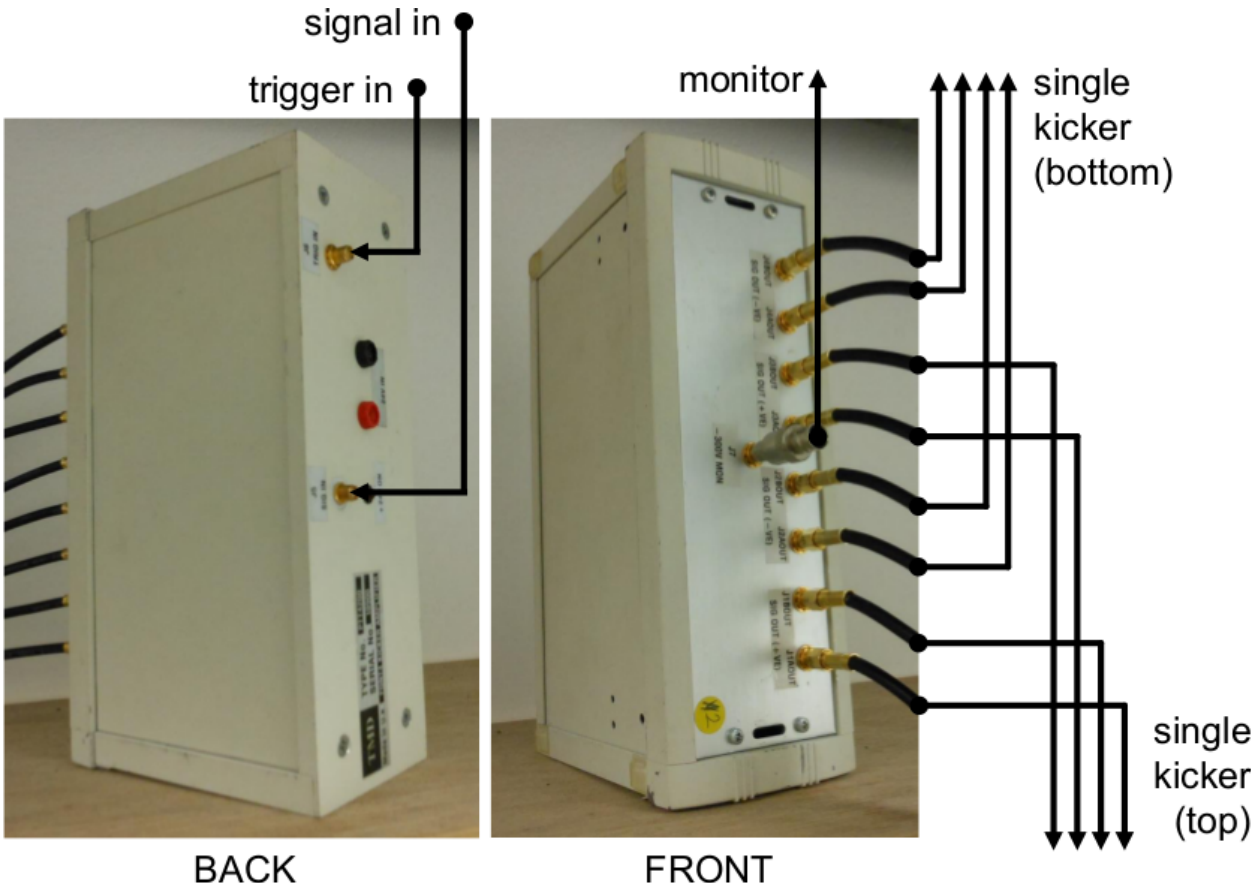

Figure 2.18: Photographs of a FONT kicker amplifier with external signals indicated.

2.4 Operational concerns for the FONT5 system

Having described the individual hardware components that together comprise the FONT feedback system, this section will describe some of the operational issues associated with the complete system. In particular, the following procedures must be performed prior to any period of data taking:

2.4.1 Eliminating the static ADC voltage offsets

Each ADC is known to have a static voltage offset; that is, the ADCs report a non-zero value even when monitoring a known zero voltage. The trim DACs are used to remove these

offsets by applying a constant corrective voltage to the input of each ADC. The required trim DAC value for each channel may be determined by fitting a straight line to the plot of ADC sample value as a function of trim DAC value; the intercept of this line with the x-axis gives the value required to zero the ADC. The required value of the trim DAC settings for each board are given in Table 2.1; as a trim DAC count is roughly equivalent to 5 μV, the static ADC voltage offsets fall approximately in the range of 7 mV to 10 mV.

	Trim DAC setting (counts)	
ADC no.	Board 1	Board 2
1	1739	1909
2	1896	1758
3	1582	1841
4	1592	1716
5	1385	1350
6	1498	1335
7	1578	1722
8	1780	1609
9	1556	1691

Table 2.1: Value in trim DAC counts required to zero the voltage of all ADCs across both FONT5 boards.

2.4.2 Optimization of the LO phase

The LO distribution supplies each processor module with the 714 MHz signal required for the RF mixers. The layout is illustrated schematically in Figure 2.19. The LO source first passes through a Mini-Circuits MTS-18B [49] transfer switch configured to provide a universal 90° phase shift when active. This master LO signal is then amplified and connected to one channel of a variable phase shifter module consisting of four Mini-Circuits JSPHS-1000 [50] phase shifters. A Mini-Circuits ZN8PD1-53+ [51] eight way power splitter is then used to provide eight copies of the LO signal (as the optimum number of processors that can be instrumented using two FONT5 boards is six, two of these signals are redundant). Four of these copies are input to a second variable phase shifter module while channels 5, 6 and 7 use the remaining three inputs of the module that the master LO signal passed through. All but one of the channels are then amplified and are ready to be sent via heliax cables to the BPM processor modules in the beamline.

The variable phase shifters are required as optimal operation of the mixers on the BPM processors occurs when the two inputs (the LO and the signal derived from the stripline outputs) are in phase [24]. This results in a maximum of the absolute value of the Σ_I output and hence the best possible signal-to-noise ratio. The fixed universal 90° phase shift is used so that the problem of finding the phase that maximizes the Σ_I is reduced to the easier problem of finding the phase that returns a zero value for the Σ_I. Once this phase is located, the 90° phase shift is removed to deliver the maximized Σ_I signal required.

Control of the phase shifters is exercised through a Data Translation DT9854 [52] USB voltage output module connected to a PC. This unit provides 8 independent analogue signals with DC voltages in the range 0-10 V with a 16-bit resolution (along with 8 digital outputs, one of which is used to activate the 90° phase shift). These signals are scaled via a custom control unit to match the 0-15 V input range of the phase shifters [53]. The total phase shift that can be achieved is over 180°; as the sign of the processor output signals is not important, this is sufficient range to allow the maximal absolute value of the Σ_I to be achieved. A calibrated (Figure 2.20) interface for control of the phase shifters was developed in LabVIEW [54].

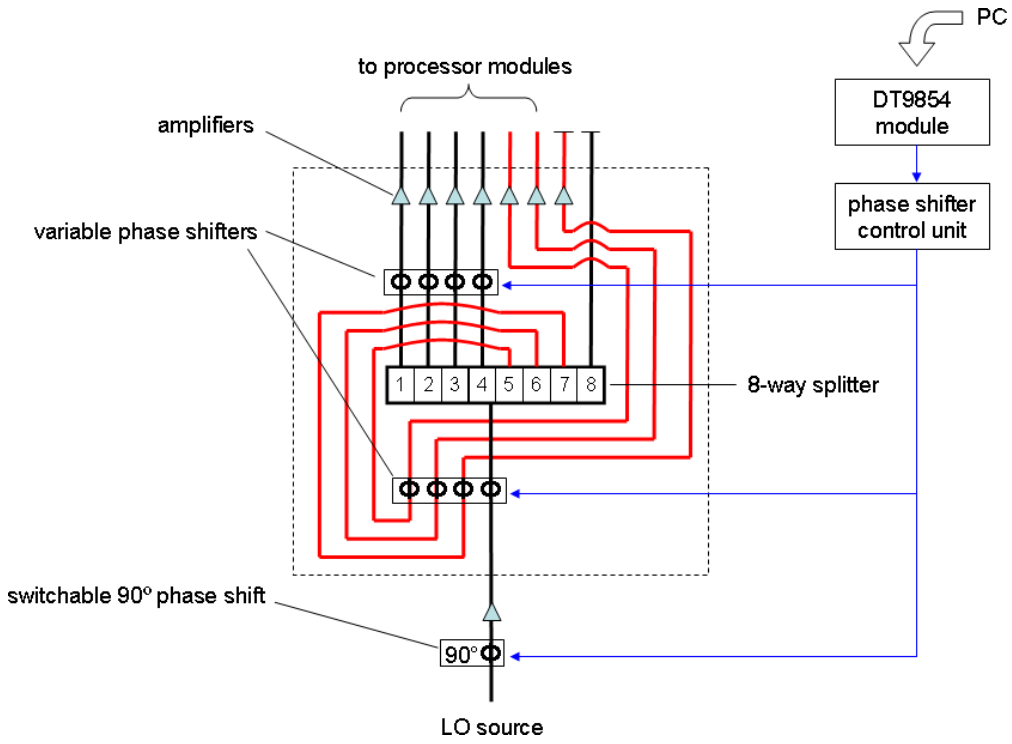

Figure 2.19: Schematic diagram of the LO distribution for the FONT processor modules.

2.4.3 Optimization of the ADC sample timing

The nine ADC channels arranged in three groups of three allow all the output signals (Δ, Σ_I and Σ_Q) from three processor modules to be monitored using one board. It is highly desirable that the assignment of processor output signal to ADC channel is constructed such that the signals from a single processor share a common ADC clock. This eliminates the possibility of timing offsets due to a conflict between the optimum ADC clock delay setting for signals from different processors.

The forms of the digitized signals are illustrated in Figure 2.21. This data is sufficient for the measurement of the position of one train of two bunches at a single BPM. The Σ_I signal in Figure 2.21(e) represents the charge. It contains two clear peaks, each of which corresponds to the passage of a bunch, at sample numbers 37 and 104. These samples may be termed the bunch samples. The bunch samples, as defined by the peaks of the Σ_I signal,

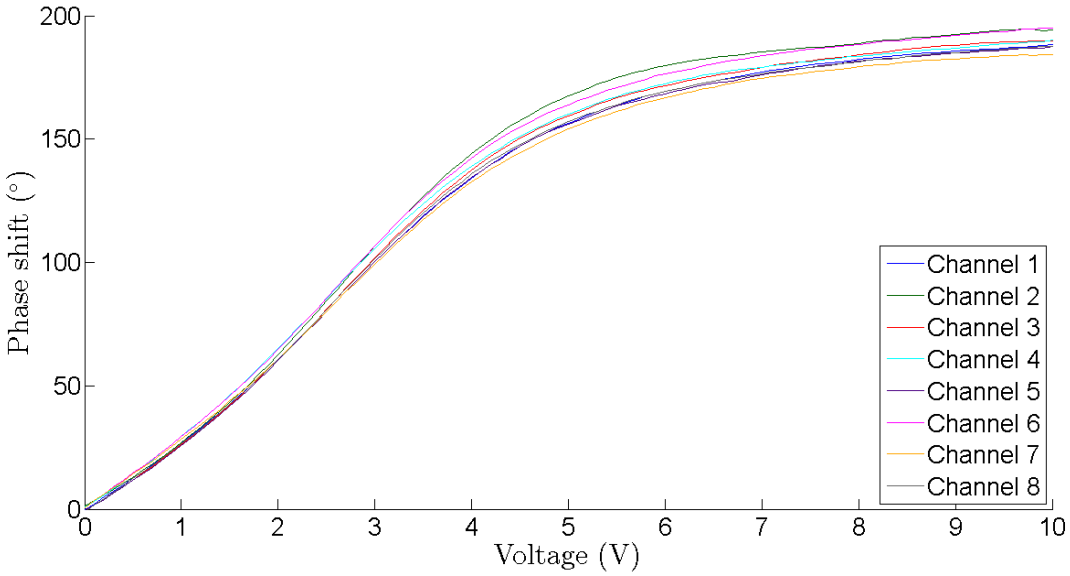

Figure 2.20: Change in phase of a 714 MHz sinusoidal input signal as a function of the voltage applied to each phase shifter.

are the only samples from each waveform used to determine the position of the bunch; they are indicated by the red circles on each waveform. As the bunch samples are separated by 67 samples and the sample period is one cycle of the 357 MHz clock, it may be deduced that the waveforms of Figure 2.21 were produced from a train of two bunches with a spacing of $67 \times 2.8 = 187.6$ ns.

The values of the Σ_Q bunch samples in Figure 2.21(a) are much smaller than those of the bunch samples from the Σ_I waveform; this indicates that the LO phase is close to optimum for that processor. The values of the Δ bunch samples in Figure 2.21(c) are also close to zero; this indicates the beam is passing close to the centre of that BPM. Figure 2.21(d) shows the form of the Δ waveform in more detail; the large peak that occurs in the Δ signal shortly after the bunch sample is a reflection of the Σ_I signal within the processor.

The position of each bunch is calculated from the ratio of the corresponding bunch samples from the Δ and Σ_I waveforms. For ideal operation, the scan delays that control the timing of the ADC clocks are set so that the bunch sample value of each Σ_I signal is maximized. At this point the sampling point is maximally coincident with the actual peak of the Σ_I signal and the best signal-to-noise ratio is obtained [4]. This concept is illustrated in Figure 2.22. The scan delay is an integer multiple of 70 ps and is represented in the schematic by the quantity Δt.

[4] As the same ADC clock is used for the Σ_I and Δ signals from a single processor, a reduction in the signal-to-noise ratio is the only adverse effect of a sub-optimal scan delay setting (although the calibration is different to the optimal case).

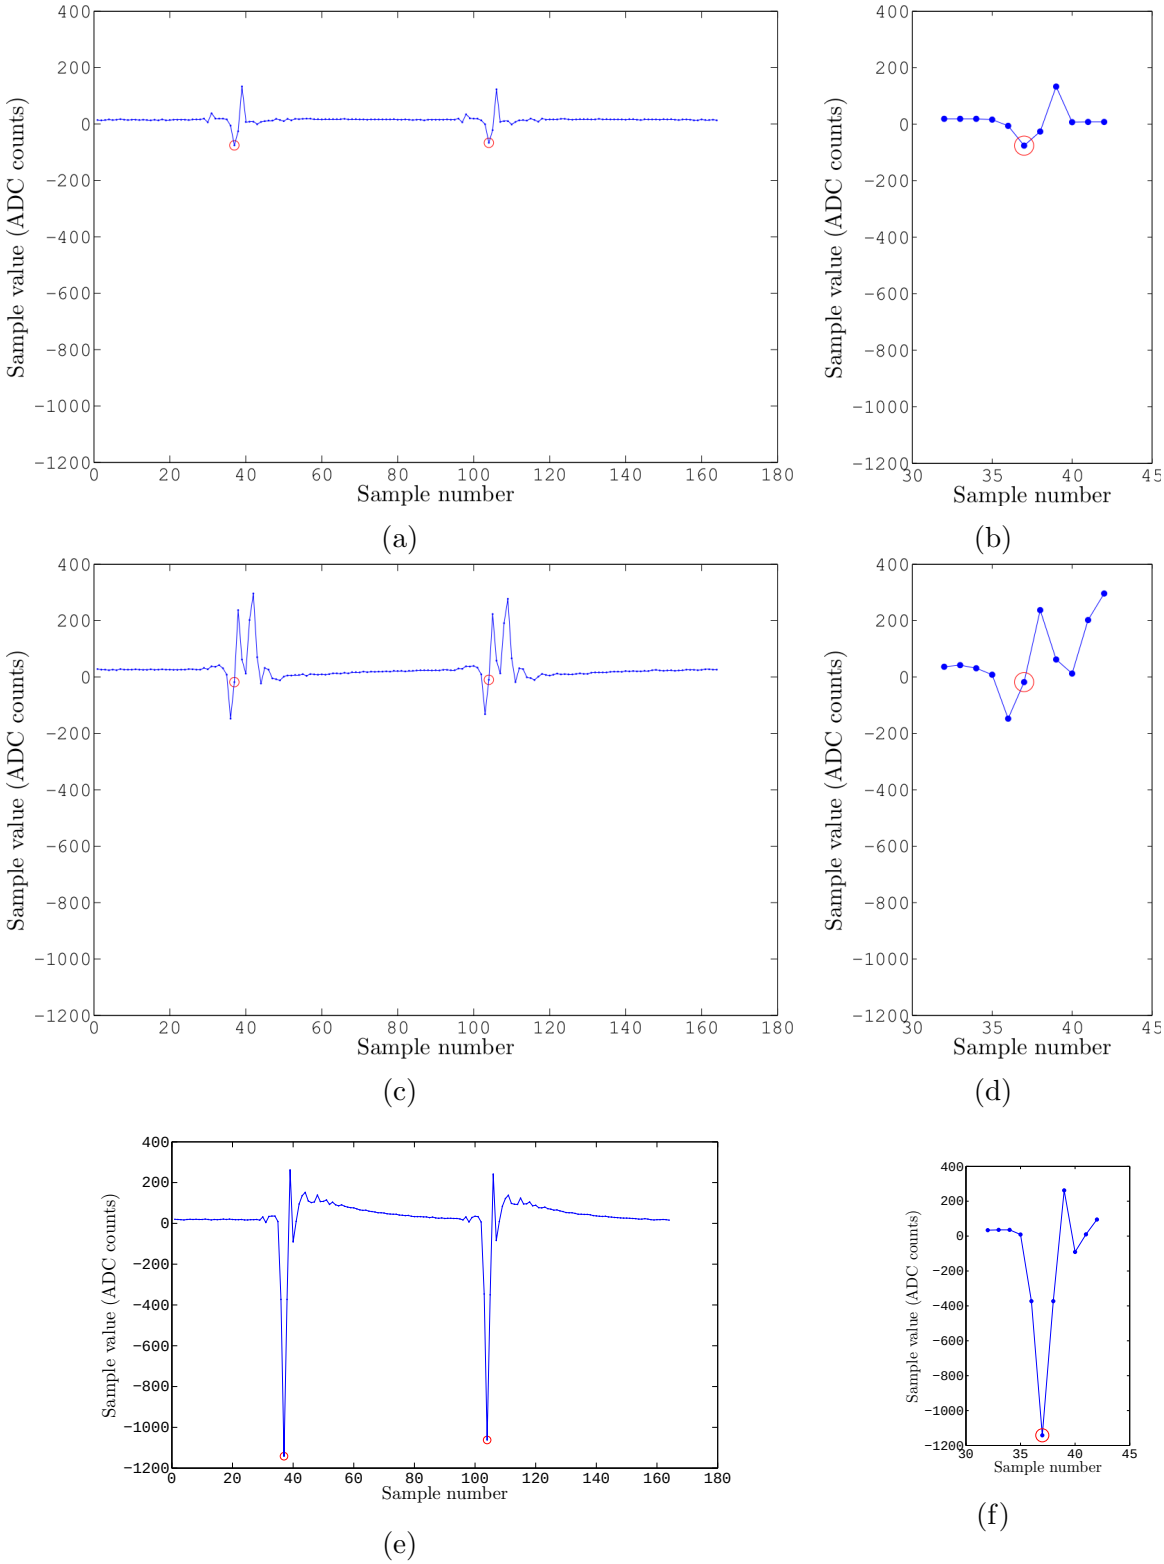

Figure 2.21: Processor output signals from a single BPM as digitized by a FONT5 board. (a) Σ_Q (entire waveform); (b) Σ_Q (detailed view of bunch 1 region); (c) Δ (entire waveform); (d) Δ (detailed view of bunch 1 region); (e) Σ_I (entire waveform); (f) Σ_I (detailed view of bunch 1 region).

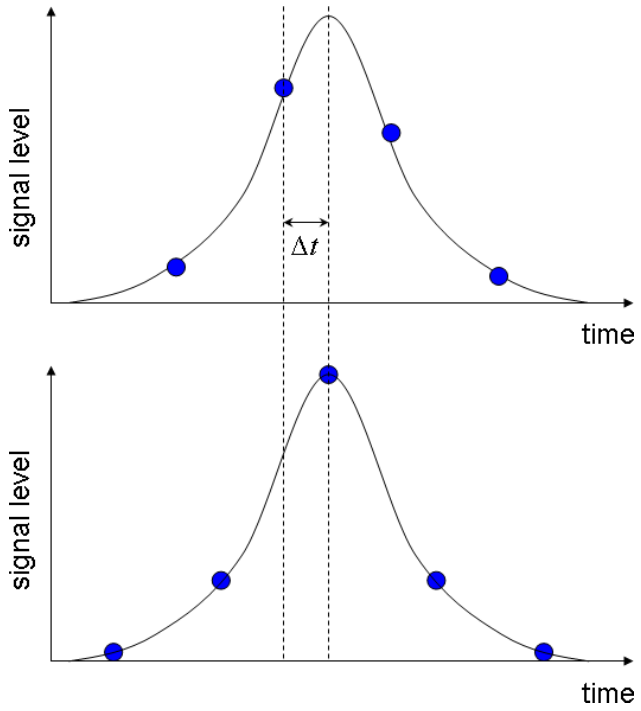

Figure 2.22: Diagram illustrating the function of the scan delay parameter. The solid black line represents the voltage of the analogue waveform; the blue circles represent the digitized samples. Top: off-peak sampling of the waveform; Bottom: optimal sampling of the waveform.

2.4.4 Centring of the beam in the BPMs

In order to satisfy the condition $R \gg y$, it is desirable to make the beam pass as close to the centre of the BPM as possible. A system utilizing piezoelectric actuator stages was developed for this purpose by a team from the Instituto de Física Corpuscular (IFIC) in Valencia, Spain [55]. The BPM movers allow motion in both the x and y directions; the dynamic range in each case is $\sim \pm 1.5$ mm and the resolution is at the micron level. Mover stages were installed on each of the dedicated FONT stripline BPMs and a LabVIEW interface was developed through which to control them. The bunch was taken to be passing through the centre of the BPM when the position of the mover was such that the average value of the Δ signal was zero [5].

2.5 Summary

The individual sub-systems that make up the FONT5 feedback system have been described, specifically: the BPMs and associated analogue processor modules that provide beam monitoring capability; the kickers and kicker amplifiers that enable control of the beam position; and the FONT5 board that acts as the crucial interface between the two. The various

[5] While the BPMs are mechanically aligned along a reference axis, the procedure used does not take into account the orbit of the beam.

procedures required to optimize the FONT system were also described.

Chapter 3

BPM Calibration and Resolution

The quality of the correction that can be achieved using the FONT feedback system is inherently limited by the performance of the beam position monitoring (BPM) system. This chapter will describe the layout of the FONT hardware at ATF and present the procedure for calibration of the BPMs that is required in order to quantify the effectiveness of the feedback correction. In addition, the identification and elimination of those factors that result in a degradation of the resolution of the FONT BPMs will be discussed.

3.1 The FONT installation at the ATF

3.1.1 Layout of beamline components

Figure 1.5 is a schematic representation of the ATF beamline that indicates with respect to the rest of the ATF the location of the FONT region. The FONT region contains the FONT stripline BPMs and the FONT feedback kickers along with all the associated processing hardware, including the FONT5 boards. The MFB1FF region is also marked; this is where the downstream stripline BPM MFB1FF resides. As the processor inputs are connected to the stripline outputs using 1 m long cables, 30 m long cables are required to connect the MFB1FF processor outputs (Δ, Σ_I and Σ_Q) to the ADC inputs of the FONT5 board in the FONT region.

More detailed diagrams labelling the beamline components in each region are given in Figure 3.1 and Figure 3.2. Each of the featured components is labelled with the name used to identify it. With the exception of the dedicated FONT hardware, where the names P1, P2 and P3 are used for the stripline BPMs and K1 and K2 are used for the kickers, the components have names that reflect both their nature and general location.

The nature of a component is given at the beginning of its name. Quadrupole magnets have names beginning with Q, corrector magnets have names beginning with Z and BPMs have names beginning with M. The general location of a component is given at the end of its name. If the name ends in X, the component is considered to be part of the extraction line. A name ending in FF identifies the component as part of the final focus. The transition

between these two sections occurs just downstream of the FONT region. In general each name also features a number that acts as an index for components of that particular type. In the extraction line, this index increases with z; the reverse is true in the final focus.

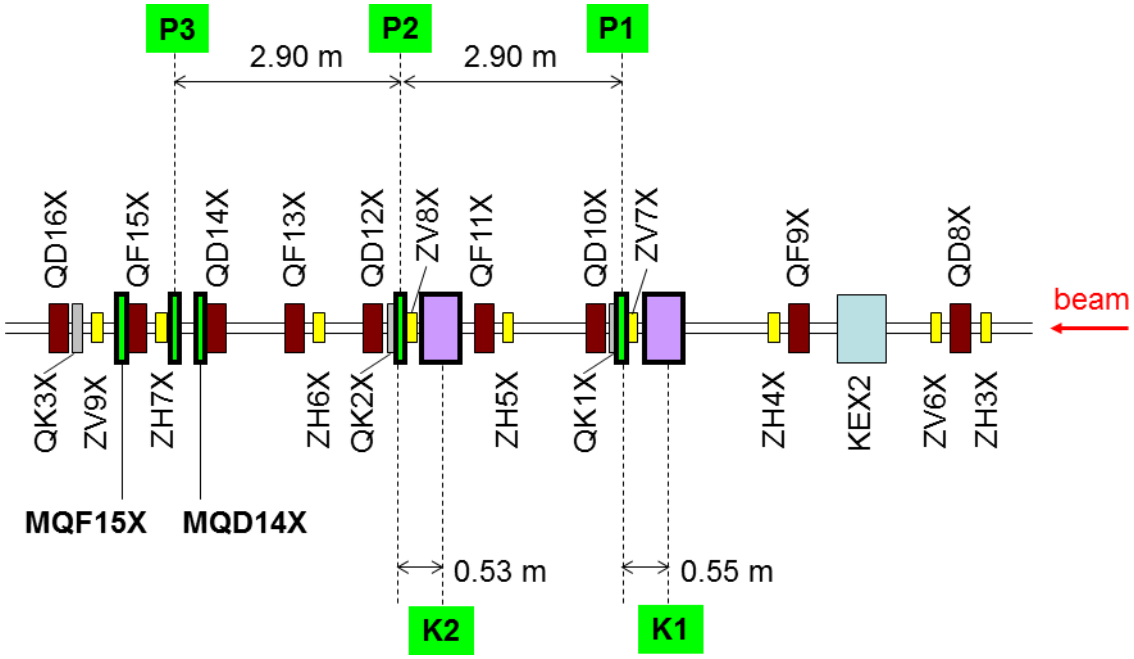

Figure 3.1: Schematic representation of the beamline in the FONT region.

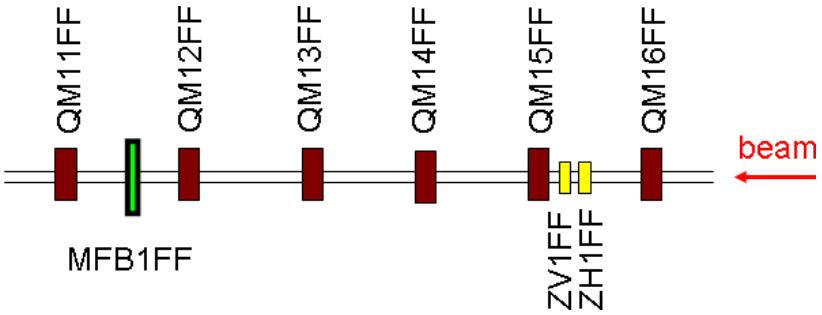

Figure 3.2: Schematic representation of the beamline in the MFB1FF region.

3.1.2 FONT processing hardware at the ATF

Figure 3.3 is a schematic indicating how the FONT hardware is connected to ATF beamline components. The nominal installation of the FONT system at the ATF uses three BPM processor modules to monitor the vertical outputs of the dedicated FONT stripline BPMs P1, P2 and P3. The eight BPM processor modules are nominally identical and so the convention in this thesis will be to refer to the output signals of the BPM processor modules using the name of the BPM that the processor module is attached to; thus, the signal $\Sigma_{I-\mathrm{P1}}$ refers to the Σ_I output of the BPM processor module whose A and B inputs are connected to the top and bottom outputs respectively of the stripline BPM P1.

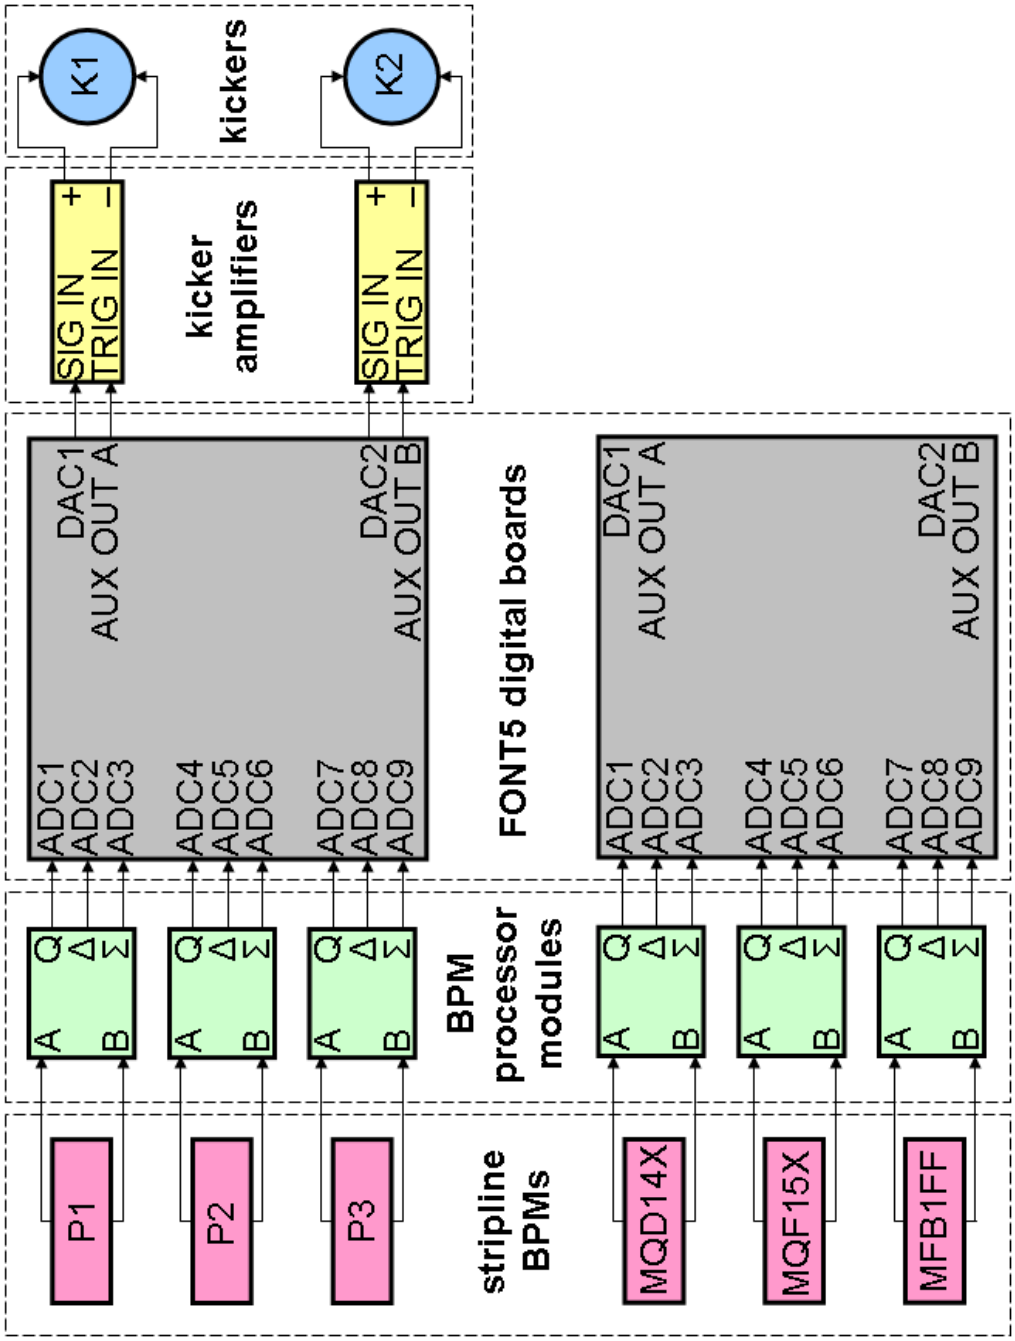

Figure 3.3: Schematic indicating how the FONT hardware is connected to ATF beamline components.

The outputs of these three analogue BPM processor modules are then digitized by a FONT5 board. The board is capable of calculating in real time the position of the bunch at P2 and P3 (from the ratio of the Δ and Σ_I bunch sample values in each case) in order to drive a pair of kicks at the nearby FONT kickers K1 and K2. In addition to this primary FONT5 board controlling the feedback, a second FONT5 board is used (along with another three BPM processor modules) to record the signals from three additional stripline BPMs. As depicted in the schematic, these BPMs are the stripline BPMs MQD14X, MQF15X and MFB1FF. MQD14X and MQF15X are both located less than 1 m from P3 while MFB1FF is located a significant distance downstream in the final focus region. The data from these secondary BPMs are primarily useful in order to study how well the increase in beam stability due to the FONT feedback system is maintained at various points in the lattice downstream of where it is applied.

The data presented in this chapter were obtained using both FONT5 boards purely as digitizers; that is, the primary FONT5 board was not used to provide a correction at the kickers K1 and K2. Such passive observation of the beam is carried out on a routine basis for the purposes of calibrating the BPMs and estimating the maximum resolution that can be achieved using the FONT hardware.

3.1.3 Analysis of FONT5 data

A FONT5 data set consists of all the digitized waveforms logged over a fixed interval of time. Only one sample per waveform is required to determine the position of a given bunch; these samples may be termed the bunch samples. A data set may be presented as the value of these bunch samples in ADC counts as a function of the extraction number. A data set from 19 April 2012, when the beam was operated in single bunch mode, is presented in this way in Figure 3.4.

The evolution with time of the value of the bunch sample from each of the three Σ_I channels on a single FONT5 board is given in Figure 3.4(a). The Σ_I values quantify the charge of the bunch. Under normal conditions of negligible beam loss the bunch charge for a given extraction is fixed (although it varies from extraction to extraction). This is confirmed by the very good agreement that exists between the Σ_I values reported at each BPM; the calculated magnitude of the correlation coefficient [1] between any pair of BPMs exceeds 0.999 in all three cases (P1 to P2, P2 to P3 and P1 to P3). The absolute value of Σ_I is plotted because it is possible to invert the sign of the BPM processor output signals by applying a 180° phase shift to the LO input. For this data set, the optimal phase of the P3 LO was such that negative values were obtained for $\Sigma_{I-\mathrm{P3}}$. The use of the absolute value thus makes clear the similarity of the Σ_I values obtained from different BPMs.

Figure 3.4(b) is the corresponding plot for the Δ signals. The Δ signals depend on both the charge and the position of the bunch; the Δ/Σ_I processing method thus amounts to charge normalization of the Δ signal. For this data set, the Δ sample values are much

[1] Pearson's product-moment correlation cofficient, defined as the covariance of a pair of variables divided by the product of their standard deviations.

smaller than the Σ_I sample values, indicating that the beam passed close to the centre of each BPM. For the FONT BPMs, this condition is trivially achieved by moving the BPMs. The plot shows that Δ_{P3} has a very strong anti-correlation to Δ_{P2}. The negative sign of the correlation coefficient simply reflects the aforementioned 180° difference in the phase of the LO between the two BPM processors; such phase differences are ultimately irrelevant as they are cancelled out when the division by Σ_I is performed.

Figure 3.4(c) shows the Σ_Q signals. Each Σ_Q signal shows approximately how the corresponding Σ_I signal would have appeared if a 90° phase shift had been applied to the input LO. The Σ_Q value, once charge normalized, provides a measure of the phase offset ϕ between the two inputs to the Σ_Q mixer: the 90° phase-shifted LO and the Σ_I output of the resistive coupler. The phase shifters in the LO distribution are set to ensure that the values of the Σ_Q bunch samples are close to zero, in which case each Σ_Q value is maximally sensitive to variations in ϕ. As this phase offset is determined primarily by the arrival time of the bunch at the BPM the Σ_Q values exhibit a high degree of correlation, although the mean value differs somewhat from BPM to BPM due to the limited accuracy of the procedure used to set the phase shifters in the LO distribution.

From the three sets of values for Δ, Σ_I and Σ_Q it is thus possible to determine for each extraction the charge normalized bunch position (Δ/Σ_I), plotted in Figure 3.4(d), and the phase offset between bunch and LO (Σ_Q/Σ_I), plotted in Figure 3.4(e).

3.2 BPM calibration

The ratio Δ/Σ_I has the units of ADC counts divided by ADC counts; therefore, in order to express the position of the bunch in a more meaningful way, it is necessary to multiply the uncalibrated position by a calibration constant k_y. The value of the calibration constant for a particular BPM may be determined by scanning the position of the beam relative to the BPM. The methods and results of this procedure will be described in this section.

3.2.1 Theoretical calibration

The theoretical expression for position in a stripline BPM (Equation 2.5) includes the ratio of $V_A - V_B$ to $V_A + V_B$. In practice, these two are replaced by the output signals of the BPM processor modules. As Δ is proportional to $V_A - V_B$ and Σ_I is proportional to $V_A + V_B$, the ratio of the two may be described as:

$$\frac{\Delta}{\Sigma_I} = \frac{g_\Delta}{g_{\Sigma_I}} \left(\frac{V_A - V_B}{V_A + V_B} \right) = \frac{g_\Delta}{g_{\Sigma_I}} \left(\frac{2y}{R} \right) \equiv k_y y \tag{3.1}$$

g_{Σ_I}, g_Δ, R and y were defined in Section 2.1.1 and k_y is the calibration constant of the BPM and processing electronics. The theoretical expression for k_y is:

$$k_y = \frac{2}{R} \left(\frac{g_\Delta}{g_{\Sigma_I}} \right) \tag{3.2}$$

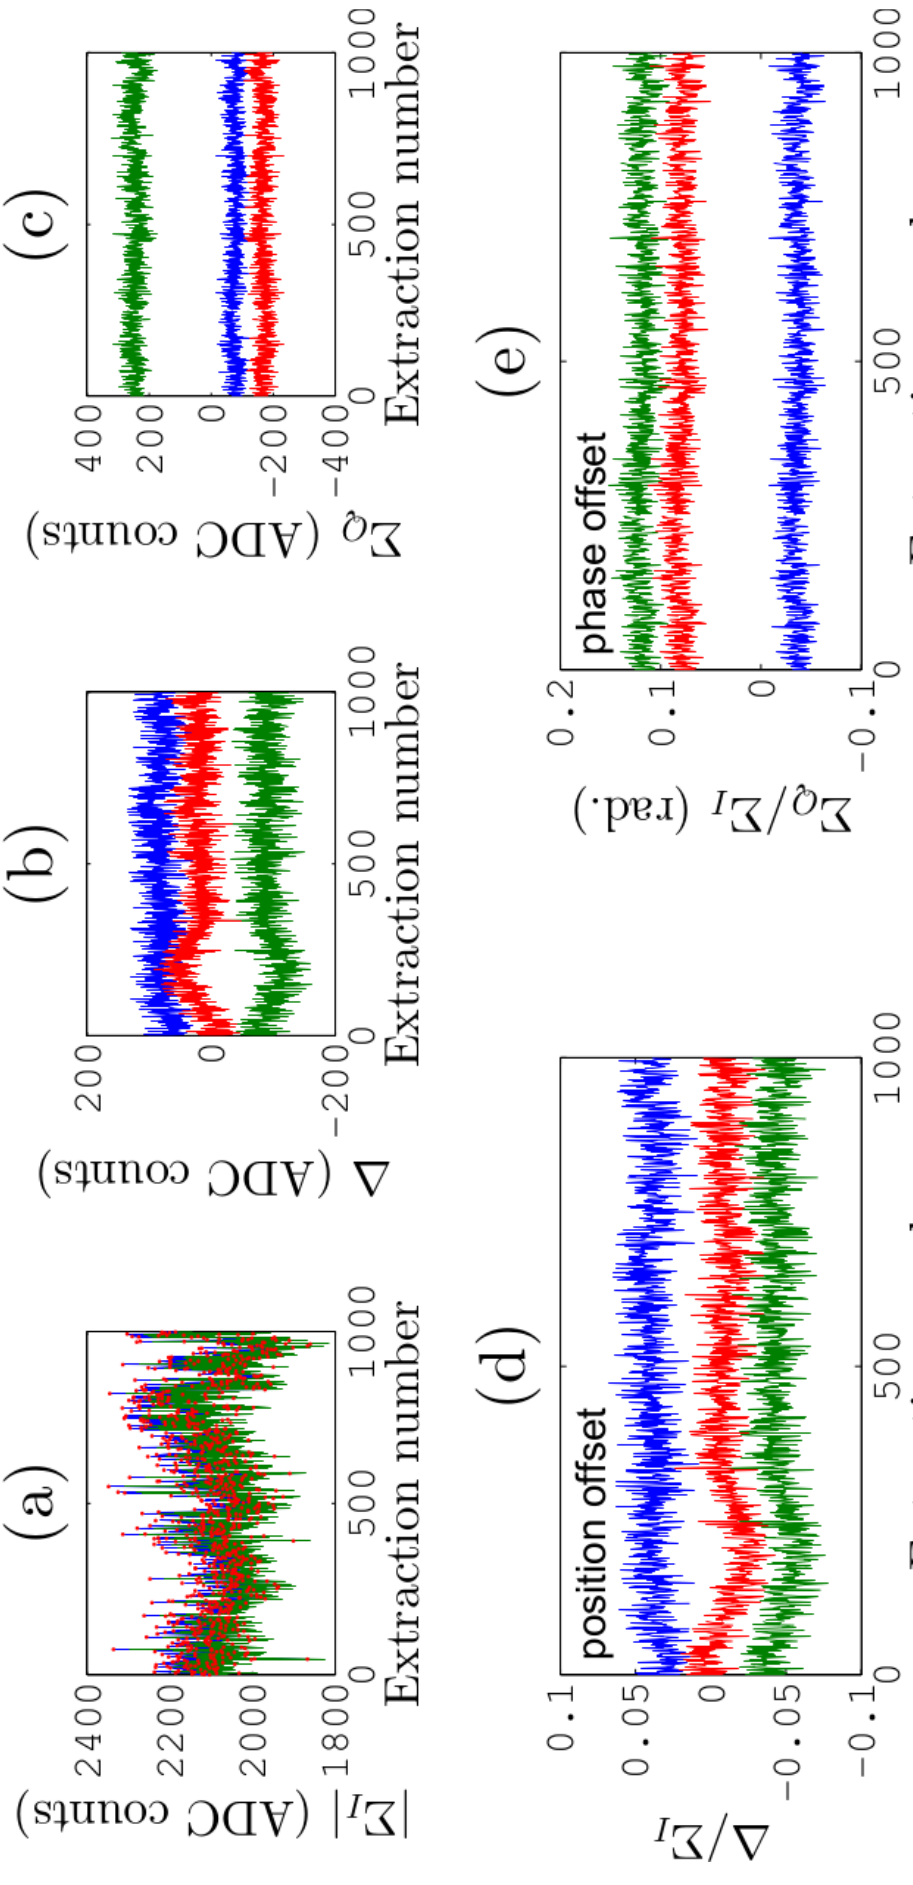

Figure 3.4: A single bunch extraction data set from 19 April 2012 represented by the value of the single bunch sample per digitized waveform along with the ratios Δ/Σ_I and Σ_Q/Σ_I as a function of extraction number. Blue: P1; Green: P2; Red: P3.

and an estimate for the calibration constant can thus be obtained from the radius of the FONT stripline BPMs ($R = 12$ mm) and the measured gains of the Δ and Σ_I channels. These are determined by connecting a 700 MHz sinuosoidal signal of 0 dBm to one of the BPM processor module inputs while the other input is terminated and then measuring the amplitude of the resulting Δ and Σ_I outputs. The results of this procedure for all 8 processor modules are given in Table 3.1 [24]. The gain ratios of the processors are found to vary considerably; the gain ratio of processor 8 is 15% lower than the mean value of $\langle \|g_\Delta / g_{\Sigma_I}\| \rangle = 16.4 \pm 1.1$. The predicted value of the mean processor calibration constant is thus $\langle \|k_{\mathrm{cal}}\| \rangle = (270 \pm 20) \times 10^{-5} \mu\mathrm{m}^{-1}$.

Processor	$\frac{g_\Delta}{g_{\Sigma_I}}$
1	16.6
2	16.1
3	17.1
4	17.8
5	16.0
7	16.5
8	14.0
10	17.0

Table 3.1: Value of the gain ratio parameter $\frac{g_\Delta}{g_{\Sigma_I}}$ for all BPM processor modules.

3.2.2 Calibration methods

The calibration procedure is performed by changing the position of the beam relative to the BPM by a known amount and observing the effect on the position reported by the FONT hardware. The change in apparent position can be brought about either by using the BPM movers or the correctors; both are assumed to provide an accurate offset of the beam relative to the BPM.

3.2.2.1 Calibration using the BPM movers

Calibration using the BPM movers is the simpler of the two methods but is only available for BPMs P1, P2 and P3. These three BPMs are mounted on stages that can be moved relative to the beam, allowing each BPM to be positioned such that the beam appears to pass close to its centre. Calibration of the BPMs is then performed by scanning the vertical positions of all three movers from -100 μm to 100 μm relative to the zero point in increments of 20 μm. Figure 3.5 shows the position of the beam over a calibration run as a function of time elapsed since the start of the run. No data is stored while the movers are in motion; each discontinuity of the line corresponds to the time taken for the movers to reach their new positions.

The calibration constant for each BPM is estimated by taking the average position of

the beam at each mover setting and then performing a linear χ^2 fit of the position as a function of mover setting. The gradient of the fit is the value for k_y. As an example, the data for a mover calibration from December 2011 are presented in Figure 3.6 (note that it was assumed that there was no error on the mover setting and the horizontal extent of the error bars is arbitrary). The resulting values of k_y are listed in Table 3.2; the error quoted is the standard error for a fitted gradient [56]. The table also includes the value of the goodness-of-fit parameter (the calculated χ^2 divided by the number of degrees of freedom ν). The plots show a small amount of saturation at the extreme negative end of the scan; as a result, the data points from the $\pm 100\mu$m mover settings are ignored for the purposes of calculating the fit.

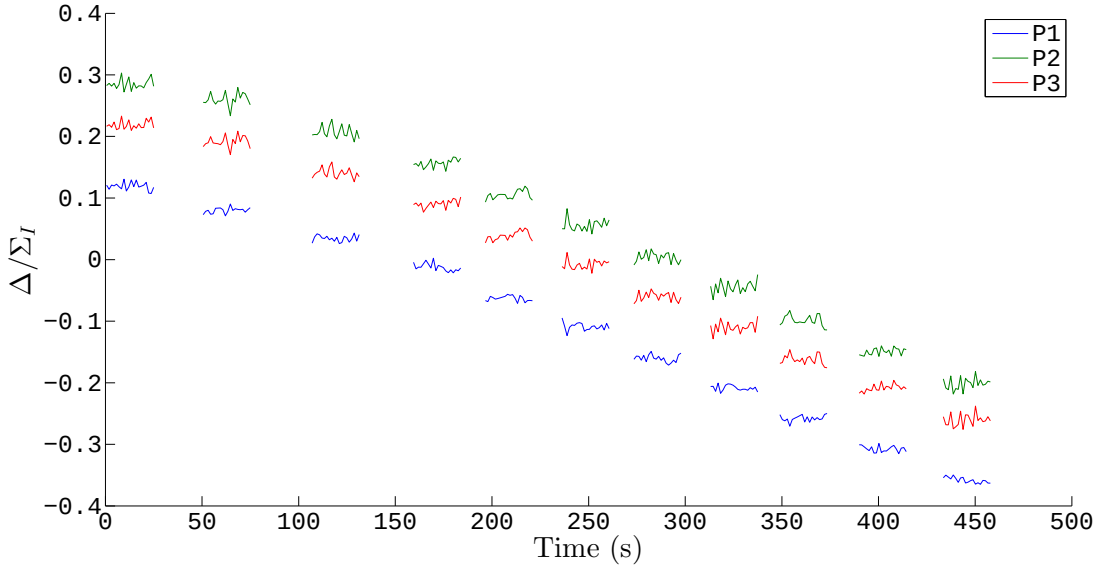

Figure 3.5: Uncalibrated vertical position of the bunch over the duration of a mover calibration run. Blue: P1; Green: P2; Red: P3.

BPM	k_y ($\frac{\Delta}{\Sigma_I}/\mu$m)	χ^2/ν
P1	$(-243 \pm 1) \times 10^{-5}$	0.3
P2	$(-255 \pm 1) \times 10^{-5}$	0.1
P3	$(-250 \pm 1) \times 10^{-5}$	0.1

Table 3.2: Value of the k_y parameter for the FONT BPMs P1, P2 and P3 obtained from a scan of the mover positions on 7 December 2011.

3.2.2.2 Calibration using the corrector magnets

Calibration using the corrector magnets involves modifying the trajectory of the beam itself by varying the current through one of several correctors located throughout the ATF

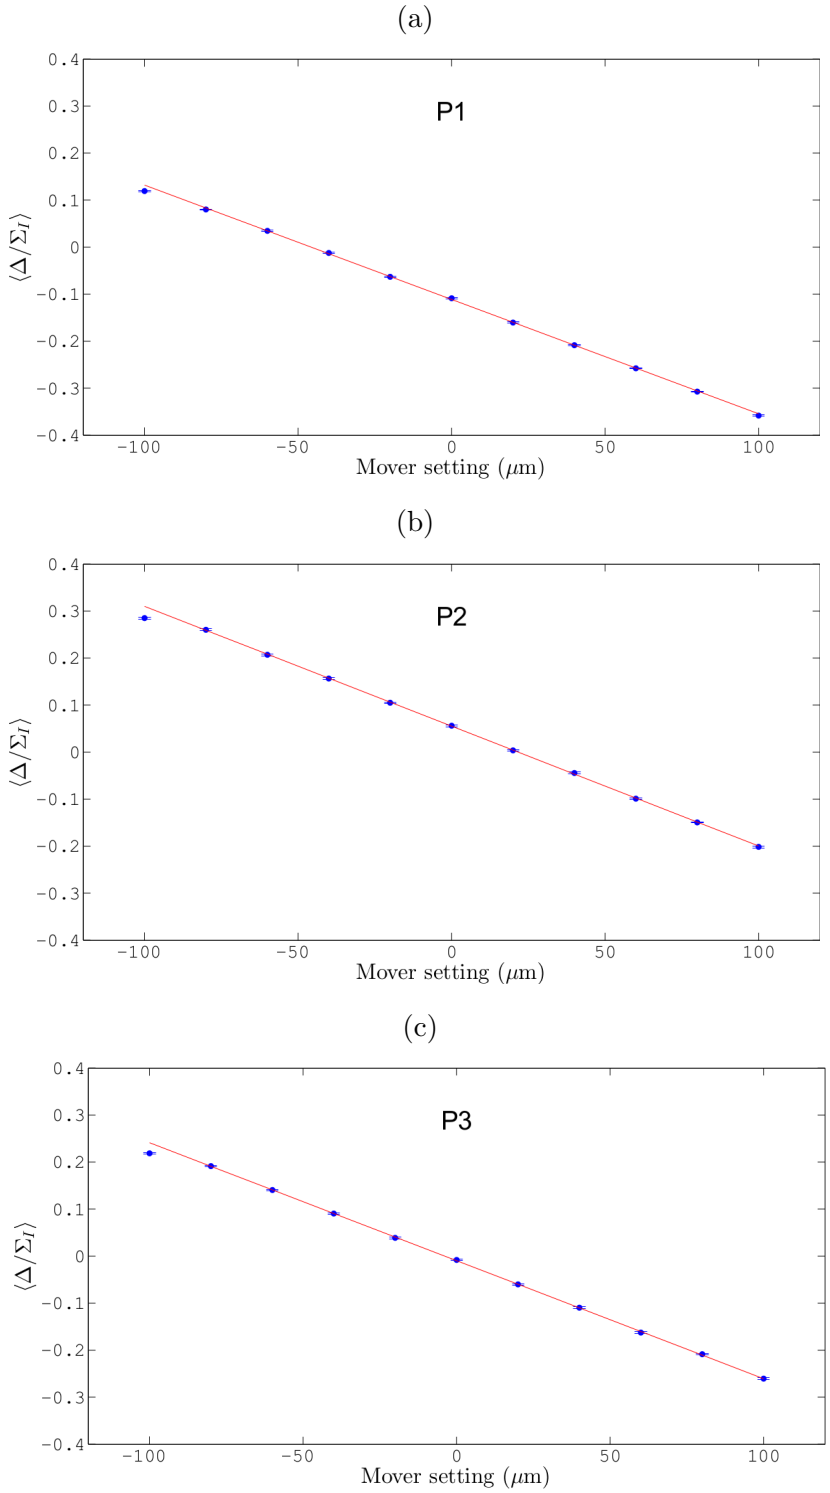

Figure 3.6: Mean value of the uncalibrated vertical position as a function of BPM mover setting (blue) and the linear χ^2 fit to these data (red) for (a) P1; (b) P2; (c) P3. Error bars represent the error on the mean.

BPM	k_y ($\frac{\Delta}{\Sigma_I}/\mu$m)	χ^2/ν
MQD14X	$(-255 \pm 1) \times 10^{-5}$	0.1
MQF15X	$(-247 \pm 1) \times 10^{-5}$	0.1
MFB1FF	$(-276 \pm 1) \times 10^{-5}$	0.6

Table 3.3: Value of the k_y parameter for the BPMs MQD14X, MQF15X and MFB1FF obtained from a scan of the ZV8X current on 7 December 2011.

extraction line. The procedure is similar to that followed when calibrating using the BPM movers, except with the added complication that the change in the current of the relevant corrector must be converted into an expected displacement of the beam at the given BPM. This requires measurement of the field as a function of power supply current.

Calibration using the corrector magnets was required for the FONT BPMs prior to the installation of the movers. It remains the only method for calibration of the other stripline BPMs commonly monitored by FONT (MQD14X, MQF15X and MFB1FF). The correctors used by FONT for BPM calibration purposes are ZV8X and ZV1FF. ZV1FF is located just upstream of MFB1FF (Figure 3.2) and can only be used to calibrate that BPM. ZV8X is located in the middle of the FONT region and can be used to calibrate MQD14X, MQF15X and MFB1FF.

3.2.2.2.1 Change in beam position due to a corrector

The change in position of the beam resulting from a change in the current setting for an upstream corrector may be predicted by considering the effect of a magnetic field on a charged particle. A particle of charge e travelling in a region of magnetic flux density B_0 follows a circular path of radius $\rho = p/eB_0$ where p is the momentum of the particle. Neglecting fringe effects, the angular deflection imparted to such a particle while traversing a dipole magnet of length L in the z-axis is:

$$\theta \sim \frac{L}{\rho} = \frac{Le}{p} B_0 \equiv \alpha_{\text{mag}} B_0 \tag{3.3}$$

where the approximation is valid for small θ. Using the ATF parameters of $p = 1.3$ GeV/c and $L = 128$ mm gives a value $\alpha_{\text{mag}} = 2.95 \times 10^{-2}$ rad/T. B_0 as a function of the current I was measured for each corrector by holding by hand a Hall probe in the vertical gap between the magnet pole and the beam pipe and recording the magnetic flux density for a number of different current settings; this data is presented in Figure 3.7. By performing a linear fit to the data, the calibration constant dB_0/dI was obtained in each case. These results are given in Table 3.4. The value of the calibration constant is found to be similar for all four of the correctors, although the value of the reduced χ^2 statistic is somewhat larger than 1 in each case. This may be because, due to time constraints, a single reading was taken for the magnetic flux density at each current setting. The error on this measurement was then taken to be the instrument error, which was 0.005 mT for the Hall probe used. The corrector current value read back from the ATF control system was taken to be perfectly accurate. For a given current setting, the change in position due to the corrector is calculated using

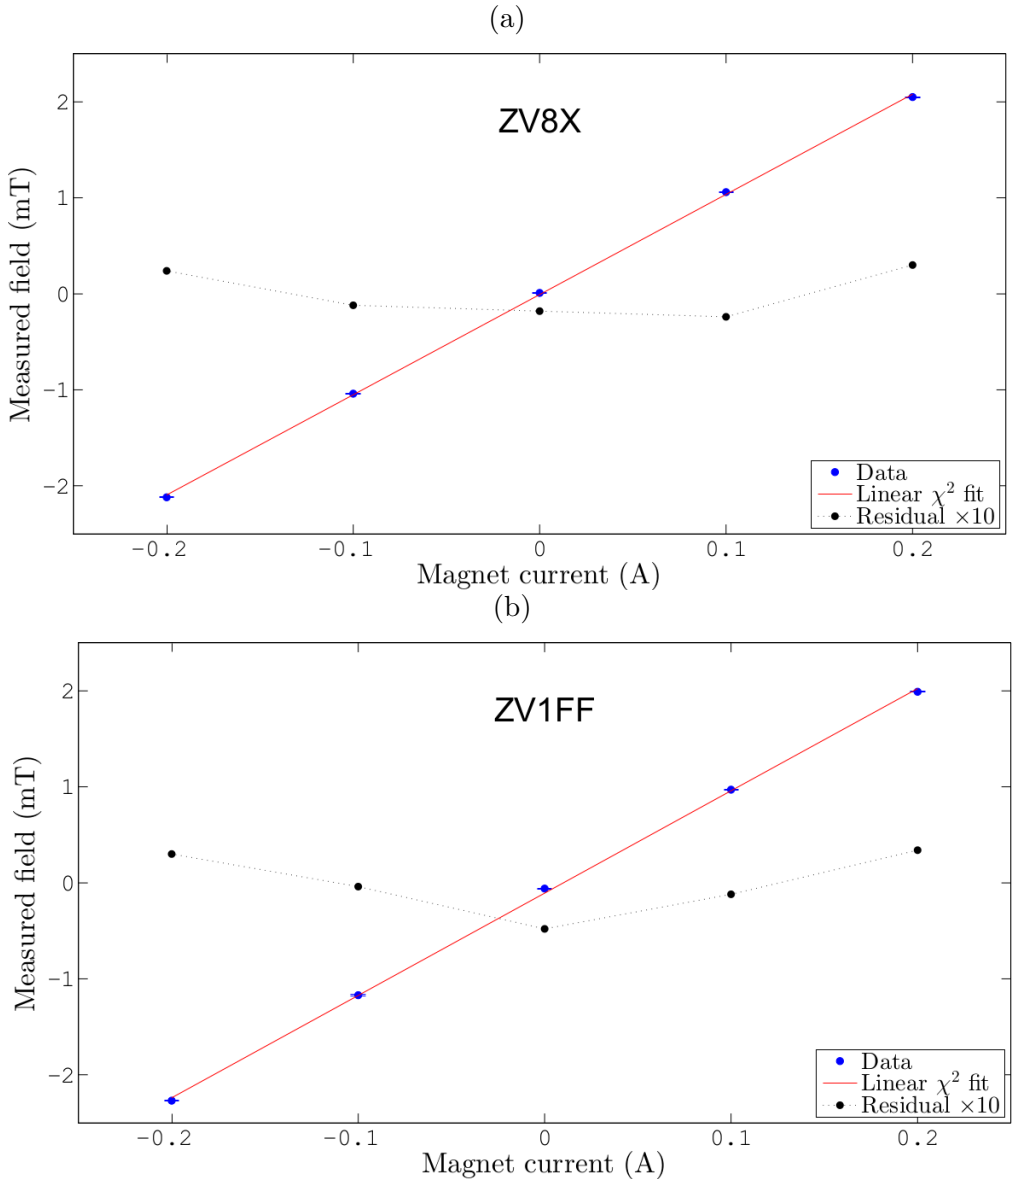

Figure 3.7: Measured magnetic flux density as a function of corrector current for: (a) ZV8X; (b) ZV1FF. Blue: measurements with the Hall probe; Red: linear χ^2 fit; Black: fit residuals multiplied by 10. Error bars represent the error on the mean.

Corrector	dB_0/dI (mT/A)	χ^2/ν
ZV8X	10.44 ± 0.02	6.7
ZV1FF	10.66 ± 0.02	12.1

Table 3.4: List of corrector calibration constants.

the appropriate elements of the transfer matrix from the corrector to the BPM of interest. Consider the beam transport from a generic corrector Z to some BPM P:

$$\begin{pmatrix} y_P \\ y'_P \end{pmatrix} = \begin{pmatrix} {}^{Z}_{P}M_{11} & {}^{Z}_{P}M_{12} \\ {}^{Z}_{P}M_{21} & {}^{Z}_{P}M_{22} \end{pmatrix} \begin{pmatrix} y_Z \\ y'_Z \end{pmatrix} \tag{3.4}$$

where y indicates the vertical position of the bunch, y' is the rate of change of y as a function of distance and ${}^{Z}_{P}M$ is the 2×2 transfer matrix for calculating y and y' at P from their values at Z. The position at P when the current through the corrector Z is set to its nominal value is:

$$y_P = {}^{Z}_{P}M_{11} y_Z + {}^{Z}_{P}M_{12} y'_Z \tag{3.5}$$

and the effect of a change in this current is to add an extra term θ to y'_Z. The displacement at P resulting from a change in current through Z is thus given by:

$$ {}^{Z}_{P}M_{12}\theta = {}^{Z}_{P}M_{12}\alpha_{\mathrm{mag}} \left(dB_0/dI\right) I \tag{3.6}$$

and the transfer matrix element is calculated from the MAD8 [57] model of the ATF extraction line using the settings for the current of the magnets. As an example, the data for a ZV8X calibration from December 2011 is given in Figure 3.8 and the resulting values of k_y in Table 3.3. Note in this case how the single range of ZV8X current maps to individual ranges of positions at each BPM and that the y axes also have different scales. As these BPMs are not mounted on movers, it is not possible to trivially force the nominal trajectory of the beam to pass close to the centre of each BPM. Note that while the current of ZV8X was being scanned the other corrector currents were left at their nominal values.

3.2.3 Calibration results

The BPM calibration procedure is performed routinely at the start of any period of data taking as part of the suite of diagnostics used to verify correct operation of the FONT hardware. As a result, there are many BPM calibration data sets incorporating many different combinations of BPM and analogue processor module, and utilizing both movers and correctors. A large selection of results from BPM calibrations in the vertical sense only and going back as far as December 2009 are collected in Figure 3.9.

By calculating the mean value of the calibration constants measured for each processor, it is possible to compare the theoretical value obtained using the corresponding gain ratio of Table 3.1 with the observed value. The results are in Figure 3.10; the general trend is for the theoretical value to overestimate the observed value typically by less than 15% (with processors 3 and 8 being exceptional in this regard). For the most part the difference appears systematic suggesting that the simple expression for the theoretical value given in Equation 3.2 is in need of refinement.

3.3 Resolution of the FONT BPM system

The resolution of the FONT BPM system is the smallest change in the position of the beam that can be accurately measured; hence, the resolution sets an upper limit on the

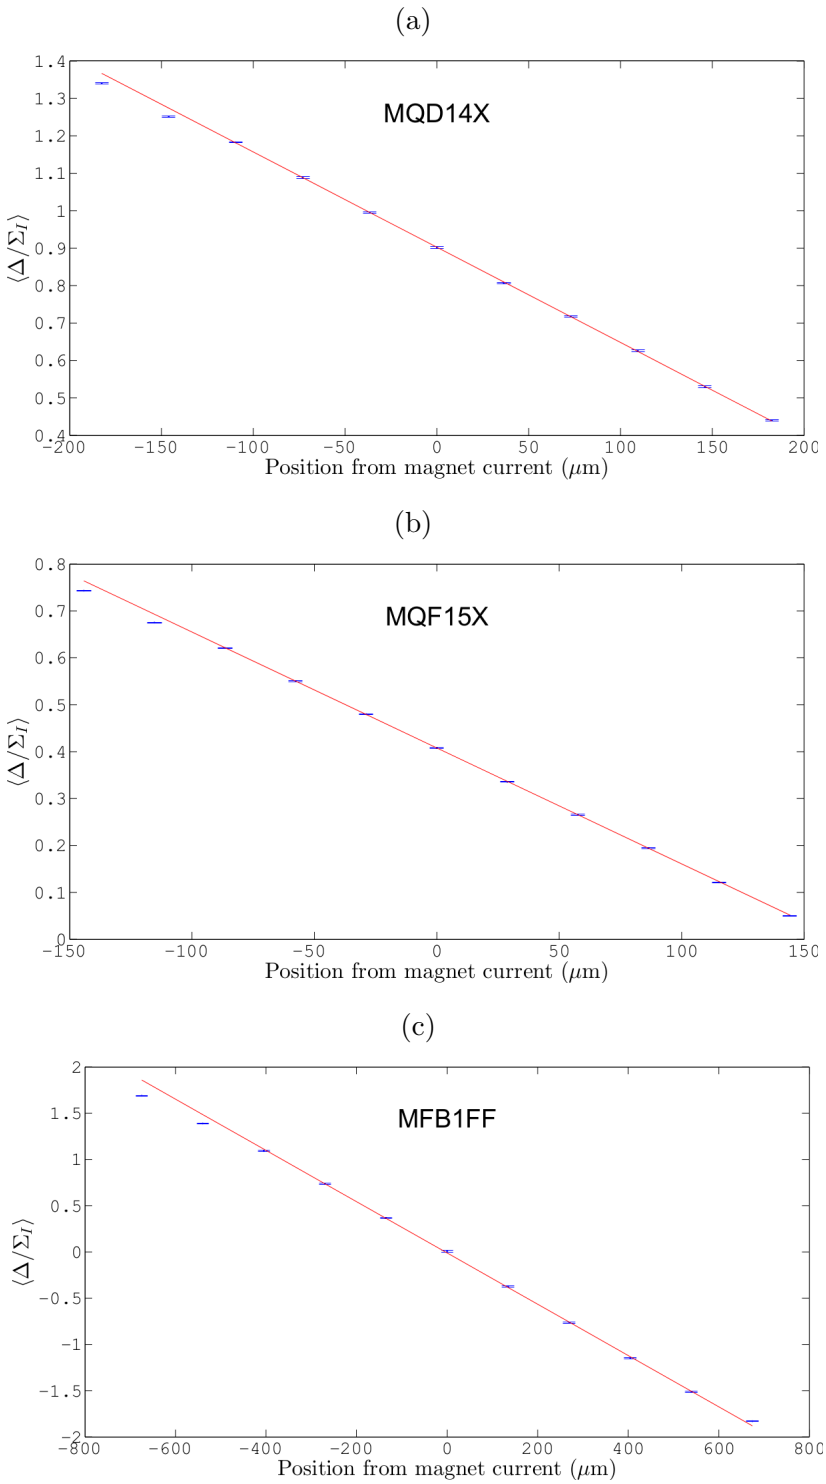

Figure 3.8: Mean value of the uncalibrated vertical position as a function of the expected displacement due to the change in ZV8X current (blue) and the linear χ^2 fit to these data (red) for (a) MQD14X; (b) MQF15X; (c) MFB1FF. Error bars represent the error on the mean.

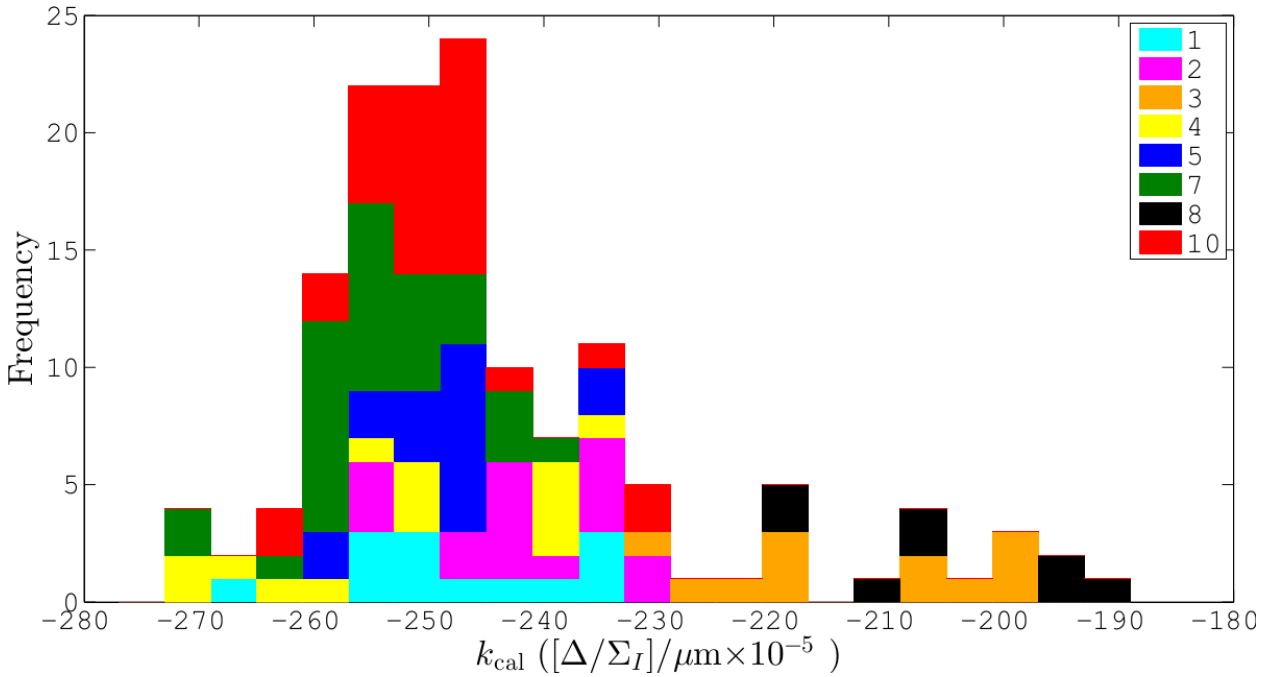

Figure 3.9: Histogram of calibration constants k_y. The bar corresponding to each bin has been coloured to reflect how many times a result within that range was obtained from a given BPM processor module, as indicated by the legend.

improvement in stability that can be achieved using the FONT feedback system.

For a given data set, the ensembles of the measured position at each BPM may be assembled into the vectors $\mathbf{y}_i$, $\mathbf{y}_j$ and $\mathbf{y}_k$, where each of the indices i, j and k denotes a specific BPM of each three-BPM system. The first three-BPM system consists of P1, P2 and P3 and the second is made up of MQD14X, MQF15X and MFB1FF.

In order to estimate the resolution of the three-BPM system, it is necessary first to remove the residual beam position offset at each BPM. For a FONT data set, this is accomplished by subtracting from the position of the nth extraction that of the $(n-1)^{\mathrm{th}}$ extraction. The result gives the change in the position of the beam in each BPM from extraction to extraction; however, assuming that the random component of the beam position follows a Gaussian distribution, this drift removal procedure will have the effect of increasing the beam jitter by a factor of $\sqrt{2}$. Consider a vector $\mathbf{N}$ consisting of n numbers drawn randomly from a normal distribution. From this vector two new vectors may be formed: the vector $\mathbf{S}$, consisting of the first $n-1$ values of $\mathbf{N}$, and the vector $\mathbf{T}$, which consists of the last $n-1$ values of $\mathbf{N}$. The standard deviation of the difference between these vectors, $\mathbf{S}-\mathbf{T}$, is given by the expression:

$$\sigma_{S-T} = \sqrt{\sigma_S^2 + \sigma_T^2 - 2\sigma_S\sigma_T\rho_{ST}} \tag{3.7}$$

where σ_S and σ_T are the standard deviations of the vectors $\mathbf{S}$ and $\mathbf{T}$ and ρ_{ST} is the correlation coefficient between the two vectors. For large values of n the two vectors will have the same standard deviation ($\sigma_S = \sigma_T = \sigma_N$) due to the $n-2$ values that are common to both; however, because of the random nature of the values, the correlation cofficient between the two vectors will be very close to zero. Thus $\sigma_{S-T} = \sqrt{2}\sigma_N$; hence, a factor of $1/\sqrt{2}$ is

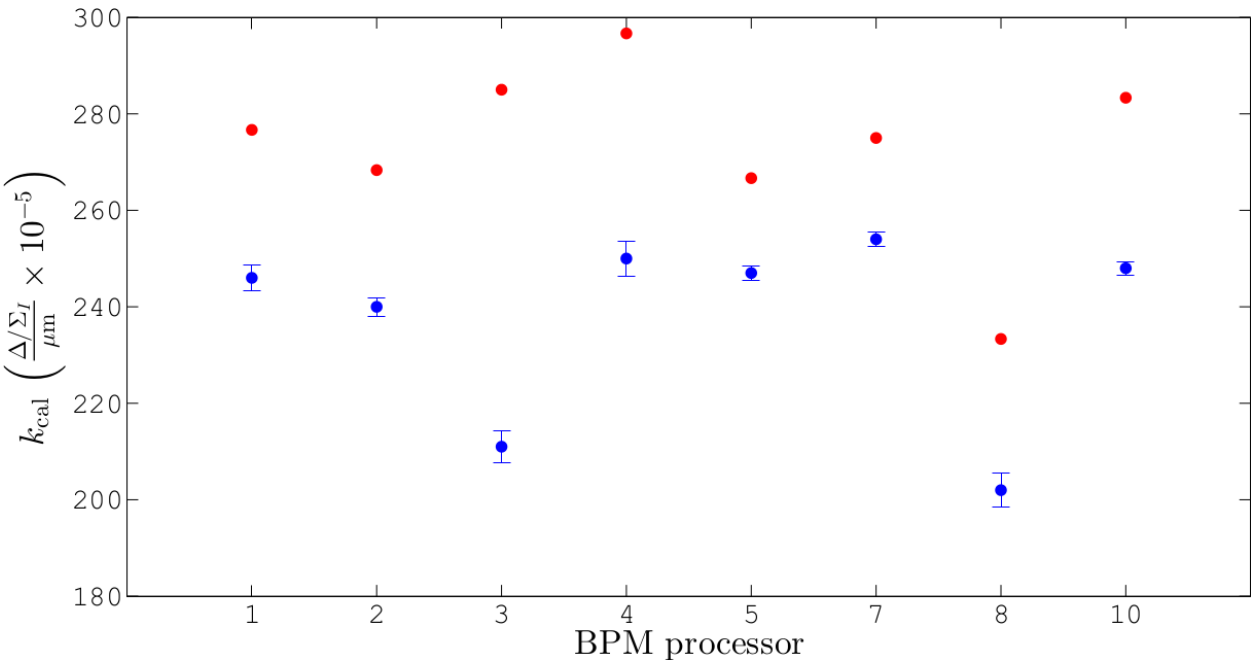

Figure 3.10: Calibration constant as a function of processor number. The mean observed value (blue) is compared to the theoretical value (red).

introduced to the definition of the drift removed position in order to compensate for the random jitter of the beam position.

The effect of the drift removal procedure is illustrated in Figure 3.11. The mean absolute position of the bunch is different for each BPM; by contrast, the drift removed position has a mean of close to zero in each case. For the data set shown, the ratio of the jitter of the drift subtracted position to the jitter of the conventional position is ~ 0.85 for all three BPMs (this ratio would be 1 only for a perfectly normally distributed position). The resolution of the BPM system may then be estimated by predicting the position at one BPM using the measured positions at the other two:

$$\mathbf{y}_i^{\mathrm{pred}} = C_{ij}\mathbf{y}_j^{\mathrm{meas}} + C_{ik}\mathbf{y}_k^{\mathrm{meas}} \tag{3.8}$$

where C_{ij} and C_{ik} are the prediction coefficients used to describe the position at BPM i as a linear combination of the positions at BPMs j and k. This vector of predicted positions is then compared with the measurements for that BPM, $\mathbf{y}_i^{\mathrm{meas}}$, to give a vector of residual position $\mathbf{R}$:

$$\mathbf{R} \equiv \mathbf{y}_i^{\mathrm{meas}} - \mathbf{y}_i^{\mathrm{pred}} = \mathbf{y}_i^{\mathrm{meas}} - (C_{ij}\mathbf{y}_j^{\mathrm{meas}} + C_{ik}\mathbf{y}_k^{\mathrm{meas}}) \tag{3.9}$$

The standard deviation of the vector $\mathbf{R}$, σ_R, can be related to the standard deviation of the measured positions:

$$\sigma_R^2 = \sigma_{y_i}^2 + C_{ij}^2\sigma_{y_j}^2 + C_{ik}^2\sigma_{y_k}^2 \tag{3.10}$$

If the resolutions of the BPMs are assumed to be identical ($\sigma_i = \sigma_j = \sigma_k = \sigma_{\mathrm{res}}$), the expression for the resolution σ_{res} is:

$$\sigma_{\mathrm{res}} = \frac{\sigma_R}{\sqrt{1 + C_{ij}^2 + C_{ik}^2}} \tag{3.11}$$

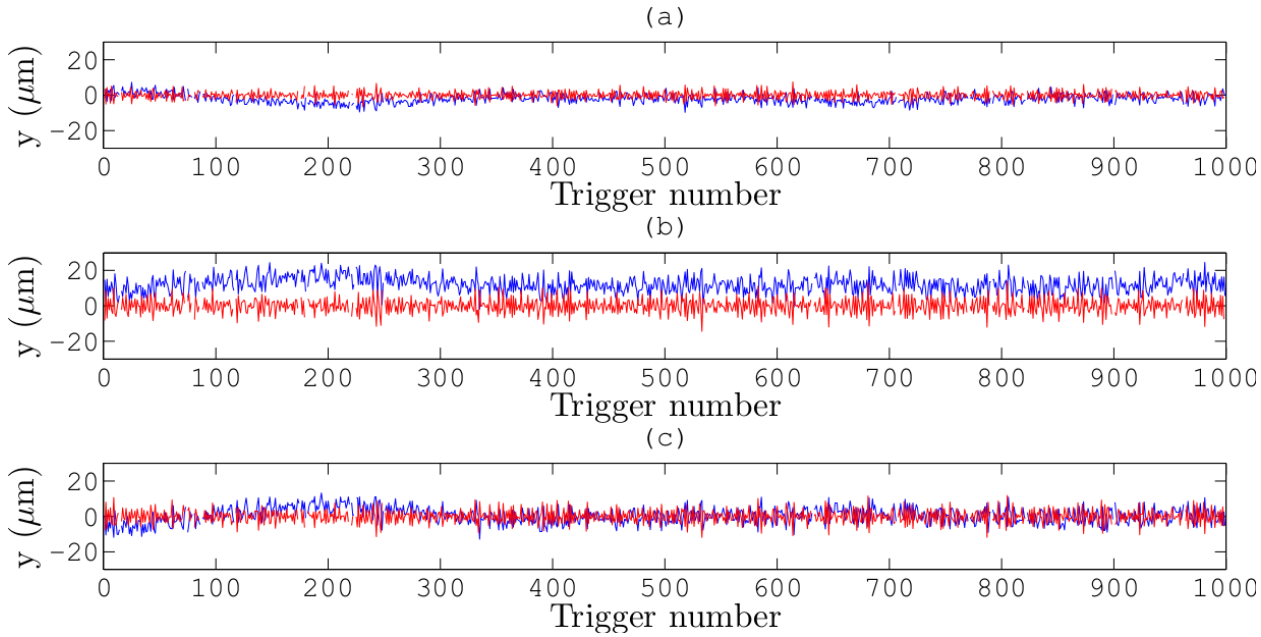

Figure 3.11: The absolute position (blue) and the drift removed position (red) for the BPMs (a) P1; (b) P2; (c) P3.

As any assigment of the three BPMs to the indices i, j and k is equally valid, it is possible to obtain three unique estimates for the resolution of the system from any single set of observations. However, as the BPMs were assumed to have identical resolution, in each case the answer is an estimate for the resolution of the system as a whole, not of that BPM for which the position is being predicted.

It is now necessary to provide a description of the prediction coefficients. For each assignment of BPMs to indices, two prediction coefficients describe how the position at one BPM may be expressed as a linear combination of the positions at the other two. Two methods for determining the values of the prediction coefficients will be described here: the first, the model method, relies solely on the beam transfer matrices. The second, the fit method, defines the prediction coefficients in order to minimize the length of the vector $\mathbf{R}$ (Equation 3.9); in this case, their values are calculated by performing a least squares fit to the residuals of the observed position data.

3.3.1 Model method

In the model method, the prediction coefficients are functions of the transfer matrix elements. In principle, knowledge of the appropriate transfer matrix allows the beam conditions at any location to be extrapolated to any other location along the beamline: the beam position at a given location is a linear function of the position and angle of the beam at some other location. In this expression, the position and angle coefficients are given by the appropriate elements of the transfer matrix between the two locations (cf. Equation 3.5). It is thus possible to use the positions y_{P1}, y_{P2} and y_{P3} and the transfer matrices ${}^{\mathrm{P1}}_{\mathrm{P2}}M$ and ${}^{\mathrm{P2}}_{\mathrm{P3}}M$ to

construct an equation of the type:

$$A_{\mathrm{P1}} y_{\mathrm{P1}} + A_{\mathrm{P2}} y_{\mathrm{P2}} + A_{\mathrm{P3}} y_{\mathrm{P3}} = 0 \tag{3.12}$$

which, with suitable modification, will describe the position at any one BPM as a function of the position at the other two. The A parameters take the values:

$$\begin{aligned} A_{\mathrm{P1}} &= {}^{\mathrm{P1}}_{\mathrm{P2}}M_{21} - \frac{{}^{\mathrm{P1}}_{\mathrm{P2}}M_{22}{}^{\mathrm{P1}}_{\mathrm{P3}}M_{11}}{{}^{\mathrm{P1}}_{\mathrm{P3}}M_{12}} \\ A_{\mathrm{P2}} &= \frac{{}^{\mathrm{P2}}_{\mathrm{P3}}M_{11}}{{}^{\mathrm{P2}}_{\mathrm{P3}}M_{12}} \\ A_{\mathrm{P3}} &= \frac{{}^{\mathrm{P1}}_{\mathrm{P2}}M_{22}}{{}^{\mathrm{P1}}_{\mathrm{P3}}M_{12}} - \frac{1}{{}^{\mathrm{P2}}_{\mathrm{P3}}M_{12}} \end{aligned} \tag{3.13}$$

As an example, when using the model method to predict the position at P1 it is found that the prediction coefficients are given by $C_{\mathrm{P1P2}} = -A_{\mathrm{P2}}/A_{\mathrm{P1}}$ and $C_{\mathrm{P1P3}} = -A_{\mathrm{P3}}/A_{\mathrm{P1}}$. Using these, and similar expressions obtained for the two other BPM permutations, the distributions of position for each FONT BPM for the data set of Figure 3.4 are given in Figure 3.12. These results show that the distribution at P1 as predicted using measurements at P2 and P3 is considerably narrower than that which is actually observed, suggesting that there is some mechanism causing a degradation of the resolution at P1. Ultimately, such a mechanism is most likely due to a construction error inherent to that BPM. As the model method is completely constrained, a large spurious component of the position at P1 causes equivalent broadening of the predicted distributions at P2 and P3. Thus the model method provides a single unique estimate of the resolution of the system and the BPM used to create the vector of residuals makes no difference to the result; the residual distributions of plots (d), (e) and (f) are identical once the numerical factors calculated from the prediction coefficients are taken into account (Figure 3.14).

The residual is thus scaled and then binned in order to produce a histogram to which a Gaussian curve is fitted. The width of this Gaussian gives an estimate for the resolution of the BPM system and the error is taken as the standard deviation of the fit coefficients divided by the square root of the number of bins. For this data set, the resolution as predicted by the model method is $\sigma_{\mathrm{res}} = 2.62 \pm 0.91$ μm. Given that the feedback system is known to reduce the jitter on the position down to the sub-micron level [46], it is apparent that the model method predicts a worse resolution for the BPM system than is actually the case.

3.3.2 Fit method

The problem with the model method arises because the weighting given to each BPM is solely determined by the transfer matrices. As there is good reason to believe that P1 has worse resolution than the other two BPMs, an alternative method that allows the positions at P2 and P3 to take a more significant role in determining the predicted position is desired.

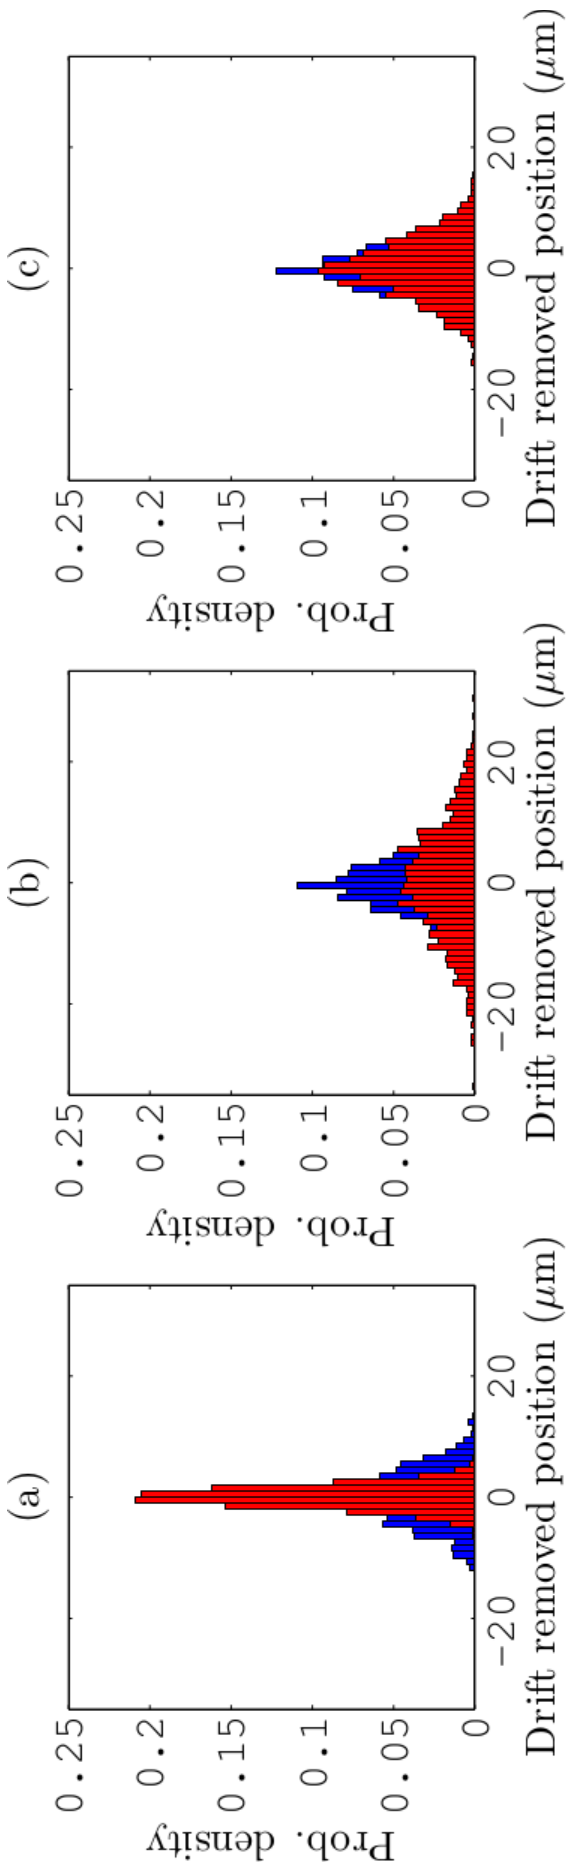

Figure 3.12: Comparison of the measured distribution of the position (blue) with the prediction obtained using the model method (red) for the BPMs (a) P1; (b) P2; (c) P3. The distribution of the residuals obtained by subtracting each prediction from each measurement is plotted underneath (green) for (d) P1; (e) P2; (f) P3.

The fit method accomplishes this by obtaining values for the prediction coefficients from the solution in the least-squares sense to the equation:

$$\mathbf{y}_i = \begin{bmatrix} \mathbf{y}_j & \mathbf{y}_k \end{bmatrix} \begin{bmatrix} C_{ij} \\ C_{ik} \end{bmatrix} \tag{3.14}$$

and the expected effect of the least-squares fit is to limit the contribution of the P1 position to the predicted positions in the other two. The minimization is performed using the MATLAB implementation of matrix left division [58]. A comparison of the prediction coefficients obtained for the two methods is given in Table 3.5. As expected, the fit method reduces

Prediction coefficient	Model method	Fit method
C_{P1P2}	0.56	1.39
$-C_{\mathrm{P1P3}}$	1.07	1.77
C_{P2P1}	1.77	0.07
C_{P2P3}	1.90	1.12
$-C_{\mathrm{P3P1}}$	0.93	0.07
C_{P3P2}	0.53	0.86

Table 3.5: Prediction coefficients as calculated using both model and fit methods.

the magnitude of C_{P2P1} and C_{P3P1} to close to zero, reflecting that the position measured at P1 is a poor predictor for the location of the beam in the two downstream BPMs. The distributions of the predicted positions for each BPM using the fit method are given in Figure 3.13. It is evident that the fit method results in a dramatic improvement of the predicted positions at P2 and P3. The same cannot be said for the prediction at P1; however, once the residuals are scaled by the appropriate numerical factor, all three estimates of the resolution obtained using the fit method represent an improvement on the value given by the model method. The scaled residuals used to estimate the resolution are plotted in Figure 3.14. Plots Figure 3.14(a), Figure 3.14(b) and Figure 3.14(c) are, modulo inversion of the x-axis, identical as previously noted. On the other hand, the fit method produces three different distributions and those of plots Figure 3.14(e) and Figure 3.14(f) are significantly narrower than that of Figure 3.14(d). The position measured at P1 has a large component that is not linearly related to the actual bunch positions at the other BPMs and as a result it is difficult to accurately predict it from the position at the other two BPMs alone. The estimates for the resolution using both methods are given in Table 3.6. For the case where

	Estimated resolution (μm)	
BPM of prediction	Model method	Fit method
P1	2.62 ± 0.91	1.58 ± 0.55
P2	2.62 ± 0.91	0.63 ± 0.20
P3	2.62 ± 0.91	0.62 ± 0.20

Table 3.6: Resolution predicted using both model and fit methods.

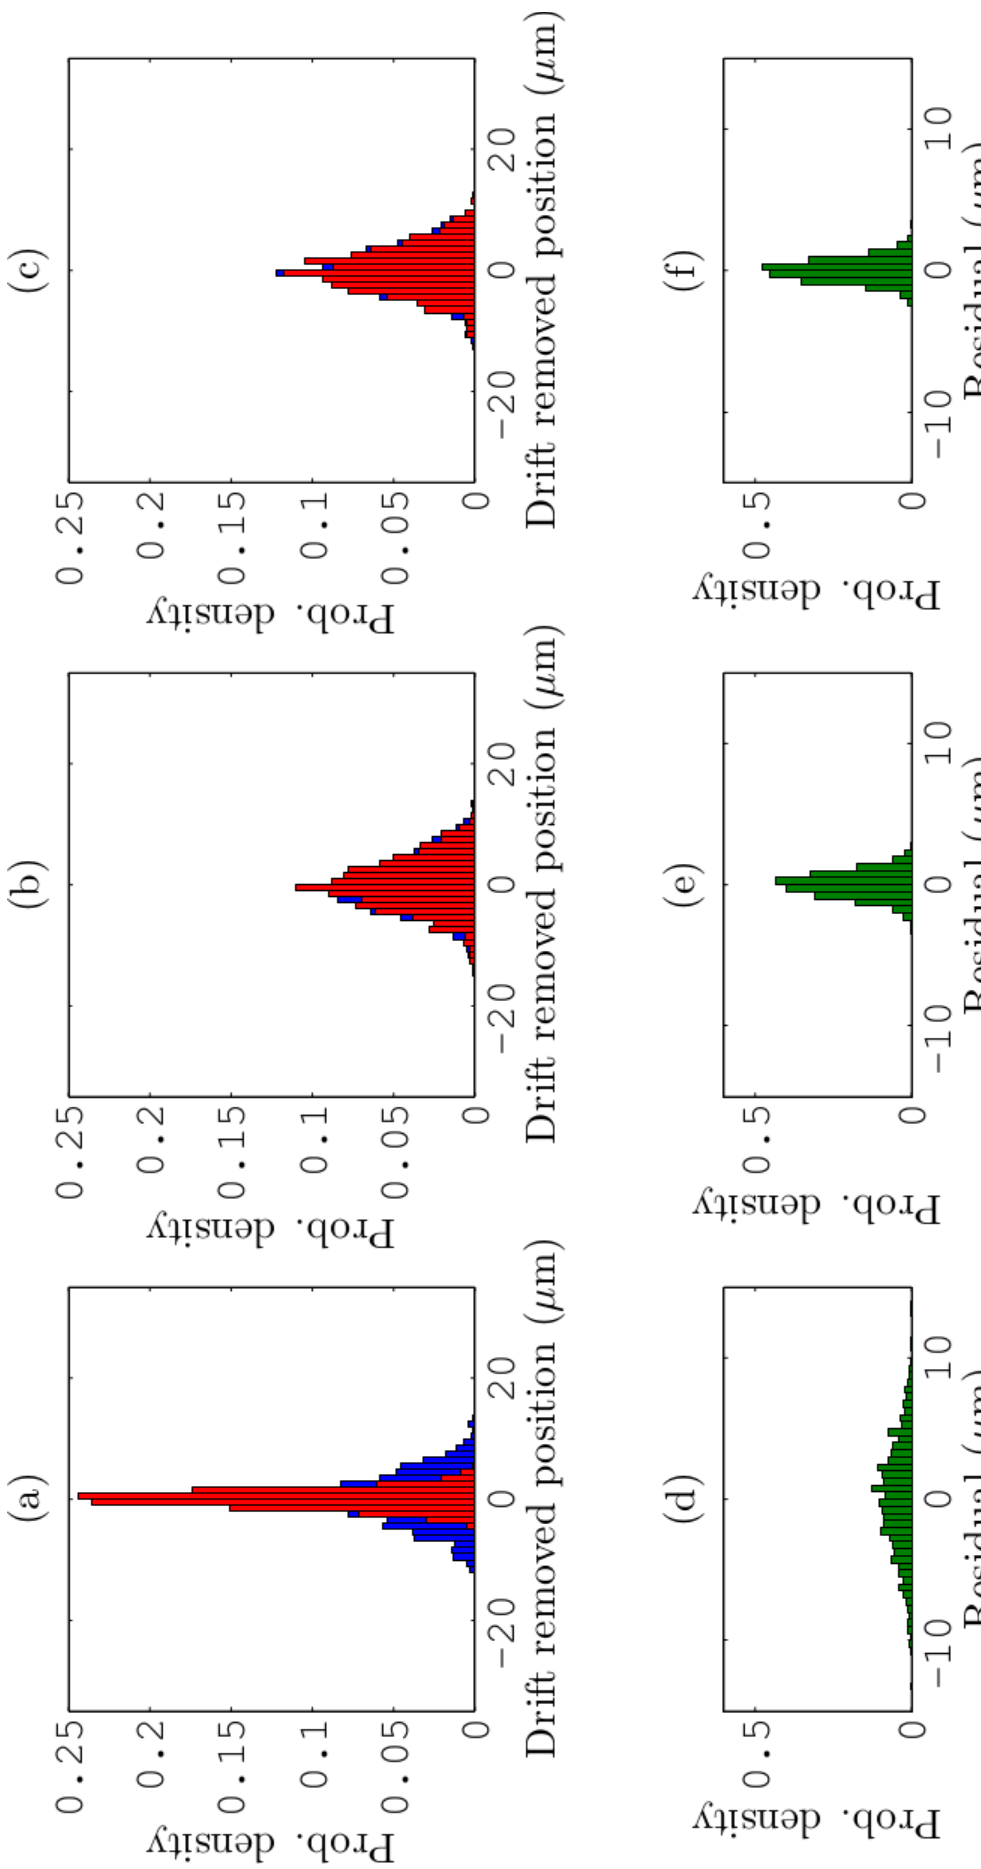

Figure 3.13: Comparison of the measured distribution of the position (blue) with the prediction obtained using the fit method (red) for the BPMs (a) P1; (b) P2; (c) P3. The distribution of the residuals obtained by subtracting each prediction from each measurement is plotted underneath (green) for (d) P1; (e) P2; (f) P3.

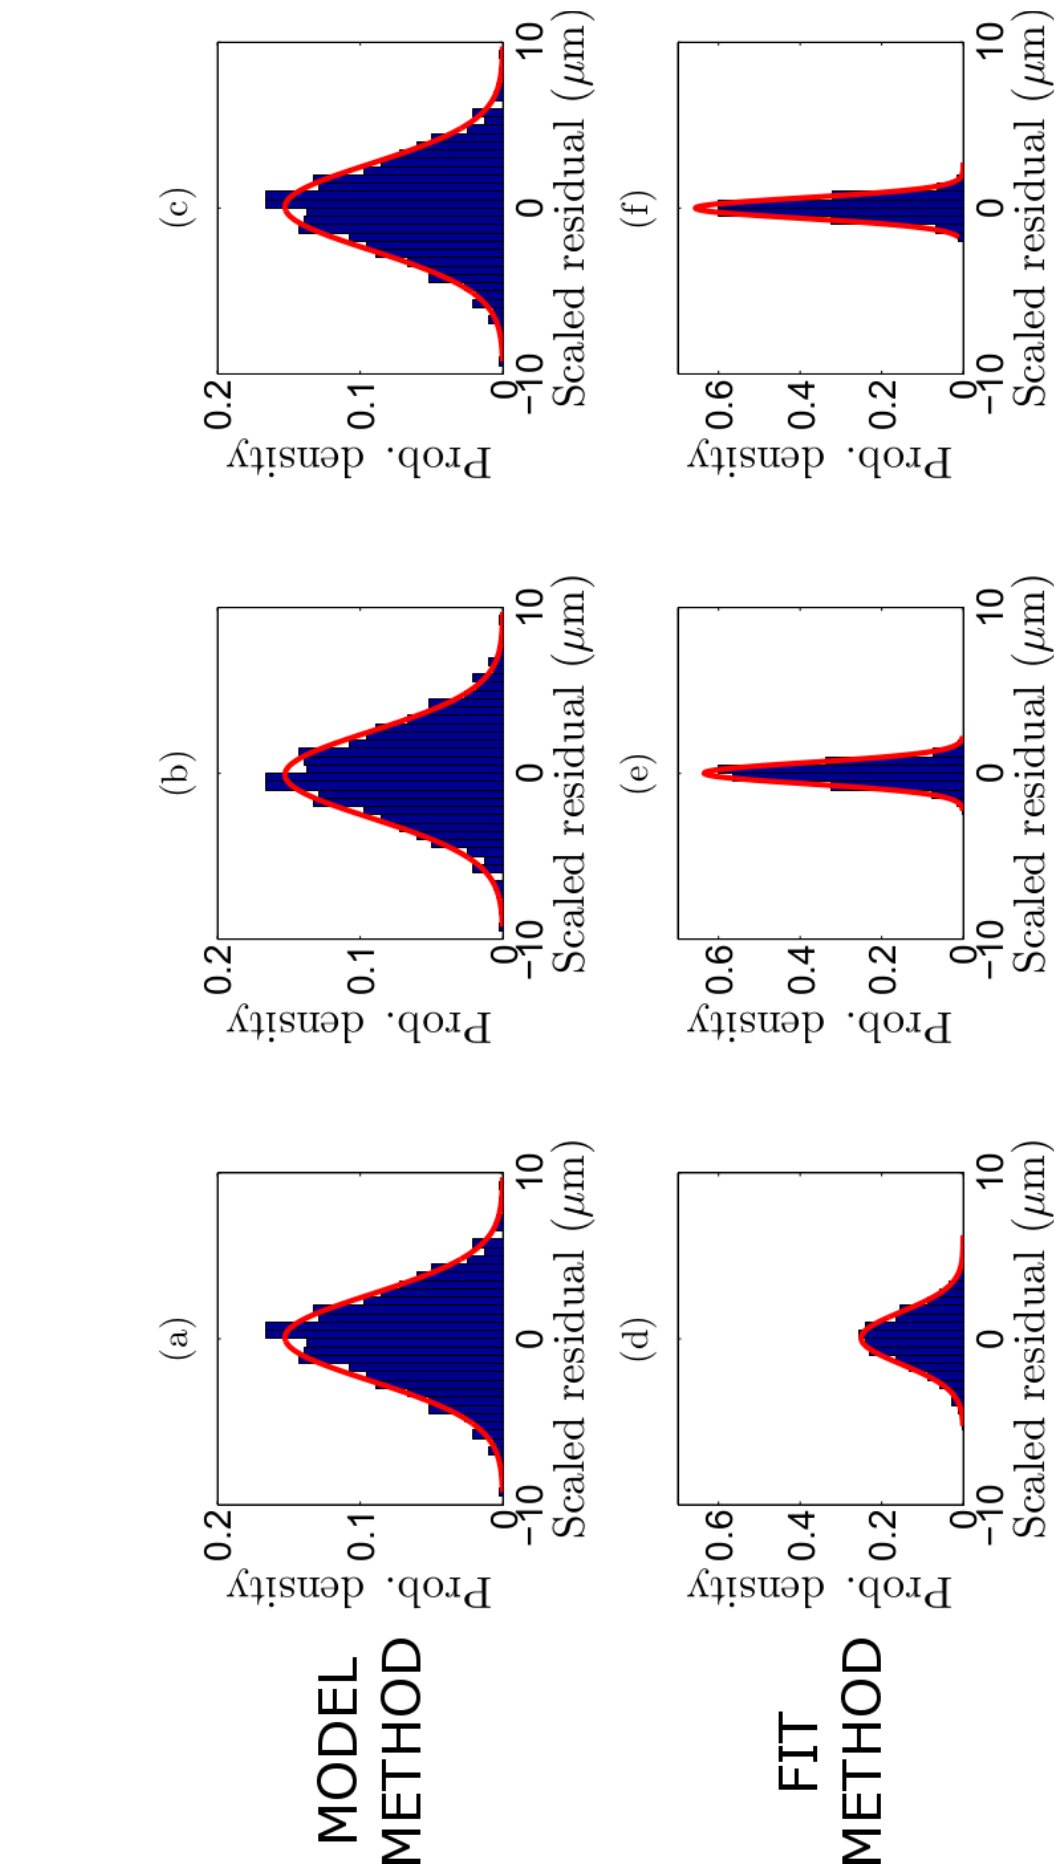

Figure 3.14: Distribution of the scaled residuals (blue) and the Gaussian fit used to estimate the BPM resolution (red). The model method was used to predict the position for (a) P1; (b) P2; (c) P3 and the fit method for (d) P1; (e) P2; (f) P3.

all three BPMs have the same low resolution, both the model method and the fit method should return similar results. As this is not the case, it is necessary to consider possible causes for the poor resolution at P1.

3.4 Effect of LO phase offset on measured position

The beam of the ATF damping ring radiates synchrotron light each time it is deflected leading to a loss in beam momentum. This momentum loss is primarily in the direction of beam travel but the distribution of the trajectories of emitted photons is such that there is a turn-by-turn reduction of the transverse momentum. The loss of transverse momentum is the goal of the damping ring; to maintain the beam orbit, the longitudinal loss of momentum must be compensated for using an accelerating cavity.

Consider a particle that arrives at the accelerating cavity later than the synchronous particle which is always accelerated by the same amount every turn. Such a particle will receive less energy from the cavity and as a result will take a shorter path around the machine due to the smaller radius of curvature it experiences in the arc sections. This particle will arrive back at the cavity earlier than it did for the previous turn until eventually it arrives earlier than the synchronous particle, at which point it will receive more energy from the cavity and follow a longer path, and so on. The difference in arrival time of the particle relative to the synchronous particle will oscillate from positive to negative; this longitudinal oscillation is known as the synchrotron oscillation [7].

As the mean phase of a bunch of particles is not exactly equal to the synchronous phase, the synchrotron oscillation causes the arrival time of the extracted bunch at the FONT measurement BPMs to vary from trigger to trigger. As mentioned in Section 2.4.2, the phase of the LO input to each processor module is set to match the phase of the bunch at that BPM in order to maximize the amplitude of the Σ_I output signal. As the LO is derived from the machine RF but the bunch phase oscillates relative to this, the LO phase can only be matched to the bunch on average. It is thus worth considering what the effect of a non-optimal LO phase is on the position measured according to the FONT system.

3.4.1 Effect of LO phase on position for an ideal processor

Ideally, a phase offset between the bunch and the LO would have no effect on the measured position (although as a large offset would negatively impact the signal levels it would be expected to have a detrimental effect on the resolution). As any arbitrary signal can be represented using Fourier analysis as a sum of trigonometric functions, consider a single frequency component of the BPM processor input signals $V_A(t)$ and $V_B(t)$, denoted $v_A(\omega)$ and $v_B(\omega)$:

$$\begin{aligned} v_A(\omega) &= a_A(\omega)\sin(\omega t) \\ v_B(\omega) &= a_B(\omega)\sin(\omega t) \end{aligned} \tag{3.15}$$

where ω is the angular frequency and $a_A(\omega)$ and $a_B(\omega)$ are the respective amplitudes of that frequency in the Fourier decompositions of $V_A(t)$ and $V_B(t)$. As this is an ideal system, there is no phase difference between the two signals. The ideal resistive coupler would add together these two inputs to form $v_{\text{coupler}}(\omega)$; the ideal 180° hybrid would first shift the phase of one input by 180° to form the intermediate signal $v'_B(\omega)$ and then add this result to the other input, thereby forming $v_{\text{hybrid}}(\omega)$. Using the inputs of Equation 3.15:

$$\begin{aligned} v_{\text{coupler}}(\omega) &= v_A(\omega) + v_B(\omega) = [a_A(\omega) + a_B(\omega)]\sin(\omega t) \\ v_{\text{hybrid}}(\omega) &= v_A(\omega) + v'_B(\omega) = [a_A(\omega) - a_B(\omega)]\sin(\omega t) \end{aligned} \tag{3.16}$$

These intermediate signals then travel to the mixers. A frequency mixer has two inputs, designated here as the LO input and the RF input, which is either of the $v_{\text{coupler}}(\omega)$ or $v_{\text{hybrid}}(\omega)$ signals. The output of the mixer is the product of the two inputs. The LO input to the processor is given by:

$$V_{\text{LO}} = a_{\text{LO}}\sin(\omega_{\text{LO}}t + \psi_{\text{LO}}) \tag{3.17}$$

where a_{LO} is the amplitude of the LO signal, ω_{LO} is the angular frequency, and the phase offset between the LO and the RF inputs is represented as ψ_{LO}. The outputs of the three mixers are thus given by:

$$\begin{aligned} v_{\Sigma_I-\text{mixer}}(\omega) &= v_{\text{coupler}}(\omega) \times a_{\text{LO}}\sin(\omega_{\text{LO}}t + \psi_{\text{LO}}) \\ v_{\Sigma_Q-\text{mixer}}(\omega) &= v_{\text{coupler}}(\omega) \times a_{\text{LO}}\cos(\omega_{\text{LO}}t + \psi_{\text{LO}}) \\ v_{\Delta-\text{mixer}}(\omega) &= v_{\text{hybrid}}(\omega) \times a_{\text{LO}}\sin(\omega_{\text{LO}}t + \psi_{\text{LO}}) \end{aligned} \tag{3.18}$$

where a 90° phase shift has been applied to the LO for the Σ_Q mixer. Each signal is the product of two sine functions and so they may be rewritten using standard trigonometric relations as the sum of two sine functions whose frequencies are given by the sum and difference of the original frequencies. The low-pass filters in the FONT processors suppress the high frequency terms, in which case the outputs become:

$$\begin{aligned} v_{\Sigma_I}(\omega) &= \tfrac{1}{2}a_{\text{LO}}[a_A(\omega) + a_B(\omega)]\cos([\omega_{\text{LO}} - \omega]t + \psi_{\text{LO}}) \\ v_{\Sigma_Q}(\omega) &= \tfrac{1}{2}a_{\text{LO}}[a_A(\omega) + a_B(\omega)]\sin([\omega_{\text{LO}} - \omega]t + \psi_{\text{LO}}) \\ v_{\Delta}(\omega) &= \tfrac{1}{2}a_{\text{LO}}[a_A(\omega) - a_B(\omega)]\cos([\omega_{\text{LO}} - \omega]t + \psi_{\text{LO}}) \end{aligned} \tag{3.19}$$

For this ideal processor the ratios of the Δ and Σ_Q signals to the Σ_I signals are expected to behave as:

$$\begin{aligned} \frac{v_{\Delta}(\omega)}{v_{\Sigma_I}(\omega)} &= \frac{a_A(\omega) - a_B(\omega)}{a_A(\omega) + a_B(\omega)} \\ \\ \frac{v_{\Sigma_Q}(\omega)}{v_{\Sigma_I}(\omega)} &= \tan([\omega_{\text{LO}} - \omega] + \psi_{\text{LO}}) \sim \tan\psi_{\text{LO}} \end{aligned} \tag{3.20}$$

For a given frequency component the ratio Δ/Σ_I has no dependence on ψ_{LO} while the ratio Σ_Q/Σ_I goes as the tangent. In practice, the band-pass filters on V_{hybrid} and V_{coupler} ensure that only a narrow band of frequency components with $\omega \sim \omega_{\text{LO}}$ survive to the inputs of the mixers and so the behaviour of the processor can be roughly approximated as the behaviour of the frequency component $\omega = \omega_{\text{LO}}$.

3.4.2 Effect of a path length difference at the processor inputs

The definition of the second stripline input to the processor in Equation 3.15 can be modified to include a constant phase offset Ψ between the two inputs. Such a phase shift arises as a consequence of the difference in path length for the signals $V_A(t)$ and $V_B(t)$ as they travel from the striplines to the processor inputs; this could be due to an error in construction of the processor itself or if the processor were attached to the BPM with a pair of poorly matched cables. In this case:

$$v_B(\omega) = a_B(\omega)\sin(\omega t + \Psi) \equiv a_B(\omega)\left[\cos(\Psi)\sin(\omega t) + \sin(\Psi)\cos(\omega t)\right] \tag{3.21}$$

and the phase shift is seen to introduce a $\cos(\omega t)$ term, or quadrature component, to the expression for $v_B(\omega t)$. The intermediate signals then become:

$$\begin{aligned} v_{\text{coupler}}(\omega) &= \left[a_A(\omega) + \cos(\Psi)a_B(\omega)\right]\sin(\omega t) + \sin(\Psi)a_B(\omega)\cos(\omega t) \\ v_{\text{hybrid}}(\omega) &= \left[a_A(\omega) - \cos(\Psi)a_B(\omega)\right]\sin(\omega t) - \sin(\Psi)a_B(\omega)\cos(\omega t) \end{aligned} \tag{3.22}$$

Assuming that the phase shift between the two inputs is small (the cables for each BPM processor were selected so as to ensure this), $\cos(\Psi) \sim 1$ and $\sin(\Psi) \sim \Psi$. Thus the amplitude of the quadrature component depends linearly on Ψ.

$$\begin{aligned} v_{\text{coupler}}(\omega) &\sim \left[a_A(\omega) + a_B(\omega)\right]\sin(\omega t) + \Psi a_B(\omega)\cos(\omega t) \\ v_{\text{hybrid}}(\omega) &\sim \left[a_A(\omega) - a_B(\omega)\right]\sin(\omega t) - \Psi a_B(\omega)\cos(\omega t) \end{aligned} \tag{3.23}$$

For the situtation where the beam is well centered, $a_A(\omega) \sim a_B(\omega)$. The quadrature component is thus expected to be the dominant contribution to $v_{\text{hybrid}}(\omega)$ and a negligible component to $v_{\text{coupler}}(\omega)$. Neglecting the quadrature component, the expression for $v_{\text{coupler}}(\omega)$ is identical to the ideal case and therefore the behaviour of the ratio Σ_Q/Σ_I post-mixer is not expected to change significantly. However, including this component in the expression for $v_{\text{hybrid}}(\omega)$ results in a new expression for $v_\Delta(\omega)$:

$$v_\Delta(\omega) = \tfrac{1}{2}a_{\text{LO}}\left\{\left[a_A(\omega) - a_B(\omega)\right]\cos(\left[\omega_{\text{LO}} - \omega\right]t + \psi_{\text{LO}}) - \Psi a_B(\omega)\sin(\left[\omega_{\text{LO}} - \omega\right]t + \psi_{\text{LO}})\right\} \tag{3.24}$$

The expected dependence of the ratio Δ/Σ_I is then:

$$\frac{v_\Delta(\omega)}{v_{\Sigma_I}(\omega)} = \frac{a_A(\omega) - a_B(\omega)}{a_A(\omega) + a_B(\omega)} - \frac{\Psi a_B(\omega)}{a_A(\omega) + a_B(\omega)}\tan(\left[\omega_{\text{LO}} - \omega\right]t + \psi_{\text{LO}}) \tag{3.25}$$

For a well-centred beam, then, the measured position as a function of LO phase is expected to vary as $\tan(\psi_{\text{LO}})$. Furthermore, the variation with phase is expected to increase linearly with the path length difference associated with a particular combination of BPM and analogue processor module along with the cables connecting the two.

3.4.3 LO phase scans

The optimal LO phase setting for a given BPM processor is that which maximizes the Σ_I bunch sample value. Equation 3.19 states that Σ_I is proportional to the cosine of ψ_{LO}. It

is thus difficult to determine the optimal phase by simple observation of Σ_I while ψ_{LO} is varied as Σ_I becomes increasingly insensitive to changes in phase as the optimum phase is approached. Conversely, the sensitivity of Σ_Q to a change in phase increases as the optimum phase is approached. Historically, this was the purpose of the Σ_Q output of each processor: to provide a convenient way to phase each LO signal. However, this use of Σ_Q is deprecated as the phase shift applied to the LO for the Σ_Q mixer is known to differ slightly from the ideal value of 90°.

Instead, the optimal phase is determined by first activating the switchable 90° phase shift at the start of the LO distribution. The correct LO phase setting for each channel is then determined either by manual adjustment of each phase shifter while monitoring the corresponding Σ_I waveform on the DAQ in real time ("by eye") or by using an advanced feature of the DAQ to scan all the phase shifters across their entire range and then fit the average Σ_I bunch sample value as a function of phase shifter setting. In either case, once the Σ_I signals have been minimized, the switchable 90° phase shift is removed so that they obtain their maximal values. Figure 3.15 shows the result that such a scan of the individual

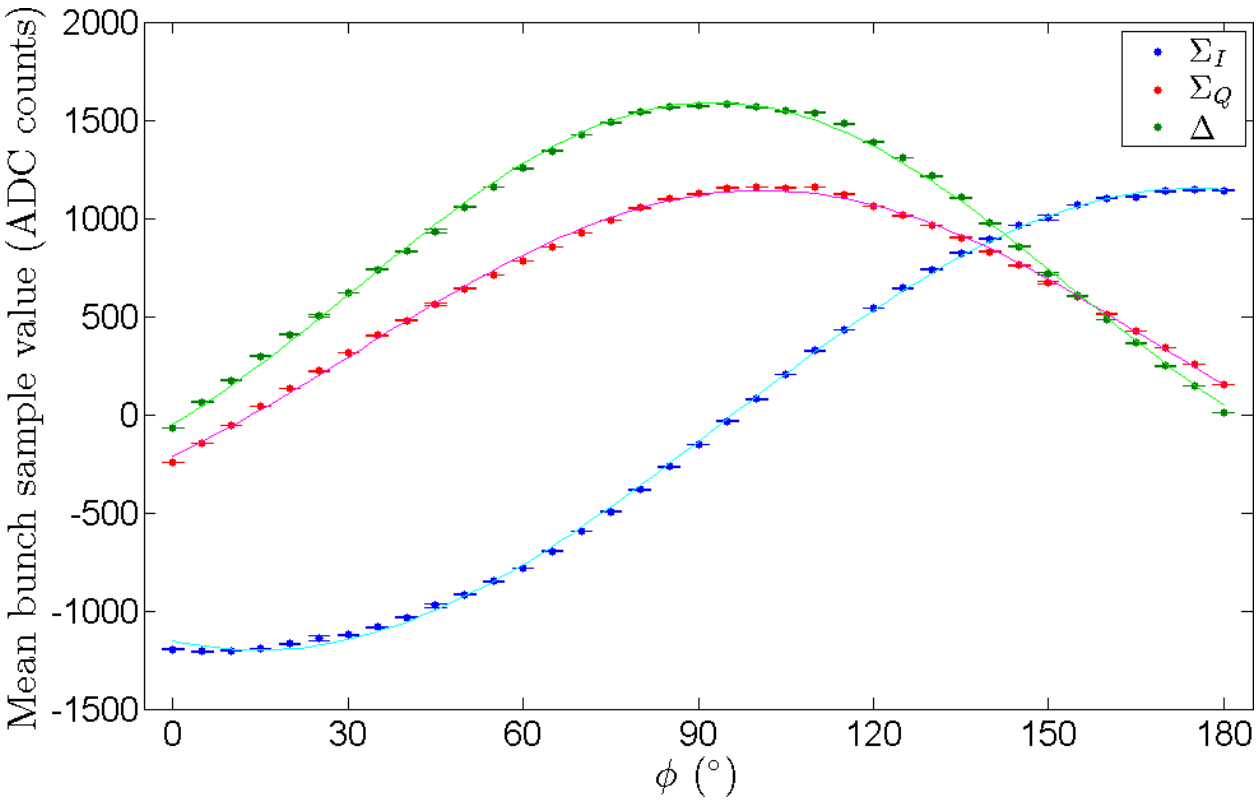

Figure 3.15: Value of the bunch samples for P1 as a function of master LO phase. Blue: $\langle\Sigma_I\rangle$; cyan: sinusoidal fit to Σ_I data; red: $\langle\Sigma_Q\rangle$; magenta: sinusoidal fit to Σ_Q data; dark green: $\langle\Delta\rangle$; light green: sinusoidal fit to Δ data.

channel LO phase setting has on the bunch sample values from a single BPM (MFB1FF). The LO phase is scanned over the entire range of the phase shifter: from 0° to 180°. Each point represents the average bunch sample value over 100 triggers. The behaviour of each signal is as predicted in the previous section; all three signals exhibit a sinusodial dependence

on the phase setting. The maximum Σ_I value is found in the region $\phi \sim 175°$; the Σ_Q data has equal amplitude to the Σ_I data but leads it by 90° and the Δ data displays the expected behaviour for a signal dominated by the quadrature component as it is also 90° out of phase with the Σ_I data. In this case, the amplitude of the Δ signal conveys the magnitude of the path length difference for the processor. In the absence of such a path length (and hence the quadrature component) the Δ data would be expected to be in phase with the Σ_I data with an amplitude determined by the position of the beam. Figure 3.16 shows the data from the same scan but now considering the uncalibrated position (Δ/Σ_I) and the phase offset (Σ_Q/Σ_I) reported at each setting by all six of the BPMs monitored using FONT hardware. In both cases a function of the form:

$$p_1 \tan(p_2 \Psi_{\mathrm{LO}} + p_3) + p_4 \tag{3.26}$$

has been fitted to the data from each BPM. Figure 3.16(b) shows that Σ_Q/Σ_I provides a very precise measure of the phase offset between the bunch and the LO, although the amount of variation in where the fit crosses the x-axis points to a phase offset between the Σ_I and Σ_Q outputs (Ψ_{Σ_Q}) that varies for each processor. These offsets are calculated from each fit and listed in Table 3.7. Ideally the offset would be 90° in each case; the primary motivation for

BPM	Processor	Ψ_{Σ_Q} (°)
P1	5	87.31 ± 0.05
P2	7	95.69 ± 0.06
P3	10	93.43 ± 0.05
MQD14X	2	90.09 ± 0.05
MQF15X	4	95.19 ± 0.05
MFB1FF	1	86.42 ± 0.05

Table 3.7: Phase offset between the Σ_I and Σ_Q outputs for the processor on each BPM.

the switchable 90° phase shift on the master LO is that it provides a more accurate 90° phase shifted version of the Σ_I signal for the phase optimization procedure than the Σ_Q signal from the processor itself. The phase offset between Σ_I and Σ_Q is not related to the phase offset between the inputs to the 180° hybrid. Figure 3.16(a) shows the effect of a phase offset on the position measured for each BPM and suggests a possible cause of the poor resolution at P1: a phase offset at P1 has a large effect on the Δ signal. The degree to which each BPM reports a phase offset as a change in position can be calculated by performing a linear least-squared fit to the region $|\Psi_{\mathrm{LO}}| < 10°$. The gradient of the fit gives the phase sensitivity expressed in units of $(\frac{\Delta}{\Sigma_I})/°$ which can then be divided by the calibration constant for that BPM to give the result in units of μm/°; these results are given in Table 3.8. Using a data set from the 19th April 2012 taken after the system had been optimized, the standard deviation of the phase offset according to the ratio Σ_Q/Σ_I is found to be around 0.6°. The phase sensitivity of Table 3.8 enables this phase jitter to be converted to a jitter in the measured position, $\sigma_{y_{\Psi_{\mathrm{LO}}}}$. The results are given in Table 3.9 and suggest that much of the observed jitter in the measured position at P1 arises as a result of jitter in the bunch phase. This provides a very reasonable explanation for why it is so difficult to predict the position at P1 using just the measured positions at P2 and P3.

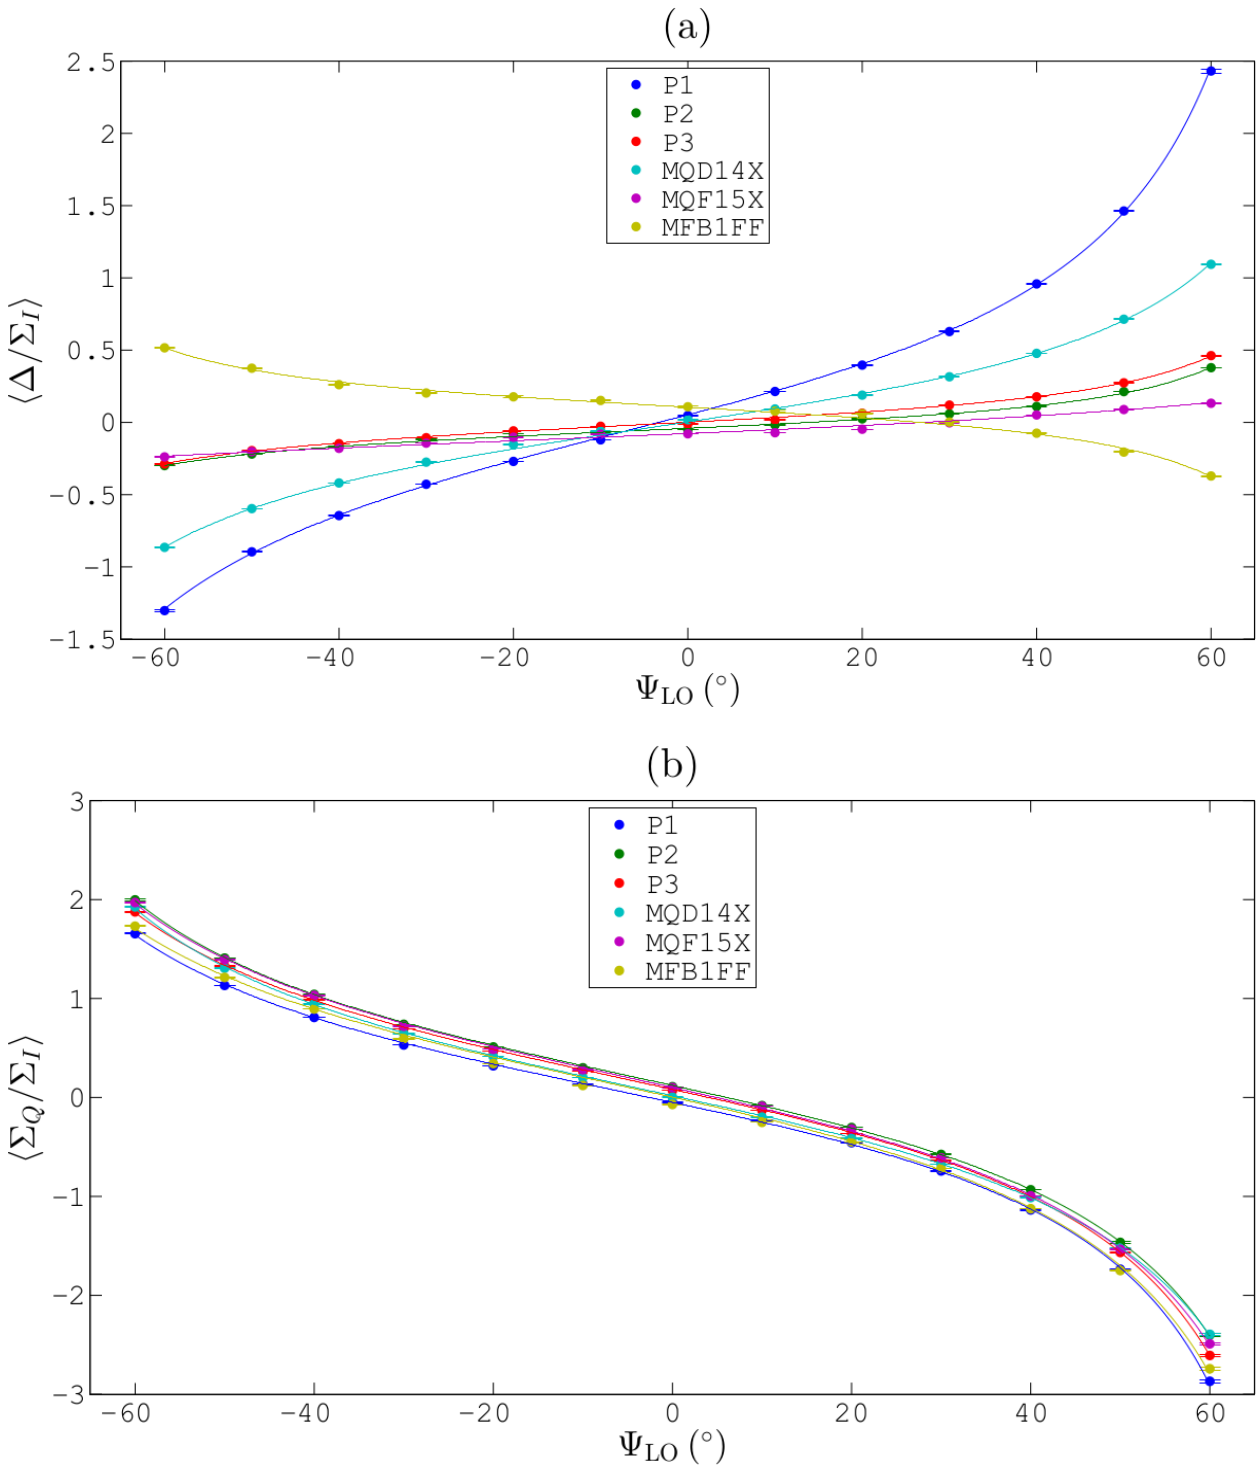

Figure 3.16: Behaviour of the ratios (a) Δ/Σ_I and (b) Σ_Q/Σ_I as a function of master LO phase for the six BPMs monitored using FONT hardware. The points give the mean value of the ratio recorded at each phase setting while the line represents the fit to this data.

BPM	Processor	$\frac{\delta y}{\delta \Psi_{\mathrm{LO}}}$ (μm/$^\circ$)
P1	5	-6.85 ± 0.05
P2	7	-1.05 ± 0.04
P3	10	-0.90 ± 0.04
MQD14X	2	-3.27 ± 0.07
MQF15X	4	-0.20 ± 0.02
MFB1FF	1	1.55 ± 0.15

Table 3.8: Phase sensitivity reported for each combination of processor and BPM.

BPM	$\sigma_{\Psi_{\mathrm{LO}}}$ ($^\circ$)	$\sigma_{y_{\Psi_{\mathrm{LO}}}}$ (μm)	$\sigma_{y_{\mathrm{measured}}}$ (μm)
P1	0.60 ± 0.01	4.09 ± 0.10	4.48 ± 0.10
P2	0.62 ± 0.01	0.65 ± 0.03	4.80 ± 0.11
P3	0.61 ± 0.01	0.54 ± 0.03	4.55 ± 0.10

Table 3.9: Jitter due to phase predicted from the measured sensitivity to phase at each BPM.

3.4.4 Phase compensation

The preferred definition of position, Δ/Σ_I, has been shown to depend to a variable extent on the phase offset between LO and bunch, Ψ_{LO}. The goal is thus to arrive at a new definition of position which is insensitive to such a change. Consider a linear combination of Δ/Σ_I and Σ_Q/Σ_I:

$$B_\Delta \left(\frac{\Delta}{\Sigma_I}\right) + B_{\Sigma_Q} \left(\frac{\Sigma_Q}{\Sigma_I}\right) = \kappa \tag{3.27}$$

where B_Δ, B_{Σ_Q} and κ are constants. The linear dependence of both Δ/Σ_I and Σ_Q/Σ_I on the phase offset Ψ_{LO} for small values of Ψ_{LO} has already been established and so Equation 3.27 may be rewritten as:

$$B_\Delta \left(m_\Delta \Psi_{\mathrm{LO}} + c_\Delta\right) + B_{\Sigma_Q} \left(m_{\Sigma_Q} \Psi_{\mathrm{LO}} + c_{\Sigma_Q}\right) = \kappa \tag{3.28}$$

where m_Δ and m_{Σ_Q} represent respectively the gradient of the straight line fit to each of Δ/Σ_I and Σ_Q/Σ_I and similarly c_Δ and c_{Σ_Q} correspond to the y-intercept in each case. As the objective is a linear combination of the two observables which is not a function of Ψ_{LO} we first differentiate with respect to Ψ_{LO} and then impose the condition $d\kappa/d\left(\Psi_{\mathrm{LO}}\right) = 0$ in order to arrive at:

$$B_\Delta m_\Delta + B_{\Sigma_Q} m_{\Sigma_Q} = 0 \tag{3.29}$$

Defining $B_\Delta = 1$ it follows that $B_{\Sigma_Q} = -m_\Delta/m_{\Sigma_Q}$. Substituting these into Equation 3.27 results in the expression for the phase compensated position, ξ:

$$\xi = k'^{-1}_{\mathrm{cal}} \left[\left(\frac{\Delta}{\Sigma_I}\right) - \frac{m_\Delta}{m_{\Sigma_Q}} \left(\frac{\Sigma_Q}{\Sigma_I}\right)\right] \tag{3.30}$$

where k'_{cal} is the calibration constant obtained from measurement of how ξ varies as a function of mover or beam position.

3.4.4.1 Effect of phase compensation

Figure 3.17 illustrates the effect of the phase compensation procedure by applying it to the data from the scan of Ψ_{LO}. Each plot shows that the phase compensated position is approximately flat over the linear region $|\Psi_{\mathrm{LO}}| < 10^{\circ}$; thus, the phase compensated position will be approximately independent of phase for phase offsets of the magnitude typically encountered at the ATF. Further evidence of the success of the phase compensation procedure can be seen in the effect it has on the BPM to BPM correlations. Figure 3.18 contains the correlation plots for both the position y and the phase compensated position ξ; in each plot, the vector of measurements at one BPM of the first three-BPM system is plotted against the vector of measurements at one of its companion BPMs. Figure 3.18(a) and Figure 3.18(b) express the conclusion of Section 3.3; namely, that the position at P1 bears little relation to that at P2 and P3, while Figure 3.18(c) shows the convincing linear relationship that exists between the latter two BPMs. Figure 3.18(d) and Figure 3.18(e) show that, in contrast to the ordinary position, the phase compensated position at any one BPM is strongly related to that measured at the other two. This is highly suggestive of a phase compensated position that more accurately reflects the trigger-to-trigger variation in the true position of the beam. The improvement to the "BPM to BPM" correlations is given numerically in Table 3.10.

	Correlation coefficient	
BPM pair	y	ξ
P1-P2	−0.16	−0.62
P1-P3	−0.26	−0.77
P2-P3	0.95	0.95

Table 3.10: Correlation coefficients for each pair of y or ξ measurements.

3.4.4.2 Phase compensated calibration

Figure 3.19 compares the behaviour of the uncalibrated y_{P1} position to that for ξ_{P1} over a scan of the P1 mover position. As far as calibration is concerned, the effect of phase compensation is mostly confined to a uniform 10 μm position offset; k'_{cal} differs from k_y by just 3×10^{-5}: a little over 1%. Given that the calibration constant changes by such a small amount even for the BPM which is most sensitive to variations in LO phase, the approximation $k'_{\mathrm{cal}} \sim k_{\mathrm{cal}}$ will be considered valid in most cases.

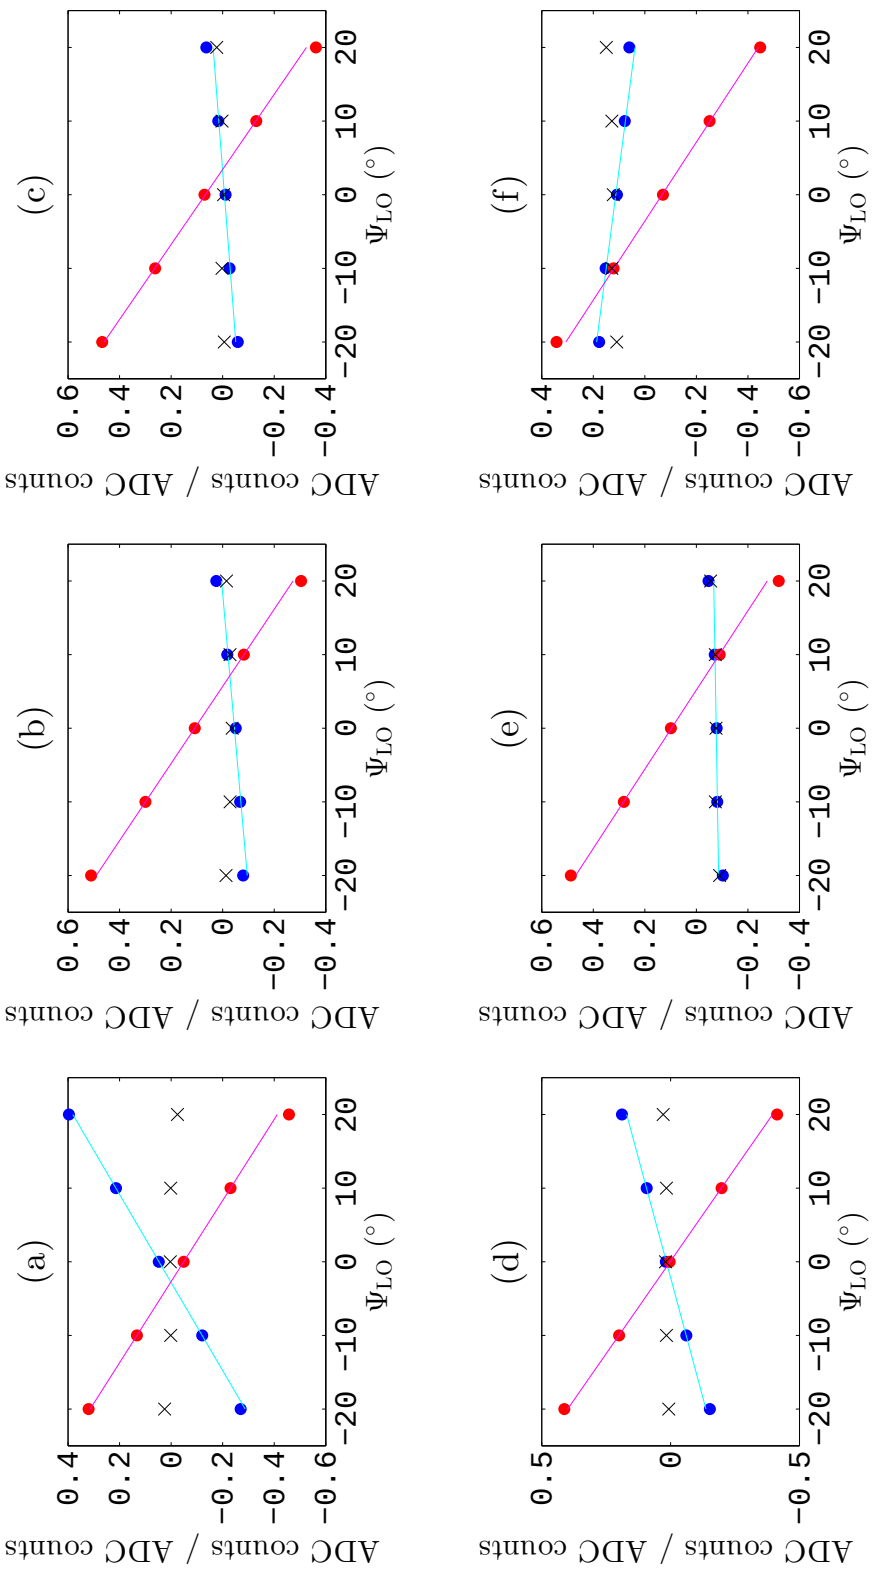

Figure 3.17: Effect of the phase compensation procedure in the linear range of the LO phase scan. Blue: $\langle \Delta/\Sigma_I \rangle$; cyan: linear fit to Δ/Σ_I data; red: $\langle \Sigma_Q/\Sigma_I \rangle$; magenta: linear fit to Σ_Q/Σ_I data; black: phase compensated $\frac{\Delta}{\Sigma_I}$ for the BPMs (a) P1; (b) P2; (c) P3; (d) MQD14X; (e) MQF15X; (f) MFB1FF.

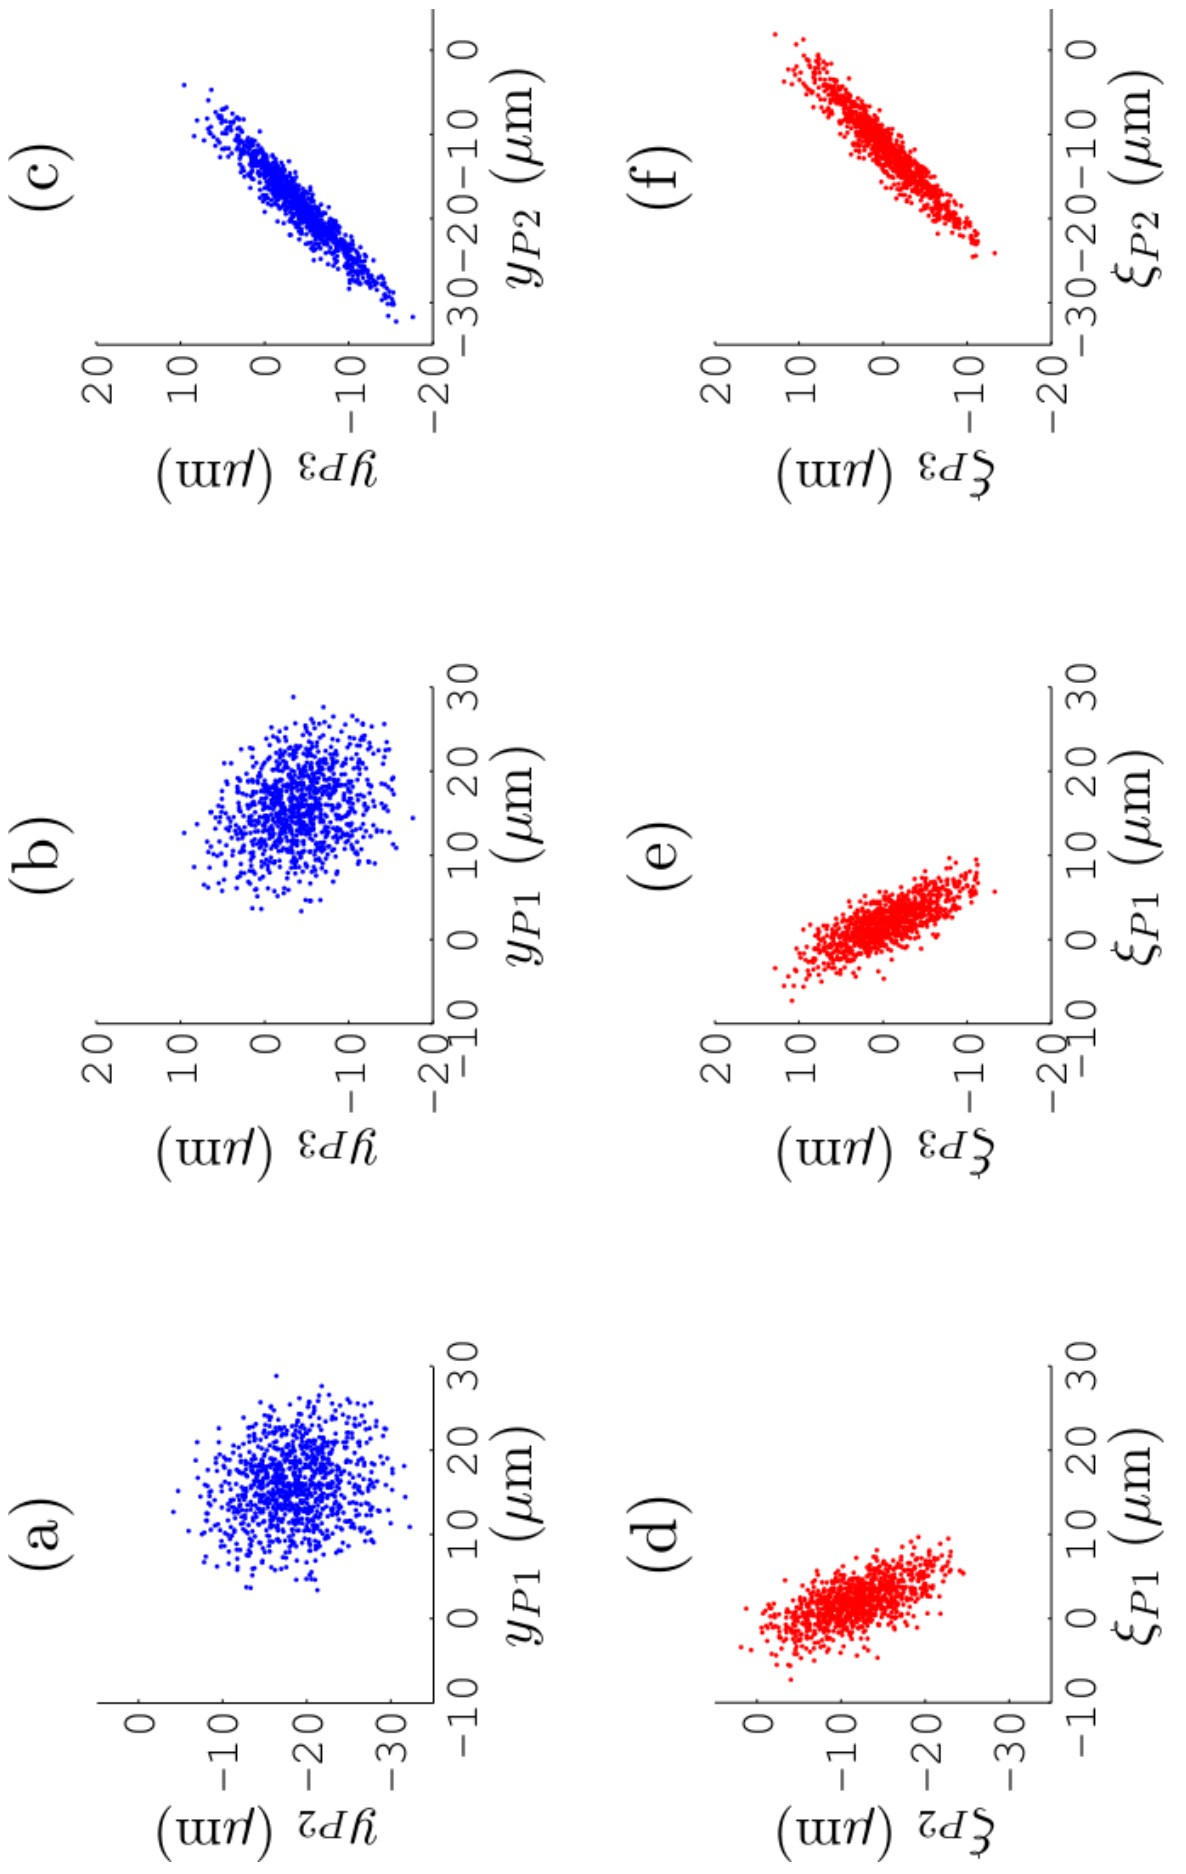

Figure 3.18: Correlation of the position y or the phase compensated position ξ for each pair of BPMs. (a) y_{P1}-y_{P2}; (b) y_{P1}-y_{P3}; (c) y_{P2}-y_{P3}; (d) ξ_{P1}-ξ_{P2}; (e) ξ_{P1}-ξ_{P3}; (f) ξ_{P2}-ξ_{P3}.

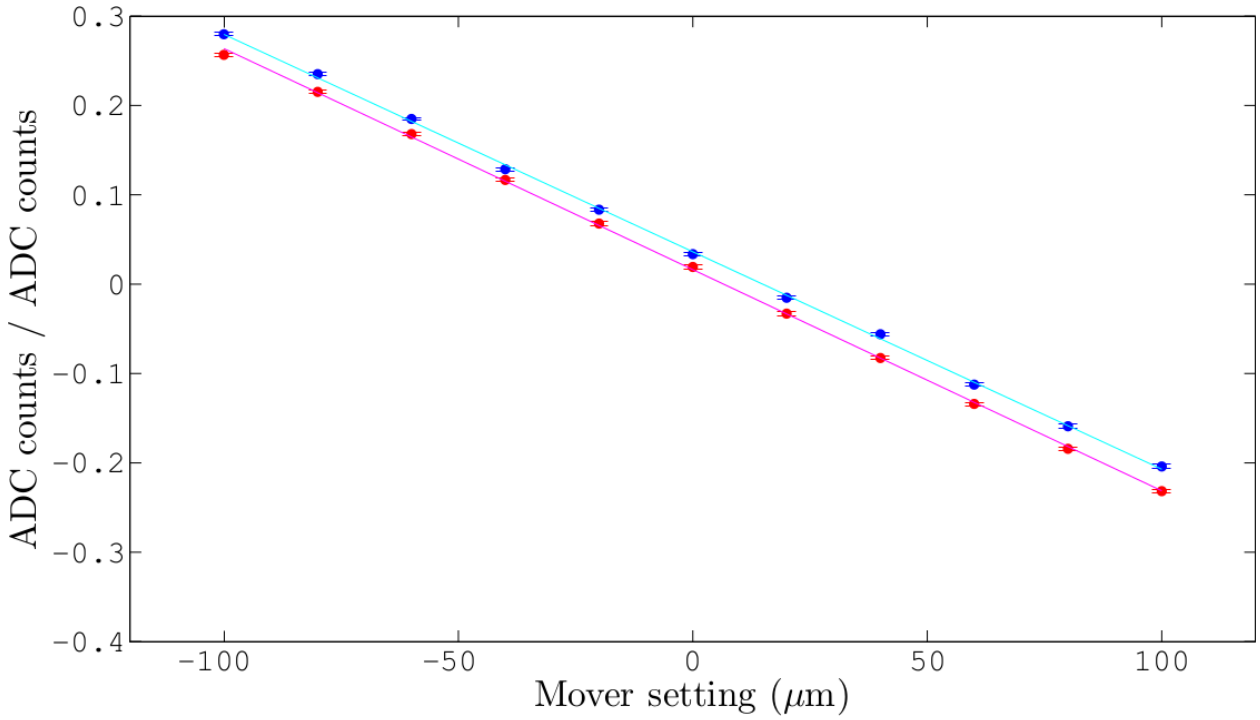

Figure 3.19: Effect of the phase compensation procedure on data for a calibration using the P1 mover. Blue: $\langle\left(\frac{\Delta}{\Sigma_I}\right)\rangle$; cyan: linear fit to $\frac{\Delta}{\Sigma_I}$ data; red: $\langle\left(\frac{\Delta}{\Sigma_I}\right) - \frac{m_\Delta}{m_{\Sigma_Q}}\left(\frac{\Sigma_Q}{\Sigma_I}\right)\rangle$; magenta: linear fit to $\left(\frac{\Delta}{\Sigma_I}\right) - \frac{m_\Delta}{m_{\Sigma_Q}}\left(\frac{\Sigma_Q}{\Sigma_I}\right)$ data. The vertical error bars represent the standard error on each mean position.

3.4.4.3 Phase compensated resolution

The resolution analysis of Section 3.3 can be repeated, replacing the ordinary position y with the phase compensated position ξ. The results are given in Figure 3.20. Compared with Figure 3.14, the model and fit methods are in much better agreement and the results are consistent in suggesting a resolution at the 0.5 μm level. However, even after phase compensation the resolution at P1 appears to be worse than its companion BPMs; this can be traced to the fact that the correlation coefficients calculated using ξ_{P1} are low compared to that for ξ_{P2}-ξ_{P3}. Ultimately, this suggests that an additional mechanism is working to degrade the resolution at P1.

	Estimated resolution (μm)	
BPM of prediction	Model method	Fit method
P1	0.84 ± 0.31	0.63 ± 0.20
P2	0.84 ± 0.31	0.45 ± 0.14
P3	0.84 ± 0.31	0.45 ± 0.14

Table 3.11: Resolution predicted from the phase compensated position using both model and fit methods.

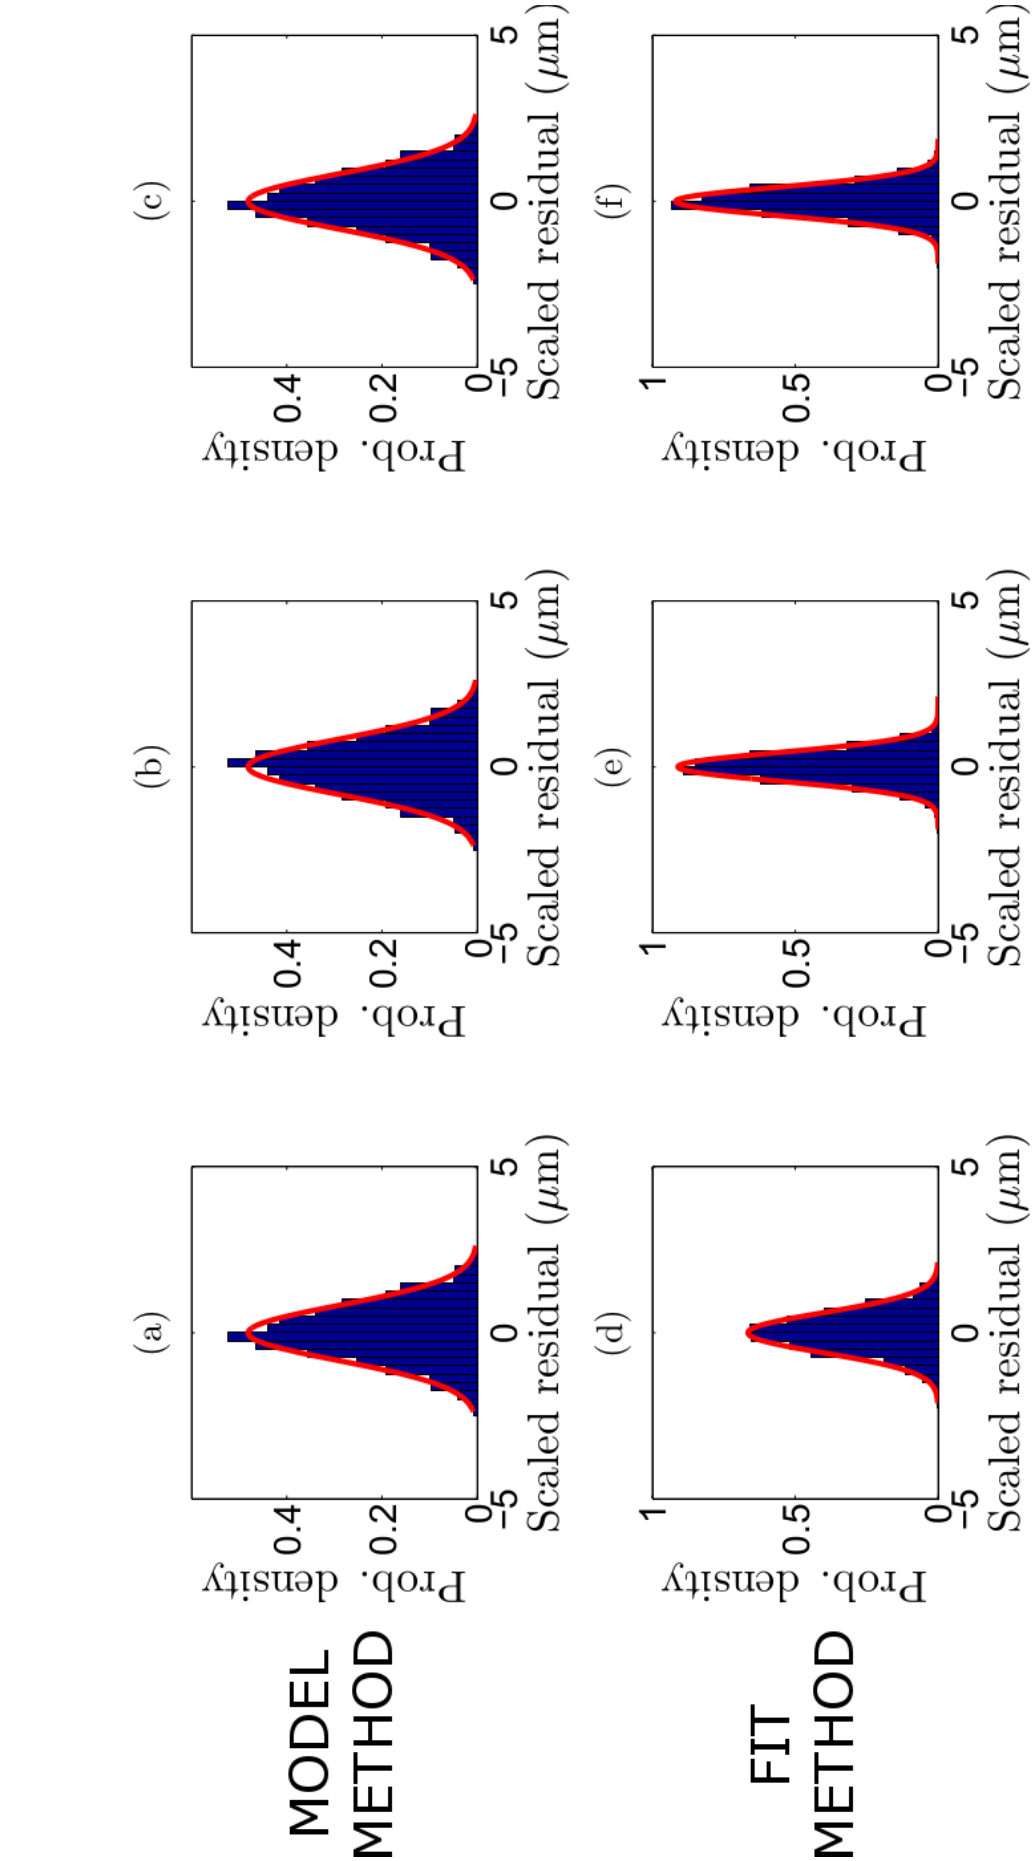

Figure 3.20: Distribution of the scaled residuals (blue) and the Gaussian fit used to estimate the BPM resolution (red) using the phase compensated position, ξ. The model method was used to predict the position for (a) P1; (b) P2; (c) P3 and the fit method for (d) P1; (e) P2; (f) P3.

3.5 Summary

Two FONT5 boards are used to take measurements at the ATF. The primary FONT5 board is responsible for feedback and uses the dedicated FONT BPMs P1, P2 and P3. A secondary board is used to monitor additional stripline BPMs downstream of the kickers. The BPMs are calibrated by moving either the beam or the BPMs themselves by a known amount and this is performed on a routine basis. The aggregated results yield a median calibration constant $k_{\text{cal}} = -0.0025\ \mu\text{m}^{-1}$.

The methods used to estimate the resolution of the BPM system were described and the results used to identify the detrimental effect of the dependence of the FONT BPM processors on the phase offset between the input LO and the bunch signals. Analysis techniques were developed to compensate for this effect and when used are found to provide positions that more accurately reflect the true location of the beam. The lowest estimate for the resolution of the three-BPM system comprised of P1, P2 and P3 is $0.45 \pm 0.14\ \mu\text{m}$.

Chapter 4

Performance of the FONT5 Feedback System

This chapter will present the derivation of the FONT5 feedback algorithm and explain how it is implemented in the firmware of the FONT5 board. The results will then be presented and the performance of the FONT5 feedback system evaluated.

4.1 The FONT5 feedback algorithm

The FONT5 feedback algorithm converts the measured position at the feedback BPMs P2 and P3 into a pair of kicks to be applied at K1 and K2 in order to correct the position of the next bunch.

4.1.1 Derivation

Consider a beamline containing a pair of kickers, K1 and K2, and a pair of BPMs, P2 and P3, where both BPMs are downstream of the two kickers. The beam itself consists of trains of two bunches; the first bunch is not kicked but the second bunch receives a kick at each of K1 and K2. The corrected position of the second bunch at P2, ${}^{\mathrm{b2}}Y_{\mathrm{P2}}$, can be described by the expression:

$$ {}^{\mathrm{b2}}Y_{\mathrm{P2}} = {}^{\mathrm{b2}}y_{\mathrm{P2}} + {}^{\mathrm{K1}}_{\mathrm{P2}}Hv_{\mathrm{K1}} + {}^{\mathrm{K2}}_{\mathrm{P2}}Hv_{\mathrm{K2}} \tag{4.1} $$

where ${}^{\mathrm{b2}}y_{\mathrm{P2}}$ is the position that bunch 2 would have had in the absence of the kicks at K1 and K2, v_{K1} and v_{K2} represent the magnitude of the kick at K1 and K2 respectively and ${}^{\mathrm{K}i}_{\mathrm{P}j}H$ are the kicker calibration constants that describe how a kick at Ki translates to a change in position at BPM Pj. Four such constants may therefore be defined, each one describing the extent to which a kick at either K1 or K2 affects the position measured at either BPM P2 or P3. A similar expression is obtained for the corrected position of the second bunch at P3:

$$ {}^{\mathrm{b2}}Y_{\mathrm{P3}} = {}^{\mathrm{b2}}y_{\mathrm{P3}} + {}^{\mathrm{K1}}_{\mathrm{P3}}Hv_{\mathrm{K1}} + {}^{\mathrm{K2}}_{\mathrm{P3}}Hv_{\mathrm{K2}} \tag{4.2} $$

The goal of the FONT feedback system is to stabilize the position of the second bunch at the two feedback BPMs, to which end the condition ${}^{\mathrm{b2}}Y_{\mathrm{P2}} = {}^{\mathrm{b2}}Y_{\mathrm{P3}} = \text{constant}$ must be imposed. For convenience, the (arbitrary) position selected as the goal of the feedback system is zero. By expressing Equation 4.1 and Equation 4.2 in the form of a matrix equation, the necessary values of the kicks are obtained:

$$\begin{bmatrix} {}^{\mathrm{K1}}_{\mathrm{P2}}H & {}^{\mathrm{K2}}_{\mathrm{P2}}H \\ {}^{\mathrm{K1}}_{\mathrm{P3}}H & {}^{\mathrm{K2}}_{\mathrm{P3}}H \end{bmatrix} \begin{pmatrix} v_{\mathrm{K1}} \\ v_{\mathrm{K2}} \end{pmatrix} = - \begin{pmatrix} {}^{\mathrm{b2}}y_{\mathrm{P2}} \\ {}^{\mathrm{b2}}y_{\mathrm{P3}} \end{pmatrix} \tag{4.3}$$

$$\begin{pmatrix} v_{\mathrm{K1}} \\ v_{\mathrm{K2}} \end{pmatrix} = - \begin{bmatrix} {}^{\mathrm{K1}}_{\mathrm{P2}}H & {}^{\mathrm{K2}}_{\mathrm{P2}}H \\ {}^{\mathrm{K1}}_{\mathrm{P3}}H & {}^{\mathrm{K2}}_{\mathrm{P3}}H \end{bmatrix}^{-1} \begin{pmatrix} {}^{\mathrm{b2}}y_{\mathrm{P2}} \\ {}^{\mathrm{b2}}y_{\mathrm{P3}} \end{pmatrix} \tag{4.4}$$

The second bunch would be located exactly on axis at both feedback BPMs if the kicks described above could be delivered to it. In practice the kicks are calculated from the measured position of bunch 1 by implicitly assuming that this is identical to the uncorrected position of bunch 2. The validity of this assumption will be investigated later in this chapter; for now, it is noted that the higher the degree of correlation between the two uncorrected bunches, the more stable the position of the second bunch will be following a correction calculated from the first bunch.

4.1.2 Implementation

In practice, Equation 4.4 is implemented in firmware via:

$$\begin{pmatrix} v_{\mathrm{K1}} \\ v_{\mathrm{K2}} \end{pmatrix} = \begin{bmatrix} {}^{\mathrm{P2}}_{\mathrm{K1}}G & {}^{\mathrm{P3}}_{\mathrm{K1}}G \\ {}^{\mathrm{P2}}_{\mathrm{K2}}G & {}^{\mathrm{P3}}_{\mathrm{K2}}G \end{bmatrix} \begin{pmatrix} {}^{\mathrm{b1}}(\frac{\Delta}{\Sigma_I})_{\mathrm{P2}} \\ {}^{\mathrm{b1}}(\frac{\Delta}{\Sigma_I})_{\mathrm{P3}} \end{pmatrix} \tag{4.5}$$

where G is the gain matrix and the uncorrected position ${}^{\mathrm{b1}}y$ is represented by the ratio Δ/Σ_I. The elements of the gain matrix will henceforth be referred to as the gain parameters; each gain parameter couples the measured position at one of the BPMs into a kick at one of the kickers. Thus, the gain parameters are in units of DAC counts per μm and are functions of the elements of the H matrix of Equation 4.3. Each of these kicker calibration constants are themselves determined by a combination of the beam transport from kicker to BPM and certain properties of the FONT feedback system (such as the attenuation present in the signal path from DAC to kicker amplifier). The kicker calibration constants have reciprocal units to the gain parameters (μm per DAC count) and may be determined by observing how the position of the beam at the BPMs varies as a function of the kick applied at each kicker.

By considering a single row of Equation 4.5 it is apparent that calculating each kick requires three arithmetic operations: division by Σ_I, multiplication by both Δ and the gain parameter and, finally, the addition of the contributions from the two BPMs to form the final kick $v_{\mathrm{K}i}$. The basic operation of the feedback module is depicted schematically in Figure 4.1.

The inputs of the feedback module are the bunch sample values from the Σ_I and Δ signals for the BPMs P2 and P3. Each Σ_I value is used as the address for a lookup table

(LUT) that contains at each address the value found by dividing the gain parameter by the 13-bit 2's complement interpretation of that address. The LUT stage thus performs both the division by Σ_I and the multiplication by the gain parameter. A multiplier computes the product of the Δ value and the output of the LUT for each BPM and then an adder combines the outputs of the two multipliers to deliver the calculated kick $v_{\mathrm{K}i}$. As the gain parameters are the only source of difference between the kicks that are calculated for K1 and K2, the feedback logic is duplicated in the firmware by including two instances of the feedback module (one for each kicker) that execute in parallel.

The LUT scheme described above is implemented in the firmware instead of a division operation on the FPGA itself because the time taken for the feedback algorithm to return a result is of critical concern. The system latency forms the subject of the next section; for now it is noted that, compared to multiplication, division by a factor that is not a power of two is an inherently slow operation to implement in a digital processor. Each LUT thus replaces a slow division in logic with the much faster operation of reading data from a RAM. The FONT feedback logic utilizes a total of four LUTs - one for each combination of BPM and kicker. The more advanced features of the feedback algorithm are discussed in detail in [46].

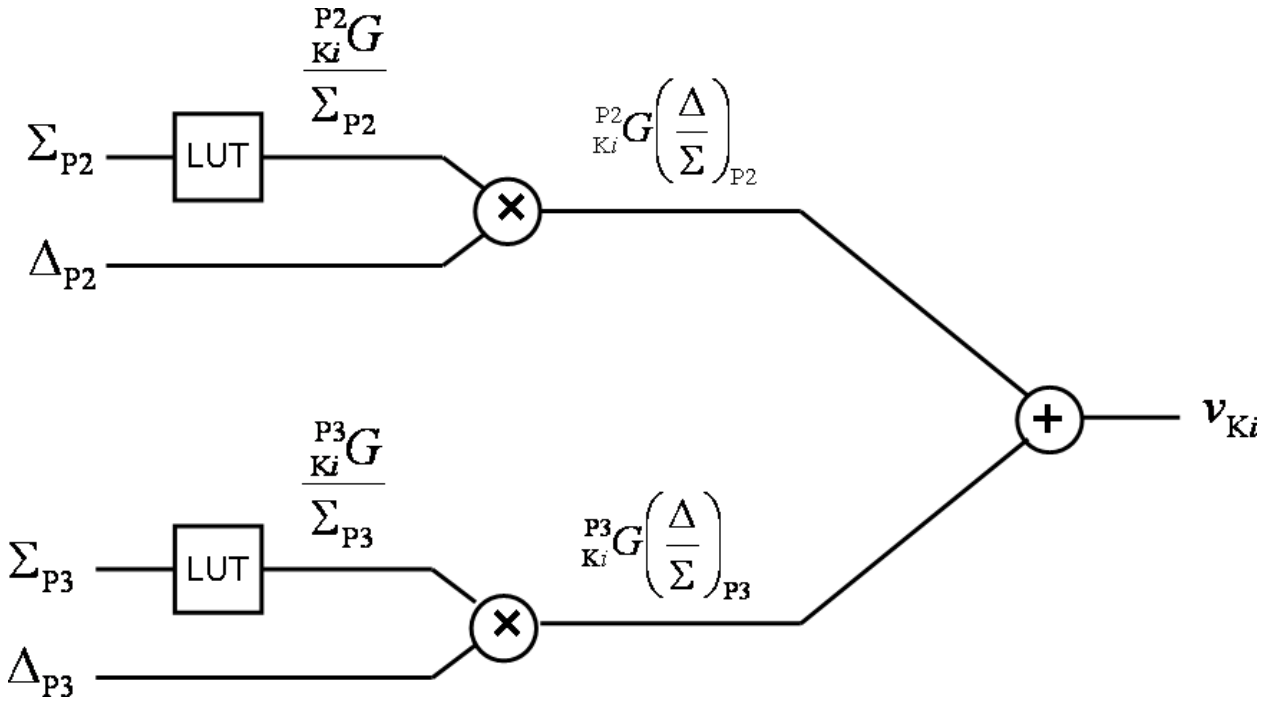

Figure 4.1: Schematic of the simplified feedback module illustrating the value at all stages of the calculation.

4.2 Latency of the feedback system

The latency of the FONT5 feedback system is defined here as the time between arrival of the first bunch at the feedback BPMs and the delivery of the necessary drive signals at the

FONT kickers in order to correct the position of the second bunch. This latency must be less than the bunch spacing in order for feedback on the bunch to bunch timescale to be possible; it is thus a critical parameter of the FONT system and one of the primary design criteria when it came to the development of the processing electronics and the kicker amplifiers. As the FONT feedback system consists of two feedback BPMs and two kickers a different figure for the latency could be obtained for each combination of BPM and kicker.

For a coupled system the most apposite measurement is that which gives the largest value, i.e. the measurement of the latency between the BPM and the kicker which each have the longest signal path from the FONT5 board. The optimal location of the FONT5 board is close to P2 (and hence K2) as this minimizes the latency due to the path length of the cables. The maximum signal path in this case consists of these three stages: first, the BPM signals travel from P3 to the FONT5 board via the analogue processor module. Second, the processor output waveforms are digitized by the FONT5 board and the results used to calculate the appropriate kicker drive signal, which is clocked on to the DAC of the FONT5 board. Third, this signal travels to K1 via the kicker amplifier.

Once clocked, the DAC maintains a constant voltage unless the value is updated. The maximum duration of the output pulse is thus equal to the length of the sampling window as a zero is clocked on to the DAC after the last ADC sample is gathered to ensure that the DAC is held at zero voltage in the time between bunch trains. The combined latency of the three stages is estimated by deliberately delaying the time at which the kicker drive signal is clocked to the DAC. As this delay is increased, at some point the delay must become so long that the signal arrives at the kicker later than the bunch it was supposed to kick. Estimating the latency is thus a matter of determining at which delay setting the second bunch ceases to be fully kicked; by convention, this condition is considered met when the kick applied to the second bunch is 90% that of the undelayed case. The method is illustrated schematically in Figure 4.2. The figure contains a pair of diagrams. In each diagram, the two blue circles represent the two bunches. The bunch arrival time at the kicker is represented by the location of the circle along the x-axis and the position of the bunch at a particular BPM is represented by the location of the circle along the y-axis. The voltage across the kicker as a function of time is also plotted in dark blue. The top diagram corresponds to the case where feedback on the bunch to bunch timescale is possible as the system latency is less than the bunch spacing. In this case, the voltage across the kicker takes the form of a pulse that rises some time after the first bunch arrives, is flat for a time sufficient for the kicker to have the maximum effect on the position of the second bunch, and then falls some time after the arrival of the second bunch. Thus, the first bunch receives no kick and is depicted as being on axis and the second bunch has been displaced by an amount proportional to the voltage across the kicker when it arrived there.

The latency of the system was earlier defined as the time it takes from the arrival of the first bunch for the kicker voltage pulse to achieve 90% of its final value. All that can be deduced from a data set with the timing of the top diagram is an upper limit on the latency (it must be less than the bunch spacing). Thus in the second diagram the kick has been delayed by a known amount such that the rising edge of the voltage pulse across the kicker now coincides with the arrival of the second bunch. This results in a change in the position

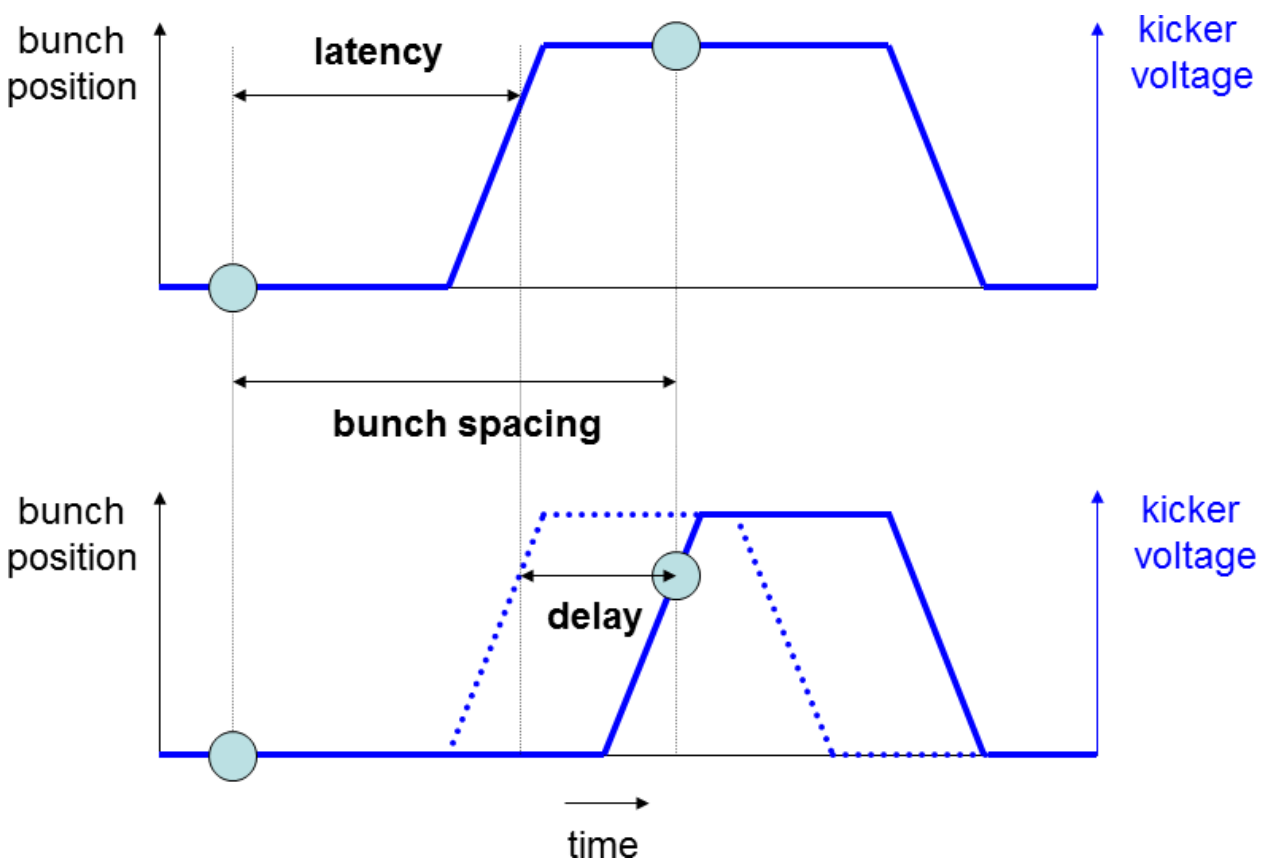

Figure 4.2: Schematic illustrating the principle behind the latency measurement.

of the second bunch relative to the first diagram; the known amount of added delay can then be subtracted from the bunch spacing to obtain an estimate of the latency of the system.

The latency of the system has been measured on many occasions; the results presented here are from a measurement performed on 7 December 2011. The K1 kick was delayed in steps of 11.2 ns from 0 ns up to 89.6 ns; for each setting, the kick was continually toggled on and off so that only half of the recorded triggers featured a kick. The position data gathered at each data setting may thus be represented as the vectors ${}^{bi}\mathbf{y}_{\mathrm{P}x}$ (for when the kick was off) and ${}^{bi}\mathbf{Y}_{\mathrm{P}j}$ (for when the kick was on), where i represents the bunch index (1 or 2) and j the BPM index (1, 2 or 3). The quantity $\Lambda_{\mathrm{P}j}$ is then defined for each delay setting as:

$$\Lambda_{\mathrm{P}j} = \left\langle \left({}^{\mathrm{b2}}\mathbf{Y}_{\mathrm{P}j} - {}^{\mathrm{b1}}\mathbf{Y}_{\mathrm{P}j}\right) - \left({}^{\mathrm{b2}}\mathbf{y}_{\mathrm{P}j} - {}^{\mathrm{b1}}\mathbf{y}_{\mathrm{P}j}\right) \right\rangle \tag{4.6}$$

and represents the average displacement due to the kick calculated by comparing consecutive kick on/kick off triggers and adjusted for the position of the first bunch. The results are given for the position in both P2 (Figure 4.3) and P3 (Figure 4.4). In each case the horizontal black dotted line indicates the 90% value of the maximum displacement and the vertical black dotted line indicates the delay setting at which it is estimated to occur. Both plots share the same basic features: the magnitude of $\Lambda_{\mathrm{P}j}$ is initially large and remains relatively constant up to about 25 ns of added delay. At this point the arrival of the second bunch at K1 starts to coincide with the rising edge of the kicker pulse and as the delay increases further the magnitude of the kick decreases until after about 55 ns of added delay the bunch

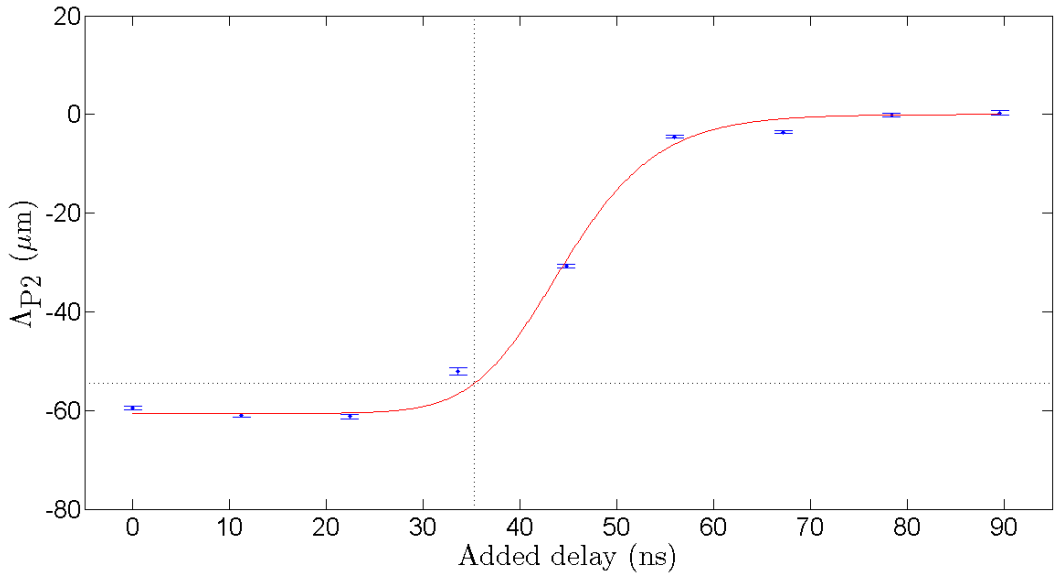

Figure 4.3: Λ_{P2} as a function of the delay applied to the constant kick at K1. The error bars represent the standard error on the mean and the red line is is obtained by fitting a generalized logistic function to the data (Equation 4.7).

ceases to be kicked altogether. The difference in the magnitude of the kick at P2 compared to that at P3 is also notable: a kick that produces a 60 µm displacement at P2 has a much less pronounced effect at P3, where the displacement is only about 15 µm. This is a result of the difference in betatron phase between the locations of the two BPMs. The latency of the system can be estimated by approximating Λ_{Pj} as a function of the added delay, t, by the generalized logistic function:

$$\Lambda_{Pj}(t) = \Lambda_{Pj}(-\infty) + \frac{\Lambda_{Pj}(+\infty) - \Lambda_{Pj}(-\infty)}{(1 + \eta e^{-\kappa(t-\tau)})^{1/\nu}} \tag{4.7}$$

where $\Lambda_{Pj}(-\infty)$, $\Lambda_{Pj}(+\infty)$, η, κ, τ and ν are parameters describing the shape of the curve. $\Lambda_{Pj}(-\infty)$ determines the asymptotic value of Λ_{Pj} as $t \to -\infty$ and is taken as the mean of Λ_{Pj} over the initial flat region. Similarly, $\Lambda_{Pj}(+\infty)$ gives the value of Λ_{Pj} as $t \to +\infty$ which is trivially taken to be zero (at $t = +\infty$, ${}^{bi}\mathbf{Y}_{Pj} = {}^{bi}\mathbf{y}_{Pj}$). ν is taken to be 0.5, forcing the region of maximum gradient to be midway between the two asymptotes. The remaining three parameters were then determined from a χ^2 fit to the data.

The value of the added delay that results in the second bunch receiving 90% of the full kick is estimated from the fit to the data. It is then subtracted from the bunch spacing of 187.6 ns to give a measure of the latency of the system in each case (Table 4.1). This latency represents the sum of the delays of all the components that make up the FONT system, as well as the intrinsic delays due to the propagation time of both beam and BPM and kicker signals. The estimated latency of these individual sources are presented in Table 4.2. The measurement of the latency of the BPM processor module was presented in Section 2.1.3, while the latency of the FONT5 board is given by the known number of clock cycles it takes for the feedback algorithm to produce a result. The remaining components were either

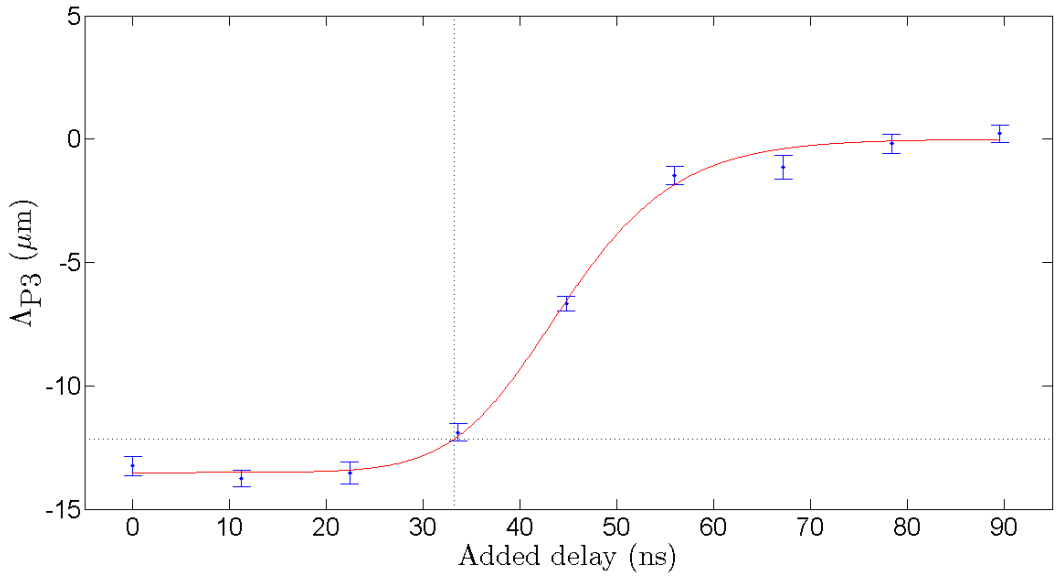

Figure 4.4: Λ_{P3} as a function of the delay applied to the constant kick at K1. The error bars represent the standard error on the mean and the red line is obtained by fitting a generalized logistic function to the data (Equation 4.7).

Feedback path	Added delay (ns)	Latency (ns)
P2→K1	35.3 ± 1.1	152.3 ± 1.1
P3→K1	33.2 ± 2.4	154.4 ± 2.4

Table 4.1: Added delay and the resultant measured latency for the feedback paths P2→K1 and P3→K1.

measured in the lab or estimated from the beamline geometry and the cable lengths [24]. The estimates from the data are both consistent with a latency of $\sim$ 153 ns which is in good agreement with the expected value of 155 ns.

4.3 Performance of the feedback system

The results presented here are from the FONT operations that took place in December 2011. For these data runs, one FONT5 board was used to monitor the BPMs P1, P2 and P3 and control the kickers K1 and K2. A second FONT5 board monitored the signals from BPMs MQD14X and MQF15X.

4.3.1 Kicker calibration

The gain parameters for feedback were calculated by scanning the kick provided by each of the kickers while the other is inactive and observing the effect on the position in the two

Component	Measured latency (ns)
Analogue processor	15
FONT5 board	39
Kicker amplifier	35
Kicker	3
Total signal / beam travel time (P3 $\rightarrow$ K1)	63
All	155

Table 4.2: Contributions to the latency from the individual sources of delay within the FONT5 system.

feedback BPMs, P2 and P3 (Figure 4.5). The results reproduce the effect previously seen in the latency measurement: a kick at K1 produces a large displacement at P2 but has little impact at P3 while a kick at K2 has a large effect at P3 and a more modest effect at P2. This behaviour merely reflects the nature of the lattice in the FONT region. The effect of each kicker is a change in the angle of the beam but the stripline BPMs only measure the position. In order to maximize the BPM response to a kicker it is therefore necessary to locate the BPM at a point in the lattice where the angle of the beam at the kicker has been largely converted to position (and vice versa); such a point is described as having a betatron phase advance $\mu = 90^\circ$ with respect to the kicker. The phase advances between pairs of FONT beamline components have been calculated from the nominal ATF optics in the past [24] and are presented here in Table 4.3. The kicker-BPM pairs K1-P2 and K2-P3 are seen to be out of phase ($\mu \sim 90^\circ$), while the pairs K1-P3 and K2-P2 are at the same phase ($\mu \sim 0^\circ$ or 180°). However, the deviation from the ideal case is significant (a fact supported by the kicker calibrations of Figure 4.5). A linear χ^2-fit to the uncalibrated position Δ/Σ_I in each

Transition	$\mu(^\circ)$
K1-K2	94.5
K1-P2	115.2
K1-P3	178.5
K2-P2	20.9
K2-P3	84.2

Table 4.3: The betatron phase advance μ between selected pairs of FONT beamline components [24].

of the feedback BPMs as a function of the strength of each kicker gives the parameters of the H matrix; the entire matrix can then be inverted to obtain the individual elements of the G matrix for use in the feedback algorithm on the FONT5 board (Table 4.4).

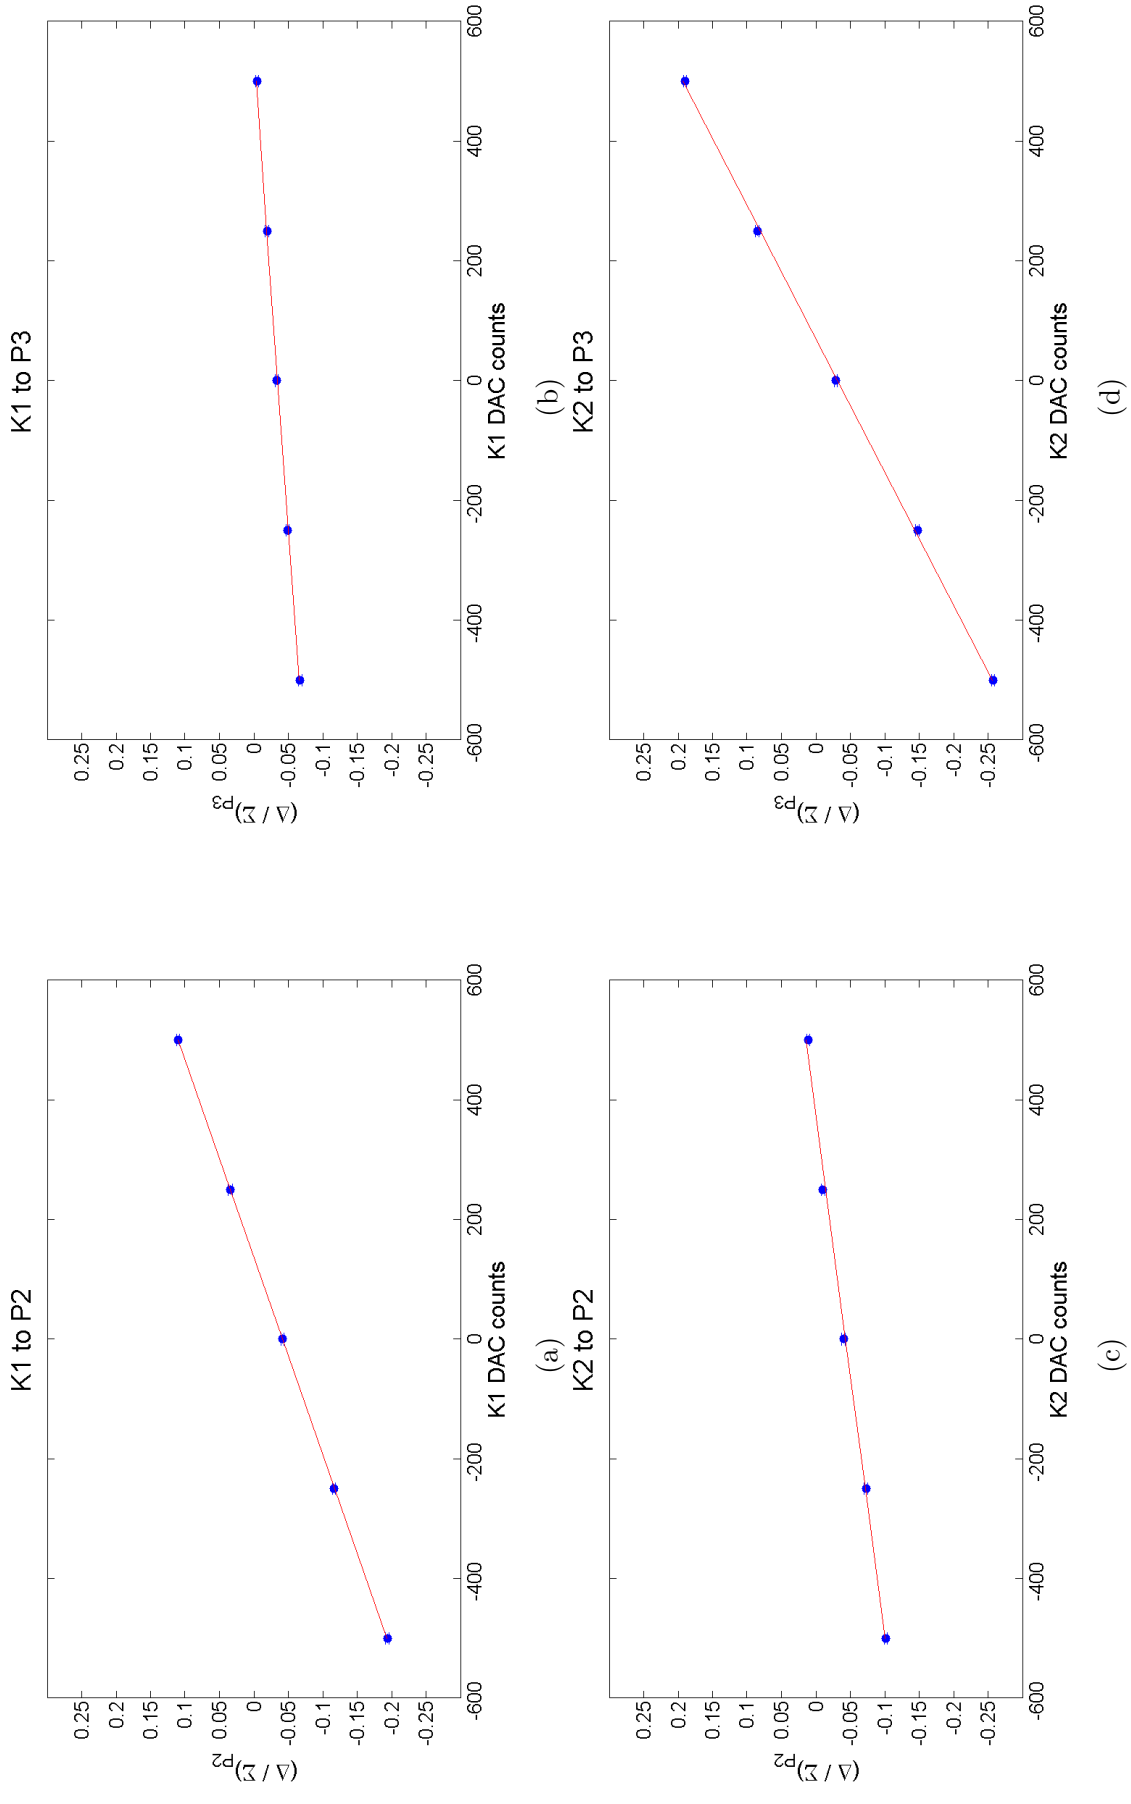

Figure 4.5: In each graph the mean value of the ratio Δ/Σ_I for bunch 2 is plotted as a function of the kick strength in DAC counts (blue dots); also shown is the result of the linear χ^2-fit to the data (red line). The four graphs correspond to the four combinations of active kicker and BPM of observation: (a) K1→P2; (b) K1→P3; (c) K2→P2; (d) K2→P3.

Kicker-BPM pair	$H(\left(\frac{\Delta}{\Sigma_I}\right)$ /DAC count)	G(DAC counts/ $\left(\frac{\Delta}{\Sigma_I}\right)$
K1→P2	$(30.3 \pm 1.3) \times 10^{-5}$	-3480 ± 210
K1→P3	$(6.2 \pm 1.3) \times 10^{-5}$	890 ± 110
K2→P2	$(11.5 \pm 1.3) \times 10^{-5}$	480 ± 100
K2→P3	$(45.0 \pm 1.2) \times 10^{-5}$	-2350 ± 170

Table 4.4: H values and derived G values.

4.3.2 Observed effect of the feedback system

Having calculated from the kicker calibrations the appropriate gain parameters for the feedback algorithm and then used this result to fill the LUTs, the FONT5 system is prepared to perform feedback. During a feedback data run, the software interface of the FONT5 board is set to alternate feedback on and off from train to train in order to reduce the impact of any change in beam conditions that may occur over the duration of a data run, as well as provide knowledge of the undisturbed beam conditions (particularly the bunch to bunch correlations). The effect of the FONT feedback system may be illustrated by comparing the distribution of the bunch position for the feedback on and feedback off populations in each BPM. The relevant metrics are the mean position of the bunch and the standard deviation of the position of the bunch, hereafter referred to as the jitter.

Figure 4.7 presents the effect of the feedback system at the BPM P1. The first bunch in the two bunch train is used to determine the kick for the second bunch but is not kicked itself; thus, any apparent difference between the feedback on and feedback off populations for bunch 1 is not caused by the feedback system. Nevertheless, the data for both bunches are presented as a cross-check. For all the BPMs monitored by FONT, the position of bunch 1 is clearly unaffected by whether the feedback system is active or not.

As P1 is upstream of K2 and only a short distance downstream of K1, the difference between the feedback on and feedback off populations for bunch 2 is expected to be small; indeed, as only one of the kicks has been supplied at this point (and the distance from that kicker is too small for the kick to have a significant effect at the BPM) it is not expected that the beam position will be correctly stabilized at this location in any case. The results confirm this as Figure 4.7(d) and Figure 4.7(e) both suggest that the effect of the feedback system at P1 was a small displacement of 2-3 μm for bunch 2 relative to bunch 1. A comparison of Figure 4.7(c) and Figure 4.7(f) reveals that the feedback system made the jitter of bunch 2 larger than that of bunch 1 as the red histogram of feedback on data is slightly wider than the blue histogram of feedback off data in the latter plot.

By comparing the feedback off distributions from Figure 4.7(a) and Figure 4.7(d) it is apparent that there was a pre-existing displacement of bunch 2 relative to bunch 1 of about -5 μm. Bunch offsets such as these are common and often arise as a result of the structure of the voltage pulse driving the extraction kicker; if the pulse is not completely stable over the entire duration of beam extraction, the two bunches receive slightly different kicks when

they leave the damping ring which causes them to have different trajectories as they travel the extraction line.

Figure 4.8 contains the equivalent plots for the BPM P2. The effect of the feedback is immediately evident in Figure 4.8(d). The distribution of bunch 2 positions for the feedback on case is narrow compared to the feedback off case and displaced by about 10 μm; however, due to the bunch offset the corrected distribution does not lie about the zero point. A correction of similar quality can be seen for P3, the other feedback BPM, in Figure 4.9.

The results for MQD14X and MQF15X, the witness BPMs, are given in Figure 4.10 and Figure 4.11 respectively. Figure 4.10(a) and Figure 4.10(d) along with Figure 4.11(a) and Figure 4.11(d) reveal that the beam is significantly off axis in both these BPMs. This is unavoidable due to the lack of mover stages for the non-FONT stripline BPMs. In both cases, the initial jitter and magnitude of the kick at each BPM are in line with the expected behaviour given the lattice. The beta function in the FONT region is plotted in Figure 4.6.

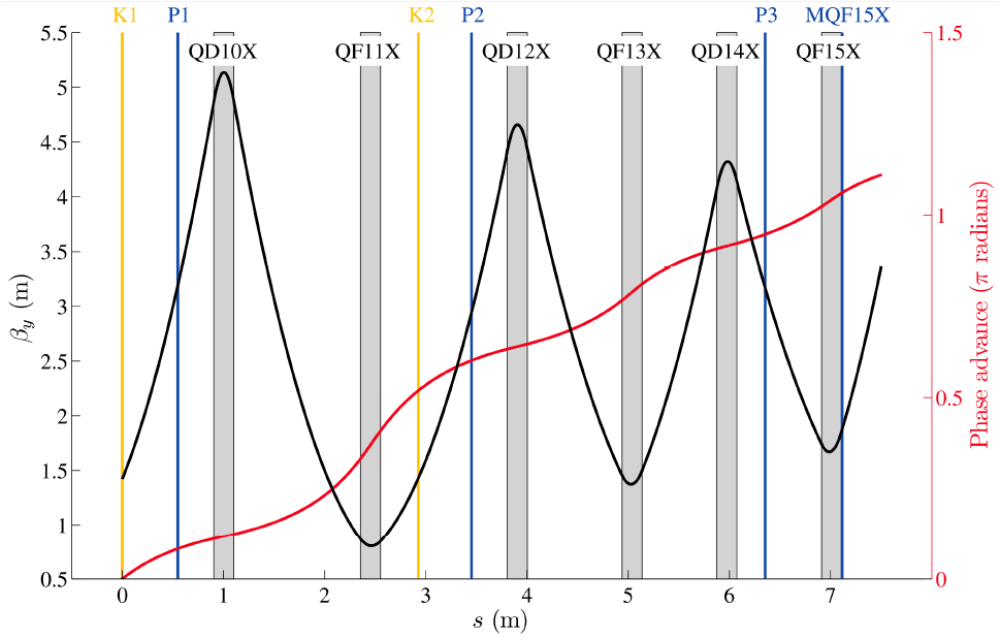

Figure 4.6: Plot of the vertical beta function in the FONT region [46].

In the extraction line of the ATF there is observed to be a strong correlation between the beam size (determined by the beta function) and the beam jitter [59]. MQD14X is located at a quadrupole that is defocusing in the x plane and consequently a large jitter in the y plane is observed. The reverse is true for MQF15X as it resides close to a quadrupole that is focusing in the x plane.

The mean and the standard deviation of the bunch position at each BPM are supplied respectively in Table 4.5 and Table 4.6. The phase compensated position of bunch i at a given BPM when no corrective kick is supplied is defined as ${}^{\mathrm{b}i}\xi$ (Equation3.30); the quantity ${}^{\mathrm{b}i}\Xi$ is introduced here as the phase compensated position when a corrective kick is supplied. Note that the beam has been stabilized at P2 to the level of the resolution estimate from Chapter 3. The effect of the feedback system can then be characterized in terms of two

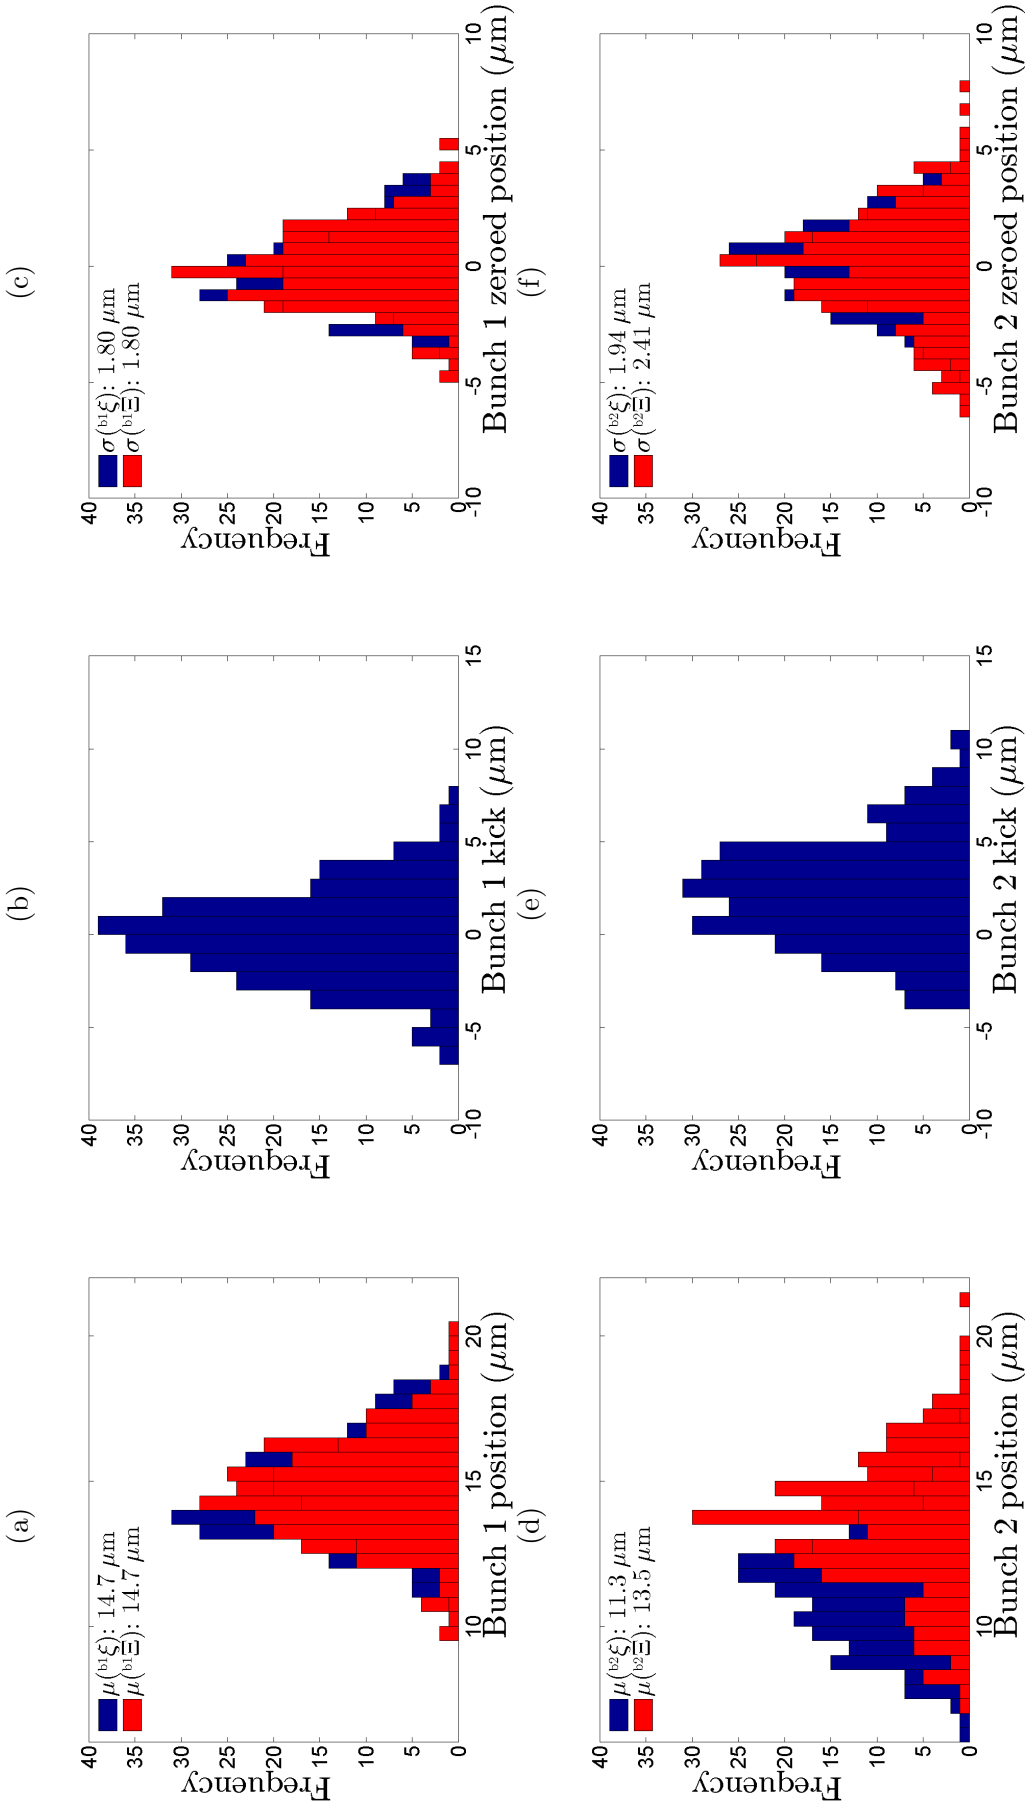

Figure 4.7: Results of the feedback run for BPM P1. (a),(d): histograms of the bunch position for feedback off (blue) and feedback on (red); (b),(e): histogram of the difference in position between each on/off pair; (c),(f): as (a) and (d), but with the mean of each distribution subtracted prior to binning. Plots (a), (b) and (c) feature the data for bunch 1 and (d), (e) and (f) that for bunch 2.

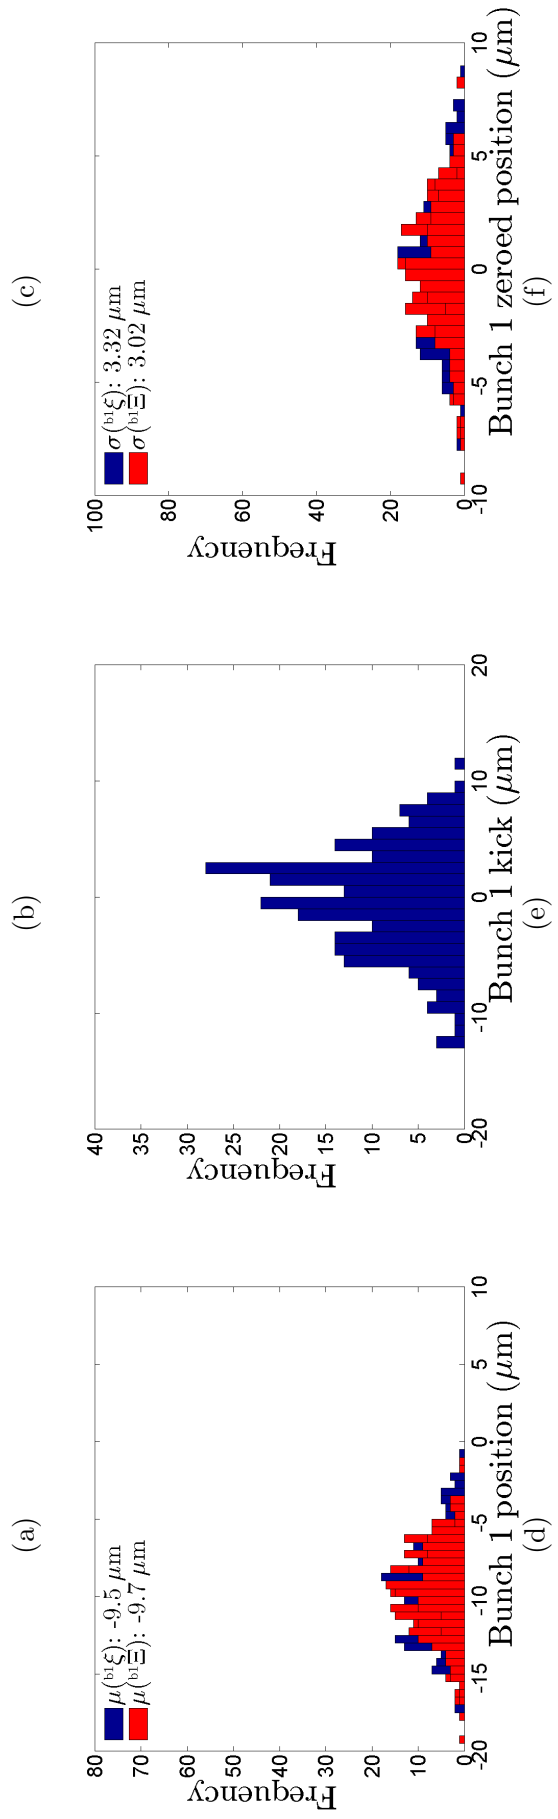

Figure 4.8: Histograms of the positions during the feedback run for BPM P2. (a),(d): histograms of the bunch position for feedback off (blue) and feedback on (red); (b),(e): histogram of the difference in position between each on/off pair; (c),(f): as (a) and (d), but with the mean of each distribution subtracted prior to binning. Plots (a), (b) and (c) feature the data for bunch 1 and (d), (e) and (f) that for bunch 2.

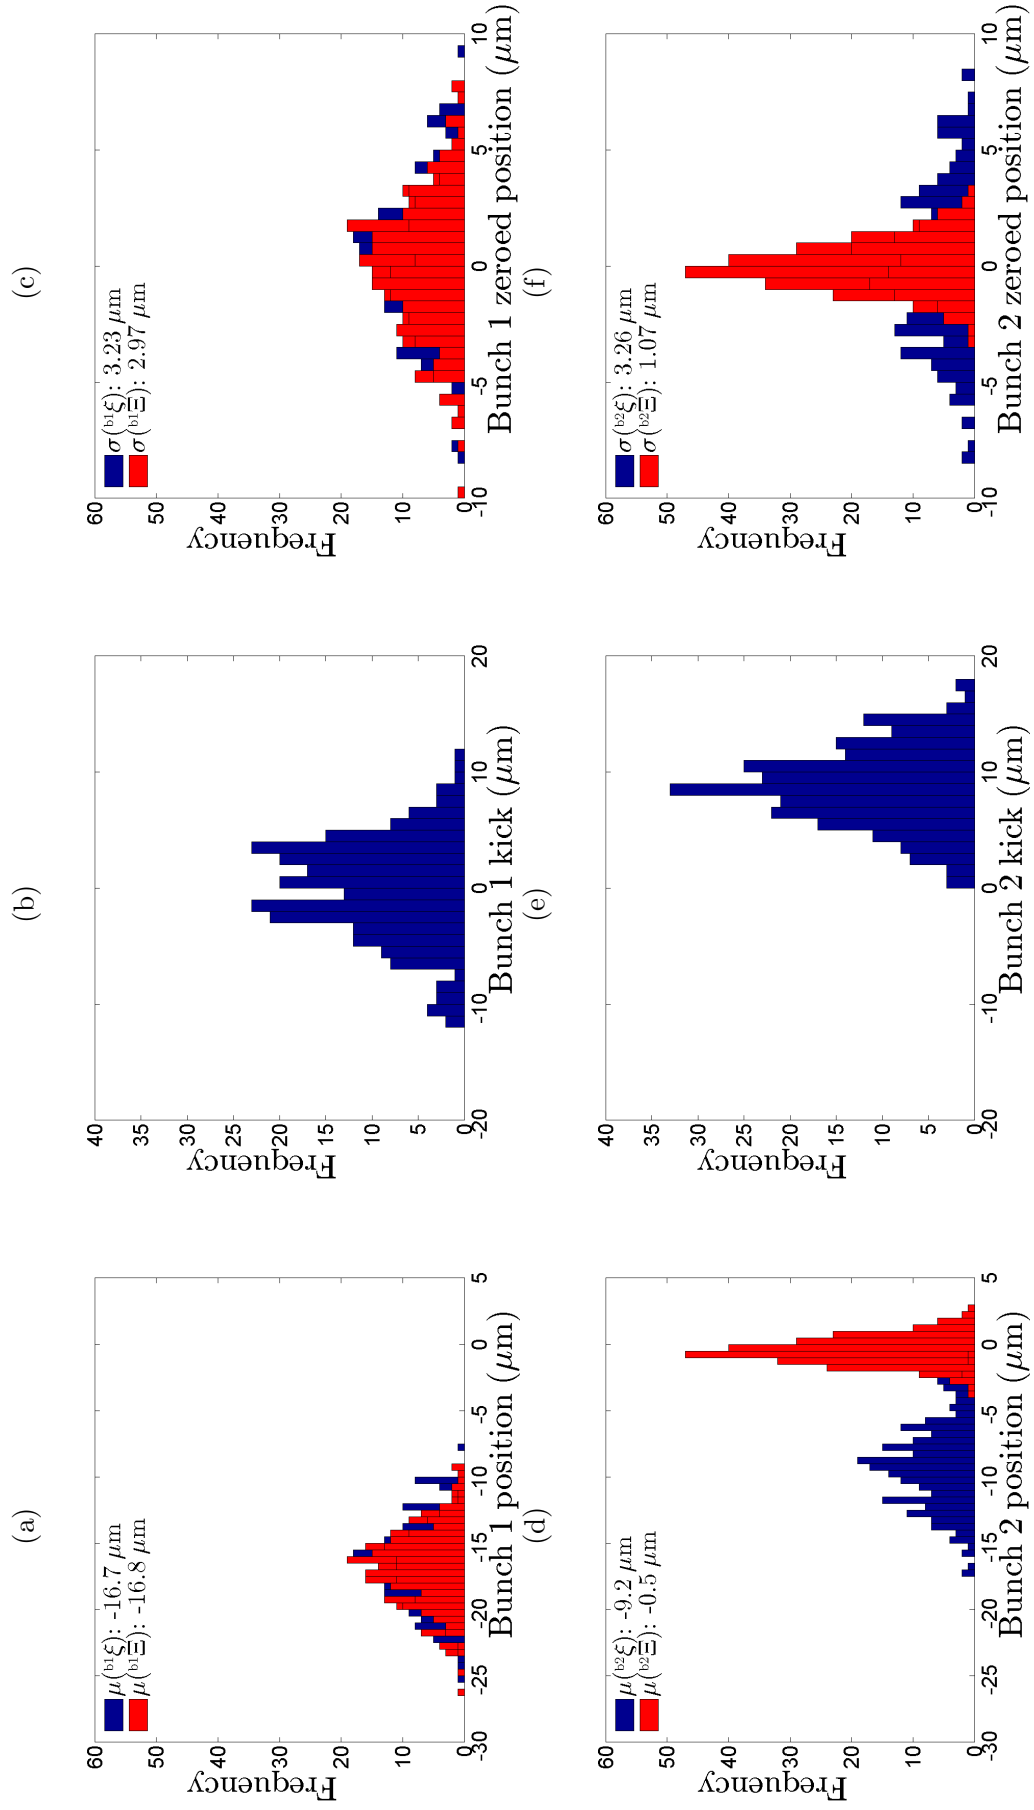

Figure 4.9: Histograms of the positions during the feedback run for BPM P3. (a),(d): histograms of the bunch position for feedback off (blue) and feedback on (red); (b),(e): histogram of the difference in position between each on/off pair; (c),(f): as (a) and (d), but with the mean of each distribution subtracted prior to binning. Plots (a), (b) and (c) feature the data for bunch 1 and (d), (e) and (f) that for bunch 2.

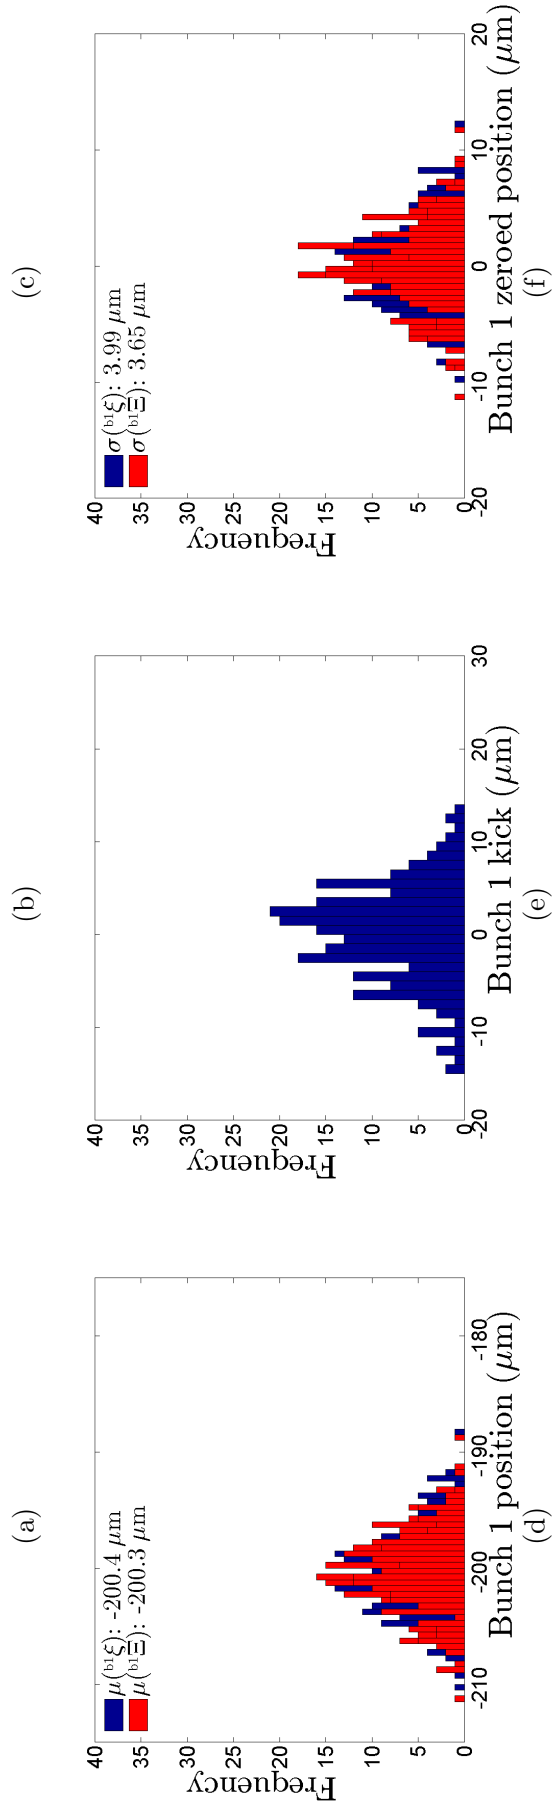

Figure 4.10: Histograms of the positions during the feedback run for BPM MQD14X. (a),(d): histograms of the bunch position for feedback off (blue) and feedback on (red); (b),(e): histogram of the difference in position between each on/off pair; (c),(f): as (a) and (d), but with the mean of each distribution subtracted prior to binning. Plots (a), (b) and (c) feature the data for bunch 1 and (d), (e) and (f) that for bunch 2.

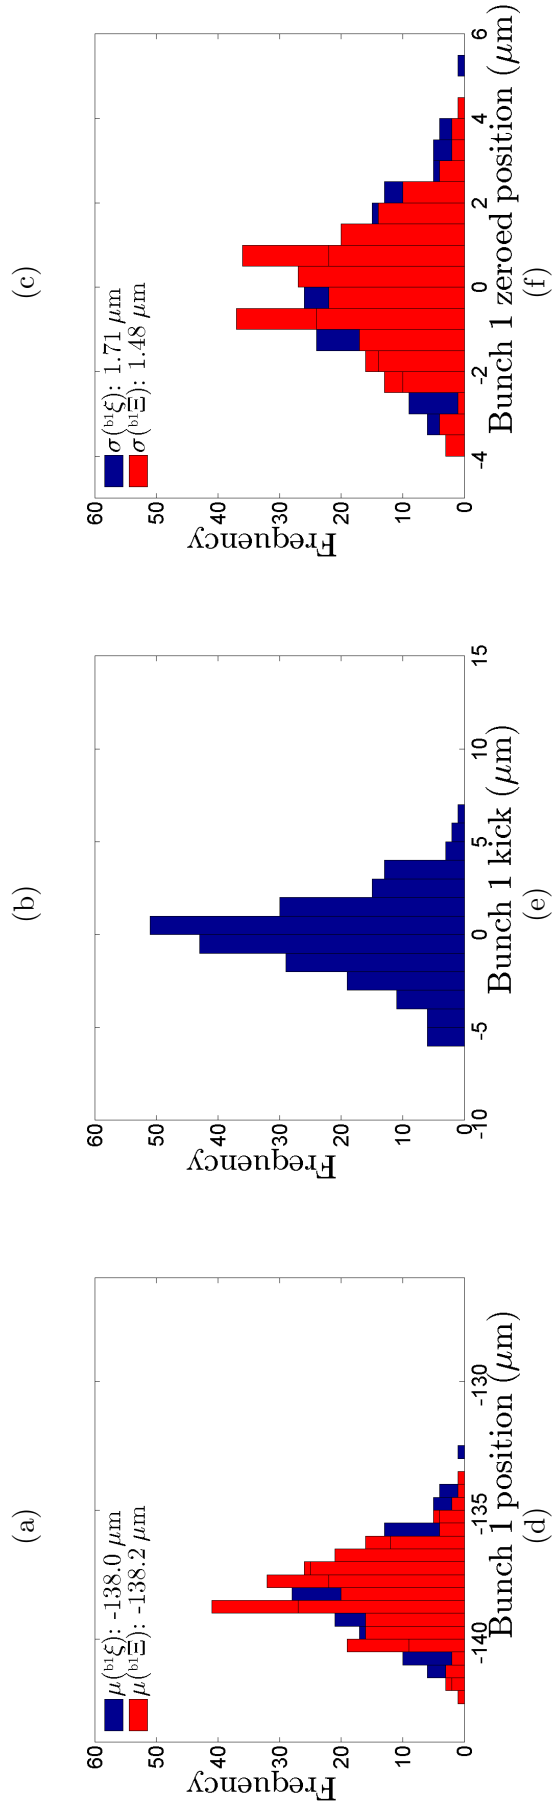

Figure 4.11: Histograms of the positions during the feedback run for BPM MQF15X. (a),(d): histograms of the bunch position for feedback off (blue) and feedback on (red); (b),(e): histogram of the difference in position between each on/off pair; (c),(f): as (a) and (d), but with the mean of each distribution subtracted prior to binning. Plots (a), (b) and (c) feature the data for bunch 1 and (d), (e) and (f) that for bunch 2.

	Bunch 1		Bunch 2	
	OFF	ON	OFF	ON
BPM	$\mu(^{b1}\xi)$ (μm)	$\mu(^{b1}\Xi)$ (μm)	$\mu(^{b2}\xi)$ (μm)	$\mu(^{b2}\Xi)$ (μm)
P1	14.7 ± 0.1	14.7 ± 0.1	11.3 ± 0.1	13.5 ± 0.2
P2	-9.5 ± 0.2	-9.7 ± 0.2	-1.8 ± 0.2	6.0 ± 0.0
P3	-16.7 ± 0.2	-16.8 ± 0.2	-9.2 ± 0.2	-0.5 ± 0.1
MQD14X	-200.4 ± 0.3	-200.3 ± 0.2	-192.5 ± 0.2	-182.3 ± 0.1
MQF15X	-138.0 ± 0.1	-138.2 ± 0.1	-134.3 ± 0.1	-130.0 ± 0.1

Table 4.5: Calculated values of the mean position for both bunches for the feedback on and feedback off cases.

	Bunch 1		Bunch 2	
	OFF	ON	OFF	ON
BPM	$\sigma(^{b1}\xi)$ (μm)	$\sigma(^{b1}\Xi)$ (μm)	$\sigma(^{b2}\xi)$ (μm)	$\sigma(^{b2}\Xi)$ (μm)
P1	1.80 ± 0.08	1.80 ± 0.08	1.94 ± 0.09	2.41 ± 0.11
P2	3.32 ± 0.15	3.02 ± 0.14	3.25 ± 0.15	0.59 ± 0.03
P3	3.23 ± 0.15	2.97 ± 0.14	3.26 ± 0.15	1.07 ± 0.05
MQD14X	3.99 ± 0.18	3.65 ± 0.17	3.93 ± 0.18	1.45 ± 0.07
MQF15X	1.71 ± 0.08	1.48 ± 0.07	1.68 ± 0.08	1.07 ± 0.05

Table 4.6: Calculated values of the position jitter for both bunches for the feedback on and feedback off cases.

metrics, the first of which is the change in the mean position of the second bunch due to the feedback system, d:

$$d = \mu(^{b2}\Xi) - \mu(^{b2}\xi) \tag{4.8}$$

and the second is the corresponding change in the jitter of the position expressed as a ratio, f:

$$f = \frac{\sigma(^{b2}\xi)}{\sigma(^{b2}\Xi)} \tag{4.9}$$

The values of these metrics are presented in Table 4.7. Note that P1 is upstream of K2; the position is only fully corrected downstream of both kickers and so the f value for this BPM is not required to be greater than one.

4.4 Factors influencing feedback performance

Although the analysis that produced the results presented in this chapter included the phase compensation procedure developed in Chapter 3, it should be noted that the firmware at the time used the non-phase compensated definition of the position ($y = \frac{\Delta}{\Sigma_I}$) in the feedback calculation. It is therefore this definition that should be considered when making predictions of the performance of the feedback algorithm.

BPM	d (μm)	f
P1	2.2 ± 0.2	0.8 ± 0.1
P2	7.8 ± 0.2	5.5 ± 0.8
P3	8.7 ± 0.2	3.1 ± 0.5
MQD14X	10.2 ± 0.3	2.7 ± 0.5
MQF15X	4.3 ± 0.1	1.6 ± 0.1

Table 4.7: Effect of the feedback system characterized in terms of the mean displacement due to the feedback system d and the jitter reduction factor f.

4.4.1 Expected jitter given observed bunch to bunch correlations

Random variation in the position of bunch 2 relative to bunch 1 (i.e. intra-train jitter) will negatively impact the minimum jitter that can be achieved for bunch 2 using the feedback system. Under the approximation that the corrected position at each BPM can be written as ${}^{\mathrm{b2}}Y = {}^{\mathrm{b1}}y - {}^{\mathrm{b1}}y$, the expected performance of the feedback system in each BPM can be estimated using the following equation:

$$\sigma({}^{\mathrm{b2}}Y) = \sqrt{\sigma({}^{\mathrm{b1}}y)^2 + \sigma({}^{\mathrm{b2}}y)^2 - 2\sigma({}^{\mathrm{b1}}y)\sigma({}^{\mathrm{b2}}y)\rho\left({}^{\mathrm{b1}}y{}^{\mathrm{b2}}y\right)} \tag{4.10}$$

where $\rho\left({}^{\mathrm{b1}}y{}^{\mathrm{b2}}y\right)$ is the correlation coefficient describing the strength of the linear relationship between the uncorrected, non-phase compensated position of bunch 1 and the uncorrected, non-phase compensated position of bunch 2. Zero bunch to bunch correlation implies that no information about the position of bunch 2 can be inferred from the position of bunch 1; under these conditions, the feedback system simply provides a random kick to the second bunch.

Table 4.8 compares the non-phase compensated jitter achieved with that predicted using the non-phase compensated jitters for bunch 1 and bunch 2 along with the measured non-phase compensated bunch to bunch correlation. The two values are found to be in good agreement.

		$\sigma\left({}^{\mathrm{b2}}Y\right)$ (μm)	
BPM	$\rho\left({}^{\mathrm{b1}}y{}^{\mathrm{b2}}y\right)$	Predicted	Observed
P2	0.99	0.5	0.6
P3	0.97	0.9	1.0

Table 4.8: Comparison of the non-phase compensated jitter achieved at the feedback BPMs P2 and P3 with that predicted from the beam conditions.

4.4.2 Phase sensitivity

Given the known issue of BPM phase sensitivity, it is interesting to consider the performance that could have been achieved had the FONT5 feedback algorithm been using the phase compensated position $\xi = \left[\left(\frac{\Delta}{\Sigma_I}\right) - \frac{m_\Delta}{m_{\Sigma_Q}}\left(\frac{\Sigma_Q}{\Sigma_I}\right)\right]$. By analogy to Equation 4.10, the following equation may be used to predict the performance of such a system:

$$\sigma(^{\mathrm{b2}}\Xi) = \sqrt{\sigma(^{\mathrm{b1}}\xi)^2 + \sigma(^{\mathrm{b2}}\xi)^2 - 2\sigma(^{\mathrm{b1}}\xi)\sigma(^{\mathrm{b2}}\xi)\rho\left(^{\mathrm{b1}}\xi^{\mathrm{b2}}\xi\right)} \tag{4.11}$$

where the non-phase compensated positions, denoted by y and Y, have been replaced by their phase compensated equivalents, denoted by ξ and Ξ. Table 4.9 presents the equivalent figures to those from Table 4.8 once the effect of the LO phase offset on the measured position at each BPM has been taken into account. Including phase compensation in the analysis

		$\sigma\left(^{\mathrm{b2}}\Xi\right)$ (μm)	
BPM	$\rho\left(^{\mathrm{b1}}\xi^{\mathrm{b2}}\xi\right)$	Predicted	Observed
P2	0.99	0.5	0.6
P3	0.98	0.7	1.1

Table 4.9: Comparison of the phase compensated jitter achieved at the feedback BPMs P2 and P3 with that predicted from the beam conditions.

results in an increase in the bunch to bunch correlation at P3; this in turn decreases the jitter expected to remain after correction to 0.7 μm. Thus, a version of the firmware that included phase compensation would be expected to deliver improved stability of the beam as long as the feedback BPMs remain sensitive to phase offsets.

4.4.3 Alternative feedback algorithm

The algorithm of the FONT feedback system is meant to stabilize the position of the beam at the two feedback BPMs and thus all points downstream. However, even under the ideal conditions of perfect BPM resolution and perfect bunch 1 to bunch 2 correlation, the algorithm implemented on the FONT5 board was found not to deliver such a perfect correction of both position and angle.

Given that the effect of a kicker is limited to a change in the angle of the passing bunch, it follows that in order to correct both position and angle to zero at a single point the kick at K1 must direct the beam on to the axis at K2 where a second kick can then remove the angle. Defining the position and angle of the bunch immediately before the kick at K2 as y_{K2} and θ_{K2} respectively, for a perfectly corrected beam it must be the case that:

$$\begin{pmatrix} y_{\mathrm{K2}} \\ \theta_{\mathrm{K2}} \end{pmatrix} + \begin{pmatrix} 0 \\ \delta\theta_{\mathrm{K2}} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \tag{4.12}$$

where $\delta\theta_{\mathrm{K2}}$ is the change in angle caused by K2. The position of the bunch immediately before K2 may be expressed in terms of the position of the bunch immediately before K1

(y_{K1}), the kick given at K1 ($\delta\theta_{\mathrm{K1}}$) and the transfer matrix from K1 to K2 ($^{\mathrm{K1}}_{\mathrm{K2}}M$):

$$\begin{pmatrix} y_{\mathrm{K2}} \\ \theta_{\mathrm{K2}} \end{pmatrix} = {}^{\mathrm{K1}}_{\mathrm{K2}}M \left[\begin{pmatrix} y_{\mathrm{K1}} \\ \theta_{\mathrm{K1}} \end{pmatrix} + \begin{pmatrix} 0 \\ \delta\theta_{\mathrm{K1}} \end{pmatrix} \right] \tag{4.13}$$

By substituting Equation 4.12 into Equation 4.13 it is apparent that in order to perfectly correct the beam the kicks must satisfy the equation:

$${}^{\mathrm{K1}}_{\mathrm{K2}}M \begin{pmatrix} y_{\mathrm{K1}} \\ \theta_{\mathrm{K1}} + \delta\theta_{\mathrm{K1}} \end{pmatrix} = \begin{pmatrix} 0 \\ -\delta\theta_{\mathrm{K2}} \end{pmatrix} \tag{4.14}$$

From the first row of this equation, the condition on $\delta\theta_{\mathrm{K1}}$ is obtained:

$$\delta\theta_{\mathrm{K1}} = -\frac{{}^{\mathrm{K1}}_{\mathrm{K2}}M_{11}}{{}^{\mathrm{K1}}_{\mathrm{K2}}M_{12}} y_{\mathrm{K1}} - \theta_{\mathrm{K1}} \tag{4.15}$$

And a similar condition on $\delta\theta_{\mathrm{K2}}$ is obtained from the second row:

$$\delta\theta_{\mathrm{K2}} = -{}^{\mathrm{K1}}_{\mathrm{K2}}M_{21} y_{\mathrm{K1}} - {}^{\mathrm{K1}}_{\mathrm{K2}}M_{22} \left(\theta_{\mathrm{K1}} + \delta\theta_{\mathrm{K1}}\right) \tag{4.16}$$

At this stage the two kicks are expressed in terms of the position and angle at K1. These can not be measured directly but may be expressed in terms of the position at P2, the position at P3, the elements of the transfer matrix from K1 to P2 and the elements of the transfer matrix from P2 to P3. This means that this algorithm can also be expressed in the form of Equation 4.5 with each gain parameter now given by a complicated combination of transfer matrix elements.

The effectiveness of this algorithm can be determined using a very simple model of the ATF beamline developed in MATLAB. An ensemble of 100,000 bunches is generated with normally distributed and uncorrelated position and angle at K1. Each bunch is then propagated downstream as far as MQF15X with the position at P2 and P3 being used to calculate the change in angle at K1 and K2 according to both the conventional algorithm implemented on the FONT5 board (described in Section 4.1.1) and the alternative algorithm that has been outlined here. The trajectories of bunch 1 and bunch 2 are set to be identical and the effect of BPM resolution (the BPMs are assumed to have the same resolution) is investigated by adding a random term to each position used in the feedback calculations. The results are given in Figure 4.12. Under the ideal condition of perfect bunch to bunch correlations the alternative algorithm is seen to provide better performance than the conventional algorithm; however, the improvement is at the 2% level for those resolutions routinely achieved from the stripline BPMs at the ATF (0.5 μm) and such a modest reduction in jitter would be difficult to observe. As the resolution of the BPM system improves the alternative algorithm becomes more attractive as the performance of the conventional algorithm is limited to 90 nm even for perfect BPMs; this small amount of residual jitter when using the conventional algorithm may be thought of as a result of the failure to explicitly modify the kick at K2 to take into account the effect of the kick at K1. The alternative algorithm does take this into account and is therefore able to deliver a perfectly stable beam at K2which in this simple model must then propagate as such ad infinitum.

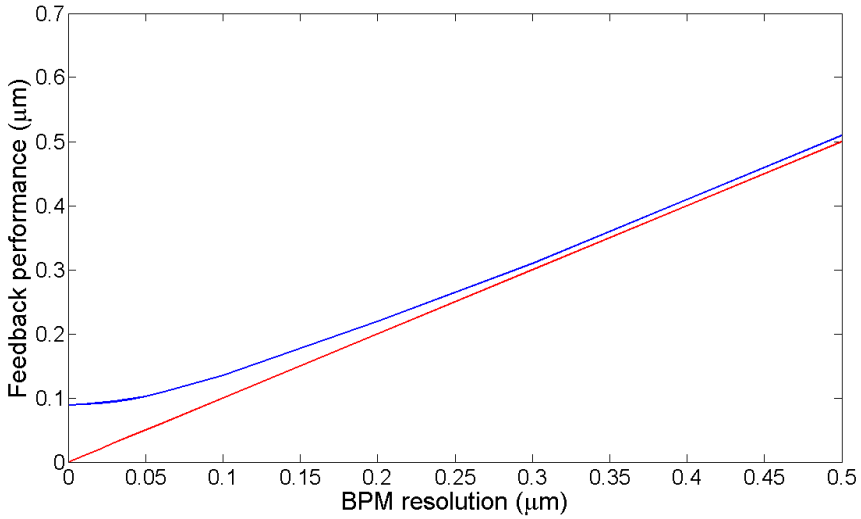

Figure 4.12: Simulated performance of the conventional feedback algorithm (blue line) and the alternative feedback algorithm (red line) as a function of BPM resolution.

4.5 Summary

The FONT feedback algorithm is constructed to deliver position stability of the second bunch at the feedback BPMs P2 and P3. The system has been demonstrated to achieve reduction in the position jitter of the second bunch by a factor of 5.5 at P2 and a factor of 3.1 at P3. The empirical relation between beam size and beam jitter in the ATF extraction line [59] is such that were this factor of 5.5 maintained to the nominal interaction point, goal 2 of the ATF2 collaboration would have been achieved. For a constant BPM resolution, the correction factors are heavily dependent on the bunch to bunch correlation at each feedback BPM. The latency of the system has been measured to be 156 ns and is in line with expectations.

Chapter 5

Design and Simulation of a Long Train Feedback System

The FONT feedback system discussed thus far is capable of performing feedback on short bunch trains consisting of either two or three bunches. This chapter will discuss how the existing FONT hardware could be used to perform feedback on longer bunch trains. The scenario that will be considered is a bunch train of the type that will be delivered by a future linear collider; that is, consisting of several thousand bunches. This scenario required a modified design for the FONT5 firmware and, in order to investigate the performance that could be achieved using these designs, it was also necessary to develop a flexible simulation of a feedback system installed in an accelerator beamline.

5.1 Simulation of a feedback system using FONT hardware

The principle of simulation of a feedback system using FONT hardware is illustrated in Figure 5.1. If the timing signals necessary for the operation of the standard FONT5 firmware can be provided, all that is required to simulate the behaviour of a BPM is a pair of analogue signals to approximate the Δ and Σ_I outputs of a FONT BPM processor module. If the amplitude of these pseudo-BPM signals can be modified in response to the kicker drive signal generated by the FONT5 board on the bunch to bunch timescale then a virtual kicker can be realised and position feedback can be simulated. A suitable source for both the FONT5 timing signals and the pseudo-BPM signals is the FONT4 digital board (Figure 5.2) [60].

5.1.1 FONT4 digital board

The FONT4 board was developed as part of the first incarnation of the FONT feedback system to incorporate digital components. Originally designed to monitor the signals from a single BPM each, a total of four were produced. Each FONT4 board features a Xilinx Virtex-4 FPGA, a pair of ADCs and a pair of DACs along with several clock inputs and

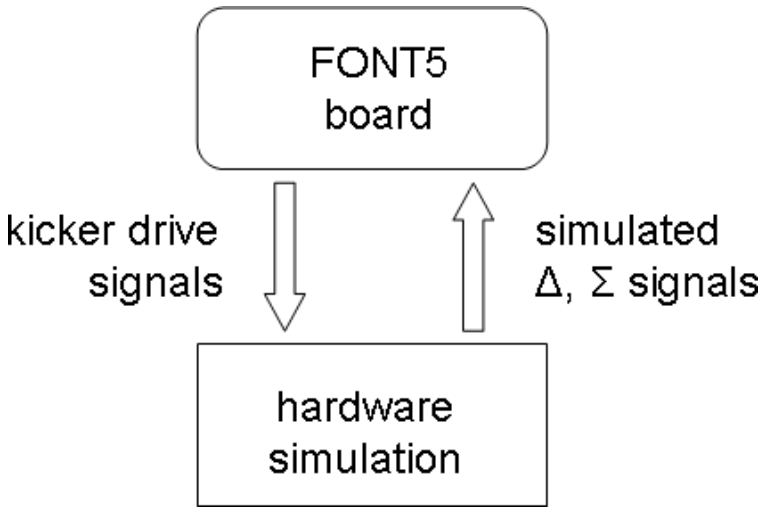

Figure 5.1: Schematic illustrating the principle of simulation of a feedback system using FONT hardware.

outputs. This arrangement of ADCs and DACs controlled by an FPGA makes a FONT4 board ideal for use as part of the simulation: the four clock outputs allow a single FONT4 board to provide all the timing signals needed to operate the FONT5 firmware, the two DACs can be used to simulate the Δ and Σ_I signals from a single BPM and the pair of ADCs allow the measurement of both analogue outputs of the FONT5 board. A single FONT4 board is thus able to represent a beamline that consists of a single BPM and a single kicker, and hence approximate the situation envisaged at the ILC. Alternatively, a pair of FONT4 boards may be used in order to reproduce the behaviour of a dual BPM, dual kicker system and serve as a replica of the FONT system as deployed at the ATF.

In the dual BPM, dual kicker scheme each FONT4 board produces either Δ or Σ_I signals. Figure 5.3 shows how a common source of 357 MHz is used to drive all three digital boards. The DAC outputs from one FONT4 board are used as the Σ_{P2} and Σ_{P3} ADC inputs to the FONT5 board. The other board provides the corresponding Δ signals, as well as the trigger and 2.16 MHz clock necessary for FONT5 board operation. The ADCs of this FONT4 board are set to monitor the DAC outputs of the FONT5 board, allowing the Δ outputs of the FONT4 board to be modified based on the kick signals provided by the FONT5 board. The development of the FONT4 firmware to deliver this behaviour forms the subject of the rest of this section. The design specification for the FONT4 feedback simulation firmware was that it be capable of performing the following tasks:

- Generation of FONT5 timing signals
- Simulation of BPM signals
- Communication over serial connection

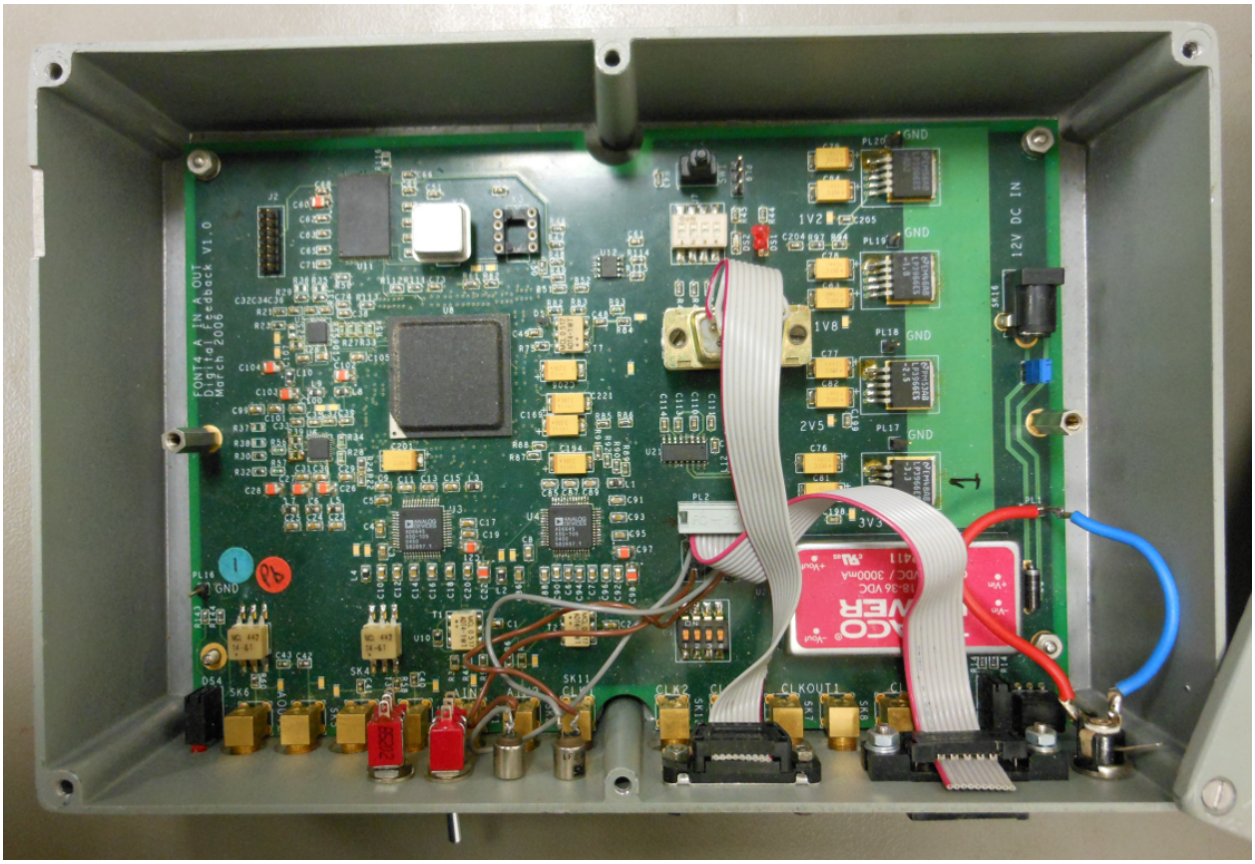

Figure 5.2: Photograph of the FONT4 digital board.

5.1.2 Generation of FONT5 timing signals

The FONT5 firmware requires a 357 MHz clock, a 2.16 MHz clock and a trigger signal in order to function. In the laboratory, the 357 MHz clock input of the FONT5 board is most conveniently driven by a dedicated clock generator. The output of this generator is split in order that it may also be used as a clock input for the FONT4 board, ensuring timing stability between the simulated signals and the device observing them. As with the FONT5 feedback firmware, the bulk of the logic of the FONT4 feedback simulation firmware is driven by this input 357 MHz clock.

The ADCs of the FONT4 board have a maximum sampling frequency of 90 MHz and thus it is necessary to frequency divide the 357 MHz clock by a factor of 5 to produce a 71.4 MHz clock; the factor of 5 was chosen so that there would be an integer number of 71.4 MHz cycles in a cycle of the 2.16 MHz clock [1] to aid in timing stability. This is then used to clock both the ADCs and the DACs; furthermore, it is itself frequency divided by a factor of 33 to produce the 2.16 MHz clock required for operation of the FONT5 board. The remaining signal, the trigger, is generated by the FONT4 board by counting cycles of the 2.16 MHz clock. Two parameters are used to determine the temporal structure of the output trigger. The first parameter, the trigger period, is the number of 2.16 MHz clock

[1] A single cycle of the 2.16 MHz clock is equivalent to 33 cycles of the 71.4 MHz; 2.16 MHz is an approximation to the actual frequency.

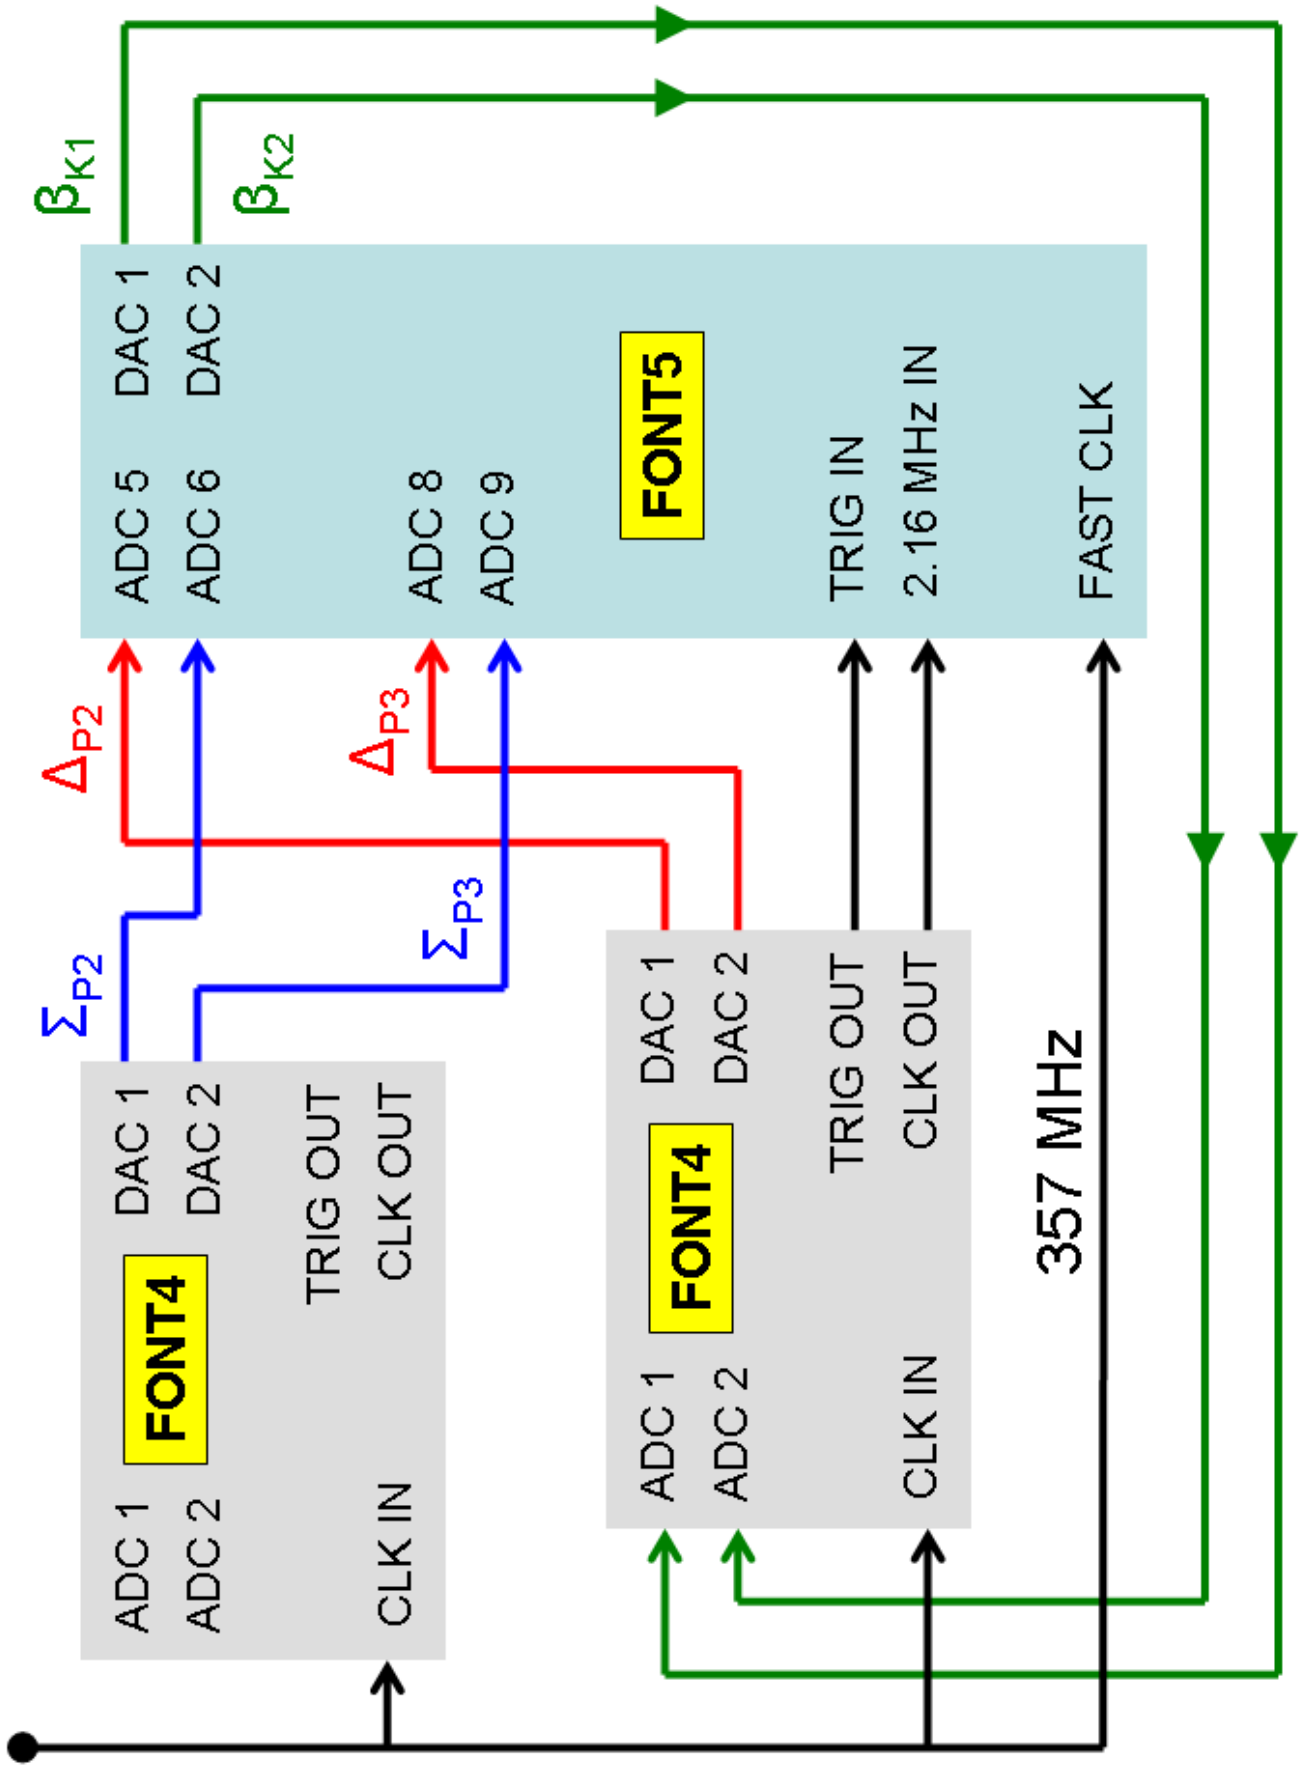

Figure 5.3: Schematic of the simulated dual BPM, dual kicker system.

cycles the counter achieves before it is reset to zero. The other parameter, the trigger select, is the value of the 2.16 MHz cycle counter that actually results in a logical one of the trigger signal.

5.1.3 Simulation of BPM signals

The simulation of the BPM signals consists of two separate tasks performed independently. The temporal structure of the simulated bunch train refers to the times at which the FONT4 DAC outputs have non-zero values. It is analagous to the bunch arrival time and is determined by a number of parameters that will be described below. The amplitude structure of the simulated bunch train refers to exactly what these non-zero values of the FONT4 DAC outputs are. These values are analagous to the bunch position and are determined by an algorithm that effectively describes the nature of the simulated beamline.

5.1.3.1 Temporal structure of the bunch train

Figure 5.4 defines the parameters that are relevant to the temporal structure of the simulated bunch train. The trigger signal, as well as being output for use by the FONT5 board, is used internally to initiate a counter of ring clock cycles designated the bunch strobe enable generation counter. After this counter reaches the value set by the trigger delay parameter, the bunch strobe enable is asserted and remains high for the number of ring clock cycles specified by the train length parameter. The bunch strobe enable signal defines the extent of the simulated bunch train; while the bunch strobe enable is low, the values on the DACs are held at zero. While the bunch strobe enable is high, a counter of 71.4 MHz cycles called the bunch strobe generation counter is active. This counter is used to control the state of the bunch strobe, the logical one of which determines the point at which the value of the DAC outputs is updated. The bunch strobe generation counter is reset to zero when it reaches the value specified by the bunch spacing parameter; as the bunch strobe has a value of 1 only while the bunch strobe generation counter has a value of zero, the resulting form of the bunch strobe is a series of pulses of duration 1 cycle of 71.4 MHz spaced regularly by the number of cycles of 71.4 MHz specified by the bunch spacing parameter.

5.1.3.2 Amplitude structure of the bunch train

Having described the parameters that determine exactly when the FONT4 DACs will have non-zero values, the method of determining exactly what this non-zero value should be will now be discussed.

5.1.3.2.1 Beamline model for simulation of a feedback system using FONT hardware

For maximum utility, the beamline to be modelled is a minimal version of that used for feedback at the ATF; that is, the model beamline consists of four locations: two kickers, K1

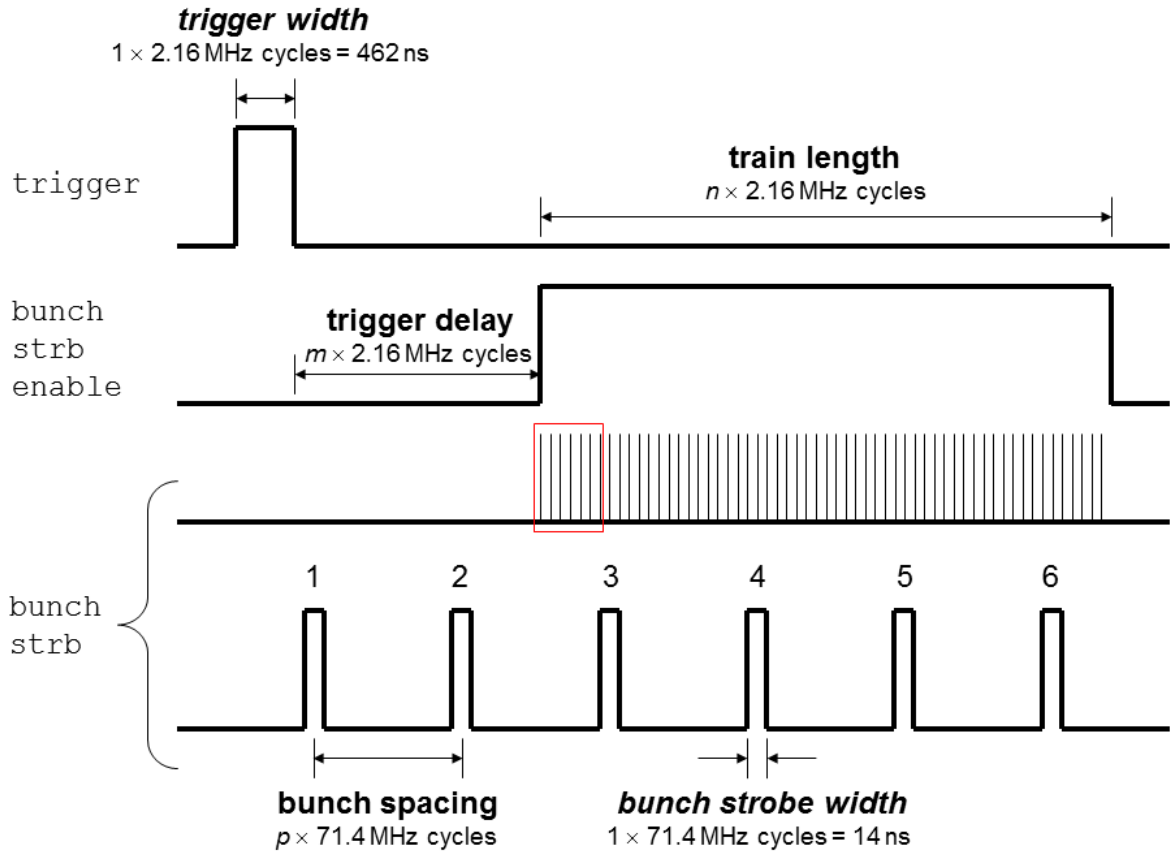

Figure 5.4: Diagram illustrating the significance of the timing parameters. The thick black lines describe the binary state of the named signals as a function of time. The bottom two lines both represent the bunch strobe; the red rectangle indicates the region that is presented on a shorter timescale below.

and K2, and two BPMs, referred to as P2 and P3 in order to conform to the convention established by the FONT system while deployed at ATF. Recalling the equations of beam transport from Chapter 4, it is possible to write the only directly observable quantities of the system (the position of the beam at the two BPMs, y_{P2} and y_{P3}) as a function of the initial properties of the beam (i.e. the position and angle at K1) and the kicks imparted at K1 and K2. Consider the position at P2:

$$y_{\mathrm{P2}} = {}^{\mathrm{K2}}_{\mathrm{P2}}M_{11}y_{\mathrm{K2}} + {}^{\mathrm{K2}}_{\mathrm{P2}}M_{12}\left(\theta_{\mathrm{K2}} + \delta\theta_{\mathrm{K2}}\right) \tag{5.1}$$

The beam properties at K2 are given by the following pair of equations:

$$\begin{aligned} y_{\mathrm{K2}} &= {}^{\mathrm{K1}}_{\mathrm{K2}}M_{11}y_{\mathrm{K1}} + {}^{\mathrm{K1}}_{\mathrm{K2}}M_{12}\left(\theta_{\mathrm{K1}} + \delta\theta_{\mathrm{K1}}\right) \\ \theta_{\mathrm{K2}} &= {}^{\mathrm{K1}}_{\mathrm{K2}}M_{21}y_{\mathrm{K1}} + {}^{\mathrm{K1}}_{\mathrm{K2}}M_{22}\left(\theta_{\mathrm{K1}} + \delta\theta_{\mathrm{K1}}\right) \end{aligned} \tag{5.2}$$

By making these substitutions it is possible to express the position at P2 as a function of the four variables y_{K1}, θ_{K1}, $\delta\theta_{\mathrm{K1}}$ and $\delta\theta_{\mathrm{K2}}$:

$$y_{\mathrm{P2}} = {}^{\mathrm{P2}}k_{y}y_{\mathrm{K1}} + {}^{\mathrm{P2}}k_{\theta}\theta_{\mathrm{K1}} + {}^{\mathrm{P2}}k_{\delta\theta_{\mathrm{K1}}}\delta\theta_{\mathrm{K1}} + {}^{\mathrm{P2}}k_{\delta\theta_{\mathrm{K2}}}\delta\theta_{\mathrm{K2}} \tag{5.3}$$

where:

$$
\begin{aligned}
{}^{\mathrm{P2}}k_y &= {}^{\mathrm{K2}}_{\mathrm{P2}}M_{11}{}^{\mathrm{K1}}_{\mathrm{K2}}M_{11} + {}^{\mathrm{K2}}_{\mathrm{P2}}M_{12}{}^{\mathrm{K1}}_{\mathrm{K2}}M_{21} \\
{}^{\mathrm{P2}}k_\theta &= {}^{\mathrm{K2}}_{\mathrm{P2}}M_{11}{}^{\mathrm{K1}}_{\mathrm{K2}}M_{12} + {}^{\mathrm{K2}}_{\mathrm{P2}}M_{12}{}^{\mathrm{K1}}_{\mathrm{K2}}M_{22} \\
{}^{\mathrm{P2}}k_{\delta\theta_{\mathrm{K1}}} &= {}^{\mathrm{P2}}k_\theta \\
{}^{\mathrm{P2}}k_{\delta\theta_{\mathrm{K2}}} &= {}^{\mathrm{K2}}_{\mathrm{P2}}M_{12}
\end{aligned}
\tag{5.4}
$$

A corresponding set of expressions is obtained for P3 by substituting the elements of ${}^{\mathrm{K2}}_{\mathrm{P2}}M$ with those of ${}^{\mathrm{K2}}_{\mathrm{P3}}M$. The positions at P2 and P3 are thus a simple function of the initial beam conditions, the angular deflections at K1 and K2, and two sets of coefficients that are derived from the beam transfer matrices. By way of illustration, the values of these coefficients as calculated from the transfer matrices corresponding to the nominal state of the ATF are given in Table 5.1.

Parameter	Value
${}^{\mathrm{P2}}k_y$	1.10
${}^{\mathrm{P2}}k_\theta$	2.04
${}^{\mathrm{P2}}k_{\delta\theta_{\mathrm{K1}}}$	2.04
${}^{\mathrm{P2}}k_{\delta\theta_{\mathrm{K2}}}$	0.53
${}^{\mathrm{P3}}k_y$	-1.32
${}^{\mathrm{P3}}k_\theta$	-1.32
${}^{\mathrm{P3}}k_{\delta\theta_{\mathrm{K1}}}$	-1.78
${}^{\mathrm{P3}}k_{\delta\theta_{\mathrm{K2}}}$	0.63

Table 5.1: Values of the k parameters for simulation of the FONT region at ATF.

5.1.3.2.2 Implementation of the model in the FONT4 firmware

In order to put this model into practice three elements are required: first, a source of the values y_{K1} and θ_{K1} that represent the initial condition of each bunch in each train; second, a source of the values $\delta\theta_{\mathrm{K1}}$ and $\delta\theta_{\mathrm{K2}}$ that represent the kicks applied to each bunch by the FONT5 board and finally, some means of including the eight coefficients derived from the beam transfer matrices of the model (Table 5.1). The former four values are expected to vary on the bunch-to-bunch timescale and will thus be referred to as the variable inputs to the FONT4 simulation calculation; the latter eight will be referred to as the coefficient inputs.

For maximum flexibility, the values y_{K1} and θ_{K1} should be capable of assuming arbitrary values dynamically. That is, it should be possible for them to vary both from bunch to bunch and from trigger to trigger. Varying the values from trigger to trigger is necessary to simulate train-to-train jitter (variation of the mean position of the train). The ability

for the values to vary from bunch to bunch enables the description of train structure (i.e. a predictable, static dependence of bunch position on location within the train) in addition to more exotic effects such as a train structure that varies from trigger to trigger.

As the FPGA includes onboard memory blocks, each y_{K1} and θ_{K1} pair is conveniently represented in firmware as the output of a pair of memory blocks. These memory blocks are designated RAM1 and RAM2, respectively. Both the ADCs and the DACs of the FONT4 board have 14-bit resolution and so this is the logical width of the data bus for each RAM. The FONT5 board does not include external memory modules and so the size of RAM1 and RAM2 is limited by the available resources on the FPGA; ultimately the width of the address bus was set at 12-bit to allow for trains consisting of up to 4096 unique bunches. A signal called the bunch counter is used as the address input of both memory blocks. This counter is set to zero by the rising edge of the trigger and then incremented for every rising edge of the bunch strobe thereafter; as a result, the values presented at the data outputs of RAM1 and RAM2 cycle through their preloaded values over the length of the simulated bunch train. The method of populating these RAMs will be detailed in a later section.

$\delta\theta_{\mathrm{K1}}$ and $\delta\theta_{\mathrm{K2}}$ are the kicks applied to each bunch at K1 and K2 respectively. During operation of the FONT feedback system at ATF, the actual kicks are determined by the two output signals from the DACs on the FONT5 board. In the case of a test-bench capable of simulating feedback these FONT5 DAC outputs must be used as the FONT4 ADC inputs. As the FONT4 ADCs sample continuously at a rate of 71.4 MHz it is necessary to implement a method of selecting those samples that correspond to the kicks calculated by the FONT5 board; this will be considered in the next section.

The final element of the model concerns the coefficient inputs that describe the model beamline. Recall that the position at each BPM is defined in Equation 5.3. In firmware, this equation appears as:

$$V = {}^{\mathrm{DAC1}}k_{\mathrm{RAM1}}P + {}^{\mathrm{DAC1}}k_{\mathrm{RAM2}}Q + {}^{\mathrm{DAC1}}k_{\mathrm{ADC1}}A + {}^{\mathrm{DAC1}}k_{\mathrm{ADC2}}B \tag{5.5}$$

where V denotes the value that will be applied to DAC1 on the FONT4 board, P and Q are the 14-bit values stored in RAM1 and RAM2 respectively and A and B are the values selected from ADC1 and ADC2 of the FONT4 board respectively. Depending on how the FONT4 DACs are connected to the FONT5 ADCs, V will be interpreted either as the Σ_I or Δ output of a stripline BPM processor. In the case that V is used as a Δ signal, the parameter ${}^{\mathrm{DAC1}}k_{\mathrm{RAM1}}$ represents the coefficient ${}^{\mathrm{P2}}k_y$. As these coefficients (Table 5.1) fall approximately in the range -2 to 2, they are conveniently represented in the firmware as seven-bit two's complement numbers stored in the form of control registers (seven-bit control registers are preferred for the following reason: as the FONT5 board transmits data as a series of bytes, in each case the most significant bit is used to mark the framing bytes that facilitate interpretation of the bytestream by the DAQ software). Each parameter is thus expressed as a number in the range -64 to 63 and a division by 32 is required in the FONT4 simulation calculation to scale the parameters to their correct range.

5.1.3.2.3 Role of the bunch strobe

The bunch strobe was introduced earlier as the signal that determines the temporal structure of the bunch train. Ultimately, all the tasks involved in producing the simulated output are controlled by signals derived from the bunch strobe. These signals are listed below, along with a brief description of their function:

- **ram read strobe**
  Clock data from the current address of RAM1 and RAM2 for use in the simulation calculation.
- **adc strobe**
  Register the current samples from ADC1 and ADC2 for use in the simulation calculation.
- **dac strobe**
  Clock the outputs of the simulation calculation on to the DACs.
- **ram write strobe**
  Clock the registered ADC values and DAC output values in to the diagnostic RAMs.
- **count strobe**
  Increment the bunch counter used to address the RAMs.

All five strobes are delayed versions of the master bunch strobe; crucially, it is possible to exercise individual control over each delay. This is achieved by counting 71.4 MHz cycles after the arrival of the bunch strobe and asserting each strobe only for a specific value of this 71.4 MHz counter. The motivation for such a scheme is illustrated in Figure 5.5. The diagram depicts the situation where a FONT5 board is used to perform feedback based on simulated Δ and Σ_I signals from a FONT4 board.

First, consider just the signals (e) and (f). Signal (e) represents the DAC output of the FONT5 board as a function of time, while signal (f) represents the state of the FONT5 bunch strobe during the same interval (the FONT5 bunch strobe will be described in more detail in the next section; for now we note that it determines which FONT5 ADC samples are input to the feedback calculation and that it is set by the user via the software interface). As indicated in the diagram, a fixed time elapses between the logical high of the FONT5 bunch strobe and the appearance on the FONT5 DAC of the result of the FONT5 feedback calculation; this delay is labelled the FONT5 board latency and corresponds to the time it takes to execute the FONT5 feedback algorithm.

Now, consider signal (d). This signal represents the DAC output of the FONT4 board, which is then used as an ADC input of the FONT5 board (i.e. as a Δ or Σ_I signal). Note how the FONT5 bunch strobe is set to coincide with the non-zero values of the FONT4 DAC output; the vertical blue lines indicate the FONT5 ADC samples that are input to the FONT5 feedback calculation. On the other hand, signal (c) is by definition coincident with the FONT4 DAC output; this signal, the FONT4 dac strobe, determines when the result of the FONT4 simulation calculation is supplied to the FONT4 DAC outputs. The FONT4 dac strobe is a delayed version of signal (a), the FONT4 bunch strobe, the generation of

which was described in Section 5.1.3.1. The FONT4 dac strobe may be manipulated at will in order to alter the arrival time of the simulated bunches, with the proviso that it must always follow signal (b), the FONT4 adc strobe.

The FONT4 adc strobe is used to capture a particular value of the FONT4 ADC input for use in the FONT4 simulation calculation. In the diagram the FONT4 adc strobe is set to coincide with the non-zero value of the FONT5 DAC output; the vertical red lines indicate the FONT4 ADC samples that are input to the FONT4 simulation calculation. As the diagram portrays the first two bunches of a train, the first value captured by the FONT4 ADC should be zero in order to reflect the lack of a correction applied to the first bunch. The requirement for the FONT4 adc strobe to precede the FONT4 dac strobe is thus simply a consequence of the time taken to operate on the FONT4 ADC inputs in order to produce the result of the FONT4 simulation calculation. In summary the FONT4 simulation calculation (Equation 5.5) proceeds as follows:

- The ram read strobe and the adc strobe get the variable inputs to the FONT4 simulation calculation (the coefficient inputs are constantly available as control register outputs).
- All variables and coefficients are sign extended to 28-bit. Sign extension increases the number of bits of a binary number while preserving the number's sign and value and is necessary in order to perform the required digital calculations.
- Each variable is multiplied by the corresponding coefficient to produce four 28-bit products.
- The four products are added together.
- The most significant seven bits and least significant five bits are discarded to give a sixteen-bit result; the former operation counteracts the sign extension and the latter operation scales the coefficients to their correct range (-2 to 2).
- The result of the calculation is saturated at fourteen bits in order to match the range of the DACs.
- The dac strobe loads the result of the calculation on to the FONT4 DACs.
- The ram write strobe clocks both the registered ADC values and the output DAC values into diagnostic RAMs.
- The count strobe increments the bunch counter in order to address the next elements of RAM1 and RAM2 at the next pulse of the bunch strobe.

5.1.3.2.4 Form of the FONT4 DAC outputs

Figure 5.6 shows FONT5 digitized examples of a FONT4 DAC output signal as well as the Σ_I output signal corresponding to actual beam data; the standard FONT5 firmware was used in both cases but only a portion of the waveform is presented. The low maximum clock

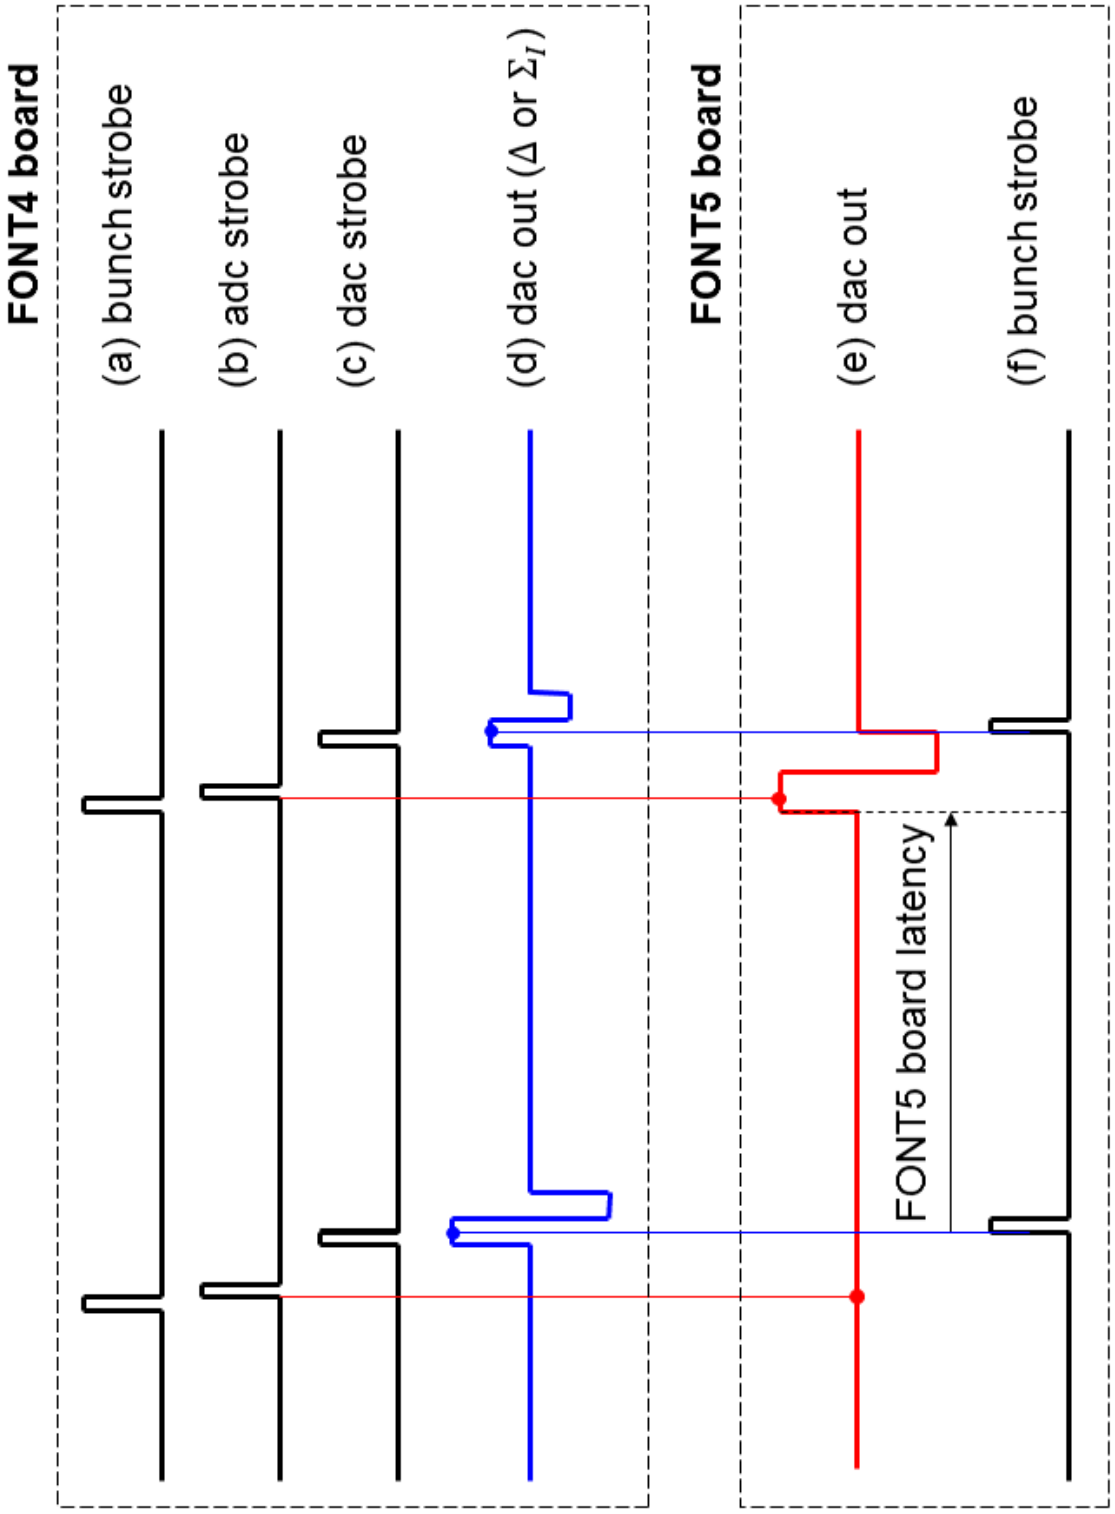

Figure 5.5: Diagram indicating the role of the bunch strobe and associated signals; each of the lines (a) to (f) represents the voltage of a signal as a function of time. Not to scale.

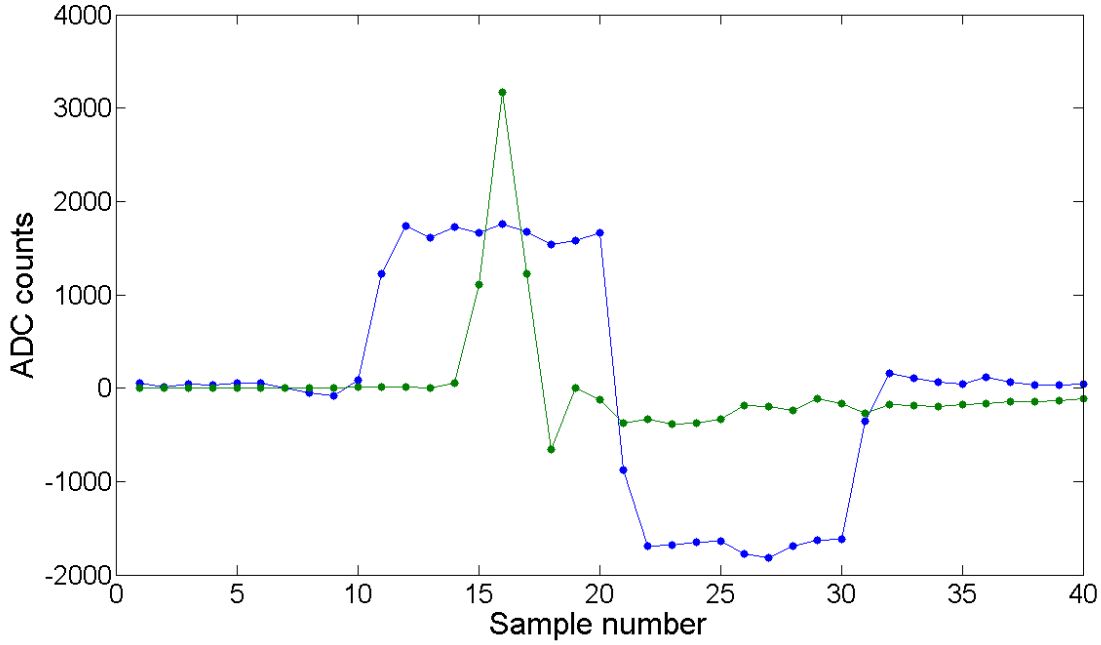

Figure 5.6: Sections of a pair of waveforms digitized by the FONT5 board (one sample = 2.8 ns). Blue: DAC output from the FONT4 testbench; Green: actual Σ_I output of a FONT BPM processor module while monitoring a BPM at ATF.

speed of the DACs on the FONT4 board results in a wide square pulse which is a poor match to the output of the actual BPM processor modules. As the FONT5 feedback calculation determines the position of a bunch using a single sample from each of the digitized Δ and Σ_I waveforms, the actual structure of the simulated signals is largely irrelevant (although it prevents the simulation from being used to study the effect of timing jitter on the analogue inputs to the FONT5 board).

The other notable feature of the plot is the bipolar nature of the FONT4 DAC output. This is necessary due to a limitation of the FONT4 board. A unipolar pulse is produced by sending a constant value to the DAC and then, after a time equal to the desired pulse width, sending a zero. When the FONT4 DAC is set to produce a train of such unipolar pulses, the actual voltage achieved by the DAC diminishes as a function of pulse number. This problem goes away if a bipolar pulse is used instead. A bipolar pulse is generated by sending a constant value to the DAC and then sending the same value with a sign change before sending the zero. By ensuring that the DAC spends equal time providing positive and negative voltages, the measured voltage of all pulses in the train is constant. The two scenarios are illustrated in Figure 5.7.

5.1.3.2.5 Communication over serial connection

The ability to communicate with a serial interface is required in order to modify the control registers and load new values into the RAMs. This logic closely follows that of the FONT5 firmware. The single control register module operates on the domain of the 40 MHz clock

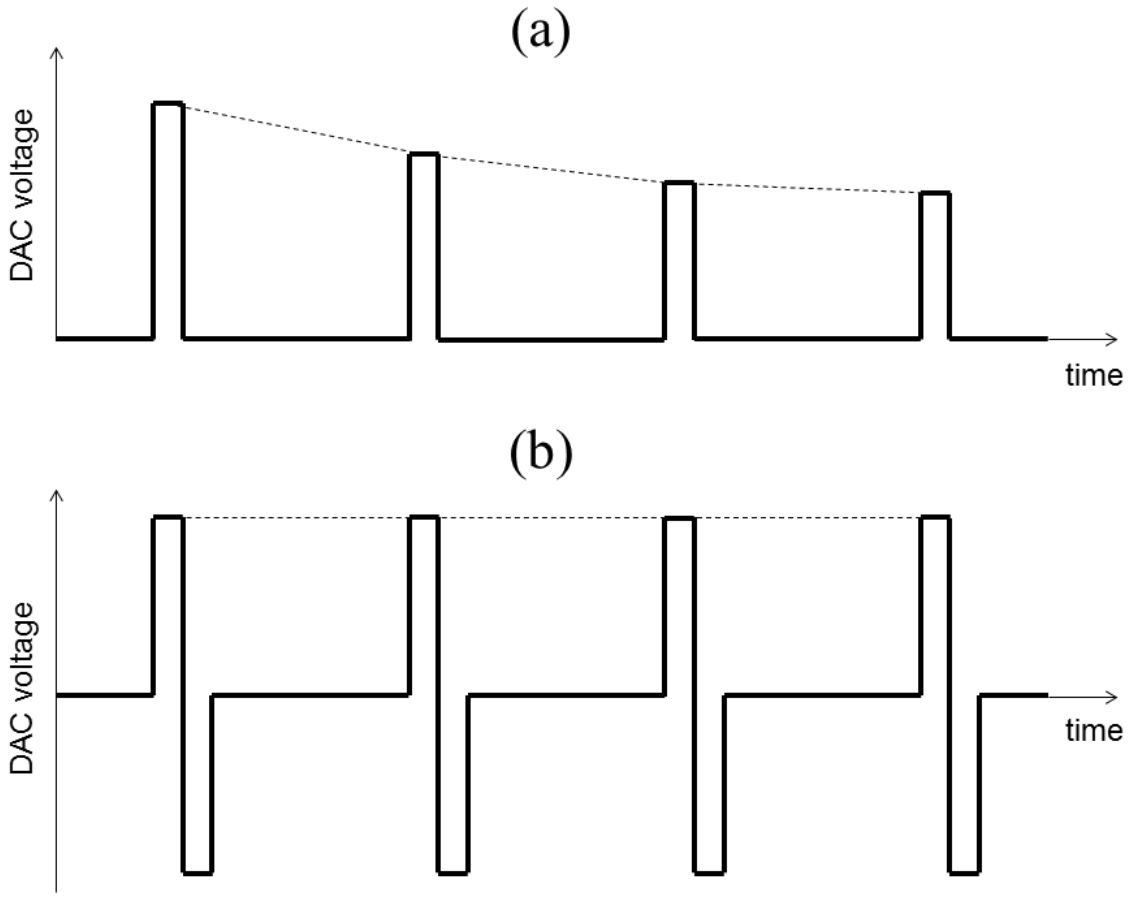

Figure 5.7: Schematic illustrating the behaviour of the FONT4 DAC when set to produce a train of: (a) unipolar pulses; (b) bipolar pulses.

provided by the onboard oscillator of the FPGA. A list of control registers is given in Table 5.2. In addition, for diagnostic purposes the FONT4 test bench firmware can be set to transmit data over a serial link every trigger. This includes the ADC values registered over the duration of the bunch train and the values sent to the DACs, i.e. the simulated BPM signal values. The data returned are described in Table 5.3. To prevent buffer overflow, the FONT4 firmware prohibits the generation of bunch trains while the board is transmitting data; this effectively limits the trigger rate, particularly when the train length parameter is large. This limit can be bypassed by disabling data transmission entirely, in which case it is possible to achieve repetition rates matching the 5 Hz design for the ILC even for a train of 2,820 bunches.

5.2 FONT5 firmware for long bunch trains

The baseline FONT5 firmware records 164 samples at a rate of 357 MS/s. This equates to a sampling window of 459.2 ns duration and is sufficient to capture the entirety of a typical ATF bunch train, consisting of two or three bunches separated in time by ∼150 ns. In order

Parameter	Size	Control register address
Trigger period	21-bit	2, 1, 0
Trigger select	21-bit	5, 4, 3
Trigger delay	8-bit	7, 6
Train length	8-bit	9, 8
Bunch spacing	8-bit	11, 10
${}^{\mathrm{DAC1}}k_{\mathrm{RAM1}}$	7-bit	13
${}^{\mathrm{DAC1}}k_{\mathrm{RAM2}}$	7-bit	14
${}^{\mathrm{DAC1}}k_{\mathrm{ADC1}}$	7-bit	15
${}^{\mathrm{DAC1}}k_{\mathrm{ADC2}}$	7-bit	16
${}^{\mathrm{DAC2}}k_{\mathrm{RAM1}}$	7-bit	17
${}^{\mathrm{DAC2}}k_{\mathrm{RAM2}}$	7-bit	18
${}^{\mathrm{DAC2}}k_{\mathrm{ADC1}}$	7-bit	19
${}^{\mathrm{DAC2}}k_{\mathrm{ADC2}}$	7-bit	20

Table 5.2: List of the parameters used for the FONT4 test bench firmware whose values are stored in control registers. Parameters that occupy more than one control register have the relevant addresses listed in descending order of significance.

to record ILC-like bunch trains consisting of ∼3,000 bunches over a period of close to 1 ms, several elements of the FONT5 firmware had to be redesigned.

The previous data storage scheme used a store strobe signal to define the period over which data should be recorded; as long as the store strobe was active, each ADC sample was stored in the corresponding RAM for that ADC channel. Such a scheme is not desirable for recording long bunch trains as the minimum clock frequency of the FONT5 ADCs makes it impractical for them to be operated at a rate other than 357 MS/s and to record data at such a high frequency for a period of 1 ms would require more memory than is available on the FPGA.

The second necessary modification relates to the generation of the bunch strobes that control which ADC samples are used to perform the FONT5 feedback calculation. The baseline firmware explicitly specified the location of each bunch within the sequence of ADC samples for all three BPMs. This is a perfectly rational choice when the number of bunches is limited to a maximum of three but becomes an increasingly inelegant solution as the number of bunches increases.

The first of these issues was resolved in the long train FONT5 firmware by abandoning the store strobe as the arbiter as to whether a given ADC sample is stored in RAM. By using the bunch strobes instead the ADCs may continue to be operated at a frequency of 357 MHz but the vast majority of the samples will be immediately discarded, with only one sample per bunch, the input of the feedback calculation, being retained for later transmission. This results in only a modest increase in the use of memory resources on the FPGA as the RAM

Description	Number of bytes
Firmware version framing byte	1
Firmware version byte	1
Board timestamp framing byte	1
Board timestamp byte	1
RAM1 (y) framing byte	1
RAM1 (y) data bytes	$2n$
RAM2 (θ) framing byte	1
RAM2 (θ) data bytes	$2n$
RAM3 (ADC1) framing byte	1
RAM3 (ADC1) data bytes	$2n$
RAM4 (ADC2) framing byte	1
RAM4 (ADC2) data bytes	$2n$
RAM5 (DAC1) framing byte	1
RAM5 (DAC1) data bytes	$2n$
RAM6 (DAC2) framing byte	1
RAM6 (DAC2) data bytes	$2n$
Control registers framing byte	1
Control registers bytes	21
Termination byte	1

Table 5.3: Data returned by the FONT4 test-bench firmware. n is the number of bunches that make up the train and is determined by the train length and bunch spacing parameters.

modules need only be enlarged in order to store a total of 4,096 14-bit values.

The second issue required the adoption of the bunch strobe logic from the FONT4 test-bench. In the long train firmware, each bunch strobe is generated from three parameters: the location of the first bunch and the bunch spacing, each expressed in cycles of 357 MHz, and a store length parameter given in units of ring clock cycles. The bunch strobes are held low outside of the window defined by the trigger input delay and the store length parameter such that increasing the store length increases the number of bunches returned; hence, unlike the baseline firmware, which can only return 164 samples at a frequency of 357 MHz, the long train firmware is able to return a variable number of samples at a frequency determined by the bunch spacing parameter. For example, a bunch spacing of 110 cycles of 357 MHz corresponding to 308 ns gives an effective sample rate of 3.25 MS/s with the exact number of samples returned depending on the store length.

A further consideration exists in the form of the potentially much greater amount of data that now has to be returned each trigger. The ADC samples account for the vast majority of the bytes returned by the FONT5 board and, as each sample consists of two bytes, the total number of ADC bytes sent every trigger is equal to $9 \times 2 \times 164 = 2,952$ for the baseline FONT5 firmware. The maximum data transmission rate over RS-232 that may be achieved with the FONT5 board is 460.8 kbps; thus, using the baseline firmware it is possible to transmit all the ADC data in just over 50 ms. Clearly if the number of ADC samples is

increased to 2,820 while the time between triggers is reduced to just 200 ms, the transmission time of the FONT5 data will be far in excess of the trigger period and thus it will no longer be possible for all the data from a trigger to be transmitted before the arrival of the next trigger. A mechanism to prevent the relevant firmware module receiving a trigger while data transmission is in progress was thus implemented in order to guard against corruption of the transmitted data. This method is not an ideal solution as it requires that a large fraction of the triggers received are simply ignored.

As increasing the transmission speed is not an option, the only possible way to successfully capture every trigger at the 5 Hz rate is to reduce the amount of data sent. The sequencer module contains the order in which the contents of the data RAMs are transmitted along with the values of the framing bytes that signal to the DAQ that the following data bytes are from a particular RAM. The sequencer module can be modified so that at each stage of transmission it first consults the value of a one-bit flag from the control registers to decide whether or not to actually send the appropriate data. Unfortunately, the only way to reduce the transmission time within the 200 ms window while storing at least the 2,820 samples required for an ILC-like bunch train requires sending at most the data from a single ADC. While this would be an acceptable, if undesirable, mode of operation on an actual machine where independent measurements could be made to determine the efficacy of the feedback correction, for the purposes of this demonstration the decision was made to increase the trigger period in order to record more data each trigger. The trigger period used was three times that of the ATF trigger period of 640 ms.

Finally, the DACs on the FONT5 board are subject to the same requirement as those on the FONT4 board: the average DAC output must be zero in order to prevent unacceptable droop while trying to maintain a constant output. The previous FONT5 DAC clocking scheme, which saw the result of the feedback calculation clocked on to the DAC a fixed number of cycles after the bunch arrived (and remain there until replaced by the value calculated for the subsequent bunch) had to be extended to include a second event where the feedback calculation multipled by -1 was clocked on to the DAC and a third event which zeroed the DAC in between bunches.

5.3 Results of simulated feedback system

With a design for the FONT5 firmware permitting the observation of long bunch trains and a suitable substitute for the BPM processor signals, it is possible to verify correct operation of the long train firmware.

5.3.1 Simulated bunch trains

The results presented here used a single FONT4 board loaded with the test-bench firmware in conjunction with the FONT5 board loaded with the long train firmware. The setup is depicted schematically in Figure 5.8; it represents a single BPM, single kicker beamline. The FONT4 test-bench was set to provide an ILC-like bunch train; that is, a total of 2,820

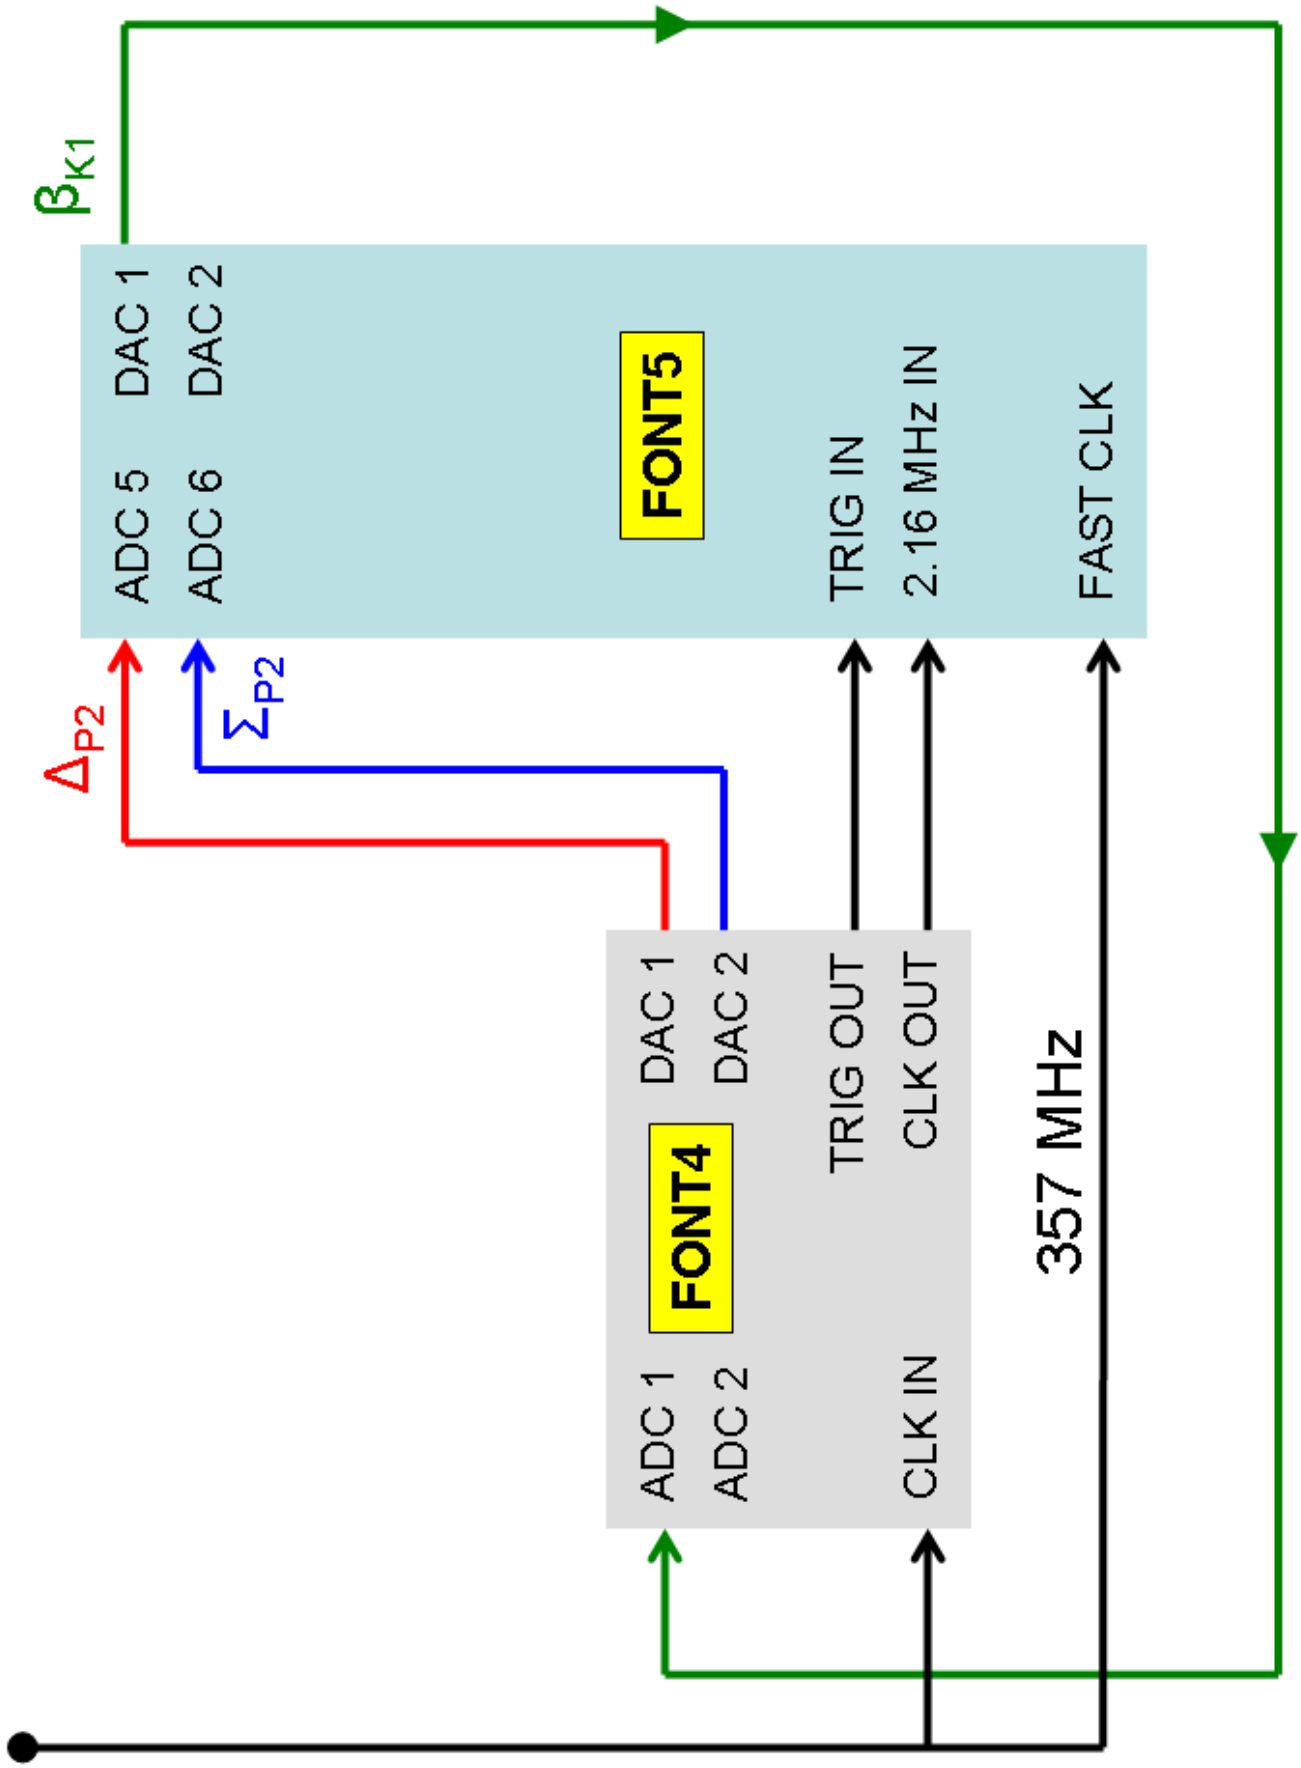

Figure 5.8: Schematic of the actual setup used to investigate performance of the long train FONT5 firmware.

bunches with a spacing of 308 ns. The Σ_I output of the FONT4 board was set to give the largest positive value it could (8191 FONT4 DAC counts $\equiv$ 500 mV). Figure 5.9 illustrates the signal as recorded by the ADCs of the FONT5 board. The FONT5 bunch strobe is configured such that the first stored ADC sample corresponds to the first bunch of the simulated train while the store length is set slightly higher than necessary (3,000 ring clock cycles or just under 1.4 ms) in order to capture some of the inter-train zero voltage region. The plot is consistent with a FONT5 board that recorded exactly 2,820 non-zero samples

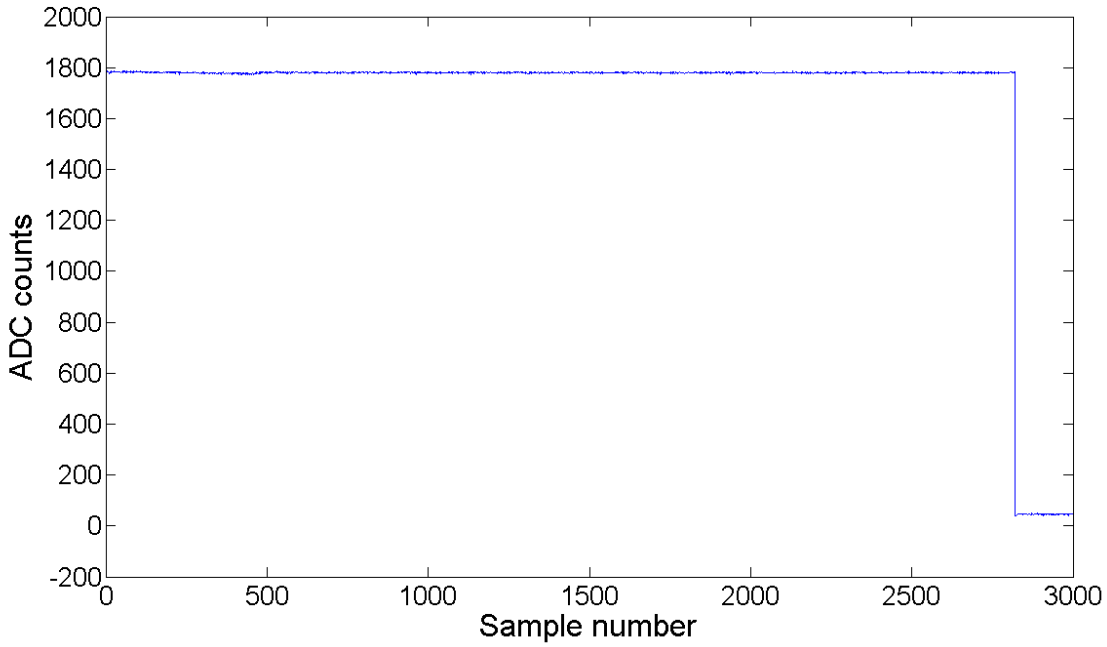

Figure 5.9: Σ_I output of the FONT4 test-bench digitized using the long train FONT5 firmware.

with a value of $\sim$1800 ADC counts and that the remainder of samples had a value close to zero, consistent with the expectation of baseline noise, thus demonstrating both FONT4 test-bench and long train firmware working correctly.

5.3.2 Simulated kicker calibration

The beamline model chosen in order to demonstrate feedback is defined by the parameters in Table 5.4. The DAC2 output of the FONT4 test-bench provides the Σ_I signal for the FONT5 board. Its value is simply the output of RAM2, which was loaded with 8191 in all locations. The DAC1 output represents the Δ signal; its value has a contribution from both RAM1 and ADC1. This corresponds to a train with zero angular offset, and a position offset determined by the value stored in RAM1, receiving a kick according to the value registered on ADC1. The value of ${}^{\mathrm{DAC1}}k_{\mathrm{ADC1}}$ determines the gain parameter of the feedback system; in this case, the largest seven-bit two's complement value has been selected in order to maximize the effect of the simulated kick. To perform the kicker calibration, the contents of RAM1 were first set to zero to represent a beam travelling on-axis at K1. After appropriate configuration

Parameter	Value
${}^{\mathrm{DAC1}}k_{\mathrm{RAM1}}$	32
${}^{\mathrm{DAC1}}k_{\mathrm{RAM2}}$	0
${}^{\mathrm{DAC1}}k_{\mathrm{ADC1}}$	63
${}^{\mathrm{DAC1}}k_{\mathrm{ADC2}}$	0
${}^{\mathrm{DAC2}}k_{\mathrm{RAM1}}$	0
${}^{\mathrm{DAC2}}k_{\mathrm{RAM2}}$	32
${}^{\mathrm{DAC2}}k_{\mathrm{ADC1}}$	0
${}^{\mathrm{DAC2}}k_{\mathrm{ADC2}}$	0

Table 5.4: Beamline model parameter values.

of the various strobes of the FONT4 test-bench, the constant DAC output of the FONT5 board was then scanned from -4000 up to 4000 in increments of 2000 (units of FONT5 DAC counts) and data was saved at each setting. The results are given in Figure 5.10. Only the first ten bunches are plotted in order to emphasize the change in position of the second bunch relative to the first bunch. The plot shows that the position of the first bunch does not change as a function of the DAC setting i.e. the constant kick is not applied to the first bunch. Each subsequent bunch in the train shares a common position that differs from that of the first bunch in a manner proportional to the FONT5 DAC setting.

From these data it is possible to obtain the gain parameter of the feedback system through the usual method of fitting bunch position as a function of FONT5 DAC setting. By taking the mean of the reciprocal of the fitted gradient over all 2,820 bunches, the P2-K1 gain parameter is found to be -5400; the mean is taken in order to remove the contribution of ADC and DAC noise.

5.3.3 Simulated feedback

As alluded to earlier, the FONT4 board is able to simulate jitter through dynamic reloading of, in this case, RAM1. This process is performed by the FONT4 software interface for the test-bench that was written in LabVIEW. A vector of 14-bit values is generated according to the desired pattern and then transmitted to the FONT4 board in the interval between triggers; thus, the position of each bunch can be made to change on the trigger-to-trigger timescale. For the following data runs, each train was assumed to be completely flat; that is, there was no change in position on the bunch-to-bunch timescale beyond the natural variation caused by noise on the ADCs and the DACs. For each trigger, the position of the train was drawn from a Gaussian distribution with a mean of zero and a standard deviation that varied from a minimum of 100 up to a maximum of 2,000 FONT4 DAC counts.

Figure 5.11 shows the degree to which the FONT5 board is able to accurately measure the output of the FONT4 board. For a data run consisting of 50 triggers, the blue line

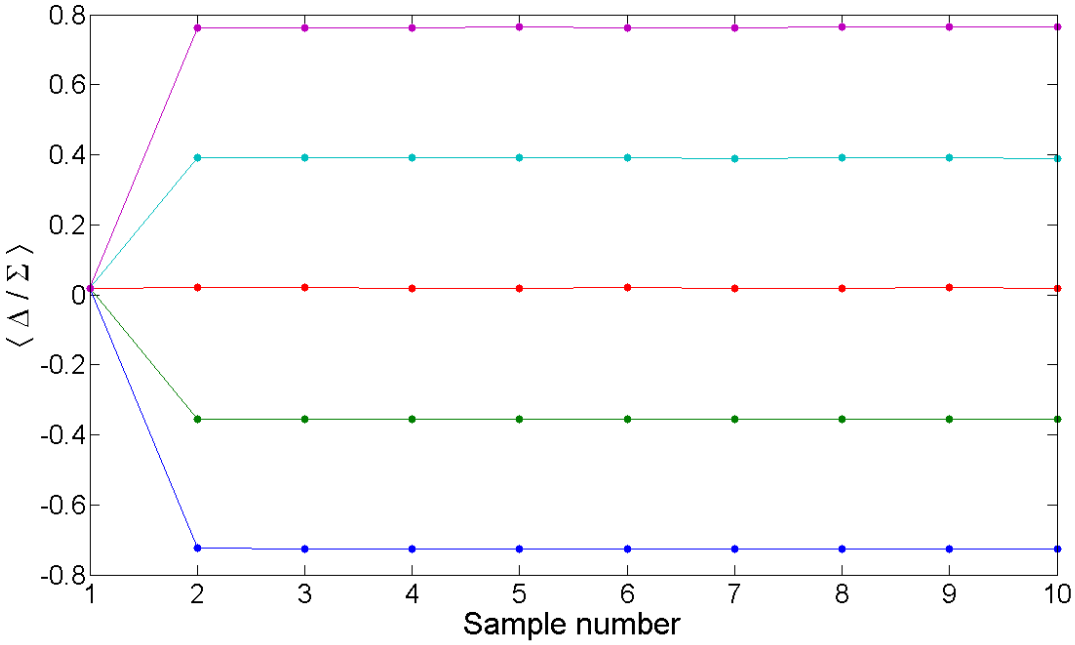

Figure 5.10: Mean recorded bunch position for the first ten bunches of the ILC-like train during a scan of the FONT5 constant DAC out value. Blue: -4000; Green: -2000; Red: 0; Cyan: 2000; Purple: 4000.

represents the Δ value of the first bunch measured by the FONT5 board. As the voltage scales of the FONT5 ADCs and the FONT4 DACs are different, the value in FONT5 ADC counts has been normalized to have a mean and a variance of one. The red line represents the value that the FONT4 board reported as having applied to the DAC; since this value has units of FONT4 DAC counts, it has been similarly normalized. The Δ value output by the FONT4 board consists of a RAM1 term, which is known perfectly in advance, and an ADC1 term, required when performing feedback, that in this case serves only to couple in random noise from that ADC channel. Thus, the deviation of the blue line from the red line is down to the combined effect of noise on the FONT4 ADC channel, noise on the FONT4 DAC channel and noise on the FONT5 ADC channel (for simulated feedback runs the noise on the FONT5 DAC channel also plays a role). Nevertheless, the correlation coefficient relating the measured Δ value to the set Δ value, expressed as an average over all bunches in the train, exceeds 0.99. The jitter of the Δ output of the FONT4 board is observed to be 3.0 FONT5 ADC counts; given that the maximum output of the FONT4 board is $\sim$ 1780 FONT5 ADC counts and this corresponds to the maximum value for Σ_I, the effective "noise floor" of the simulation is $\sigma\left(\Delta/\Sigma_I\right) = 1.7 \times 10^{-3}$. Figure 5.12 shows the feedback algorithm operating on a long train data set consisting of 100 triggers by plotting the position profile of each train (the bunch position as a function of bunch index). In this example, the standard deviation of the train position is 2,000 FONT4 DAC counts. For each train the FONT5 board applies the correction calculated from the position of bunch one to bunch two. As a result, the position of bunch two is approximately zero. The residual position of bunch two is then used to calculate the correction for bunch three. This correction

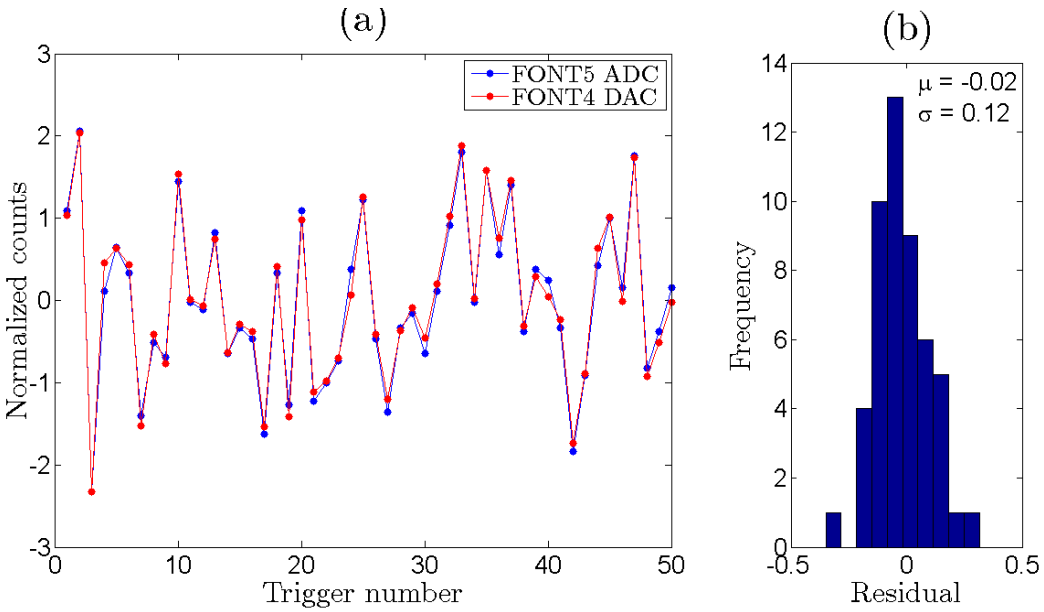

Figure 5.11: (a) Δ value for the first bunch plotted as a function of trigger number; the blue line gives the value measured by the FONT5 board in normalized FONT5 ADC counts and the red line the value as defined at source on the FONT4 board in normalized FONT4 DAC counts. (b) the distribution of the residual Δ (the mean μ and standard deviation σ of the distribution are indicated on the plot.)

is then added to the last value applied to the DAC in order that the position remains close to zero for the duration of the train. Figure 5.13 illustrates the performance of the feedback algorithm for different values of the train jitter. As usual, only the first few bunches are plotted to ensure that the behaviour of bunch one is clearly visible. The plot shows that in each case the standard deviation of the bunch position is reduced from its initial value down to the minimum possible value as determined by the resolution of the system. The standard deviation can be calculated separately for the corrected position of each of the 2819 corrected bunches and the mean of these values provides an estimate of the resolution of the FONT5 feedback board when operated with the FONT4 test-bench: $\sigma(\Delta/\Sigma_I) = 2 \times 10^{-3}$.

5.3.4 Simulated gain scan

By way of verification of the FONT4 test-bench, a pair of gain scans were performed. In each case the gain was scanned from a minimum of 0.2 up to a maximum of 1.8 times the nominal value of -5400. In the first case, each address of RAM1 was loaded with the same non-zero value such that the beam appeared to have zero jitter but a large constant offset. The result is given in Figure 5.14. The plot shows quite clearly the effect a non-nominal gain has on the early region of the bunch train. Gains less than nominal result in the corrected position slowly converging towards zero while gains greater than nominal result in a bunch position that performs damped oscillations about zero. The second gain scan used a beam centred at zero but with a jitter of 100 FONT4 DAC counts in order to observe the effect of changing the gain on the mean and standard deviation of the beam position. Figure 5.15

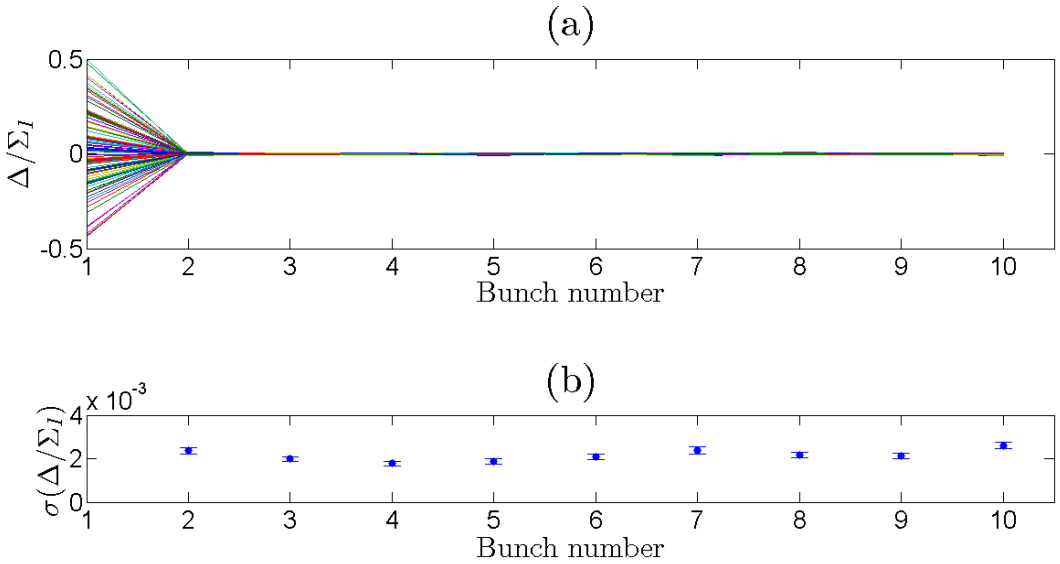

Figure 5.12: For a nominal gain feedback run: (a) position profile for the first ten bunches from 100 trains; (b) jitter profile (bunch jitter as a function of bunch index) for the first ten bunches.

shows that bunch two behaves as expected; using the correct value for the gain parameter results in the mean position being as close to zero as possible and the standard deviation of position being at a minimum. However, when the final bunch of the train is examined it is evident that the mean position of the bunch no longer depends on the gain parameter (the feedback correction has converged on the correct value in all cases) and the minimum bunch jitter is actually to be achieved using the minimum gain parameter. This is a consequence of the inherent noise in the system; any variation in position on the bunch-to-bunch timescale is purely coming from the ADCs and the DACs of the combined system. The bunch-to-bunch correlation of this random component is effectively zero and thus feedback is not able to improve the situation.

5.4 Summary

Using the specially developed FONT4 test-bench to simulate the Δ and Σ_I outputs of a FONT BPM processor, a design for the FONT5 firmware that is capable of performing feedback on an ILC-like bunch train consisting of 2,820 bunches separated in time by 308 ns has been demonstrated for a model beamline consisting of a single BPM and a single kicker.

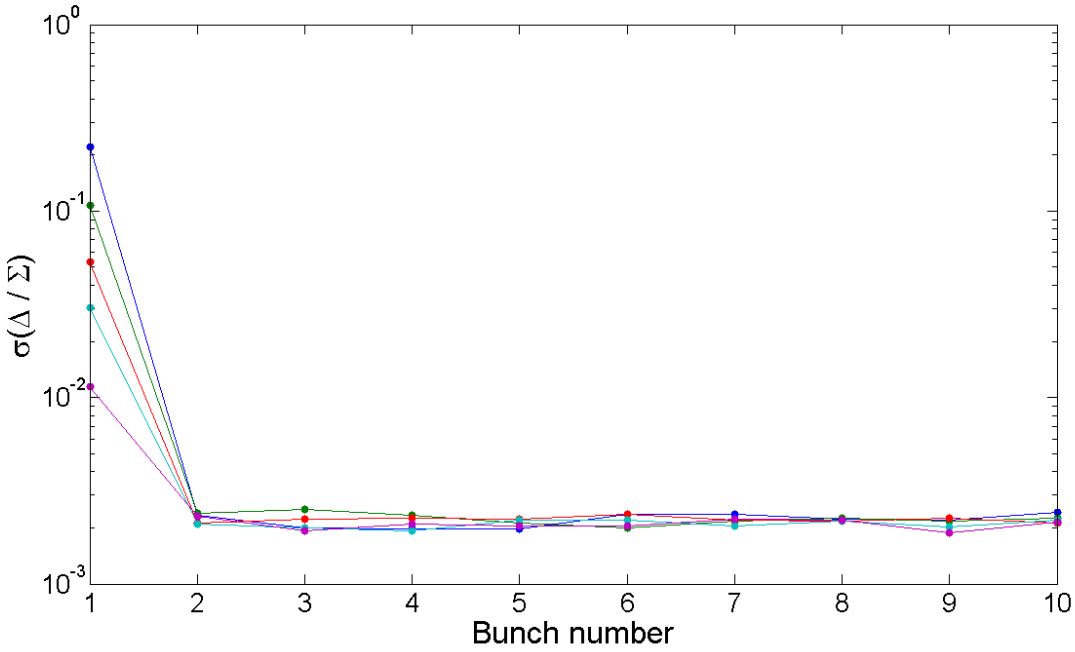

Figure 5.13: Standard deviation of bunch position as a function of bunch number for various values of the train jitter (in FONT4 DAC counts). Blue = 2000; Green = 1000; Red = 500; Cyan = 250; Purple = 100.

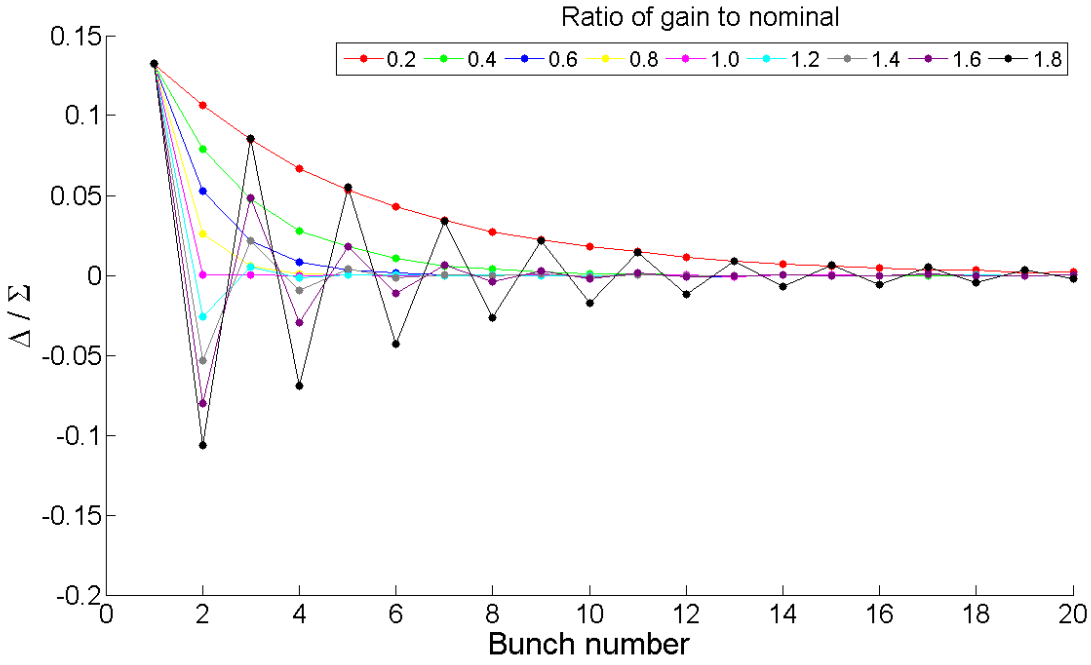

Figure 5.14: Mean bunch position as a function of bunch number for various values of the feedback gain parameter, as indicated in the legend.

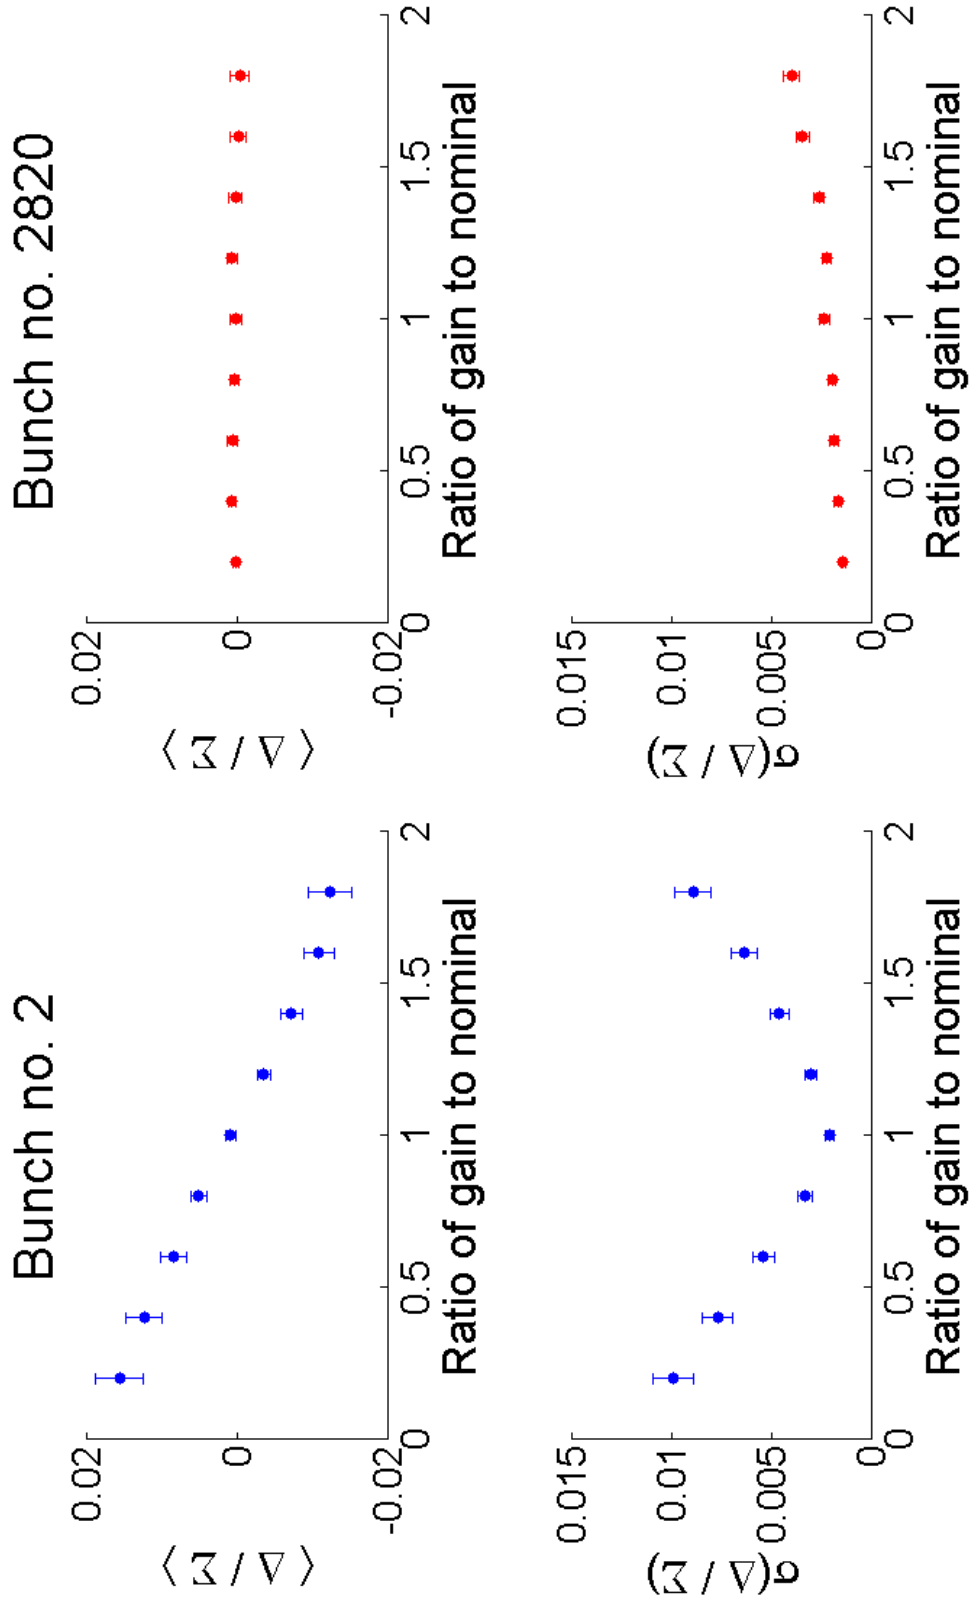

Figure 5.15: Mean (top) and standard deviation (bottom) of bunch 2 (left, blue) and bunch 2820 (right, red) as a function of the feedback gain parameter.

Chapter 6

Conclusions

6.1 Summary

With the discovery at the LHC of the final particle predicted by the Standard Model, the Higgs boson, the case for a linear lepton collider capable of performing the precision measurements necessary for the determination of its properties is demonstrated. The design for such a machine exists in the form of the ILC, which represents the culmination of several decades of research and development covering multiple aspects of the field of accelerator physics.

One such aspect is the delivery of a high luminosity at the IP with beams measuring just 640×5.7 nm. In order to maximize the overlap of the beams in the presence of ground motion on the intra-train timescale, a feedback system is necessary. The FONT feedback system aims to use the deflection due to the Coulomb interaction of one of the beams downstream of the IP in order to provide a corrective kick to the other beam upstream of the IP.

To this end low latency BPM processing electronics and a high voltage, short rise time amplifier were developed along with a digital feedback processing unit for the prototype FONT feedback system (Section 2). Supplemented with a second digital board enabling the monitoring of additional striplines downstream of the feedback kickers, the FONT5 system was installed at the ATF and operated during a number of user shifts.

The calibration constants for the BPM system were obtained and the observed variation shown to be due to differences in the analogue processor modules (Section 3.2). The resolution of the three-BPM system P1, P2 and P3 was estimated and found to be 2.6 ± 0.9 μm if the resolutions at all BPMs were assumed to be identical and the transfer matrices were used to predict the position at any one BPM from the other two (Section 3.3.1). An alternative method that allowed the prediction to be weighted in favour of the more reliable BPMs estimated the resolution of the system to be 0.9 ± 0.3 μm and provided evidence that the position measured at P1 had a large component that was uncorrelated to the positions at P2 and P3 (Section 3.3.2).

By deliberately misphasing the LO signal input to each processor module, the depen-

dence of the reported position on a phase offset between the LO and the bunch signals was determined for each BPM and found to be much more severe for P1 (Section 3.4.3). By developing the analysis to take into account the Σ_Q signal that quantifies the LO phase offset, the resolution prediction was vastly improved and the lower limit for the resolution of the three-BPM system found to be 0.45 ± 0.14 μm (Section 3.4.4.3). The system was then operated in feedback mode using the positions measured at P2 and P3 to correct both the position and angle of the beam using the kickers K1 and K2. The latency was found to be 156 ns (Section 4.2), well within the latency budget determined by the bunch spacing of the ILC, and the stability that the feedback system could provide at the feedback BPMs found to be at the level predicted by the bunch to bunch correlations of the ATF beam. The best performance achieved was a reduction in the jitter by a factor of 5.5 at P2 (Section 4.3).

A new design for the FONT5 board firmware was created to work with ILC-like bunch trains consisting of several thousand bunches (Section 5.2) and, using a hardware simulation running on the FONT4 digital board of the previous generation of the FONT feedback system as a substitute for an accelerator beamline consisting of a single BPM and a single kicker (Section 5.1), verified to be capable of operation under such conditions (Section 5.3).

6.2 Outlook

The elimination of the phase sensitivity of the FONT BPM system has been shown to be expected to deliver an improvement to the performance of the FONT5 feedback system. One method of accomplishing this is to install remote controlled phase shifters on the inputs to the BPM processor modules; the phase sensitivity can then be minimized with the equipment in situ. Besides this specific issue, future work with the FONT system at ATF is likely to use radically different configurations to the one described in this thesis. The potential presence of jitter sources downstream of the FONT region would limit the performance that the upstream FONT system could achieve at the IP. Thus, a kicker has been installed in the final focus region of the ATF immediately upstream of the nominal IP; the positions measured at P2 and P3 could be used to drive a correction at this IP kicker instead of at K1 and K2 [1]. Going further, it is envisaged that the FONT5 board will be used to digitize the output signals from the high resolution cavity BPMs residing in the IP chamber itself [61] as part of a local feedback system at the IP. This work is expected to feature in the forthcoming theses of other members of the FONT group [62, 63].

[1] As the IP kicker is a mere 595 mm upstream of the nominal IP it is only able to change the position by a few microns at that location; however, this is sufficient given that the jitter is generally less than 100 nm at the nominal IP.

Bibliography

[1] M. Herrero. "The Standard Model" (1998). ArXiv:hep-ph/9812242.

[2] M. Franklin C. Campagnari. *Rev. Mod. Phys.*, **69** (1997) 137. ArXiv:hep-ex/9608003.

[3] K. Kodama et al. *Phys. Lett. B*, **504** (2001) 218–224. ArXiv:hep-ex/0012035.

[4] A. Cho. *Science*, **337** (2012) 141–143.

[5] P. Higgs. *Phys. Rev. Lett.*, **13**, 16 (1964) 508–509.

[6] A. Hoffman. "Measurement and Control of Charged Particle Beams". Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology. Cambridge University Press (2004).

[7] E. Wilson. "An Introduction to Particle Accelerators". Oxford University Press (2001).

[8] O. Brüning. "Turnaround Time in Modern Hadron Colliders & Store-length Optimization". In "Proceedings of BEAM07", (2007). CARE-HHH-APD Workshop on Finalizing the Roadmap for the Upgrade of the CERN & GSI Accelerator Complex.

[9] "Physics at the International Linear Collider". `http://lcsim.org/papers/DBDPhysics.pdf` (2013). Physics Chapter of the ILC Detailed Baseline Design Report.

[10] R. Lipton. "Muon Collider: Plans, Progress and Challenges" (2012). ArXiv:1204.3538.

[11] "International Linear Collider Reference Design Report: Volume 3 - Accelerator". `http://ilcdoc.linearcollider.org/record/6321/files/ILC_RDR_Volume_3-Accelerator.pdf?version=4` (2007).

[12] W. Kilian et al. *Phys. Lett. B*, **373** (1996) 135–140. ArXiv:hep-ph/9512355.

[13] M. Palmer. "Lecture A3 Part 2 - Low Emittance Ring Design" (2010). Lectures from 5th International School for Linear Colliders, Villars-sur-Ollon, Switzerland.

[14] "International Linear Collider Technical Design Report: Volume 3.II: Accelerator - Baseline Design". `http://www.linearcollider.org/ILC/Publications/Technical-Design-Report` (2013).

[15] "The International Linear Collider - Gateway to the Quantum Universe". `http://www.linearcollider.org/pdf/ilc_gateway_report.pdf` (2007).

[16] F. Tecker. "CLIC (and room temperature RF)" (2010). Lectures from 5th International School for Linear Colliders, Villars-sur-Ollon, Switzerland.

[17] R. Amirikas et al. "Ground motion & comparison of various sites". EUROTeV-Report-2005-023-1 (2005).

[18] J. Resta-López. "Luminosity Performance Studies of Linear Colliders with Intra-train Feedback Systems" (2009). ArXiv:0902.2915.

[19] `http://legacy.kek.jp/intra-e/index.html`. KEK homepage.

[20] F. Hinode, S. Kawabata, H. Matsumoto, K. Oide, K. Takata, S. Takeda & J. Urawaka. "Accelerator Test Facility: Design and Study Report" (1995). KEK Report 95-4.

[21] R. Shafer. "Beam position monitoring". Technical report, Los Alamos National Laboratory (1990). AIP Conference Proceedings.

[22] V. Sargsyan. "Comparison of Stripline and Cavity Beam Position Monitors". Technical report, DESY (2004). TESLA Report 2004-03.

[23] B. Grishanov et al. "ATF2 Proposal Vol. 1" (2005). KEK Report 2005-2.

[24] R. Apsimon. "The Development and Implementation of a Beam Position Monitoring System for use in the FONT Feedback System at ATF2". D.Phil. thesis, University of Oxford (2011).

[25] R. Apsimon et al. "Development of Fast Micron-Resolution Beam Position Monitors for Linear Collider Beam-Based Feedback Systems". In "Proceedings of 2008 IEEE Nuclear Science Symposium and Medical Imaging Conference", (2008).

[26] Mini-Circuits. "Datasheet for LFCN-1000 Low Pass Filter". `217.34.103.131/pdfs/LFCN-1000.pdf`.

[27] Mini-Circuits. "Datasheet for SCPJ-2-9 Power Splitter/Combiner". `217.34.103.131/pdfs/SCPJ-2-9.pdf`.

[28] Mini-Circuits. "Datasheet for SYM-2 Frequency Mixer". `217.34.103.131/pdfs/SYM-2.pdf`.

[29] Mini-Circuits. "Datasheet for JPS-2-4 Power Splitter/Combiner". `217.34.103.131/pdfs/JPS-2-4.pdf`.

[30] Mini-Circuits. "Datasheet for HPQ-08 Power Splitter/Combiner". `217.34.103.131/pdfs/HPQ-08.pdf`.

[31] Mini-Circuits. "Datasheet for ZX60-4016E Amplifier". `217.34.103.131/pdfs/ZX60-4016E+.pdf`.

[32] Mini-Circuits. "Datasheet for VAT-3+ SMA Fixed Attenuator". 217.34.103.131/pdfs/VAT-3+.pdf.

[33] Mini-Circuits. "Datasheet for VLF-120 Low Pass Filter". 217.34.103.131/pdfs/VLF-120.pdf.

[34] Avtech. "Datasheet for AVM Series Pulse Generators". www.avtechpulse.com/catalog/page018_cat11_avm_rev14.pdf.

[35] C. Clarke. "The Interaction Point Collision Feedback System at the International Linear Collider and its Sensitivity to Expected Electromagnetic Backgrounds". D.Phil. thesis, University of Oxford (2008).

[36] R. Apsimon et al. "Latest beam test results from ATF2 with the FONT ILC prototype intra-train beam feedback systems". In "Proceedings of the International Particle Accelerator Conference, 2010", (2010).

[37] Xilinx, Inc. "Virtex-5 FPGA User Guide" (2009). www.xilinx.com/support/documentation/user_guides/ug190.pdf.

[38] Xilinx, Inc. "Virtex-5 FPGA Data Sheet: DC and Switching Characteristics" (2009). www.xilinx.com/support/documentation/data_sheets/ds202.pdf.

[39] Xilinx, Inc. "Platform Flash PROM User Guide" (2009). www.xilinx.com/support/documentation/user_guides/ug161.pdf.

[40] Texas Instruments. "Datasheet for ADS5474 14-bit, 400 MSPS Analog-to-Digital Converter" (2007). www.ti.com/lit/ds/symlink/ads5474.pdf.

[41] Analog Devices. "Datasheet for AD9744 14-bit, 210 MSPS TxDAC D/A Converter" (2005). www.analog.com/static/imported-files/data_sheets/AD9744.pdf.

[42] Linear Technology. "Datasheet for LTC2624 Quad 16-bit Rail-to-Rail DACs in 16-Lead SSOP" (2004). cds.linear.com/docs/en/datasheet/2604fd.pdf.

[43] standards.ieee.org/findstds/standard/1149.1-2013.html.

[44] S. Palnitkar. "Verilog HDL". Prentice Hall (1996).

[45] Xilinx, Inc. "Xilinx ISE 10.1 Design Suite Software Manuals and Help" (2008). www.xilinx.com/itp/xilinx10/books/manuals.pdf.

[46] B. Constance. "Design and Beam Testing of a Fast, Digital Intra-train Feedback System and its Potential for Application at the International Linear Collider". D.Phil. thesis, University of Oxford (2011).

[47] Mini-Circuits. "Datasheet for ZPUL Pulse Amplifiers". 217.34.103.131/pdfs/ZPUL-21.pdf.

[48] www.tmd.co.uk.

[49] Mini-Circuits. "Datasheet for MTS-18B Transfer Switch". 217.34.103.131/pdfs/MTS-18XL-B.pdf.

[50] Mini-Circuits. "Datasheet for JSPHS-1000 Phase Shifter". 217.34.103.131/pdfs/JSPHS-1000.pdf.

[51] Mini-Circuits. "Datasheet for ZN8PD1-53+ Power Splitter/Combiner". 217.34.103.131/pdfs/ZN8PD1-53+.pdf.

[52] Data Translation. "Datasheet for DT9854 Analog Output Module". www.datatranslation.com/products/dataacquisition/usb/DT9853-DT9854/.

[53] Mini-Circuits. "Typical Performance Data for JSPHS-1000 Phase Shifter". 217.34.103.131/pages/s-params/JSPHS-1000_VIEW.pdf.

[54] www.ni.com/labview/.

[55] A. Faus-Golfe. "Accelerator R&D for Future Large Colliders at IFIC" (2009). Private communication.

[56] G. Strang. "Introduction to Applied Mathematics". Wellesley-Cambridge (1986). P. 398.

[57] http://mad8.web.cern.ch/mad8/.

[58] http://www.mathworks.co.uk/help/matlab/math/systems-of-linear-equations.html.

[59] Glen White. Private communication.

[60] P. Burrows et al. "The FONT4 ILC Intra-train Beam-based Digital Feedback System Prototype". In "Proceedings of PAC07", (2007). PAC07-MOPAN108.

[61] Y. Kim et al. *Phys. Rev. ST Accel. Beams*, **15**, 042801.

[62] M. Davis. Unpublished D.Phil. thesis.

[63] N. Blaskovic Kraljevic. Unpublished D.Phil. thesis.