



Design and beam testing of a fast, digital intra-train
feedback system and its potential for application at the
International Linear Collider

Ben Constance
Brasenose College

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

Trinity Term 2011

Abstract

Documented in this thesis is the development of the FONT5 beam-based, intra-train feedback system. It is designed to provide the beam stability at the interaction point of the International Linear Collider necessary to maintain a high luminosity. It has been possible to take advantage of the relatively large bunch spacing to implement a flexible system based around fast digital electronics, enabling additional features such as bunch train flattening and kicker droop compensation.

A prototype feedback system has been installed at the Accelerator Test Facility in KEK, Japan, comprising three beam position monitors, two electromagnetic kickers and associated processing electronics. The installation is arranged to stabilise both the position and angle jitter of the beam in the vertical plane. Firmware for a Field Programmable Gate Array which runs the feedback algorithms is described, as are the results of a set of experiments to characterise the system behaviour. The feedback latency is measured to be 135 ns for position feedback alone, and 147 ns when operating in both degrees of freedom.

Using a three-bunch train with spacing of 154 ns, a vertical position stability of 430 nm in the feedback BPM is achieved. This is seen to fulfil the requirements of the International Linear Collider intra-train interaction point feedback system, even if the machine were to operate with the proposed ‘Low-N’ parameter set where the bunch spacing is 50% of nominal.

A statistical framework is developed to help understand the operation of the feedback upon trains with imperfect bunch-to-bunch position correlation, and is applied to the results of the FONT5 experiments. In this way, the prospect of using the prototype FONT5 installation to help meet the goals of the ATF2 final-focus test beamline is considered.

Acknowledgements

Much of the physics done today requires complex, often vast, apparatus beyond the means of any one individual, sometimes even of one nation, to design and build. Similarly, producing a thesis in a field such as accelerator physics is an enterprise made possible by the assistance and support of a great number of colleagues and collaboration partners, plus many more people besides, and I shall take this opportunity to thank them.

It has been my fortune to work in a brilliant group on the FONT experiments, with a great professor, Phil Burrows. Thanks foremost to him for making Oxford a welcoming place, for his continual guidance and giving me the chance to do such interesting work. Thanks to Glenn Christian, whose work laid the foundation for this thesis, and to Javier Resta-Lopez for many helpful chats. This work would not have been possible without my fellow students, Robert Apsimon and Douglas Bett, both on and off shift. Thanks too to Colin Perry, without whom these might have been thought experiments, and to Christina Swinson, Christine Clarke, Tony Hartin, Hamid Dabiri Khah and everyone else whose effort has gone into the FONT project.

I am very grateful to Junji Urakawa and the entire ATF collaboration for their support and hospitality, and to those many colleagues out at KEK who were always quick to offer help and advice. My time in Japan was all the better for the company of those who joined us for evenings in the dormitory smoking room. Above all, Tsukuba would not have been the same were it not for Jolanta Brodzicka, whom I thank for everything.

Throughout the many years that I have directed toward this goal, my parents, Simon and Dorothy Constance, could not have been more supportive, and I have them to thank more than anyone else. Also, I owe Nichola Taylor for all the encouragement and, along with David Woods, for the many great times in both London and Oxford.

There is a beginning to most paths, and in this respect I would like to thank old mentors Mrs Chidlow, for helping set me upon mine, and Dr Joyner, who by giving me a laboratory ensured I would never want to leave it.

Contents

1	Introduction	1
1.1	The Standard Model of particle physics	1
1.1.1	The Higgs boson	2
1.2	Extension of the Standard Model	4
1.2.1	Supersymmetry	5
1.3	Experiments at the Terascale	6
1.3.1	$\mu^+\mu^-$ collider rings	8
1.3.2	Linear e^+e^- colliders	8
2	The accelerator physics of the International Linear Collider	9
2.1	Transverse beam dynamics	9
2.1.1	The nominal orbit	9
2.1.2	Focussing with quadrupole magnets	10
2.1.3	Linear betatron motion	11
2.1.4	Transfer matrix formalism	13
2.1.5	Beam size, emittance and jitter	15
2.1.6	Off-momentum orbits	17
2.1.7	Chromatic aberration	19
2.2	Overview of the International Linear Collider	20
2.2.1	Electron source	22
2.2.2	Positron source	22
2.2.3	Damping rings	23
2.2.4	Main linear accelerators	26
2.2.5	Beam delivery systems	29
2.2.6	Detectors	30
2.2.7	Extraction lines and dumps	31
2.3	The ILC interaction point	31
2.3.1	Luminosity	31
2.3.2	IP crossing angle and crab cavities	32

2.3.3	The hourglass effect	33
2.3.4	Beam disruption	33
2.4	Degradation of the luminosity	35
2.4.1	Effects of ground motion	37
2.4.2	Damping ring extraction kicker	38
2.4.3	Jitter in the main linacs	38
2.4.4	Interaction point feedback	39
3	The FONT5 prototype digital intra-train feedback system	40
3.1	The Accelerator Test Facility	40
3.1.1	Electron source	43
3.1.2	Linear accelerator	43
3.1.3	Damping ring	44
3.1.4	The ATF2 extraction line and final focus	46
3.1.5	The physics goals of the ATF2 project	48
3.2	The objectives of the FONT5 feedback system installation at ATF2	48
3.2.1	ILC interaction point intra-train feedback technology demonstration .	49
3.2.2	Contribution to the physics goals of ATF2	49
3.3	The FONT5 feedback system installation	49
3.3.1	Stripline beam position monitors	51
3.3.2	Fast analogue BPM processors	53
3.3.3	BPM mover system	55
3.3.4	Stripline kickers	56
3.3.5	Kicker amplifiers	56
3.3.6	The FONT5 digital feedback board	58
3.4	FONT5 BPM performance	60
3.4.1	Charge normalisation	60
3.4.2	BPM calibration using corrector magnets	62
3.4.3	BPM calibration with movers	64
3.4.4	Estimation of the BPM resolution	66
3.5	Feedback algorithms	68
3.5.1	Vertical position feedback	68
3.5.2	Vertical position and angle feedback	69
3.5.3	Bunch train flattening	69
3.5.4	Kicker pulse droop compensation	70
3.5.5	Constant DAC output mode and calibration of the kickers	72

4	FPGA firmware implementation of a flexible feedback system	75
4.1	Control register interface	75
4.2	ATF timing system	78
4.2.1	357 MHz clock input electronics	78
4.2.2	Digital input thresholds	79
4.2.3	Synchronisation with the ATF2 beam	80
4.3	BPM signal digitisation	81
4.3.1	ADC synchronisation	82
4.3.2	ADC data phase feedback for sampling stability	83
4.3.3	Fine adjustment of the ADC sample time	84
4.3.4	ADC inputs and pedestal subtraction with trim DACs	85
4.4	Pulse-by-pulse acquisition of BPM data	87
4.4.1	Dual-port block RAM	87
4.4.2	Data stream format	89
4.5	Feedback signal generation	92
4.5.1	Virtex-5 DSP48E slices	92
4.5.2	Look-up tables for charge normalisation and gain application	92
4.5.3	Digital signal processing for a single feedback loop	93
4.5.4	Digital signal processing for coupled feedback loops	96
4.6	RS232 interface for control of the FONT5 system	99
4.6.1	RS232 control protocol	99
4.6.2	FONT5 graphical user interface	101
5	Performance of the FONT5 feedback system at ATF	103
5.1	Feedback gain	103
5.1.1	Nominal system gain for a single feedback loop	103
5.1.2	Nominal system gain in a coupled feedback system	104
5.1.3	Betatron phase advance and feedback system location	105
5.1.4	Optimal gain for an imperfect beam	106
5.2	Latency measurements	108
5.3	Effects of LO phase jitter on the BPM resolution	111
5.3.1	Local oscillator phase jitter	111
5.3.2	Obtaining limits on the BPM resolutions	113
5.4	Measurement of the jitter covariance matrix	114
5.4.1	Ideal case with perfect resolution	114
5.4.2	Correcting for degradation by known resolutions	115
5.5	Modelling the FONT5 system	115
5.5.1	Modelling of P2-to-K1 single loop feedback	116

5.5.2	Modelling of P3-to-K2 single loop feedback	117
5.5.3	Modelling the coupled feedback system	117
5.6	Feedback performance in single-loop mode	119
5.6.1	Vertical position jitter reduction at P2	120
5.6.2	Vertical position jitter reduction at P3	123
5.7	Performance of two coupled feedback loops	124
5.7.1	Effect on vertical position and angle jitter	125
5.8	Effects of gain variation on feedback performance	129
5.8.1	Scan of the P2-to-K1 loop gain	130
5.8.2	Scan of the P3-to-K2 loop gain	131
5.8.3	Scans of the coupled loop gains	132
5.8.4	Expected future performance and extrapolation to the final focus . .	134
6	Conclusions	136
6.1	Summary	136
6.2	Demonstration of an ILC interaction point feedback system	137
6.3	Jitter reduction in two degrees of freedom	138
6.4	Outlook	138
 Appendices:		
A	Statistical modelling of intra-train feedback systems	140
A.1	Statistical transport of beam jitter	140
A.2	Treatment of multi-bunch trains	141
A.3	Feedback transformations - the two-bunch case	142
A.4	Generalisation of feedback transformations to n bunches	143
A.5	Modelling of a single intra-train feedback loop	144
A.6	Modelling of coupled intra-train feedback loops	145
A.7	Treatment of BPM resolution in a single feedback loop	145
A.8	Treatment of BPM resolution in coupled feedback loops	148
B	Deconvolution of resolution effects from BPM measurements	150

Chapter 1

Introduction

1.1 The Standard Model of particle physics

Modern scientific endeavour has delivered to us a wondrous array of advancements over the centuries. It is impossible to overstate the revolutionary way in which it has changed our perceptions and understanding of, as well as our interaction with, the environment we found ourselves in and which constitutes us. From the astrophysical surveys which illuminate the universe over the grandest of scales, to the developments in medicine which have enriched our lives and the communication technologies having so great an impact on our societies, science and reason have changed our world like no other force could.

Standing amongst these great achievements is the Standard Model of particle physics. Based upon two of the most important discoveries of the 20th century, special relativity and quantum mechanics, the Standard Model is surely one of physics' most resounding successes. Using its mathematical framework, the structure of all the visible matter in the universe may be described in terms of twelve fundamental particles, fermions, which include the six leptons and the six quarks (see table 1.1). For each fermion exists an anti-matter equivalent with opposite charge. In the Standard Model, all interactions between the fermions take place via three fundamental forces: the strong, the weak and the electromagnetic forces. As a quantum theory, energy exchanged in these interactions is in discrete packets, and it is carried by force mediating particles, bosons. The strong force is mediated by the gluon, the weak by the W^\pm and Z^0 bosons and the electromagnetic by the photon (see table 1.2).

Stable, visible matter is built from only the first generation of fermions, namely the up and down quarks and the electron. The heavier, higher generation fermions produced in experiments are short-lived and soon decay. Bound into pairs (mesons) or triplets (baryons) by the strong force, the quarks are found only as composite particles known collectively as hadrons. The gluons responsible for binding the hadrons together are self-interacting, such that the two or three valence quarks are joined by a sea of quarks, anti-quarks and gluons

which become apparent in high energy collisions [1].

Generation	Leptons		Quarks	
1	Electron e^-	Electron neutrino ν_e	Up u	Down d
2	Muon μ^-	Muon neutrino ν_μ	Charmed c	Strange s
3	Tau τ^-	Tau neutrino ν_τ	Top t	Bottom b

Table 1.1: The fundamental fermions of the Standard Model.

Boson	Function
Photon γ	Electromagnetic force
Weak bosons W^\pm, Z^0	Weak force
Gluon g	Strong force
Higgs boson h	Electroweak symmetry breaking

Table 1.2: The fundamental bosons of the Standard Model.

Stable baryons, the proton and the neutron, are able to bind together under the right circumstances, held by residual strong forces to form atomic nuclei of various size and charge. The electromagnetic force binds electrons to nuclei, resulting in atoms, which come together due to electromagnetic forces to form molecules. Then chemistry can take effect, and driven by evolution, structures of breathtaking complexity may arise.

Appealing as it may be, this frankly beautiful and comfortingly reductionist view of nature is not complete. Particle colliders have created all the fundamental fermions of the Standard Model in high energy collisions, which have subsequently been observed in particle physics experiments before going on to decay. The strength and behaviour of the three forces have been likewise studied, and in all cases, the predictions of the Standard Model have been borne out. However, there remain pieces of the puzzle which refuse to fit perfectly, phenomena yet to be observed, and important physical truths the Standard Model alone cannot explain. Over the following sections, some important pointers to physics beyond the Standard Model are briefly discussed.

1.1.1 The Higgs boson

There is a particle predicted by the Standard Model that is yet to appear in experiments. In the theory's framework, a single force, the electroweak, takes the place of the electromagnetic and weak forces. The mediators of this unified force are massless, as are the leptons on which

it acts. Clearly the world we observe does not behave in this way. The way in which theory and observation are reconciled in the Standard Model is known as the Higgs mechanism. A field is postulated to exist whose unexcited state is unstable, and whose infinitely degenerate stable ground states possess a non-zero potential. All particles of the Standard Model may interact with this field, known as the Higgs field. Though the overall theory is symmetric, this symmetry is hidden, or broken, when the behaviour of particles is viewed with respect to a ground state of the Higgs field. In particular, viewed from this perspective, the massless electroweak bosons give way to a massless photon plus the massive W^\pm and Z^0 bosons. The fermions gain their mass through interaction with the same field [2].

As well as giving particles their mass, the Higgs mechanism predicts that a further particle should exist, the Higgs boson, thus far unobserved. The masses gained by the fermions and the weak bosons are not directly predicted by the Standard Model, nor is the mass of the Higgs boson itself. Masses measured experimentally are found to be consistent with the Standard Model framework however, and may be used to predict the expected mass of the Higgs. Figure 1.1 shows the quality of fits of the Standard Model to precision electroweak measurements as a function of Higgs boson mass [3].




Figure 1.1: Quality of global fits of precision electroweak measurements to the Standard Model vs. Higgs boson mass. The yellow area indicates the mass range excluded by direct searches at CERN’s LEP [3].

Experiments at CERN’s Large Electron-Positron (LEP) collider and Fermilab’s Tevatron proton-antiproton collider have made searches for the Higgs boson at various energies. These experiments did not uncover evidence of a Higgs boson, and so may be used to exclude potential masses of the particle. These excluded ranges are shown in figure 1.2, along with

the higher masses excluded indirectly by measurement of electroweak effects which involve the Higgs boson [3, 4].

Taken together, these studies suggest that a Standard Model Higgs boson would most likely be found in the mass range of 114–158 GeV, with masses at the lower end of this range being more probable. Experiments at CERN’s Large Hadron Collider (LHC), which began colliding protons with a centre-of-mass energy of 7 TeV in March 2010 [5], are designed in part to confirm or refute the presence of a Standard Model Higgs in this mass range.

The Higgs mechanism in the Standard Model provides a means for breaking the electroweak symmetry, but it is not the only way in which nature could have conspired to do so. Alternative theories abound [6], and until the particle landscape at Terascale energies is more fully mapped no firm conclusions may be drawn. The LHC may well produce the Standard Model Higgs boson, or perhaps it will be found at far higher mass than predicted. It is possible that multiple Higgs bosons may be seen, or that none at all will appear. All these potential outcomes would be equally profound in their impact on our understanding of the universe.




Figure 1.2: Mass ranges for the Higgs boson excluded by direct searches at LEP and the Tevatron, as well as indirectly by precision electroweak measurements [4].

1.2 Extension of the Standard Model

Despite the Standard Model’s unparalleled success in the energy regimes so far explored, its deficiencies are well known. Some important shortcomings may be summarised as follows:

Gravity

The Standard Model completely ignores the existence of the fourth fundamental force of nature.

Neutrino oscillations

Experiments have looked at solar neutrinos produced in the sun’s nuclear reactions (e.g. SNO [7]), neutrinos produced as cosmic rays strike the Earth’s atmosphere (e.g.

Super-Kamiokande [8]) and also man-made beams of neutrinos sent from particle accelerators for hundreds of kilometres through the Earth's crust (e.g. MINOS [9]). All have reached the same conclusion: the various flavours of neutrino are linear combinations of quantum mechanical states of definite, non-zero mass, clearly in contradiction to the Standard Model which assumes each flavour of neutrino is massless.

Dark matter

Astronomical observations of the rotation curves of galaxies, the radial dependence of their angular velocities, have shown sharp deviations from the predictions of Newtonian gravity. A currently favoured explanation is that of cold dark matter [10]. In this theory, the visible matter making up the stars of a galaxy is accompanied by a vast halo of massive particles which interact only by the weak and gravitational forces. There is no particle in the Standard Model consistent with the proposed properties of these Weakly Interacting Massive Particles (WIMPs).

The Higgs hierarchy problem

As the energy to which the Standard Model is assumed to be valid is increased, the size of the radiative corrections to the mass of the Higgs diverge quadratically [1]. To keep the Standard Model valid at high energy and the Higgs mass consistent with the experimentally preferred range requires cancellation of the radiative corrections. This in turn means that the Standard Model parameters must be very finely tuned indeed, an undesirable quality of the Model.

Unification

As discussed previously, the Standard Model unifies the electromagnetic and the weak forces. The strong force is not part of this unification. However, the search for beauty and simplicity begets hopes of a Grand Unified Theory, in which all the fundamental forces may be considered as low-energy expressions of a single unified force.

1.2.1 Supersymmetry

Supersymmetric theories [11] are a class of theory that propose various modifications to the Standard Model. In such theories, a new symmetry between bosons and fermions is postulated. A new particle is introduced for every particle in the Standard Model, such that each boson has a fermionic superpartner and vice versa. When the symmetry of these theories is broken, the superpartners acquire different masses. Since no supersymmetric particles have been detected in experiments, these new particles must have masses in excess of their Standard Model counterparts.

Although no experimental evidence for these theories exists, they have theoretical properties which make them very attractive. Firstly, the radiative corrections to the Higgs

boson mass from each pair of superpartners cancel, providing a natural means of avoiding the quadratic divergence in the hierarchy problem [11]. Secondly, in many supersymmetric theories, the lightest neutralino would fulfil the requirements of the dark matter WIMPs predicted by astrophysical measurements [10].

Another tantalising prospect hinted at by supersymmetric theories is the unification of the strong force with the electroweak. The measured coupling strengths of the fundamental forces vary with energy, or equivalently distance, in both the Standard Model and supersymmetric theories. As shown in figure 1.3, supersymmetry (in this case a theory known as the Minimal Supersymmetric Model [11]) can cause the expected couplings to converge at a single value for energies of order 10^{16} GeV, whereas in the Standard Model they fail to do so [12].




Figure 1.3: Running of the inverse coupling strengths of the electromagnetic, weak and strong interactions in the Standard Model (left) and Minimal Supersymmetric Model (right) [12].

Although only a theoretical construct, supersymmetry predicts new particles at the Terascale. As such, evidence for or against the theories will soon emerge from data gathered at the LHC.

1.3 Experiments at the Terascale

Already, experiments at the LHC have begun gathering data at an energy of 7 TeV in the centre of mass, the first step in charting what new physics lies waiting at Terascale energies. The machine will eventually produce collisions at 14 TeV, and as statistics accumulate so too will evidence of the new phenomena which will guide attempts to overcome the Standard Model's shortcomings. It is impossible to predict what will be found, but numerous simulations of the theories presented in the previous sections, and many more besides, have shown the immense discovery power of the LHC [13, 14].

Once the Terascale landscape has been illuminated, the next step is to begin making precision measurements of the properties of the particles discovered; it will be essential to

determine accurately their masses, couplings and spins. This will enable different models to be distinguished from each other, and so determine the nature of the grander truth behind the Standard Model. This enterprise will require a machine that collides not hadrons, but leptons [6].

The protons at the LHC are composite particles, with collisions taking place between their constituent quarks and gluons. As only a portion of the proton energy goes into any given event, the precise initial conditions of the interaction cannot be controlled. This is in contrast to an e^+e^- collider, which may produce collisions at a specific and variable centre of mass energy. The particles can also be produced with a high degree of polarisation. In addition, protons interact via the strong force, leading to large amounts of background particles. These must be separated from interesting events by the detector, whereas a detector for an e^+e^- collider may be designed primarily for precision [15].

Historically, there has been great synergy between hadron and lepton colliders. Hadron machines have pushed back the high energy frontier, with complimentary, precision measurements being made at e^+e^- colliders. After the first observation of the weak bosons in the UA1 experiment at CERN's Super Proton Synchrotron, it was precision measurements from e^+e^- experiments at LEP and the Stanford Linear Collider (SLC) which accurately determined their properties and provided a powerful confirmation of the Standard Model's predictions [16]. Previously, the energy reach of fixed target proton collisions at Fermilab had enabled the discovery of the bottom quark [17], which are now produced in copious amounts for detailed study by the e^+e^- B-factories KEKB and PEP-II [18, 19].

A lepton collider in the 500 GeV – 1 TeV range would be capable of undertaking precision measurements of the new particles produced at the LHC, as well as top quarks, whilst enabling many other important studies [6]. There is world-wide consensus, embodied in reports from the International Committee for Future Accelerators [20], that a Terascale e^+e^- collider should be pursued vigorously by the international community.

Accelerating electrons and positrons to Terascale energies presents a number of technical challenges. In a collider ring, the power radiated as synchrotron radiation scales as [21]:

$$P \propto \frac{1}{R} \left(\frac{E}{m} \right)^4 \quad (1.3.1)$$

for beam energy E , particle mass m and ring radius R . The power lost in this way must be replaced using RF cavities, limiting the energy which can be reached using a ring of given radius. Indeed, for electrons, the 209 GeV achieved at LEP is considered to be near the practical limit for a circular machine [22]. Two possibilities remain: either use heavier particles in a ring, or else accelerate electrons and positrons in a straight line and thus minimise the emission of synchrotron radiation. The state of research in these fields is briefly summarised in the following sections.

1.3.1 $\mu^+\mu^-$ collider rings

While the possibility of a $\mu^+\mu^-$ collider ring has been considered [23], the particles' short lifetime of around 2 μs means that current technologies for cooling beams are unusable. Novel approaches would be necessary to cool the muon beams to a level at which they could be stored. Though research into ionisation cooling is ongoing, the technology is years away from reaching the demonstration stage [24]. Cooling is just one of the obstacles to overcome before a muon collider becomes feasible.

1.3.2 Linear e^+e^- colliders

Linear e^+e^- colliders represent the only near-term solution for producing Terascale lepton collisions. The concept was proven at the SLC, which provided collisions at the 91.2 GeV Z^0 resonance. Currently, two international collaborations are working on two designs for Terascale linear colliders: the International Linear Collider (ILC) and the Compact Linear Collider (CLIC).

The International Linear Collider

In 2004, the International Technology Review Panel recommended [25] that a linear collider based around superconducting RF cavities be designed to reach 500 GeV in the centre-of-mass, with the option of extending this energy reach to 1 TeV in the future. The international collaboration behind the ILC has responded by releasing a detailed Reference Design Report [26] for a machine fulfilling these criteria, with a Technical Design Report expected at the end of 2012 [27]. The ILC is the most advanced of the Terascale lepton collider options, with a great deal of international research ensuring the principles of its operation are well understood. This thesis is a contribution to that research, and the ILC will be discussed in more detail in chapter 2.

The Compact Linear Collider

While the ILC aims for 500 GeV – 1 TeV collisions, the CLIC design is for a linear collider with an even greater energy reach. Nature is, if one thing, unpredictable, and it will be results from the LHC which indicate the required energy range of the next linear collider. Through the use of room temperature accelerating cavities and a novel Two Beam Acceleration technique, it is anticipated that CLIC could reach energies of up to 3 TeV [28]. At this stage the design of CLIC is less mature than that of the ILC, but it too is being actively pursued with a strong international effort. There are many parallels between the design of ILC and CLIC, with the close collaboration between the two groups yielding many benefits [29].

Chapter 2

The accelerator physics of the International Linear Collider

2.1 Transverse beam dynamics

Throughout this thesis the formalism summarised in this section will be used to describe the beam dynamics as anticipated for the ILC, and to understand the operation of the intra-train feedback system developed by the Feedback on Nanosecond Timescales (FONT) group. The field of accelerator physics is broad and well studied, and a number of papers and texts will be referenced. The discussion is restricted to the transverse dynamics of beams in transfer lines (as opposed to storage rings), which has most relevance to the work presented here.

2.1.1 The nominal orbit

The nominal trajectory, or orbit, followed by a beam through an accelerator lattice is defined by the location and strength of dipole (or bending) magnets. The left-handed, curvilinear coordinate system adopted uses s to represent the longitudinal distance along this nominal orbit, with x and y being the perpendicular horizontal and vertical axes respectively. For ultra-relativistic electrons, $s = ct$. Due to the Lorentz force, such a particle travelling in a dipole field of uniform magnitude B will traverse a circular arc with the equation of motion:

$$\frac{dp_T}{ds} = eB \quad (2.1.1)$$

where p_T is the transverse component of the particle momentum and e the electronic charge. A particle on the nominal orbit is bent by an angle θ , where:

$$\frac{d\theta}{ds} = \frac{1}{\rho} = \frac{eB}{p_0} \quad (2.1.2)$$

for a bending radius ρ and a nominal beam momentum p_0 . The quantity $p_0/e = B\rho$ is referred to as the magnetic rigidity of the beam. It will be seen that dipole magnets also have a weak focussing effect on beams passing through them.

2.1.2 Focussing with quadrupole magnets

In the strong focussing technology employed in many modern accelerators, particle beams are contained using magnetic lenses. Quadrupole magnets, with hyperbolic poles arranged as shown in figure 2.1, give rise to a transverse magnetic field which increases linearly with displacement from the magnet centre:

$$B_x = \frac{p_0}{e}ky, \quad \text{and} \quad B_y = \frac{p_0}{e}kx \quad (2.1.3)$$

where the normalised quadrupole strength k has been defined as:

$$k = \frac{e}{p_0} \frac{\partial B_y}{\partial x} \quad (2.1.4)$$

The field provides a restoring force in one plane. In the perpendicular plane, the quadrupole acts as a defocussing lens. Therefore two families of quadrupoles are required to confine the beam in both planes, referred to as (horizontally) focussing and (horizontally) defocussing. For the focussing quadrupole, the equation of motion from the Lorentz force gives the differential deflection of a particle as:

$$\frac{d\theta_x}{ds} = -kx, \quad \text{and} \quad \frac{d\theta_y}{ds} = +ky \quad (2.1.5)$$

which, for a short quadrupole of length l , is the equation of a lens with focal length $1/kl$. Using short quadrupoles to represent accelerator components is commonly referred to as the thin lens approximation.




Figure 2.1: Arrangement of the poles in a quadrupole magnet [30].

2.1.3 Linear betatron motion

A transport line's lattice is the arrangement of components along the beamline, including the magnetic elements which define the beam orbit. Particles travelling along the nominal orbit undergo transverse oscillations, known as betatron oscillations, due to the influence of magnetic lenses. Using x' to denote the derivative of a particle's horizontal coordinate with respect to s , the equation of motion of a particle with moving through the lattice is [21]:

$$x'' - \frac{\rho + x}{\rho^2} = \frac{eB_y p_0}{p_0 p} \left(1 + \frac{x}{\rho}\right)^2 \quad (2.1.6)$$

and similar in the vertical plane. Assuming a linear machine, with pure dipole and quadrupole fields only, then the magnetic field may be written as:

$$B_y = -B + \frac{p_0}{e} kx, \quad \text{and} \quad B_x = \frac{p_0}{e} ky \quad (2.1.7)$$

where it is understood that both the dipole field B and the quadrupole gradients k vary with s depending on the layout of the magnets, and are all zero for much of the lattice. Substituting the fields into 2.1.6 and keeping only terms of first order in x and y gives, for a particle with nominal momentum $p = p_0$:

$$x'' + K_x(s)x = 0 \quad (2.1.8)$$

with a similar equation in the vertical, and where:

$$K_x(s) = k(s) - \frac{1}{\rho(s)^2}, \quad \text{and} \quad K_y(s) = -k(s) \quad (2.1.9)$$

The second order differential equation is an example of Hill's equation, describing a motion reminiscent of simple harmonic motion but with a variable restoring force. Since there is little to no bending in the vertical plane, the $1/\rho^2$ term corresponding to weak focussing from dipole magnets is neglected. The solution is well known [31], and may be written for one transverse plane as:

$$x = \sqrt{\varepsilon\beta(s)} \cos[\psi(s) + \psi_0] \quad (2.1.10)$$

with the constraint that:

$$\frac{d\psi}{ds} = \frac{1}{\beta(s)} \quad (2.1.11)$$

Substituting the solution back into Hill's equation and collecting sine and cosine terms separately reveals:

$$[\omega'' - \omega(\psi')^2 + \omega K_x] \cos(\psi + \psi_0) - [2\omega'\psi' + \omega^2\psi''] \sin(\psi + \psi_0) = 0 \quad (2.1.12)$$

Here, $\omega = \sqrt{\beta}$ has been introduced to simplify the formula. All of β , ω and ψ are functions of s . Since the sine and cosine functions are orthogonal, their respective coefficients must each separately vanish for the solution to be generally valid. From the sine terms:

$$2\omega\omega'\psi' + \omega^2\psi'' = \frac{d}{ds}(\omega^2\psi') = 0 \quad (2.1.13)$$

which is satisfied due to equation 2.1.11. For the cosine term to vanish, it is required that:

$$\beta^2 K + \frac{1}{2}\beta\beta'' - \frac{1}{4}\beta'^2 = 1 \quad (2.1.14)$$

which is known as the betatron envelope equation. The function $\beta(s)$ describes the amplitude modulation of the particle's betatron oscillations as it traverses the lattice, while $\psi(s)$ is the oscillation phase. The particular solution of $\beta(s)$ is determined by the initial condition at the start of the lattice. Often, a dedicated set of quadrupoles at the start of the lattice is used to match the initial beta function to that at the end of the injecting lattice or particle source.

For later convenience, a further two functions are defined:

$$\alpha(s) = -\frac{\beta'(s)}{2}, \quad \text{and} \quad \gamma(s) = \frac{1 + \alpha(s)^2}{\beta(s)} \quad (2.1.15)$$

with α , β and γ being collectively referred to as the Twiss parameters. By substitution

of the solution to Hill's equation, the following quantity, known as the Courant-Snyder invariant [31], is seen to be constant:

$$\gamma x^2 + 2\alpha x x' + \beta x'^2 = \varepsilon \quad (2.1.16)$$

This equation describes an ellipse of area $\pi\varepsilon$ in the $x - x'$ phase space of the particle. The particle's betatron phase determines its location on the ellipse. The area of this ellipse is conserved throughout the lattice, although its orientation and aspect will vary with s . The relationship between the phase space ellipse and the Twiss parameters is shown in figure 2.2.




Figure 2.2: Phase space ellipse from transverse betatron oscillation [21].

2.1.4 Transfer matrix formalism

A particularly useful method of tracing the state of a particle through the lattice makes use of transfer matrices. It is possible to rewrite Hill's equation (2.1.8) in matrix form as:

$$\frac{d\mathbf{x}}{ds} = \begin{pmatrix} 0 & 1 \\ -K(s) & 0 \end{pmatrix} \mathbf{x}(s) = A\mathbf{x}(s) = 0 \quad (2.1.17)$$

where the betatron state vector of the particle in one transverse plane is:

$$\mathbf{x}(s) = \begin{pmatrix} x(s) \\ x'(s) \end{pmatrix} \quad (2.1.18)$$

In most cases, the value of $K(s)$ with a magnetic element is constant. For constant K , this equation has the solution:

$$\begin{aligned}
\mathbf{x}(s) &= e^{As} \mathbf{x}_0 \\
&= \left[I \cos \sqrt{K}s + A \frac{1}{\sqrt{K}} \sin \sqrt{K}s \right] \mathbf{x}_0 \\
&= \begin{pmatrix} \cos \sqrt{K}s & \frac{1}{\sqrt{K}} \sin \sqrt{K}s \\ -\sqrt{K} \sin \sqrt{K}s & \cos \sqrt{K}s \end{pmatrix} \mathbf{x}_0 \\
&= M(s|0) \mathbf{x}_0
\end{aligned} \tag{2.1.19}$$

such that the particle's state vector at the start of the lattice may be transformed into the state at any other point by application of the transfer matrix $M(s|0)$. Using the matrix Ω , where:

$$\Omega = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \tag{2.1.20}$$

the transfer matrix is seen to satisfy the symplectic condition:

$$M^\dagger \Omega M = \begin{pmatrix} 0 & \sin^2 \sqrt{K}s + \cos^2 \sqrt{K}s \\ -\sin^2 \sqrt{K}s - \cos^2 \sqrt{K}s & 0 \end{pmatrix} = \Omega \tag{2.1.21}$$

By substituting equation 2.1.10 into the matrix equation, with the initial conditions $\psi_0 = 0$ and $\psi_0 = \pi/2$ providing orthogonal solutions [32], all four matrix elements may be solved simultaneously to give the general form of the transfer matrix:

$$M(s_2|s_1) = \begin{pmatrix} \sqrt{\frac{\beta_2}{\beta_1}} (\cos \Delta\psi + \alpha_1 \sin \Delta\psi) & \sqrt{\beta_1 \beta_2} \sin \Delta\psi \\ \frac{1}{\sqrt{\beta_1 \beta_2}} [(\alpha_1 - \alpha_2) \cos \Delta\psi - (1 + \alpha_1 \alpha_2) \sin \Delta\psi] & \sqrt{\frac{\beta_1}{\beta_2}} (\cos \Delta\psi + \alpha_2 \sin \Delta\psi) \end{pmatrix} \tag{2.1.22}$$

Since the value of $K(s)$ is generally piecewise constant, transfer matrices may be defined for particular lattice elements then multiplied together to form larger sections of beamline. For the linear machine, the following matrices are required:

- Drift space of length l , where $K(s) = 0$ throughout:

$$M_{drift} = \lim_{K \rightarrow 0} M(s+l|s) = \begin{pmatrix} 1 & l \\ 0 & 1 \end{pmatrix} \tag{2.1.23}$$

- Dipole magnet with bending angle and radius θ and ρ respectively, where $K(s) = 1/\rho^2$

throughout:

$$M_{bend} = M(s + l|s) = \begin{pmatrix} \cos \theta & \rho \sin \theta \\ -\frac{1}{\rho} \sin \theta & \cos \theta \end{pmatrix} \simeq \begin{pmatrix} 1 & l \\ 0 & 1 \end{pmatrix} \quad (2.1.24)$$

- Focussing quadrupole of strength $K(s) = k$ and length l :

$$M_F = M(s + l|s) = \begin{pmatrix} \cos \sqrt{k}l & \frac{1}{\sqrt{k}} \sin \sqrt{k}l \\ -\sqrt{k} \sin \sqrt{k}l & \cos \sqrt{k}l \end{pmatrix} \simeq \begin{pmatrix} 1 & 0 \\ -kl & 1 \end{pmatrix} \quad (2.1.25)$$

- Defocussing quadrupole of strength $K(s) = -|k|$ and length l :

$$M_D = M(s + l|s) = \begin{pmatrix} \cosh \sqrt{|k|}l & \frac{1}{\sqrt{|k|}} \sinh \sqrt{|k|}l \\ \sqrt{|k|} \sinh \sqrt{|k|}l & \cosh \sqrt{|k|}l \end{pmatrix} \simeq \begin{pmatrix} 1 & 0 \\ +kl & 1 \end{pmatrix} \quad (2.1.26)$$

where the thin lens approximations given are valid in the limit $l \rightarrow 0$. As a simple example which will be of later use, the transfer matrix of the commonly used focussing-defocussing (FODO) cell may be determined by means of transfer matrices. Alternating focussing and defocussing quadrupoles of equal strength are separated by drifts of length L . The transfer matrix from the centre of one focussing quadrupole to the next is given by:

$$M_{FODO} = M_{F/2} M_{drift} M_D M_{drift} M_{F/2} \quad (2.1.27)$$

which, for thin lenses of focal length f , is:

$$M_{FODO} = \begin{pmatrix} 1 - \frac{L^2}{2f^2} & 2L \left(1 + \frac{L}{2f}\right) \\ -\frac{L}{2f^2} \left(1 - \frac{L}{2f}\right) & 1 - \frac{L^2}{2f^2} \end{pmatrix} \quad (2.1.28)$$

A string of FODO lattice cells may be joined end to end to form a transfer line.

2.1.5 Beam size, emittance and jitter

A beam consists of a large number of particles each following its own trajectory through the lattice. All particles enter the lattice with their own initial phase space coordinates, and hence betatron phase, and have their own Courant-Snyder invariant ellipse with a conserved area. The beam may be treated statistically, with the beam covariance matrix σ describing

the ensemble of particles:

$$\sigma(s) = \begin{pmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{12} & \sigma_{22} \end{pmatrix} \quad (2.1.29)$$

where the RMS beam size $\sigma_x = \sqrt{\sigma_{11}}$, the RMS divergence $\sigma_{x'} = \sqrt{\sigma_{22}}$ and σ_{12} is the covariance.

A beam is generally well modelled by a bivariate normal distribution in phase space [21]. Considering such a distribution at a point in the lattice where $\alpha(s) = 0$, which implies no correlation between the particles' positions and angles, the phase space area occupied by the beam may be related to an upright ellipse. In the literature, various arbitrary choices have been made as to what proportion of the beam should be enclosed by this ellipse; here, the favoured definition is an ellipse with semi-axes of lengths σ_x and $\sigma_{x'}$. The upright ellipse has the equation:

$$\frac{\sigma_{x'}^2}{\varepsilon_{rms}} x^2 + \frac{\sigma_x^2}{\varepsilon_{rms}} x'^2 = \varepsilon_{rms} \quad (2.1.30)$$

where the RMS emittance $\varepsilon_{rms} = \sigma_x \sigma_{x'}$ has been defined. The area enclosed by the ellipse is $\pi \varepsilon_{rms}$. In general, $\alpha(s) \neq 0$. At a given point in the lattice, the ellipse may be rotated to align its semi-axes with the principal axes of the beam distribution. This rotation gives an ellipse with the following equation [33]:

$$\frac{\sigma_{x'}^2}{\varepsilon_{rms}} x^2 + \frac{\sigma_x^2}{\varepsilon_{rms}} x'^2 - \frac{2\text{Cov}[x, x']}{\varepsilon_{rms}} xx' = \varepsilon_{rms} \quad (2.1.31)$$

Comparing this to the Courant-Snyder invariant in equation 2.1.16, it is possible to identify that:

$$\sigma_x = \sqrt{\beta \varepsilon_{rms}}, \quad \sigma_{x'} = \sqrt{\gamma \varepsilon_{rms}}, \quad \text{and} \quad \text{Cov}[x, x'] = -\alpha \varepsilon_{rms} \quad (2.1.32)$$

As for all Courant-Snyder invariant ellipses, in the absence of interactions between particles, radiative effects and acceleration, the bounded phase space area is conserved. It is easy to verify that:

$$\det[\sigma(s)] = \sigma_{11}\sigma_{22} - \sigma_{12}^2 = (\beta\gamma - \alpha^2)\varepsilon_{rms}^2 = \varepsilon_{rms}^2 \quad (2.1.33)$$

Finally, for a betatron state vector where $\mathbf{x}(s_2) = M(s_2|s_1)\mathbf{x}(s_1)$, the beam covariance matrix is given by:

$$\begin{aligned} \sigma(s_2) &= \left\langle (\mathbf{x}(s_2) - \langle \mathbf{x}(s_2) \rangle)(\mathbf{x}(s_2) - \langle \mathbf{x}(s_2) \rangle)^\dagger \right\rangle \\ &= \left\langle \left(M\mathbf{x}(s_1) - \langle M\mathbf{x}(s_1) \rangle \right) \left(M\mathbf{x}(s_1) - \langle M\mathbf{x}(s_1) \rangle \right)^\dagger \right\rangle \\ &= M\sigma(s_1)M^\dagger \end{aligned} \quad (2.1.34)$$

and as the transfer matrix is symplectic with $\det[M(s_2|s_1)] = 1$ (section 2.1.4), then:

$$\det[\sigma(s_2)] = \det[\sigma(s_1)] = \varepsilon_{rms}^2 \quad (2.1.35)$$

therefore the emittance is conserved under transport as required.

A concept related to the beam size is the beam jitter. When the beam receives an anomalous kick from, for example, injection imperfections or magnet vibration, coherent betatron oscillations may be excited in its constituent particles. Otherwise referred to as beam jitter, these coherent oscillations might be excited on a train-to-train basis, or from bunch-to-bunch within a train, and correspond to displacement of the beam centroid in $x-x'$ phase space (see figure 2.2). Over many trains, the distribution of the centroid displacement in x is widest at locations of maximum $\beta(s)$, whereas in x' it is widest at maximum $\gamma(s)$. Thus, the magnitude of the beam jitter is often quoted relative to the beam size:

$$\frac{\sigma_x^{jitter}}{\sigma_x} = \frac{\sigma_{x'}^{jitter}}{\sigma_{x'}} \sim \mathcal{O}(0.1) \quad (2.1.36)$$

Appendix A outlines a statistical model of bunch-to-bunch jitter, of which train-to-train jitter is simply a special case with perfect correlation between bunches.

2.1.6 Off-momentum orbits

So far only particles with nominal momentum, $p = p_0$, have been considered. A real beam contains particles with a distribution of different momenta, which will take slightly different trajectories through bending magnets. A lattice region where particles of differing momenta are on different orbits is referred to as dispersive. If the deviation of a particle's momentum from the nominal, δ , is written as:

$$\delta = \frac{p - p_0}{p_0} \quad (2.1.37)$$

and substituted into equation 2.1.6, then, to first order, Hill's equation becomes:

$$x'' + K_x(s)x = \frac{\delta}{\rho(s)} \quad (2.1.38)$$

If the solution to the homogeneous equation solved in section 2.1.3 is denoted $x_\beta(s)$, then the general solution to equation 2.1.38 may be written as:

$$x(s) = x_\beta(s) + D(s)\delta \quad (2.1.39)$$

where the dispersion function $D(s)$ is a particular solution to the inhomogeneous equation:

$$D'' + K_x(s)D = \frac{1}{\rho(s)} \quad (2.1.40)$$

A particle with $\delta \neq 0$ undergoes its betatron oscillations about a new off-momentum orbit determined by $D(s)$. Equation 2.1.40 in matrix form is:

$$\frac{d\mathbf{D}}{ds} = \begin{pmatrix} 0 & 1 \\ -K(s) & 0 \end{pmatrix} \mathbf{D}(s) + \begin{pmatrix} 0 \\ \frac{1}{\rho(s)} \end{pmatrix} = A\mathbf{D}(s) + \mathbf{b}(s) \quad (2.1.41)$$

Again $K(s)$ and $\rho(s)$ are taken to be piecewise constant. Choosing the particular solution where $\mathbf{D}(s_0) = 0$:

$$\begin{aligned} \mathbf{D}(s) &= e^{-As} \int e^{As} \mathbf{b} ds \\ &= \begin{pmatrix} \frac{1}{\rho K} (1 - \cos \sqrt{K}s) \\ \frac{1}{\rho \sqrt{K}} \sin \sqrt{K}s \end{pmatrix} \end{aligned} \quad (2.1.42)$$

which, using equation 2.1.39, gives:

$$\mathbf{x}(s) = M(s|0)\mathbf{x}_0 + \delta \mathbf{D}(s) \quad (2.1.43)$$

or equivalently the following 3×3 matrix equation:

$$\begin{pmatrix} x(s) \\ x'(s) \\ \delta \end{pmatrix} = \begin{pmatrix} m_{11} & m_{12} & \frac{1}{\rho K} (1 - \cos \sqrt{K}s) \\ m_{21} & m_{22} & \frac{1}{\rho \sqrt{K}} \sin \sqrt{K}s \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_0 \\ x'_0 \\ \delta \end{pmatrix} \quad (2.1.44)$$

The components m_{ij} are just those of the standard 2×2 transfer matrix. Again, 3×3 matrices may be constructed for particular magnetic elements, which may then be used to propagate an off momentum particle through the lattice.

Here the matrices of concern are again for quadrupoles, dipoles and drift lengths. For the quadrupoles and drifts, the dispersion vector vanishes as $1/\rho = 0$. In the case of bending magnets where $K = 1/\rho^2$ throughout and with bending angle $\theta = l/\rho$, the matrix is:

$$M_{bend}^{3 \times 3} = \begin{pmatrix} \cos \theta & \rho \sin \theta & \rho (1 - \cos \theta) \\ -\frac{1}{\rho} \sin \theta & \cos \theta & \sin \theta \\ 0 & 0 & 1 \end{pmatrix} \quad (2.1.45)$$

2.1.7 Chromatic aberration

Chromatic aberration refers to variation of a quadrupole's focussing strength with particle momenta. If p_0 is the nominal beam momentum and k_0 the nominal focussing strength of a quadrupole, then from equation 2.1.4 an off-momentum particle will experience focussing with a strength given by:

$$k = \frac{p_0}{p} k_0 \quad (2.1.46)$$

Using sextupole magnets, this first-order chromatic aberration may be compensated. Figure 2.3 shows how the six poles are arranged and the resulting field direction in such a magnet. Unlike the quadrupole magnet whose field increases linearly with displacement from the centre, the sextupole field increases quadratically. It is then a non-linear element, with strength defined as:

$$S = \frac{-1}{B\rho} \frac{\partial^2 B_y}{\partial x^2} \quad (2.1.47)$$




Figure 2.3: Arrangement of the poles in a sextupole magnet [30].

The horizontal and vertical components of the magnetic field in a sextupole are [21]:

$$B_x = -B\rho Sxy, \quad \text{and} \quad B_y = -\frac{1}{2}B\rho S(x^2 - y^2) \quad (2.1.48)$$

In order to provide chromaticity correction, the sextupole must be located in a dispersive region. To see this, consider decomposing the orbit of an off-momentum particle into components due to dispersion and betatron motion. Neglecting vertical dispersion:

$$x(s) = x_\beta(s) + D(s)\delta, \quad \text{and} \quad y(s) = y_\beta(s) \quad (2.1.49)$$

The vertical magnetic field experienced by the particle due to a sextupole is:

$$B_y = \frac{-B\rho S}{2} [2x_\beta D\delta + (x_\beta^2 - y_\beta^2) + D^2\delta^2] \quad (2.1.50)$$

and the horizontal component is:

$$B_x = -B\rho S [D\delta + x_\beta] y_\beta \quad (2.1.51)$$

By comparison with equation 2.1.4, it is clear that the first term in both planes corresponds to a momentum-dependent quadrupole field with a focussing strength of $-SD\delta$. Thus, by judicious choice of the sextupole strength, the chromatic aberration introduced by a quadrupole may be cancelled by the sextupole. However, there are two remaining terms in equation 2.1.50: the second term is independent of momentum and is referred to as a geometric aberration, while the third term is a second order chromatic aberration. In addition, the third term in equation 2.1.51 introduces coupling between the horizontal and vertical betatron oscillations. These remaining aberrations can have a significant effect on beam quality, though further measures may be taken to mitigate their impact.

A simple example involves the use of a $-I$ transformer lattice, composed of two identical FODO cells. Each cell is tuned such that the quadrupole focal lengths are related to the drift space between the lenses by $f = L/\sqrt{2}$, giving each cell a $\pi/2$ betatron phase advance in both planes. Then, by equation 2.1.28, the transfer matrix of the combined cells is the negative identity matrix. By separating two identical sextupoles with a $-I$ transformer, the geometric aberrations introduced cancel each other exactly [34].

In practice, a pair of sextupoles per plane must be used to correct both chromaticities. Optimal chromatic correction for minimal sextupole strength is obtained by placing those sextupoles for horizontal chromatic correction at locations of low β_y and high β_x , with the opposite is true for vertical correction. While ideally the two families of sextupoles would not interfere with each other, this is not always possible [32].

There is a price for introducing non-linear elements into the lattice. The requirements of different machines vary, depending whether, for example, they be circular or linear or if they require local compensation schemes. Each application requires a detailed analysis, and the resulting installations are categorised by the order to which they eliminate aberration.

2.2 Overview of the International Linear Collider

Given the complexity of the ILC, a comprehensive description of the machine is far beyond the scope of this thesis. In the following sections, the main components of the ILC will be summarised with emphasis on relevant systems. These summaries are based on the

August 2007 Reference Design Report released by the ILC Global Design Effort [26], which contains detailed descriptions of ILC systems for readers who require them. Unless otherwise indicated, ILC design parameters presented in this chapter are taken from the Reference Design Report.




Figure 2.4: Schematic layout of the International Linear Collider [26].

An overview of the proposed ILC installation is shown in figure 2.4. Electron and positron beams circulate in the two central damping rings, where their emittance is reduced, before being transported to opposite ends of the facility. Each is then accelerated by ~ 11 km of superconducting linear accelerator before they collide at the central interaction point.

Table 2.1 lists the nominal beam specifications for the ILC, as well as for three alternative parameter sets: Low N (fewer particles per bunch), Large Y (larger vertical emittance) and Low P (low beam power). All parameter sets are designed to achieve the same energy and luminosity. The alternative parameter sets will not be discussed here, other than to note that the Low N configuration has only half the nominal bunch spacing, imposing tighter latency constraints on any intra-train feedback systems.

Parameter	Nominal	Low N	Large Y	Low P
Repetition rate (Hz)	5	5	5	5
Number of particles per bunch (10^{10})	2	1	2	2
Number of bunches per pulse	2625	5120	2625	1320
Bunch interval in the Main Linac (ns)	369.2	189.2	369.2	480.0
Average beam current in pulse (mA)	9.0	9.0	9.0	6.8
Normalised x emittance at IP (mm·mrad)	10	10	10	10
Normalised y emittance at IP (mm·mrad)	0.04	0.03	0.08	0.036
Beta function in x at IP (mm)	20	11	11	11
Beta function in y at IP (mm)	0.4	0.2	0.6	0.2
RMS beam size in x at IP (nm)	639	474	474	474
RMS beam size in y at IP (nm)	5.7	3.5	9.9	3.8
RMS bunch length (μm)	300	200	500	200
Disruption parameter in x	0.17	0.11	0.52	0.21
Disruption parameter in y	19.4	14.6	24.9	26.1
Luminosity enhancement factor	1.71	1.48	2.18	1.64
Geometric luminosity ($10^{34} \text{ cm}^{-2} \text{ s}^{-1}$)	1.20	1.35	0.94	1.21
Luminosity ($10^{34} \text{ cm}^{-2} \text{ s}^{-1}$)	2	2	2	2

Table 2.1: Beam parameters for the ILC [26].

2.2.1 Electron source

Electron bunches will be generated at the ILC using a photocathode based direct current gun. Laser light will illuminate the photocathode, producing trains of 2625 bunches separated by 369 ns, each containing 2×10^{10} electrons with at least 80% polarisation. Bunching and pre-acceleration take place in room temperature cavities, bringing the electrons to 76 MeV, before a superconducting linac is used to increase the energy to 5 GeV prior to injection into the damping ring. The repetition rate of the entire machine, and hence the electron source, will be 5 Hz.

2.2.2 Positron source

Positrons are generated by illuminating a Ti-alloy target with a high power, multi-MeV photon beam. The positron bunch train has the same structure as the electron train, again with 2×10^{10} particles per bunch. Generating the requisite high intensity photon beam is done using the main ILC electron beam. After the main electron beam reaches 150 GeV, it is passed through a 147 m superconducting, helical undulator with an on-axis field strength

of 0.86 T and a 1.15 cm period. Upon exiting the undulator, the main electron beam is accelerated further to its nominal 250 GeV, while the resulting 131 kW photon beam strikes the metal target generating e^+e^- showers, which are captured and accelerated. After the electrons and photons have been separated, the positrons are accelerated to 400 MeV using a room temperature L-band linac before being transported to the damping ring complex. A superconducting linac then brings the positrons to 5 GeV prior to injection into their damping ring.

2.2.3 Damping rings

It is the essential role of the damping rings to cool the initial electron and positron beams, producing the extremely low emittances necessary for achieving the ILC's design luminosity. The mechanism by which this is achieved is synchrotron radiation damping, in which the circulating particles lose energy through the emission of photons. The instantaneous power emitted by a relativistic electron (or positron) with energy E in a magnetic field of strength B is [35]:

$$P_\gamma = \frac{e^2 c^3}{2\pi} C_\gamma E^2 B^2 \quad (2.2.1)$$

where c is the speed of light, e the electronic charge and the constant C_γ is defined as:

$$C_\gamma = \frac{4\pi}{3} \frac{r_e}{(mc^2)^3} = 8.85 \times 10^{-5} \text{ m GeV}^{-3} \quad (2.2.2)$$

for the electron rest mass m and classical radius r_e . Considering a simple damping ring, where bends are achieved with perfect dipoles and there is no dipole field outside the bends (an isomagnetic ring), the average radiated power for a particle in the design orbit is then [35]:

$$\langle P_\gamma \rangle = \frac{c C_\gamma E_0^4}{L \rho_0} \quad (2.2.3)$$

where E_0 is the design energy, ρ_0 the fixed dipole bending radius and L the ring's circumference. In a damping ring, one or more RF cavities are installed which accelerate the particles longitudinally on each turn to compensate for the energy lost by synchrotron radiation. Those particles in a bunch with higher energies emit more radiation than lower energy particles, leading to damping of the longitudinal phase space. In addition, photons tend to be emitted in a similar direction to a particle's instantaneous velocity, but is replaced by the RF cavity in the longitudinal direction alone. This ensures damping of the transverse phase space [21].

The damping of the synchrotron and betatron oscillations is exponential in nature, and

can be shown [35] to proceed with time constants as listed below:

$$\tau_s = \frac{1}{2 + \mathcal{D}} \frac{2E}{\langle P_\gamma \rangle}, \quad \tau_x = \frac{1}{1 - \mathcal{D}} \frac{2E}{\langle P_\gamma \rangle}, \quad \tau_y = \frac{2E}{\langle P_\gamma \rangle}, \quad (2.2.4)$$

where the damping partition number \mathcal{D} is given, again for an ideal ring, by:

$$\mathcal{D} = \frac{\alpha_c L}{2\pi\rho_0} \quad (2.2.5)$$

The momentum compaction factor α_c quantifies the change in orbit radius in response to a momentum deviation, and may be calculated by integrating the dispersion function around the damping ring lattice:

$$\alpha_c = \oint \frac{D(s)}{\rho} ds \quad (2.2.6)$$

The limit to which the beam emittance may be reduced is set by the quantum nature of synchrotron radiation emission. Discrete photons are emitted in random, independent events. Each emission from a particle excites small random energy and betatron oscillations, which accumulate over time. The net effect of this random disturbance of individual particle orbits is an increase in the beam transverse and longitudinal emittances. While this effect will not be analysed in detail here (see for example [21]), it is competition between phase space damping and quantum excitation which sets the equilibrium emittances to which the beam in a damping ring will eventually settle.

In comparison to the horizontal dispersion, the vertical dispersion is orders of magnitude smaller. For a perfectly aligned ring the vertical emittance would settle to a negligible level, and in practice the equilibrium vertical emittance is dominated by small vertical dispersion arising from quadrupole offsets, and coupling with the horizontal phase space through quadrupole rolls. It is often then convenient to quote the equilibrium emittance of a ring as the normalised (or invariant) natural emittance, the limit to which the horizontal emittance may be damped in the absence of space charge and other non-radiative effects.

At the ILC, the two damping rings are envisaged to sit in a single, centrally located tunnel. Figure 2.5 shows the layout of one of the rings. Both rings will have identical parameters, with the exception of initial beam emittance at injection which is smaller for the electron ring. The circumference of the rings will be 6.695 km, with the arc sections formed from theoretical minimum emittance (TME) lattice cells. These cells are designed such that the horizontal dispersion and beta functions are minimised at the mid-point of each bending dipole. As suggested by their name, TME cells give, in theory, the optimal equilibrium emittance for the ring [36]. They are composed of room temperature magnets. In addition, sextupole and skew quadrupole magnets are in place to allow chromaticity and betatron coupling correction respectively.




Figure 2.5: Layout of the ILC damping ring [26].

Between the arcs are six straight sections per ring, each composed of simple FODO lattice cells again at room temperature. Two of a ring's straight sections accommodate 18 superconducting RF cavities operating at 650 MHz, a subharmonic of the main linac frequency, to provide longitudinal acceleration. Other insertion devices in the straights include fast injection and extraction sections, and superconducting wiggler magnets with a peak field of 1.67 T. The wiggler magnets increase the amount of synchrotron radiation produced and hence reduce the rings' damping times.

Bunch trains at the ILC, with their long 369 ns spacing, would in their entirety stretch for around 290 km, two orders of magnitude larger than the damping ring circumference. In order for such a long train to be stored, the bunches must be interleaved in the ring with a much shorter (nominally 6 ns) spacing between bunches. For this scheme to be possible, the injection and extraction schemes must utilise very fast kickers such that an individual bunch may be targeted in the ring without disturbing its neighbours. Injection and extraction are interleaved, in the sense that early damped bunches are extracted while later bunches are still circulating. Tight requirements on the kicker speed require the use of fast, electromagnetic stripline kickers in conjunction with septum magnets. The development of a pulser appropriate for driving the ILC injection/extraction kickers is an ongoing research project [37].

Table 2.2 lists some of the ILC damping rings' main parameters.

Parameter	Value
Energy	5.0 GeV
Circumference	6.695 km
RF frequency	650 MHz
Harmonic number	14,516
Momentum compaction factor	4.2×10^{-4}
Nominal horizontal tune	52.40
Nominal vertical tune	49.31
Horizontal/vertical damping time	25.7 ms
Longitudinal damping time	12.9 ms
Damping partition number	2×10^{-4}
Normalised natural emittance	5.0 $\mu\text{m rad}$

Table 2.2: Parameters of the ILC damping rings [26].

2.2.4 Main linear accelerators

Before being accelerated to their design energy, the damped electron and positron beams must be transported from the central damping ring tunnel to the start of their respective linacs at opposite edges of the facility. The low emittance transport lines responsible are the Ring to Main Linac (RTML) lines. After a collimation section, each RTML brings its beam down a long (15.5 km and 14.3 km for e^- and e^+ respectively) return section before the turnaround, in which the beamline curves around 180° back on itself (see figure 2.6). After the turnaround is a double pair of superconducting solenoids designed to rotate the particles' spin vector as required for the current experiment, and a two-stage bunch compressor utilising superconducting RF cavities to achieve a compression factor of 30-45.

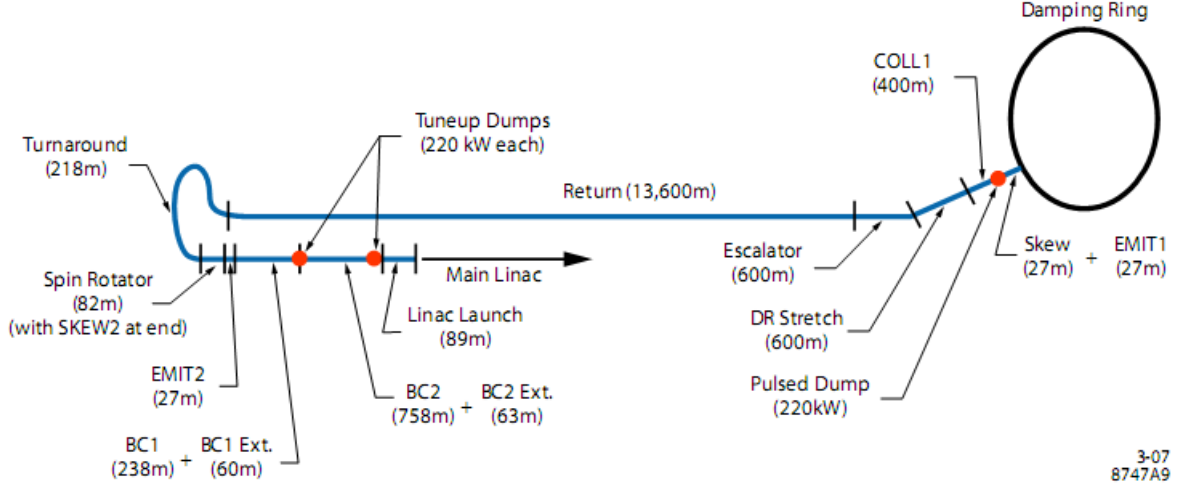


Figure 2.6: Layout of the ILC RTML [26].

The majority of the RTML magnets are at room temperature, and additional upright and skew quadrupoles are used to correct for residual dispersion and betatron coupling. Three laser wire emittance measurement stations will be used to monitor the beam quality. As well as slow trajectory feedbacks operating at the 5 Hz machine repetition rate, the turnaround allows for a fast intra-train feed-forward system in which the position of a bunch upstream of the arc may be measured, processed and used to generate an orbit correction using a fast kicker downstream of the arc. The arc's geometry allows ~ 600 ns processing time, meaning the feed-forward algorithm may be implemented in low latency digital electronics.

To bring the beams up to their design energy of 250 GeV, two superconducting linear accelerators are used. These linacs take the two beams the ~ 11 km back to the centre of the facility where they are delivered to the interaction point. Each is composed of ~ 8500 metre-long nine-cell niobium TESLA-type [38] cavities which operate at 1.3 GHz. To achieve the required energy gain, an average accelerating gradient of 31.5 MV m^{-1} will be necessary, and producing such cavities on an industrial scale is the subject of intense international research.




Figure 2.7: 1.3 GHz nine-cell niobium cavity manufactured by Mitsubishi Heavy Industries [39].

It is essential that the main linacs do not have too severe an impact on the low emittance of the beams. Being superconducting, it has been possible to design the ILC cavities with a relatively large iris (70 mm) aperture which reduces the strength of short-range transverse wakefields [40]. However, induced fields persist in the superconducting cavities meaning that long-range, bunch-to-bunch transverse resonances may lead to breakup of the bunch train. The cavities are designed with higher order mode damping ports, which along with a large bunch spacing help mitigate this effect. The expected emittance growth during acceleration is given in table 2.3 along with general parameters of the linacs.

Parameter	Value
Final energy	250 GeV
Length	11.5 km
# accelerating cavities	8,500
Accelerating gradient	31.5 MVm ⁻¹
RF frequency	1.3 GHz
Klystron repetition rate	5 Hz
Klystron output power	10 MW
Klystron pulse length	1.565 ms
RF pulse power per cavity	293.7 kW
Initial/final normalised horizontal emittance	8.4 / 9.4 μm
Initial/final normalised vertical emittance	24 / 34 nm

Table 2.3: The ILC main linac parameters [26].

2.2.5 Beam delivery systems

Once the beams have been accelerated to the full 250 GeV, they must be focussed to the extremely small sizes required to achieve the ILC's design luminosity. The Beam Delivery Systems (BDS) must match the beams from the main linacs and safely transport them to their respective final focus sections, after which they are brought into collision at the Interaction Point (IP) with a crossing angle of 14 mrad. The components of the BDS along with the design optical functions are shown in figure 2.8.

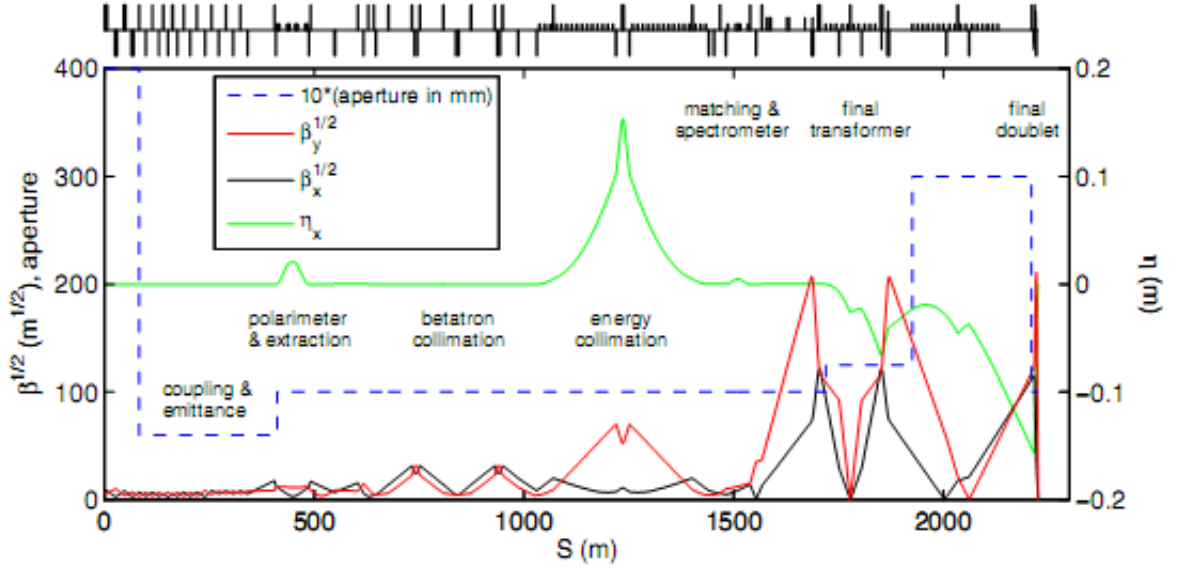


Figure 2.8: Layout of a BDS beamline along with the design optical functions [26].

At the very start of the BDS is a sacrificial collimation section which forms part of the Machine Protection System (MPS). Any beam which has developed a large trajectory error during acceleration will be prevented from doing severe damage to the rest of the BDS, although the sacrificial spoilers which intercept the beam would need to be replaced after suffering an impact. Downstream of the MPS collimation is a set of skew quadrupoles for betatron coupling correction followed by four laserwire beam size monitors for emittance measurement. A magnetic chicane then enables both energy measurement and another sacrificial collimation section forming part of the MPS, in this case to protect from off-energy beams. The MPS is completed with an emergency extraction system which contains both fast and slow kickers capable of directing the beam to a secondary beam dump if an errant bunch is detected. The fast kickers can rise within the 369 ns bunch spacing allowing a beam to be dumped mid-train. This system is also to be used during tuning, safely disposing of the beam whilst the linac trajectories are optimised.

Further collimation sections for both betatron and energy excursions are located down-

stream of the MPS installations. These are designed to control the extent of the beam halo which would otherwise generate unacceptable backgrounds in the detector at the IP. Shielded from multiple high energy bunch strikes by the MPS, the spoilers used here are intended to survive the impact of the beam halo particles. An additional source of background results from muon production in the collimators, which will be suppressed with a 5 m deep iron shield magnetised to 1.5 T in such a way as to leave the beam unperturbed.

At the end of each BDS is a Final Focus (FF) section, with the task of providing the necessary demagnification to achieve the exceptionally small spot sizes (639 nm and 5.7 nm in the horizontal and vertical respectively) required at the IP. The demagnification itself is performed by the Final Doublet (FD), a pair of strong superconducting quadrupoles. Chromatic aberration (see section 2.1.7) introduced by the FD has a strong influence on the achievable spot size, and must be compensated. For the ILC FF, the scheme to perform this compensation is based upon the Next Linear Collider compact final focus design [41], in which a sextupole is positioned beside each FD quadrupole to provide local cancellation of the chromaticity. An upstream dipole creates the necessary dispersion across the interleaved magnets. In order to cancel partially the geometric aberrations introduced by the FD sextupoles, another pair of sextupoles are positioned upstream. At the levels of focussing required at the ILC, higher order aberrations become significant and will require the use of higher order magnets (octupoles and decapoles) to minimise their impact.

2.2.6 Detectors

Detectors at the ILC will be very precise instruments able to capitalise on the unique opportunities the collider will provide. Though it is essential that multiple detectors are available to provide independent measurements, providing multiple interaction points to situate them is an expensive prospect. In particular, multiple interaction points require multiple BDS lines. The current proposal for ILC is to provide a single interaction point with two detectors mounted in a “push-pull” configuration, where the detectors are installed in such a way as to allow either one to be moved to the interaction point with minimal disruption to machine operation (a shutdown of order days [42]).

In August 2009, the International Detector Advisory Group released a report [43] validating two of the competing ILC detector designs, the Silicon Detector (SiD) and the International Large Detector (ILD). Though the detector designs will not be discussed here, the letters of intent submitted by the SiD [44] and ILD [45] groups are available online and contain detailed descriptions.

2.2.7 Extraction lines and dumps

After colliding at the interaction point, the beams must be transported safely to their dumps. A combination of the strong focussing at the FF and interactions between the two beams (discussed in the section 2.3), the exiting beams are of very poor quality. The primary beams come out of collision with large angular divergence and energy spread, accompanied by secondary particles and intense photon beams, and must be captured by the large-acceptance extraction lines. The photon beam travels alongside the primary beam and they are dumped together. Each extraction line begins with a pair of superconducting quadrupoles, after which room-temperature magnets guide the beam ~ 400 m to its dump. Prior to the dump, two energy chicanes are used to enable an energy spectrometer and a polarimeter diagnostic [46].

The dumps themselves are stainless steel cylinders containing water at a pressure of 10 bar. A magnet is used to sweep the bunch train across the titanium dump window in an arc of 3 cm radius, distributing the beam energy over a larger area. This, along with keeping the water flowing at $1 - 1.5 \text{ ms}^{-1}$, limits the rise in temperature due to a bunch train to 40°C during normal operation.

2.3 The ILC interaction point

At the interaction point of the ILC, incredibly dense beams of high energy electrons and positrons collide. Understanding the eventual performance of the machine requires the consideration of a number of factors resulting from these extreme conditions, as well as the imperfections in construction and various dynamic effects such as ground motion.

2.3.1 Luminosity

A key requirement in the design of colliders for particle physics experiments is ensuring a high rate of interactions, to both gather sufficient statistics for measurements and to maximise returns from a what is invariably a large investment. The figure of merit for the machine which governs the achievable rate of interactions is the luminosity \mathcal{L} , which is related to the interaction rate R by:

$$R = \mathcal{L}\sigma \quad (2.3.1)$$

where σ is the cross-section for the physical process of interest and the luminosity is in units of interactions per second per unit area. For perfectly aligned bunches colliding head-on, and which have Gaussian longitudinal and transverse profiles, the luminosity is given by [21, 47]:

$$\mathcal{L} = \frac{f_B N^2 H_D}{4\pi^{3/2} \sigma_s} \int_{-\infty}^{+\infty} \frac{1}{\sigma_x \sigma_y} e^{-\frac{s^2}{\sigma_s^2}} ds \quad (2.3.2)$$

where N is the number of particles per bunch, f_B the bunch-crossing frequency, σ_s the RMS bunch length and, in general, the RMS transverse bunch sizes $\sigma_{x/y}$ are functions of s . The factor H_D is known as the pinch enhancement factor and is discussed in section 2.3.4. For the nominal ILC design this factor is about 1.7.

If the simplifying assumption is made that there is no longitudinal variation in the transverse bunch widths, the integral in equation 2.3.2 is easily evaluated to give the geometric luminosity:

$$\mathcal{L}_0 = \frac{f_B N^2}{4\pi\sigma_x\sigma_y} H_D \quad (2.3.3)$$

To first order, circular collider rings achieve high luminosity by virtue of their high bunch crossing frequencies. Bunches circulate around the machine and come into collision on every revolution. The highest energy lepton collider to date, the circular Large Electron-Positron (LEP) collider, achieved a luminosity of $10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ with a bunch crossing frequency of 44 kHz [48]. Now installed in the old LEP tunnel, the Large Hadron Collider (LHC) aims to achieve $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ with a 40 MHz crossing frequency.

In a linear collider, each bunch has only a single pass through the interaction point before it is dumped, a major limit on the achievable crossing frequency. For the nominal ILC parameter set, $f_B = 5 \text{ Hz} \times 2625 \simeq 13 \text{ kHz}$, yet the design calls for a luminosity of $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. It is for this reason that such small beams are required at the ILC interaction point, as the luminosity is inversely proportional to cross-sectional area.

2.3.2 IP crossing angle and crab cavities

By bringing the beams into collision with a small horizontal crossing angle (14 mrad at the ILC), parasitic collisions outside the IP may be avoided. It also facilitates a simpler extraction of the spent beams. However, the angle also means that the bunches would not naturally collide head on, resulting in a reduction to the luminosity. When the bunch length is much smaller than the vertical beta function at the IP, the reduction in luminosity for a crossing angle ϕ is given by [49]:

$$\mathcal{L} = \frac{\mathcal{L}_0}{\sqrt{1 + \left(\frac{\sigma_s}{\sigma_x} \tan \phi\right)^2}} \quad (2.3.4)$$

In order to prevent this loss of luminosity, the ILC is designed with a crab cavity for each beam located 13.4 m from the IP. These nine-cell superconducting RF cavities operate at 3.9 GHz in the TM110 mode, generating a 5 MVm^{-1} deflecting gradient. The head and tail of the bunch receive opposite transverse deflections, while central particles pass unaffected. This has the effect of rotating the bunches en route to the IP, ensuring they collide head-on and eliminating the crossing angle's effect on the luminosity.

2.3.3 The hourglass effect

After the FDs, the beam sizes decrease, reaching a minimum at the IP. The non-zero length of the bunches means that not all collisions take place at this minimum, and the particle densities are lower than that assumed for the geometric luminosity in equation 2.3.3. The beta functions increase quadratically from the IP, giving the hourglass effect its name. Using a star to denote the values of the Twiss parameters and beam sizes at the IP then, since $\alpha^* = 0$:

$$\beta_{x/y}(s) = \beta_{x/y}^* + \frac{s^2}{\beta_{x/y}^*} \quad (2.3.5)$$

Writing the beam sizes in 2.3.2 in terms of those at the IP gives:

$$\begin{aligned} \mathcal{L} &= \frac{f_B N^2 H_D}{4\pi^{3/2} \sigma_s \sigma_x^* \sigma_y^*} \int_{-\infty}^{+\infty} [1 + (s/\beta_x^*)^2]^{-1/2} [1 + (s/\beta_y^*)^2]^{-1/2} e^{-\frac{s^2}{\sigma_s^2}} ds \\ &= \mathcal{L}_0 R_{HG} \end{aligned} \quad (2.3.6)$$

Therefore the hourglass effect reduces the geometric luminosity by a factor R_{HG} . The integral required to determine this factor may be computed numerically [47]. By using short bunches, the drop in luminosity due to the hourglass effect is reduced.

2.3.4 Beam disruption

During collision, each bunch is influenced by the collective electromagnetic field of the opposing bunch. With the dense particle bunches of the ILC, these strong fields will have a number of important effects. In general, these beam-beam effects are characterised by the horizontal and vertical disruption parameters $\mathcal{D}_{x/y}$, a pair of Lorentz-invariant, dimensionless quantities defined as [50]:

$$\mathcal{D}_{x/y} = \frac{2r_e N \sigma_s}{\gamma (\sigma_x + \sigma_y) \sigma_{x/y}} \quad (2.3.7)$$

where r_e is the classical electron radius and γ the relativistic gamma.

In the case of e^+e^- collisions, the mutual attraction between opposing bunches can lead to an increase in luminosity due to self-focussing, otherwise known as the pinch effect. The value of the disruption parameter governs the strength of the interaction: in the weak regime ($\mathcal{D}_{x/y} \ll 1$) the field of one bunch acts approximately as a lens to the other with a focal length of $\sigma_s/\mathcal{D}_{x/y}$, while in the strong regime ($\mathcal{D}_{x/y} > 1$) particles of each bunch begin to oscillate in the opposing bunch's field.

Though limited analytic models have been developed to predict the luminosity enhancement due to disruption [51], most practical studies make extensive use of computer simu-

lations. At the ILC, the design horizontal and vertical disruption parameters are 0.17 and 19.4 respectively, leading to a pinch enhancement factor (section 2.3.1) $H_D = 1.71$.

After interaction, particles exit the field of the opposing bunch having received angular deflections, the sizes of which are dependent again on the disruption parameters. Unsurprisingly, the precise angular distributions of the divergent beams are difficult to predict analytically and are again studied through simulation. The scale of the distributions is set by the nominal deflection angle given by [50]:

$$\theta_0 = \frac{\sigma_x}{\sigma_s} \mathcal{D}_x = \frac{\sigma_y}{\sigma_s} \mathcal{D}_y \quad (2.3.8)$$

which for the nominal ILC parameters is $\sim 360 \mu\text{rad}$. As an example, the results of a simulation [51] are given in figure 2.9 which shows the vertical distribution of the beam divergence in a disrupted, flat ($\sigma_x \gg \sigma_y$) beam where $\sigma_s = 0.2 \beta_y$ for a range of vertical disruption parameters.




Figure 2.9: Simulated vertical angular disruption as a function of the disruption parameter for a flat beam with $\sigma_s = 0.2 \beta_y$ [51].

In addition to an increase in the angular divergence of the particles after the interaction point, if there is an initial offset between the centroids of the incoming bunches then they will both receive an overall angular deflection. This beam-beam deflection was observed at the Stanford Linear Collider (SLC), shown in figure 2.10. Also in the figure is a simulation of the similar situation expected at the ILC, which shows that bunches with offsets of nanometre order will receive deflections in the tens of microradians.




Figure 2.10: Observed beam-beam deflection curve at the SLC (left) [52] and a simulated beam-beam deflection curve for 500 GeV collisions at the ILC (right) [53].

Though non-linear, the large amplification of the beams' relative offset at IP by the beam-beam deflection provides an important diagnostic for the ILC. By using a BPM located a few metres downstream of the IP, the tens or hundreds of microns deflection resulting from the beam-beam interaction is easily measured. With knowledge of the beam-beam deflection curve the offset at the IP may then be inferred. Indeed, this BPM signal forms the basis for the fast feedback system under development by the FONT group, which would be used to steer the beams into collision and maintain maximum luminosity.

2.4 Degradation of the luminosity

With the ILC relying on the consistent overlap of intense and highly focussed particle beams to achieve its design luminosity, alignment and stability tolerances throughout the machine are tight. Of primary concern is transporting the damped beams from the rings to the interaction point without too severe an emittance growth, for any increase in the size of the colliding beams will quickly degrade the luminosity. Static misalignment of beamline components will cause emittance growth by a number of means: quadrupole offsets and rolls introduce dispersion and betatron coupling respectively, for example. Sextupole displacements are another source of coupling, while tilted RF cavities give undesirable transverse kicks.

Standard laser alignment methods are not precise enough to ensure that the emittance will be preserved as required. Even with a very precise installation, Beam Based Alignment (BBA) algorithms must be employed to steer the beam, using dipole correctors, through the quadrupole centres on a dispersion-free path. Various BBA algorithms have been developed [54], though all are ultimately limited by the resolution of the BPMs used. A number of simulations of the ILC have shown that, in the presence of static misalignments, BBA

will allow the necessary emittance preservation [26].

There are, however, a number of dynamic effects which can still degrade the luminosity from its design value. Beam positions may drift over time due to, for example, temperature changes in magnet cooling systems or cables carrying timing signals. On shorter timescales, beam jitter can stem from a variety of sources including fluctuations in magnet power supply outputs, or in accelerating RF amplitude/phase coupled with residual dispersion. The damping ring extraction kickers might introduce significant jitter and, importantly, the motion of the ground along with other vibrations will cause motion of the beamline components over a wide frequency range.

The presence of beam jitter can impact the luminosity both indirectly, by worsening emittance growth in the low emittance transport lines and main linacs, and directly, by generating transverse offsets between the colliding beams at the IP. In order to keep emittance growth manageable, the RMS jitter should be kept below the RMS beam size before the FF. In particular, these concerns require that $\sigma_{x/y}^{jitter} \lesssim 0.5 \sigma_{x/y}$ at the exit of the main linacs, and $\sigma_{x/y}^{jitter} \lesssim \sigma_{x/y}$ at the FD entrance [55].

Vibration of the FD quadrupoles will translate directly to movement of the beam focal points. Due to the flat beam at the IP, the impact of jitter in the vertical plane is far more severe than a similar magnitude jitter in the horizontal. At the IP, and in the absence of disruption effects, a relative vertical offset Δy between the beams leads to a drop in luminosity given by [56]:

$$\mathcal{L} = \mathcal{L}_0 e^{-\frac{\Delta y^2}{4\sigma_y^2}} \quad (2.4.1)$$

This analytic expectation is plotted in figure 2.11, along with the results of a simulation including beam disruption. In order that the luminosity loss due to jitter be kept at below 10%, the relative vertical offset at the IP must be stabilised to $\lesssim 0.5 \sigma_y$ [57].




Figure 2.11: Rigid model (solid red) and simulation (dashed blue) of luminosity loss as a function of vertical offset at the ILC IP [57].

A number of feedback systems will be used to mitigate the luminosity loss due to these dynamic effects as described in the following sections.

2.4.1 Effects of ground motion

Seismometers have been used at various accelerator sites to measure the power spectral density of the vertical displacement of the ground [58]. Figure 2.12 shows the spectra from four different accelerator sites, plus that from a remote mine acting as a reference. The power tends to fall off as $1/f^4$, although other features are seen. The low frequency peak near 0.2 Hz observed at all sites is due to ocean waves. Cultural noise due to, for example, traffic begins to dominate at frequencies above 1 Hz.




Figure 2.12: Vertical ground motion power spectral density measured at various sites (left), and the same data integrated up from a range of cut-off frequencies to give the RMS ground motion as a function of frequency (right) [58].

Also shown in figure 2.12 are the same data, but integrated upwards from a cut-off frequency. The integrated spectra give the RMS vertical ground motion in the frequency band from the cut-off up to the maximum measured frequency (80 Hz). At the low end of the frequency range, micron level motion may be expected on minute timescales. Any movement of beamline components which is slow relative to the 5 Hz repetition rate may be compensated for using a pulse-to-pulse feedback system. Dipole corrector magnets located in the linac and BDS will be used to fix the local beam trajectory relative to BPM readings, with the corrections applied after each beam pulse. The exact location and type of 5 Hz feedback loops employed has yet to be finalised, but the low iteration rate means all BPM signals can be brought into the accelerator control system and processed in software, which may then vary the corrector magnet strengths. A software based system means the 5 Hz feedback is flexible, with algorithms easily modified and tested.

Figure 2.12 also shows that above 5 Hz, significant ground movement may be expected. Motion in the tens of nanometres might be experienced between machine pulses. Since the

slow feedback cannot correct for this, fast intra-train feedback and feed-forward systems will be required at the ILC.

2.4.2 Damping ring extraction kicker

As described in section 2.2.3, each bunch is extracted from its damping ring by an independent firing of the fast extraction kickers. The design of the pulsed power supplies that will drive these kickers has not yet been finalised, and the reproducibility of the kicker pulses will determine the amount of transverse bunch-to-bunch jitter introduced into the extracted trains. It is not only the magnitude of this jitter which could be troublesome, but also the potential lack of bunch-to-bunch correlation; an intra-train feedback system operating on uncorrelated jitter will serve only to exacerbate the problem.

The solution is made possible by the geometry of the RTML (section 2.2.4) with its arced turnaround section. By measuring the offset of a bunch before it enters the arc, a feed-forward signal may be generated to provide a correction to that same bunch using a kicker after the arc. Each bunch in the train will be measured and corrected independently. Approximately 600 ns are available for signal processing as the beam traverses the arc section, allowing low-latency digital electronics to be used to implement more advanced feed-forward algorithms.

2.4.3 Jitter in the main linacs

The focussing used in the main linacs is relatively weak, and beam jitter stemming from ground motion is unlikely to significantly degrade the emittance of the beam as it is accelerated. It is possible that if ground motion of around 100 nm RMS or more is experienced, however, then the jitter introduced in the main linacs may cause unacceptable emittance growth as the beam passes through the stronger focussing of the BDS [26]. In this case, an intra-train feedback system at the linac exit could be used to reduce the component of the jitter that is correlated bunch-to-bunch.

This feedback system, if necessary, would need BPMs with around 100 nm resolution, a requirement that can currently be met only by cavity BPMs. Fields induced by a passing bunch persist in cavity BPMs, and if they do so for longer than the bunch spacing then resolving individual bunches in the train is difficult or impossible. Research into operating cavity BPMs in multi-bunch mode is ongoing [59], and will determine the form any feedback system at the main linac exit will take.

In addition, residual higher order modes in the accelerating cavities can lead to static bunch-to-bunch offsets in the train. A feedback system at the main linac exit could be configured to apply additional, constant kicks on a bunch by bunch basis to eliminate this static train structure.

2.4.4 Interaction point feedback

To prevent the beam jitter at the IP from having too severe an impact on the luminosity, the beams must be stabilised at the nanometre level. Vibration of the FD quadrupoles is the dominant source of jitter at the IP, with the motion translating directly to the focal points. A dedicated intra-train feedback system will be required to steer the beams into collision.

Since the offset at the IP is so small, the feedback system will make use of the amplification provided by the beam-beam deflection described in section 2.3.4. A BPM located in one of the extraction lines, a matter of metres beyond the IP, will be used to measure the tens or hundreds of microns deflection received by the beam. Because of the large deflection, a micron-resolution stripline BPM may be used, whose signal can be processed with minimal latency and which is easily capable of resolving individual bunches. The BPM reading is used to infer the beams' offset at the IP and generate a feedback signal, which is amplified and used to drive a stripline kicker a few metres upstream of the IP in the opposing beam's FF section. This geometry ensures a minimal signal path length and thus latency for the system.

With the relatively long 369 ns spacing of the ILC bunch train, the feedback can operate in a true bunch-to-bunch sense. The first bunch in the train is measured, and the kicker rises in time to steer the second bunch of the opposing beam in an attempt to zero the beam-beam deflection and therefore the beams' offset at the IP. Then, after each subsequent bunch, the feedback signal is updated. Given the length of the bunch spacing, the necessary signal processing may be performed using low latency digital electronics. Figure 2.13 shows a schematic of an IP feedback installation.

The FONT5 feedback system described in this thesis is designed to fulfil the requirements of an interaction point feedback system for the ILC.




Figure 2.13: Schematic of an interaction point feedback installation for the ILC.

Chapter 3

The FONT5 prototype digital intra-train feedback system

Installed at the KEK Accelerator Test Facility (ATF), FONT5 is a prototype intra-train feedback system designed to meet the operational requirements of an interaction point feedback for the ILC. Micron resolution stripline beam position monitors are combined with stripline kickers to form a beam-based feedback capable of operating in a number of different modes. In particular, the FONT5 system takes advantage of the relatively long (of order 100 ns) bunch spacing planned for the ILC and available at ATF to perform digital signal processing as part of its feedback loops, a feature which allows for both high resolution single-shot beam position monitoring and flexibility, allowing the implementation of advanced feedback algorithms.

3.1 The Accelerator Test Facility

At KEK, Japan, the ATF was conceived as a test bed for the generation of ILC-like, super-low emittance electron beams. Beam operation at the facility commenced in 1997 and confirmation that the design beam parameters had been achieved came in 2001. To generate ILC-like beams, the accelerator complex contains an electron source, linear accelerator, damping ring and extraction line. These subsystems are discussed in turn over the following sections.

The ATF has a number of modes of operation, but in each case the electrons are accelerated in the linac to 1.28 GeV prior to storage in the damping ring. After reaching its equilibrium emittance the pulse is extracted from the damping ring. Originally, a short section of beamline loaded with diagnostic instruments carried the extracted beam to the beam dump. In December 2008 however, the original extraction line was replaced as part of the ATF2 project. This project is a test installation of the type of final focus doublet that

would be employed at the ILC interaction point. It is in the ATF2 extraction line that the prototype FONT5 feedback system was installed.

A summary of the main beam parameters at ATF is given in table 3.1, and figure 3.1 shows the layout of the ATF/ATF2 facility.

Beam parameter	Value
Beam energy	1.28 GeV
Bunch charge	$0.1 \times 10^{10} - 1 \times 10^{10} \text{ e}^-$
Horizontal emittance	$2 \times 10^{-9} \text{ m rad}$
Vertical emittance	$6 \times 10^{-11} \text{ m rad}$
Typical horizontal RMS beam size	70 μm
Typical vertical RMS beam size	7 μm
Single bunch pulse repetition rate	1.5 Hz
Three-bunch pulse repetition rate	0.5 Hz
Spacing of three-bunch train	$\leq 154 \text{ ns}$

Table 3.1: Nominal parameters of the Accelerator Test Facility beam [60, 61].

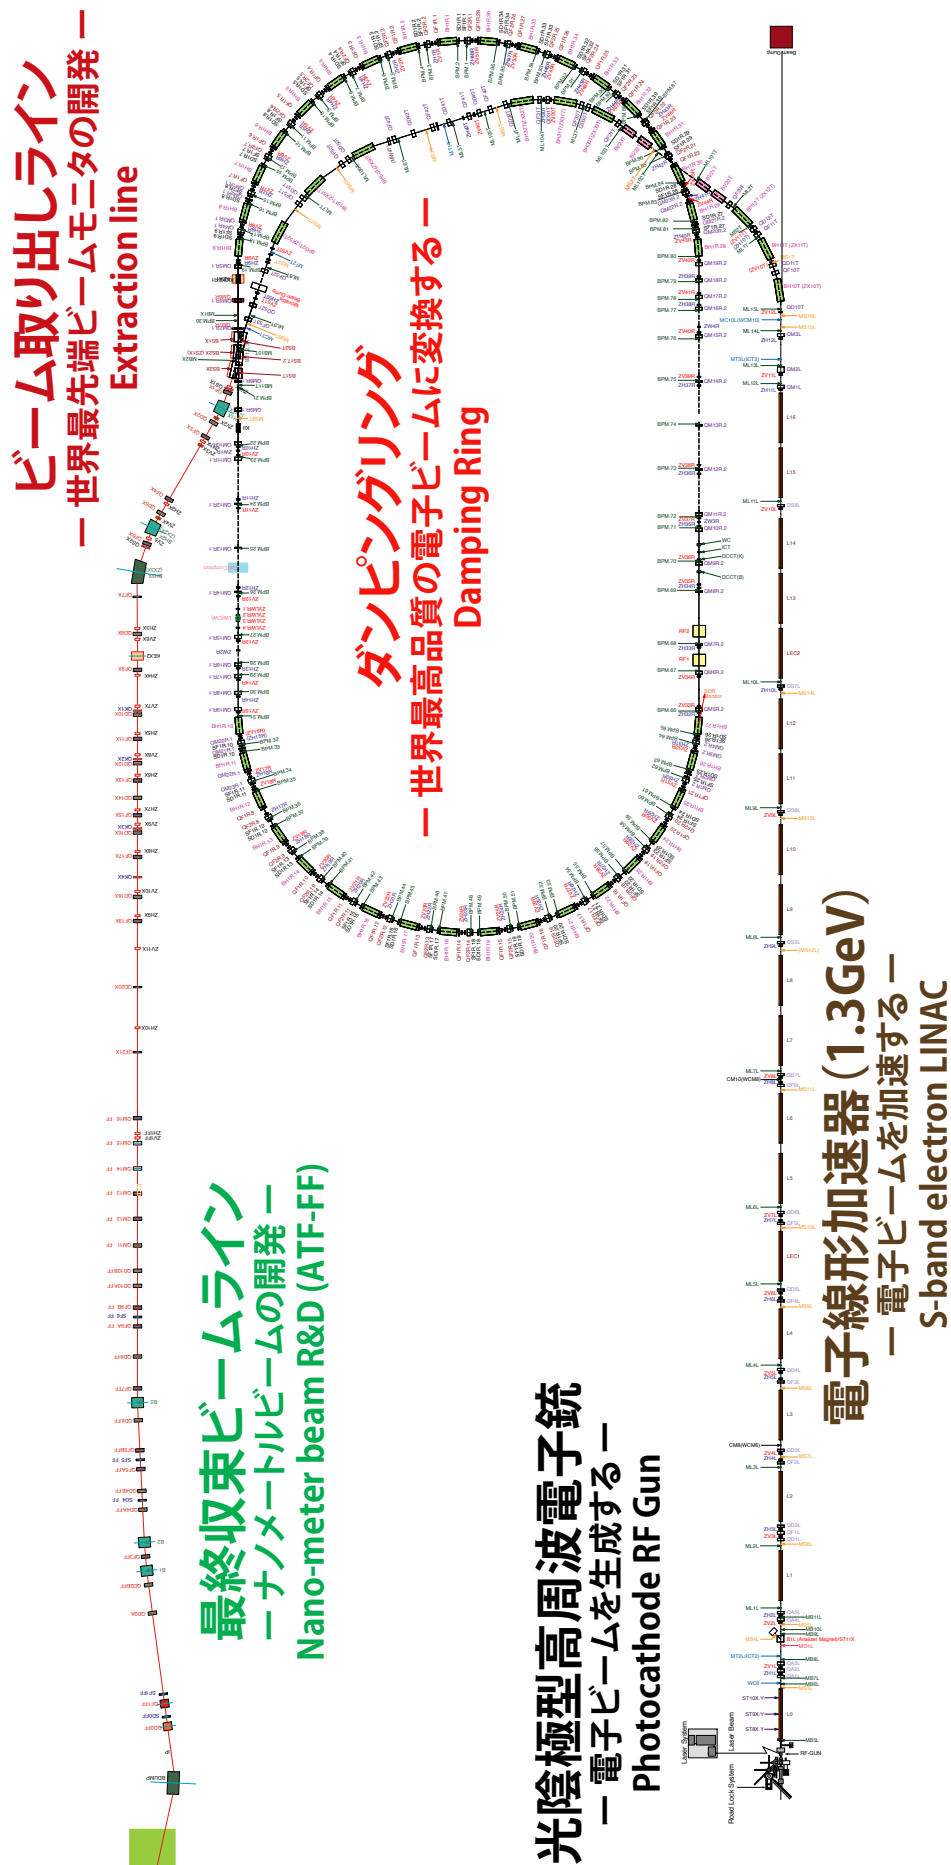


Figure 3.1: Layout of the ATF/ATF2 accelerator complex [62].

3.1.1 Electron source

While electrons at ATF were initially generated using a thermionic gun, this was removed and replaced in 2002 in a successful effort to improve beam quality. Currently, electron bunches are generated with a photo-cathode RF gun driven by a 266 nm [63] laser system. This source allows bunch trains of various structure to be produced. Typical bunches have lengths of order 10 ps and a charge of around 10^{10} electrons. Common operating modes include:

- Pulses containing a single bunch
- Trains of 20 bunches with 2.8 ns spacing

The repetition rate of the gun is variable from 0.7 – 6.2 Hz [61].

3.1.2 Linear accelerator

Electron bunches from the gun are accelerated to 1.28 GeV by a 90 m long linac consisting of 17 S-band accelerating structures. The requisite RF power, at 2856 MHz, is provided by 10 klystrons each capable of producing 4.5 μ s pulses at a rate of up to 12.5 Hz. These pulses have a peak power of 80 MW, but are compressed using the SLAC Energy Development [64] scheme, whereby RF cavities are used to reduce the pulse length to 1 μ s. After compression, the peak power available to each cavity is 200 MW. The accelerating field gradient achieved in the cavities is, on average, 26 MVm⁻¹.

At the exit of the linac the accelerated electron bunches are injected into the ATF damping ring. This is achieved with a combination of fast pulsed kicker magnet (pulse length 120 ns [61]) and a septum magnet.

Parameter	Value
Final energy	1.28 GeV
Length	90 m
# accelerating cavities	17
Accelerating gradient	26 MVm ⁻¹
RF frequency	2856 MHz
# klystrons	10
Klystron repetition rate	12.5 Hz
Klystron peak RF power	80 MW
Klystron pulse length	4.5 μ s
RF pulse compression factor	4.5
Peak RF power per cavity	200 MW

Table 3.2: The ATF linear accelerator parameters [65, 66].

3.1.3 Damping ring

Integral to the generation of ATF’s super-low emittance beams is the damping ring performance. The damping ring at ATF is racecourse shaped with a circumference of 138.6 m, with the two straight sections accommodating a number of insertion devices. These devices include:

- Septum and kicker magnets for injection and extraction of the beam.
- Wiggler magnets to reduce the ring’s damping time by increasing the amount of synchrotron radiation produced.
- An RF cavity required to provide longitudinal acceleration. The cavity operates at 714 MHz, which being a subharmonic of the linac frequency ensures that the frequency of the electron bunches and cavity RF are compatible.

In order to reduce the equilibrium horizontal emittance in the ring, the arc sections are built from low-emittance lattice cells, or FOBO cells (named in analogy to the FODO cell, but reflecting the absence of a defocussing quadrupole), designed to minimise the horizontal dispersion and beta function at the centre of each bend. Each arc consists of 18 such cells. Figure 3.2 shows a schematic of one of the cells. Each cell contains a number of magnets: one horizontal and one vertical dipole corrector; two sextupoles for chromaticity correction; two horizontally focussing quadrupoles; and a combined-function bending magnet, designed to provide vertical focussing.




Figure 3.2: Schematic of the FOBO lattice cells which constitute the ATF damping ring arc sections [67].

Listed in table 3.3 are the main damping ring parameters. The ring has a harmonic number of 330, and every other RF bucket may be filled with an electron bunch giving 165 usable buckets spaced by 2.8 ns. It is the way in which these buckets are filled that determines the structure of the extracted train. When operating in single-bunch mode, a linac pulse consisting of single bunch is injected into a damping ring bucket. In multi-bunch mode, a single firing of the injection kicker directs a 20 bunch pulse from the linac, with 2.8 ns spacing, into consecutive damping ring buckets.

Generating the ILC-like, three-bunch train of the type used in the FONT5 experiments is slightly different. In this case the linac produces three single bunch pulses with its usual, 1.5 Hz, repetition rate which are injected into the ring as usual. No bunches are extracted until all three have been injected, and after the third bunch has been damped a single extraction kicker pulse extracts all three together. In this way, a three-bunch train enters the extraction line.

It is possible to specify the precise buckets into which the three bunches are injected. This choice of buckets determines the spacing of the extracted three-bunch train, meaning the spacing can be varied in 2.8 ns steps. For example, injecting into the 0th, 55th and 110th buckets will give three equally spaced bunches in the ring and produce an extracted train with a 154 ns spacing, the maximum possible. Choosing instead the 0th, 54th and 108th means the bunches will not be spaced symmetrically around the ring, and the extracted train will have a spacing of 151.2 ns.

Parameter	Value
Circumference	138.6 m
Revolution frequency	2.16 MHz
RF frequency	714 MHz
Harmonic number	330
Momentum compaction factor	0.00214
Nominal horizontal tune	15.17
Nominal vertical tune	8.56
Horizontal damping time	17.0 ms
Vertical damping time	27.3 ms
Longitudinal damping time	19.5 ms
Normalised natural emittance	5.1 $\mu\text{m rad}$

Table 3.3: Parameters of the ATF damping ring [63, 68].

ATF’s current extraction scheme uses a double kicker configuration. This configuration is designed to minimise the impact of kicker pulse jitter on the extracted train. In addition to the extraction kicker magnet in the damping ring, a second kicker is placed in the extraction line such that they are separated by a π radian phase advance. Both kickers are driven from the same pulsed power supply, such that any variation in the pulse affects both kicks in the same way. The phase advance means that beam jitter which is introduced at the first kicker should be cancelled out at the second [69].

Note that the injection and extraction kickers’ pulse length is far too long for generating ILC-like trains with more than three bunches. Generation of such a beam requires a kicker that can select a bunch from a single damping ring bucket without disturbing the neighbouring bunches. Development of a fast stripline extraction kicker capable of bunch-by-bunch extraction is the subject of ongoing research at ATF [37]. With bunch-by-bunch extraction, trains of 20 or 60 bunches with *c.* 150 ns bunch spacing would be possible and the FONT5 system is designed with this eventuality in mind.

3.1.4 The ATF2 extraction line and final focus

At ATF, beams with emittances close to those required at the ILC have been demonstrated successfully. In order to achieve the ILC design luminosity, two other factors must be addressed: focussing the beam to produce a nanometre scale spot size in the vertical, and providing nanometre level stability at the focal point. The ATF2 [70] extraction line and final focus was designed to achieve these goals, with installation completed in December 2008. Figure 3.3 shows the ATF2 beamline, highlighting some of the installed instrumenta-

tion. The FONT5 prototype installation is located near to the start of the beamline and is discussed in detail in section 3.3.




Figure 3.3: Schematic of the ATF2 beamline [71].

The ATF2 beamline extends to around 90 metres. Beams are extracted from the ATF damping ring into a matching and diagnostic section. Upright and skew quadrupole magnets are used to correct for beta mismatch, coupling and dispersion. A number of diagnostic instruments including wire scanners, stripline BPMs, high resolution cavity BPMs, optical transition radiation monitors and a laserwire beam size monitor are used to characterise the beam as it is delivered to the final focus section.

As for the ILC, ATF2's final focus is based on the Next Linear Collider compact final focus [41], in which the final doublet and sextupoles are interleaved, but has been scaled to match the ATF energy of 1.3 GeV. The final focus section itself stretches 35 metres, and its layout is shown in figure 3.4.




Figure 3.4: Layout of the compact final focus design used for ATF2 [72].

With the two downstream sextupoles interleaved with the final doublet quadrupoles, the problematic chromaticity introduced by the quadrupoles may be locally compensated. For this to be possible, dispersion is generated across the final doublet using an upstream dipole magnet. Geometric aberrations introduced by the pair of sextupoles are cancelled by another, upstream, pair.

Synchrotron radiation emitted in the final doublet broadens the momentum distribution of the beam, which in turn limits the spot size at the focal point (Oide effect [73]). Longitudinal wake fields introduce additional momentum spread, and local chromaticity

compensation minimises the impact of these effects [41], giving the final focus a large bandwidth of $\sim 1\%$ [72].

At the focal point of the final focus, or interaction point, is located the Shintake beam size monitor [74]. This instrument is designed to make measurements of the extremely small beam sizes by analysing the profile of Compton scattered photons from the interaction of the electron beam with a laser interference fringe pattern.

3.1.5 The physics goals of the ATF2 project

The ATF2 project main physics goals are twofold and summarised below:

Goal 1: demonstration of 37 nm beam size

The first goal of the ATF2 project is to demonstrate that the electron beam may be focussed to its design value of 37 nm. Simulations have demonstrated [72] that this will require the beam to be stable at around 30% of the beam size as it enters the final focus optics.

Goal 2: demonstration of nanometre level beam stability

ATF2's second goal is to stabilise the focussed beam at the interaction point to the nanometre level. Undoubtedly this will require beam-based feedback, and work is currently being undertaken to combine a few-nanometre resolution cavity BPM [75] situated *downstream* of the interaction point with a stripline kicker to form a fast, analogue, intra-train position feedback system. To ensure the focussed beam is stable enough to allow the interaction point feedback to operate, extraction line feedbacks must reduce the magnitude of the jitter at the final focus entrance to below a micron [72].

3.2 The objectives of the FONT5 feedback system installation at ATF2

As well as a technology demonstration for an ILC interaction point feedback, the FONT5 installation at ATF2 will provide essential contributions to the realisation of ATF2 goals 1 and 2 as described in section 3.1.5. These two objectives are discussed in the following sections.

3.2.1 ILC interaction point intra-train feedback technology demonstration

As discussed in section 2.4.4, an intra-train feedback system will be required at the ILC interaction point to provide nanometre-level beam stability. The beam-beam deflection means that large angular kicks are imparted to the bunches travelling out of the interaction point (figure 2.10). Indeed, nanometre level relative offsets between the beams are expected to give deflections of many tens of microns in the feedback BPM. These large deflections are easily measured with a fast sub-micron resolution BPM, such that the overall system resolution may be considered as the error δp_T in the transverse momentum imparted by the kicker. The resulting angular error $\delta\theta$ is then:

$$\delta\theta = \frac{c}{E}\delta p_T \quad (3.2.1)$$

where E is the beam energy and c the speed of light. Thus if the feedback system can provide micron level position stability for a 1 GeV beam, it can provide stability at the nanometre level for a 1 TeV beam. As an ILC technology demonstration then, an objective of the FONT5 installation at ATF2 is to reduce the vertical position jitter at a feedback BPM to a micron or less. This type of feedback experiment requires a single feedback loop, that is, one BPM and one kicker.

3.2.2 Contribution to the physics goals of ATF2

In section 3.1.5, the jitter requirements at the entrance to the ATF2 final focus are given with regards to the ATF2 goals 1 and 2. The second objective of the FONT5 installation is to help provide this necessary stability. This aim requires a more complicated feedback; the FONT5 system is near the start of the ATF2 beamline, and reducing the position jitter alone at a single feedback BPM will not provide the necessary stability at the final focus. Instead the FONT5 installation may be used to decrease both the position and angle jitter of the beam, meaning that at every point downstream of the feedback system the jitter magnitude is reduced by a constant factor. To achieve this, two feedback loops operating in tandem are required, a system comprising two BPMs and two kickers.

3.3 The FONT5 feedback system installation

As shown in figure 3.3, the FONT5 feedback system is situated near the start of the ATF2 extraction line. The complete system consists of several beamline elements, namely three stripline BPMs and two stripline kickers, as well as various electronic components. The constituent parts of the system are discussed in the following sections.

Figure 3.5 shows the locations of the FONT5 BPMs and kickers with respect to the ATF2 extraction line quadrupoles. According to ATF2 naming conventions, the horizontally focussing quadrupoles are incrementally labelled QF[n]X and the defocussing QD[n]X. The three FONT5 stripline BPMS are labelled P1, P2 and P3, while the two stripline kickers are known as K1 and K2. Also shown is the ATF2 stripline BPM MQF15X. Although not part of the feedback system, the BPM MQF15X is occasionally disconnected from the ATF control system and used as an independent witness to the FONT5 operation. The FONT5 BPMs were made to the same design as the ATF2 stripline BPMs such as MQF15X.

Also shown in figure 3.5 are the vertical betatron function magnitude $\beta_y(s)$ and phase advance. In order to compute these functions, the nominal Twiss parameters (section 2.1.3 at the kicker K1) were taken from a Methodical Accelerator Design (MAD [76]) model maintained by the ATF collaboration. A simple finite element analysis implemented in MATLAB [77] was then used to transport the Twiss parameters through the FONT5 region of the ATF2 lattice.

The lattice locations of the BPMs and kickers were chosen to have betatron phase differences as close as possible to the ideal for two degree-of-freedom feedback. See section 5.1.3 for a discussion of the preferred phase advances.

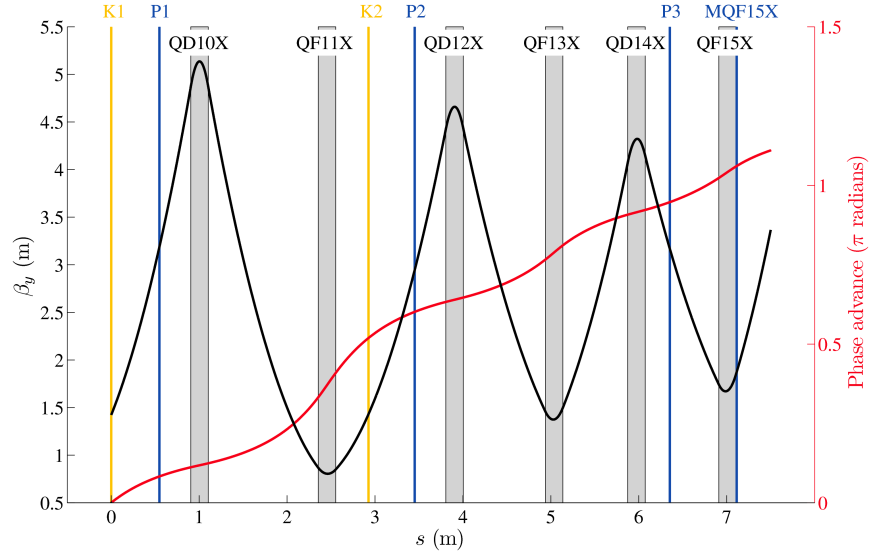


Figure 3.5: Location of the FONT5 feedback system beamline elements with respect to the ATF2 quadrupoles. The Twiss parameter $\beta_y(s)$ is shown along with the betatron phase advance. Beam propagation is from left to right.

The schematic in figure 3.6 shows how the various components of the FONT5 feedback system are connected. At the heart of the system lies the FONT5 digital feedback board. This Printed Circuit Board (PCB) is based around a Field Programmable Gate Array [78] (FPGA), a configurable silicon chip able to implement any number of digital systems. The

FONT5 board digitises the BPM readings, performs digital signal processing and generates feedback signals which may be applied at K1 and/or K2. Since the board is reprogrammable, both kicker outputs may be an arbitrary function of the combined BPM inputs. The system is therefore flexible and capable of running a wide variety of feedback algorithms.

Each BPM has dedicated fast analogue processing electronics (labelled A1-A3 in figure 3.6). The output of these processors is amplified prior to digitisation which essentially eliminates any quantisation noise from the Analogue to Digital Converters (ADCs). The outputs of the FONT5 board - that is the feedback signals - are also amplified to bring them up to the ideal amplitude to drive the final component in the system; the fast kicker amplifiers labelled TMD (see section 3.3.5). All of these components are discussed further in the following sections.

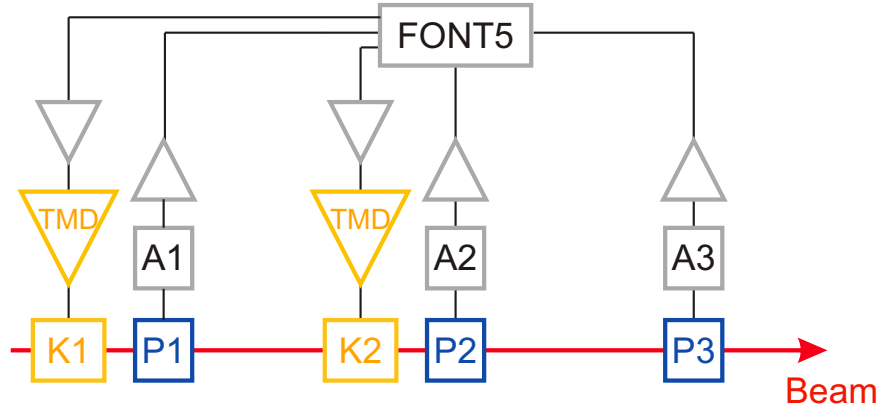


Figure 3.6: Schematic showing the connections between the various electronic and beamline components that make up the FONT5 feedback system.

3.3.1 Stripline beam position monitors

One of the beam position monitoring systems employed in the ATF2 extraction line consists of a number of stripline, or directional coupler, BPMs. The BPMs necessary for the FONT5 feedback system were for simplicity built to the same design specification as those at ATF2. Table 3.4 lists the main parameters. Consisting of four electrodes mounted within a vacuum chamber, one pair top-bottom and the other left-right, the BPMs are capable of measuring beam offsets in both the horizontal and vertical.

Electrode parameter	Value
Length	12 cm
Separation	2.4 cm
Angle subtended	52°

Table 3.4: Parameters of the ATF2 and FONT5 stripline BPM electrodes.

At the upstream end of each electrode signals are carried out via a coaxial cable and as such, the electrode forms part of a transmission line. The downstream end is terminated in the characteristic impedance of the electrode, z_0 . As the beam travels it induces an image current on the beampipe wall. Due to the ultra-relativistic velocity of the electrons, the electromagnetic fields from the bunches are Lorentz-compressed to such an extent that effectively only TEM fields remain. As such, the time-profile of the bunch train is reproduced in the induced image current.

At the edges of the electrodes there are breaks in the electrical continuity of the beampipe wall. As the field lines from a bunch make the transition from wall to electrode, a voltage is induced across the break causing TEM pulses to propagate both out via the coaxial cable and along the electrode. Assuming the resultant wave along the electrode propagates with a similar velocity to the beam, it will arrive at the downstream discontinuity at the same moment another pair of pulses, of opposite polarity, are generated by the beam. Being of opposite polarity the incoming electrode-bound wave will interfere with the pulse entering the downstream termination, all but annihilating it.

Since the remaining wave will take $t \simeq l/c$ s to reach the upstream pickoff, when the beam is dead-centre there will be a time-varying, bipolar voltage signal produced with the form:

$$v(t) = \frac{wz_0}{4\pi b} \left[I_b(t) - I_b\left(t - \frac{2l}{c}\right) \right] \quad (3.3.1)$$

with w the width of the electrode, l the electrode's length, z_0 its characteristic impedance, b the beampipe radius and $I_b(t)$ the current of the beam. The properties of the voltage signal from an electrode in response to an off-centre beam are well covered in the literature [79]. By solving Laplace's equation in two dimensions, the voltages v_T and v_B observed at the top and bottom electrodes can be shown to be proportional to:

$$\begin{aligned} v_T(t) &\propto Q(t) \left[1 + \frac{4}{\phi} \sum_{n=1}^{\infty} \frac{1}{n} \left(\frac{r}{b}\right)^n \cos(n\theta) \sin\left(\frac{n\phi}{2}\right) \right] \\ v_B(t) &\propto Q(t) \left[1 + \frac{4}{\phi} \sum_{n=1}^{\infty} \frac{1}{n} \left(\frac{r}{b}\right)^n \cos(n\theta) \sin\left(\frac{n\phi}{2} + n\pi\right) \right] \end{aligned} \quad (3.3.2)$$

where $\phi = w/b$ is the angle subtended by the electrode, $y = r \cos \theta$ is the beam's vertical offset and $Q(t)$ is the time-varying charge profile of the beam. If only the first order ($n = 1$) terms are kept, taking the sum v_Σ and difference v_Δ of the voltage signals from two opposing electrodes gives:

$$\begin{aligned} v_\Sigma(t) &= v_T(t) + v_B(t) \propto Q(t) \\ v_\Delta(t) &= v_T(t) - v_B(t) \propto Q(t) y \end{aligned} \quad (3.3.3)$$

and as such, to first order, the position of the beam may be calculated from:

$$y \propto \frac{\int v_\Delta(t) dt}{\int v_\Sigma(t) dt} \propto \frac{v_\Delta(t = t_{peak})}{v_\Sigma(t = t_{peak})} \quad (3.3.4)$$

Equality may be obtained by multiplying the right-hand side of equation 3.3.4 by some constant to be determined by calibration. The sum of the signals gives information regarding the charge of the electron bunches, and the procedure by which the difference of the signals is divided by the sum is referred to as *charge normalisation* of the BPM reading.

3.3.2 Fast analogue BPM processors

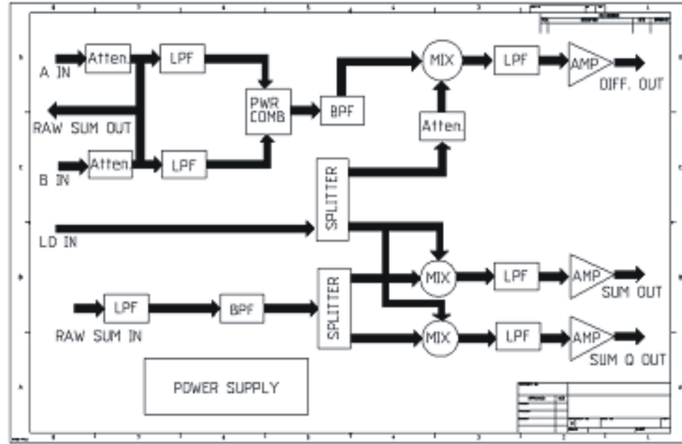


Figure 3.7: Schematic of the FONT3 analogue front-end BPM processing electronics [80].

By analysing the synchrotron light emitted by the electron bunches in the ATF damping ring using a streak camera, it is possible to determine the bunch length. Instrumentation at ATF has been used for this purpose. Typical RMS measurements give 30 ± 10 ps, where the longer lengths are observed at higher beam currents [63]. Since the temporal profile of the beam is reproduced in the raw stripline BPM signals, the observed voltage pulses will be very narrow (dependent on the bandwidth of the readout electronics), and direct digitisation is very difficult.

For the FONT3 analogue feedback experiment at ATF, an analogue PCB was developed which forms both sum and difference voltages from two opposing stripline electrodes, and downconverts these signals to baseband. In the FONT3 experiment the baseband difference signal itself was amplified to form the feedback signal in a very fast analogue position feedback system. For the purposes of the FONT5 feedback experiment, these baseband signals are suitable for digitisation. Design of the analogue processor and results of the FONT3 experiment have been documented [81].

Figure 3.7 shows a schematic of the analogue processing electronics. In brief, after high frequency noise has been removed by low-pass filters, the processor uses a 180° hybrid to subtract one stripline signal from the other. This difference as well as a direct sum of the stripline signals are passed to RF mixer circuits. The local oscillator (LO) for the mixers is a 714 MHz signal from the ATF timing system that is phase-locked to linac accelerating RF and therefore the beam. The baseband components of the mixer outputs are isolated with low-pass filters before being passed on to a pre-amplification stage.

Three mixers output three baseband signals: the difference, the sum and a quadrature sum signal produced by inverting the LO. When the phase of the LO is adjusted correctly, the sum and difference are maximised whilst the quadrature sum is minimised. Thus, the quadrature sum is a useful diagnostic for LO phase adjustment.

The components used to construct the processors were chosen to keep latency low, with the current generation achieving a signal-processing time of 10 ns [80]. As well as keeping the latency of the feedback system low, this also means it is easily capable of resolving the individual bunch positions of the ATF2's *c.*150 ns spaced, three-bunch train.

Each processor board can combine two opposing strip line signals, producing a pair of sum and difference signals. To instrument a BPM fully, two such boards are necessary for monitoring the horizontal and vertical axes.

As can be seen in figure 3.8, the main peak of the processor signal is well described by a Gaussian curve. The RMS width of the main peaks is approximately 3 ns. The design [81], detailed operational analysis [82] and role in the FONT4 experiments [62, 80] of these processors are well documented in various FONT group theses.




Figure 3.8: Example difference (left) and sum (right) signals from the FONT analogue BPM processors. Waveforms were logged with a 5 GSPS digital oscilloscope [83].

3.3.3 BPM mover system

In the summer of 2009, mover systems were installed for each of the three FONT5 BPMs. These movers were designed and manufactured at the Instituto de Física Corpuscular in Valencia, Spain [84]. They are two-axis movers capable of adjusting the BPM position in both the horizontal and the vertical, and an example photograph is shown in figure 3.9.



Figure 3.9: Photograph showing the stripline BPM P2 mounted on its associated mover.

Each mover's range is limited to ± 1.5 mm in both the horizontal and the vertical, with

step sizes as small as 1 μm possible [83]. The two motors in each mover are controlled over a single RS232 connection. With three movers, three RS232 connections are required which are all routed onto the ATF local area network via an Ethernet serial server. Thus the position of each BPM may be independently adjusted using software running in the ATF control room.

One benefit of the movers is that they provide a means of calibrating the BPMs. Whilst holding the beam at a constant position, a BPM can be stepped by a known increment through a series of positions. By recording the average BPM reading at each position, a calibration curve may be plotted. Example curves are shown in section 3.4.3.

Most importantly, the position resolution of the BPM measurements is given by [82]:

$$R_y = \frac{cR_v}{v_\Sigma} \sqrt{1 + \frac{y_0^2}{c^2}} \quad (3.3.5)$$

where R_v is the voltage resolution of the sum and difference channels, v_Σ is the magnitude of the sum signal, y_0 is the mean beam offset and c is a constant. The system resolution is degraded at large beam offsets, making it difficult to measure the relatively small beam jitter. Meeting ATF2's experimental goals requires careful tuning of the whole machine, and the FONT5 feedback system must be able to operate at various different optics where small beam offsets in the feedback BPMs are by no means guaranteed. Using the movers, it is possible to centre the BPM on the beam such that the jitter is measured with respect to a zero offset, thus maximising system resolution.

3.3.4 Stripline kickers

Control of the ATF2 beam is achieved with two stripline kickers. Similar to the stripline BPMs, each kicker consists of a pair of electrodes mounted in a vacuum chamber. The pair of electrodes is connected to coaxial cables at the downstream end and shorted together at the upstream end, forming a transmission line. In order to provide vertical beam deflections the kickers are orientated with their electrodes at the top and bottom of the beam pipe. When current pulses of opposite polarity are applied to the two electrodes, a passing electron bunch receives a transverse deflection [85]. The kickers used in the FONT5 system were provided by the SLAC National Accelerator Laboratory.

3.3.5 Kicker amplifiers

The John Adams Institute for Accelerator Science [86] worked with the UK-based RF company TMD Technologies [87] to produce an amplifier to drive the stripline kickers. The amplifier is very fast with a latency of ~ 35 ns, and is able to provide up to ± 30 A of drive

current. The output of the amplifier can be maintained for $\sim 10 \mu\text{s}$, meaning it is compatible with ATF plans to produce pulses of 20 or 60 bunches with *c.* 150 ns bunch spacing [88].

Figure 3.10 shows an example of the TMD amplifier output in the laboratory. The input was a 500 ns flat pulse, and the voltage was measured across a dummy load simulating a FONT5 kicker [89].




Figure 3.10: Example output of the TMD amplifier in response to a 500 ns flat pulse [89].

In figure 3.11, the TMD amplifier can be seen installed for initial beam testing in the original ATF extraction line.




Figure 3.11: Photograph showing the TMD kicker amplifier at ATF.

3.3.6 The FONT5 digital feedback board

All digitisation and digital signal processing for the feedback system take place on the re-programmable FONT5 digital feedback board. The PCB was designed and built at Oxford University, and is based around a Xilinx Virtex-5 [90] XC5VLX50T FPGA. Firmware running on the FPGA, designed in Verilog and stored on a Xilinx XCF32P programmable read-only memory, controls three banks of three Texas Instruments ADS5474 [91] ADCs. Each bank of three is clocked independently by the FPGA. Both the FPGA and the ADCs are clocked at 357 MHz by dividing the BPM analogue mixers' 714 MHz LO signal, with the ADCs capturing 14-bit samples at 357 MSPS. The ADS5474 were chosen in part due to their low latency; at 357 MHz their latency is just 9.8 ns

Each bank of three ADCs is assigned to one of the FONT5 BPMs as shown in table 3.5, digitising both the vertical and horizontal difference signals along with a single sum signal from each analogue front end. The ADC inputs are transformer coupled and have Linear Technology LTC2624 [92] Digital to Analogue Converters (DACs) wired in parallel, allowing the signal baseline voltage to be trimmed to zero.

ADC group	Channel	Assignment
1	1	P1 horizontal difference
	2	P1 vertical difference
	3	P1 sum
2	4	P2 horizontal difference
	5	P2 vertical difference
	6	P2 sum
3	7	P3 horizontal difference
	8	P3 vertical difference
	9	P3 sum

Table 3.5: FONT5 digital board ADC channel assignments.

Section 4.3 will deal with the FPGA-based digitisation logic in more detail. On each beam pulse, a set of 164 samples separated by 2.8 ns (357 MSPS) are logged for each of the ADC channels. This is sufficient to capture the three-bunch ATF2 train. The phasing of the clock to each bank of three ADCs is independently variable in 75 ps taps, allowing a sample to be timed such that it sits stably on the peak of the processor output. Since the sum signal has an RMS width of ~ 3 ns (section 3.3.2), there are five samples per bunch: one on peak, two at approximately $\pm 1\sigma_\Sigma$ and two at approximately $\pm 2\sigma_\Sigma$. The ratio of the peak difference to the peak sum sample gives a bunch's position according to equation

3.3.4.

Figure 3.12 shows an example of a digitised sum signal in response to a three-bunch train in the ATF2 extraction line with 154 ns spacing. The outputs of BPM P2's analogue front end electronics are amplified using 17 dB low noise amplifiers from Mini-Circuits [93] before being digitised by the FONT5 digital board. After being processed by FPGA logic, the data are transmitted over RS232 to an Ethernet serial server before being collected over a local area network from the ATF control room. Using the trim DACs the signal pedestal has been nullified, and the decaying voltage present between bunches is due to the amplification stage and does not affect the bunch measurements.




Figure 3.12: Example of a digitised sum signal logged by the FONT5 digital feedback board for a single three-bunch train in the ATF2 extraction line, as observed in the BPM P2. Each datapoint represents an ADC sample.

Output from the board is produced with four Analog Devices AD9744 [94] 14-bit DACs. The FPGA is able to use these DACs to pass feedback or other drive signals to the FONT5 kickers. In addition, the board has a JTAG [95] interface for programming the FPGA and an RS232 chip to enable control and data acquisition. An on-board 40 MHz oscillator is used to clock the parts of the firmware's logic that are not time-critical, for example the RS232 communications. Finally, a number of digital inputs and outputs allow for necessary timing signals to be monitored and generated.

All the board's inputs and outputs are made via Micro Coaxial (MCX) connectors. During operation, the FONT5 PCB is mounted in a fan-cooled case. Patch cables connect the MCX connectors to front panel BNC connectors. Figure 3.13 shows the exposed PCB.




Figure 3.13: Photograph showing the FONT5 digital board removed from its case.

3.4 FONT5 BPM performance

3.4.1 Charge normalisation

Figure 3.14 shows typical charge fluctuations in the ATF2 extraction line over a half-hour period. These data were obtained by monitoring the sum signal of BPM P2's analogue front end using the FONT5 digital board. Clearly the charge is not particularly stable and an RMS variation of $\sim 5\%$ is apparent in this instance. Conditions at the ATF strongly affect the charge profile, with temperature oscillations in the timing system and magnet cooling water along with ambient temperature changes due to weather conditions playing prominent roles. For the feedback system to operate at its full potential, these charge fluctuations must be taken into account.




Figure 3.14: Typical charge profile of a succession of three-bunch trains in the ATF2 extraction line over a half-hour period.

Charge normalisation of the BPM readings can refer to two related procedures. When the FONT5 board returns its digitised analogue processor outputs from the three BPMs, a software Data Acquisition system (DAQ) is able to identify the peak samples for each bunch (section 4.4). By dividing the horizontal and vertical difference samples by the sum samples in software, a set of six positions (bunches 1-3 horizontal and vertical positions) are obtained for each BPM. This DAQ and software processing provide a pulse-by-pulse position readout which is invaluable during operation, but is far too slow to form part of the feedback system proper.

The position signals required for feedback are generated in the FONT5 digital board FPGA. Whilst the FPGA's DAQ logic is operating, its fast feedback logic is already at work producing feedback signals with the lowest possible latency. Charge normalisation in the logic is realised using a look-up table (LUT). Firstly the ADC samples containing the peaks of the vertical difference and the sum signals from the feedback BPM(s) are identified. The vertical difference sample is delayed whilst the sum sample addresses a LUT loaded with its reciprocal. The output of the LUT is then multiplied with the vertical difference sample using a Virtex5 DSP48 [96], giving a value proportional to the position alone. This value is used to generate the feedback signal.

3.4.2 BPM calibration using corrector magnets

Obtaining accurate calibrations of the FONT BPMs' position responses is essential for understanding the effectiveness of the feedback system, even if not essential for its operation. The preferred method which has been tried and tested involves using an upstream corrector dipole magnet (in this case magnet ZV6X) to introduce vertical deflections in the beam. The properties of the magnet are well understood via modelling, magnetic field measurements and cross-calibration with ATF2 BPM systems. Transfer matrices are used to calculate what offsets at the BPM locations are expected to result from the beam deflections. By comparing the BPM readings with a series induced offsets, accurate and consistent calibrations are obtained. Example calibration curves from the 10th of May 2010 are shown below, and the procedure is discussed in more detail elsewhere [82]. Each data point is the mean bunch position over approximately 30 pulses.




Figure 3.15: Calibration curves for P1. The mean uncalibrated position measurements are plotted against the deflection introduced using the dipole corrector magnet ZV6X. Data are for bunch 1 (a), bunch 2 (b) and bunch 3 (c), where solid lines show linear fits.




Figure 3.16: Calibration curves for P2. The mean uncalibrated position measurements are plotted against the deflection introduced using the dipole corrector magnet ZV6X. Data are for bunch 1 (a), bunch 2 (b) and bunch 3 (c), where solid lines show linear fits.




Figure 3.17: Calibration curves for P3. The mean uncalibrated position measurements are plotted against the deflection introduced using the dipole corrector magnet ZV6X. Data are for bunch 1 (a), bunch 2 (b) and bunch 3 (c), where solid lines show linear fits.

Table 3.6 summarises the calibration constants obtained from linear fits to the calibration data. The dimensionless difference/sum voltage ratio is multiplied by the gradient to give a position in microns. Very good agreement, better than 1%, is observed between data from different bunches in the same BPM. The constants obtained for BPMs P2 and P3 have a similar level of agreement, though a higher constant is observed for P1. These differences may be accounted for by variation in the analogue front-end electronics [82]. The measured offsets differ from bunch to bunch since the train is not flat, with the train shape changing as the betatron phase advances between BPMs.

		χ^2/df	Gradient (μm)	Offset (μm)
P1	Bunch 1	5.46	412 ± 1	-2.44 ± 0.39
	Bunch 2	7.15	414 ± 1	-8.60 ± 0.41
	Bunch 3	5.28	415 ± 1	3.92 ± 0.35
P2	Bunch 1	39.2	396 ± 3	24.7 ± 1.08
	Bunch 2	38.9	397 ± 3	-7.93 ± 1.03
	Bunch 3	27.7	397 ± 3	-9.75 ± 1.05
P3	Bunch 1	12.8	396 ± 2	23.8 ± 0.83
	Bunch 2	9.56	397 ± 2	11.2 ± 0.73
	Bunch 3	10.7	398 ± 2	0.09 ± 0.75

Table 3.6: Summary of the parameters of linear fits to the corrector magnet calibration curves. Standard errors on the parameters are given.

3.4.3 BPM calibration with movers

The FONT5 BPM mover system installed at the ATF2 extraction line during the Spring of 2010 provides a second means of BPM calibration. By holding the beam at a constant offset and moving a BPM, the BPM's response may be calibrated. The data presented in this section describe the first operation of the mover system, and the mover position axis is simply the value input into the mover control software.

Especially in P3, some deviation from the predicted response is clear and believed to be due to inaugural quirks of the mover installation and operation. These effects are systematic, potentially due to, for example, mechanical friction, and as such the uncertainty in the obtained calibration constants is larger than the statistical errors indicate. There is still a broad agreement between the calibrations obtained using corrector magnets and those using the movers, however.

Further investigation of the mover response has been documented [82]. Again in the following figures, there are around 30 pulses averaged per datapoint. Table 3.7 summarises the parameters of the linear fits to the calibration curves.




Figure 3.18: Calibration curves for P1. The mean uncalibrated position measurements are plotted against the BPM mover's specified offset. Data are for bunch 1 (a), bunch 2 (b) and bunch 3 (c), where solid lines show linear fits.

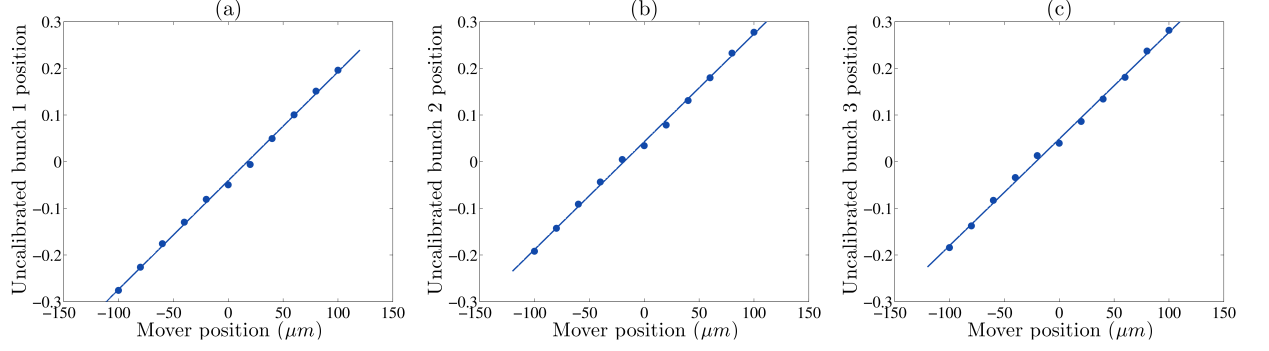


Figure 3.19: Calibration curves for P2. The mean uncalibrated position measurements are plotted against the BPM mover's specified offset. Data are for bunch 1 (a), bunch 2 (b) and bunch 3 (c), where solid lines show linear fits.

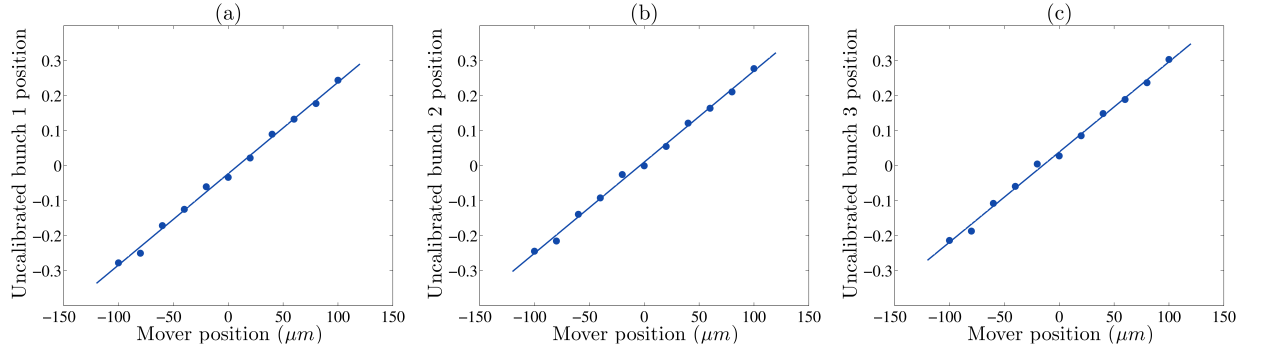


Figure 3.20: Calibration curves for P3. The mean uncalibrated position measurements are plotted against the BPM mover's specified offset. Data are for bunch 1 (a), bunch 2 (b) and bunch 3 (c), where solid lines show linear fits.

		χ^2/df	Gradient (μm)	Offset (μm)
P1	Bunch 1	5.90	427 ± 3	5.19 ± 0.47
	Bunch 2	6.10	429 ± 3	-1.11 ± 0.42
	Bunch 3	7.50	429 ± 3	12.1 ± 0.50
P2	Bunch 1	22.9	429 ± 6	17.4 ± 0.85
	Bunch 2	26.5	433 ± 6	-18.4 ± 0.88
	Bunch 3	25.7	438 ± 7	-21.2 ± 1.03
P3	Bunch 1	32.6	383 ± 7	8.74 ± 1.24
	Bunch 2	35.5	384 ± 7	-3.92 ± 1.21
	Bunch 3	35.9	388 ± 8	-15.1 ± 1.35

Table 3.7: Summary of the parameters of linear fits to the BPM mover calibration curves. Standard errors on the parameters are given.

3.4.4 Estimation of the BPM resolution

All BPMs in the FONT5 system have an intrinsic resolution, defined as the RMS width of the position measurement errors introduced on a bunch-by-bunch basis. Both thermal noise in the analogue processor and digitisation noise in the ADC contribute, although in practice the amplification of the analogue processor signals prior to digitisation make thermal noise by far the dominant source.

Estimating the BPM resolutions requires at least three simultaneous sets of position measurements. In the absence of further information, it must also be assumed that the resolutions of the three BPMs are equal. Given the three BPMs P1, P2 and P3, the position of the beam in any given BPM is a linear combination of the other two positions:

$$c_1 y(s_{P1}) + c_2 y(s_{P2}) + c_3 y(s_{P3}) = 0 \quad (3.4.1)$$

where the constants c_i are uniquely determined by the transfer matrices between the BPMs. For a given bunch within a beam pulse, the true beam position is related to the measured beam position by:

$$y^*(s_{Pi}) = y(s_{Pi}) + r^{Pi} \quad (3.4.2)$$

The random variable r^{Pi} is drawn from a normal distribution with standard deviation equal to the resolution of the i^{th} BPM, R^{Pi} . For this beam pulse, the *residual* δ^{Pi} is defined as the difference between the measured position in the i^{th} BPM and the position in that BPM as predicted from the measured positions in the other two BPMs. If y^* denotes the measured position, then:

$$\begin{aligned} \delta^{Pi} &= y^*(s_{Pi}) + \frac{c_j y^*(s_{Pj}) + c_k y^*(s_{Pk})}{c_i} \\ &= \left[y(s_{Pi}) + \frac{c_j y(s_{Pj}) + c_k y(s_{Pk})}{c_i} \right] + \left[r^{Pi} + \frac{c_j r^{Pj} + c_k r^{Pk}}{c_i} \right] \end{aligned} \quad (3.4.3)$$

Using equation 3.4.1, the first term vanishes. Considering an ensemble of many beam pulses and taking the variance of the residuals, it can be seen that if the BPM errors are uncorrelated:

$$\text{Var}[\delta^{Pi}] = \sigma_{Pi}^2 = (R^{Pi})^2 + \frac{c_j}{c_i} (R^{Pj})^2 + \frac{c_k}{c_i} (R^{Pk})^2 \quad (3.4.4)$$

The system of three equations may be written in matrix form:

$$\begin{pmatrix} 1 & c_2/c_1 & c_3/c_1 \\ c_1/c_2 & 1 & c_3/c_2 \\ c_1/c_3 & c_2/c_3 & 1 \end{pmatrix} \begin{pmatrix} (R^{P1})^2 \\ (R^{P2})^2 \\ (R^{P3})^2 \end{pmatrix} = \begin{pmatrix} \sigma_{P1}^2 \\ \sigma_{P2}^2 \\ \sigma_{P3}^2 \end{pmatrix} \quad (3.4.5)$$

Encoded in the 3×3 , rank 1 matrix is the geometry of the beamline. Since the system of three equations is rank deficient, there exists no unique solution for the three resolutions in terms of the measured residuals. A unique solution may be obtained by applying the constraint $R^{P1} = R^{P2} = R^{P3} = R$. Figure 3.21 shows the three bunch residuals from a 454-pulse data set taken at ATF2 on 16th April 2010. Each distribution has been scaled by the geometric factors from equation 3.4.5, and a Gaussian fit performed. The measured BPM resolution R is given by the Gaussians' standard deviations, which are summarised in table 3.8 along with their standard errors.

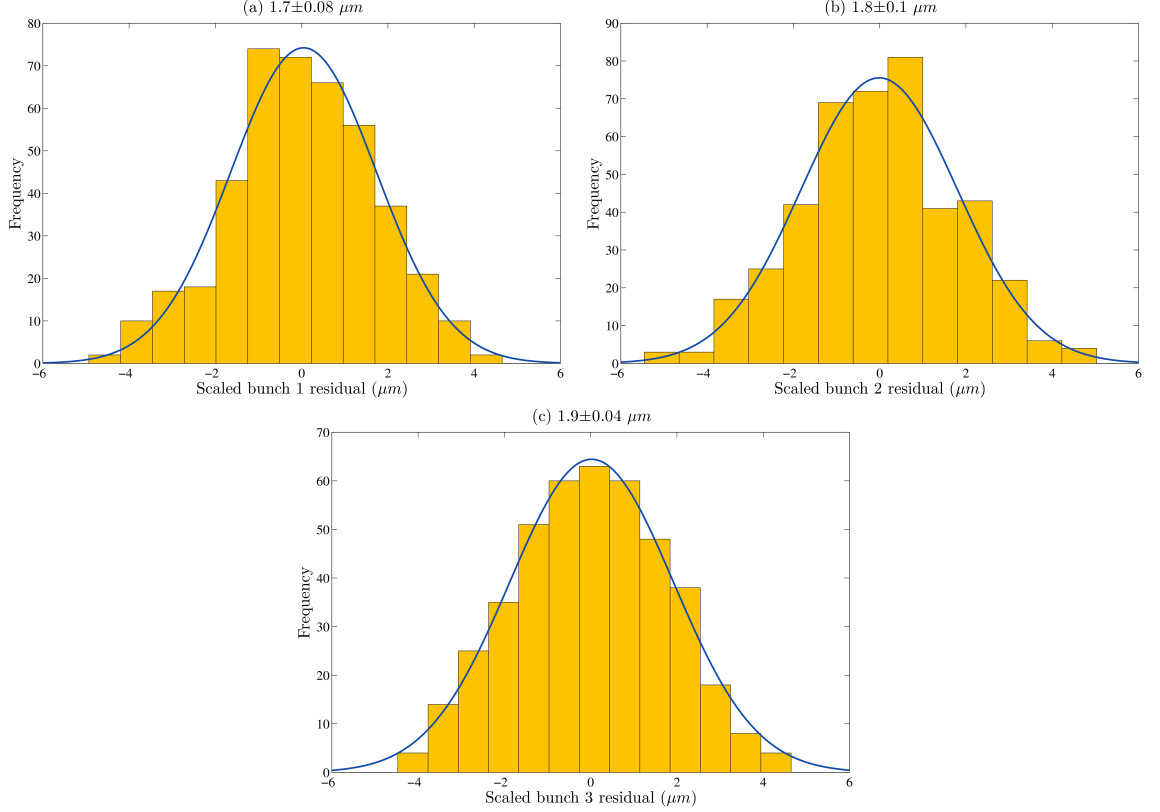


Figure 3.21: Distribution of position residuals for bunches one (a), two (b) and three (c). Each distribution has been scaled by the appropriate geometric factor from equation 3.4.5. Gaussian fits to the distributions are shown by solid lines.

Bunch	Resolution (μm)
1	1.7 ± 0.08
2	1.8 ± 0.10
3	1.9 ± 0.04

Table 3.8: Measured FONT5 BPM resolutions with standard errors using data from three-bunch trains. The resolutions of BPMs P1, P3 and P3 have been assumed to be identical.

Resolution measurements using this method consistently give results below 2 μm , which in itself is very impressive for a low latency single pass stripline BPM measurement. These results are valid only in the case that BPMs P1, P2 and P3 have equal resolution, however, and there is reason to believe that some FONT5 BPMs are operating with resolution better than 1 μm . Refinements of the resolution estimations are discussed in section 5.3.

3.5 Feedback algorithms

Before laying out the details of the FPGA logic implementation of the FONT5 feedback system, the algorithms used are outlined in this section. A more mathematical treatment of the system is presented in section 5.5.

3.5.1 Vertical position feedback

Vertical position feedback, or single loop feedback, is designed to minimise vertical position jitter at a single point in the lattice. Such a system is necessary at the ILC interaction point to maintain luminosity (section 2.4.4). A single BPM-kicker pair is required for such a system, and in the FONT5 vertical position feedback experiments at ATF2 a feedback loop is created between the BPM P2 and the kicker K1 (see figure 3.6).

Being an intra-train, or bunch-by-bunch, feedback system the algorithm followed is a simple iterative procedure:

1. The position offset of the n^{th} bunch at the feedback BPM is measured.
2. A gain factor is multiplied with the n^{th} bunch measurement to determine the kick strength necessary to correct the measured offset.
3. The value of the kick so determined is added to the $(n - 1)^{\text{th}}$ correction stored in a ‘delay loop’ register, forming the n^{th} correction.
4. The n^{th} correction is output on the DAC that drives the kicker K1, and stored in the delay loop register for the next iteration.
5. Before the $(n + 1)^{\text{th}}$ bunch arrives at K1, the kicker output reaches its specified value.

Since the first bunch in a train necessarily receives no correction, it is commonly referred to as the *pilot bunch*. Because the second bunch has been corrected before its offset is measured, the delay loop mechanism is necessary to maintain the current correction. On each iteration the correction is tweaked in response to any position offset of the current bunch. In this way, the feedback system can deal with bunch trains of any length even if the train profile is not flat.

Latency is, in an intra-train feedback system, of great importance as it determines the minimum bunch spacing it can operate on in a true bunch-to-bunch sense. This means that in a practical implementation the various stages of signal processing must proceed as efficiently as possible. With this in mind, the FONT5 digital board implements application of the gain factor in the same step as charge normalisation. Each entry of the charge normalisation LUT is simply multiplied by the gain factor, such that the gain is applied during the subsequent multiplication stage.

3.5.2 Vertical position and angle feedback

Though the FONT5 intra-train feedback for the ILC interaction region is designed to provide only position feedback, other applications require a true reduction in jitter. This is possible by reducing both the position and angle jitter using two degree-of-freedom feedback. To do this in the vertical plane requires two BPMs and two kickers, and for the FONT5 experiments in the ATF2 extraction line BPMs P2 and P3 are used along with both kickers K1 and K2. When operating on two degrees-of-freedom, the feedback system will result in reduced jitter magnitude at every downstream point in the lattice, enabling targets such as micron-level stability at the entrance to the ATF2 FF (section 3.2.2).

Two isolated single loops could, if positioned at perfect betatron phase (section 5.1.3), correct both the position and angle jitter. To achieve the best operation in practice however, two loops must be coupled together to prevent them interfering. In the coupled FONT5 feedback system, the kicks from K1 and K2 must both be linear combinations of the measured bunch positions in P2 and P3.

In essence this may be treated, for the purposes of digital processing, as a set of four gains with the kick at K1 being $\theta_1 = G_1 y^{P2} + G_2 y^{P3}$ and similarly for K2. In firmware this is implemented with four LUTs in parallel, each scaled by the appropriate gain. Again, they simultaneously perform charge normalisation of the position measurements. The sum of the output of each pair of LUTs goes out onto the DACs.

As ever, latency is a key concern. By coupling two loops together in this way, the latency is unavoidably increased. The length of the beam pipe and the return cable path between the two outermost elements, K1 and P3, sets the minimum latency.

3.5.3 Bunch train flattening

Rather than being flat, the three-bunch trains extracted from the ATF damping ring invariably have bunch-to-bunch position structure. This structure can be quite constant over long periods of time for a given machine configuration. If the feedback operates on bunches which jitter about different mean positions, the jitter of each bunch will still be reduced. However the bunches will still jitter around different means.

To remove this train structure the FONT5 digital feedback board may be used to apply a different kick to each bunch which remains constant pulse-to-pulse. This constant kick may be added to the bunch's feedback signal by the FPGA logic, thus allowing both jitter reduction and train flattening to be performed simultaneously.

Using software, constant kicks for both K1 and K2 may be specified for each bunch manually, or in the future automatically using a slow software feedback. The kicks are chosen to bring the mean positions of bunches two and three into line with the mean position of the pilot bunch. Doing so will reduce the range of kick available for feedback, but with typical beam jitters of order a micron this is rarely a concern.

Figure 3.22 shows the profile of a bunch train viewed in BPM P2 with and without application of bunch train flattening by the FONT5 system. Each datapoint represents the average bunch position over 30 pulses, with the error bars representing the RMS jitter. The individual bunch kicks were specified manually and the flattening is not perfect; a slow software feedback would give improved results.




Figure 3.22: Flattening of the bunch train in BPM P2. The blue datapoints show the mean uncorrected bunch positions, and the yellow datapoints the mean positions with a different constant kick applied to bunch two and three. Each datapoint is averaged over 30 pulses.

3.5.4 Kicker pulse droop compensation

These algorithms assume a flat kicker pulse in K1 and K2, which unfortunately is not the case. As the kicker first rises there will be some overshoot before it settles, which may catch early bunches especially where the latency is close to the bunch spacing. For long bunch trains, the kicker pulse will display an exponential decay in its output (see figure 3.10).

Another strength of digital feedback systems is the ability to correct for such a kicker pulse shape in the logic, mathematically increasing the specified kick to compensate for the droop. Of course there is a limit to how far the kicker can be driven, but with observed jitters at the few micron level it is possible to deal with long bunch trains.

The mechanism used in the FPGA logic to compensate for kicker droop over the current ATF2 three-bunch train is to apply a finite impulse response filter (FIR) to the two DAC outputs which drive K1 and K2. The FIR approach allows overshoot of an arbitrary shape to be compensated, and can be configured by calibration.

For a single feedback loop of gain G , the i^{th} kick is given by:

$$\theta_i = Gy_i + \delta_i \quad (3.5.1)$$

where $\delta_i = \theta_{i-1}$ is the current value stored in the delay loop register. With an n -tap FIR filter on the output (see figure 3.23), this becomes:

$$\theta_i = Gy_i + \delta_i + k_1Gy_{i-1} + k_2Gy_{i-2} + \dots + k_nGy_{i-n} \quad (3.5.2)$$

By appropriately specifying the tap weights k_j , the overshoot and decay of the kicker pulse is compensated. The number of taps necessary depends on the number of bunches in the train and the shape of the kicker pulse (itself dependent on the kicker impedance). This is best determined empirically using beam measurements.

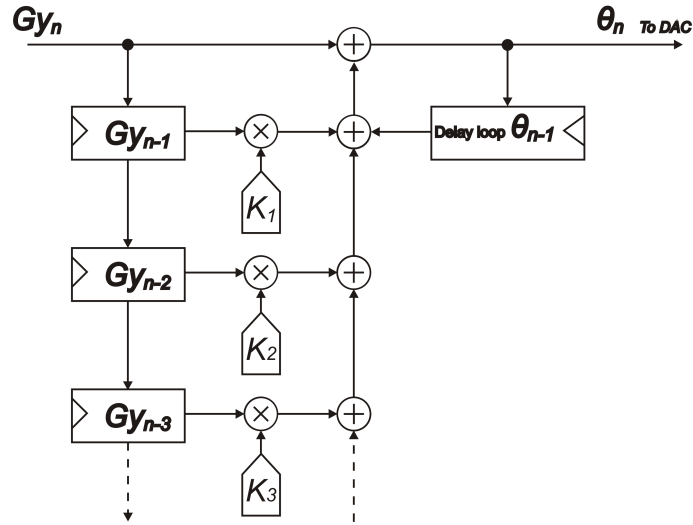


Figure 3.23: Schematic showing the implementation of the FIR filter on each FONT5 DAC output to compensate for droop of the kicker pulses.

For the current version of the FONT5 system, only a single tap is implemented as there are just three bunches per ATF2 train. The architecture of the filter makes the addition of further taps straightforward in the future. For longer trains, the initial overshoot of the kicker will affect only the first few bunches. Later in the kicker pulse its tendency is to decay exponentially, allowing a simpler infinite impulse response (IIR) filter to compensate. This would be done by scaling the delay loop by a single coefficient on each iteration of the feedback. With a combination of FIR and IIR filters, a long bunch train could be well

corrected despite an imperfect kicker pulse shape and using just a small number of filter coefficients.

3.5.5 Constant DAC output mode and calibration of the kickers

Another simple addition to a digital feedback system is the ability to specify a single, constant kick over the entire bunch train at each kicker. This causes the system to behave functionally similarly to a simple dipole corrector magnet, which in itself is only occasionally of use. In the FONT5 system however, the constant kick replaces the feedback signal but retains its timing characteristics. That is, the constant kick rises as soon as possible after the pilot bunch arrives with the delay being the system latency. In section 5.2 this mechanism is used to perform measurements of the FONT5 latency.

In addition, a series of constant kicks applied at a kicker may be used to calibrate it. The data of figures 3.24 and 3.25, taken on the 21st of April 2010, were obtained by setting the magnitude of the FONT5 DAC outputs to a range of constant values. Due to droop of the kicker pulses, the third bunch receives a smaller kick than the second. The gradients of the second bunch's response are used to calculate the feedback gains following the procedure outlined in section 5.1. The ratio of the second and third bunch gradients is used to calculate the coefficient k_1 of the FIR filter. If v^D is used to represent the DAC output voltage then:

$$k_1 = \frac{dy_2/dv^D}{dy_3/dv^D} - 1 \quad (3.5.3)$$

where y_n is the position of the n^{th} bunch. It should be noted that the position axes of the following plots have been calibrated using the BPM calibration coefficient of section 3.4.2 in order to show clearly the range of deflections. However, when determining the feedback gains, the various gradients are kept in units of uncalibrated position per DAC count in order to maintain compatibility with the FONT5 digital feedback board.

Tables 3.9 and 3.10 summarise the calibration constants measured for kickers K1 and K2 respectively.




Figure 3.24: Calibration curves for K1. The positions of bunches 2 and 3 vs. kick strength are shown in the BPMs P2 (a) and P3 (b). Linear fits are shown by solid lines.




Figure 3.25: Calibration curves for K2. The positions of bunches 2 and 3 vs. kick strength are shown in the BPMs P2 (a) and P3 (b). Linear fits are shown by solid lines.

		χ^2/df	Constant
P2	Bunch 2	2.99	75.3 ± 0.5
	Bunch 3	3.73	68.3 ± 0.6
P3	Bunch 2	23.9	29.4 ± 1.7
	Bunch 3	23.5	26.5 ± 1.7

Table 3.9: Calibration constants for the kicker K1 in the BPMs P2 and P3, extracted from linear fits to the calibration curves. The constants are measured in nanometres per DAC count output onto the kicker.

		χ^2/df	Constant
P2	Bunch 2	6.49	21.5 ± 0.8
	Bunch 3	6.66	17.3 ± 0.8
P3	Bunch 2	13.2	86.3 ± 1.5
	Bunch 3	20.5	69.0 ± 1.7

Table 3.10: Calibration constants for the kicker K2 in the BPMs P2 and P3, extracted from linear fits to the calibration curves. The constants are measured in nanometres per DAC count output onto the kicker.

Chapter 4

FPGA firmware implementation of a flexible feedback system

Firmware for the Virtex-5 FPGA on the FONT5 digital feedback board has been designed and implemented using the Verilog [78] hardware design language. It is a modular design, and its operation is outlined in this chapter. Verilog also provides an efficient means of archiving the design, containing as it does a full behavioural description of the firmware. The FONT5 project files are maintained by the FONT group at the John Adams Institute for Accelerator Science [86].

4.1 Control register interface

During operation, the timing and behaviour of the FONT5 feedback system is governed by a set of parameters. These parameters are stored in a number of 7 bit registers, which for historic reasons are split into two arrays. The first array is clocked at 357 MHz and the second at 40 MHz. All the registers may have new values stored in them by transmitting updates from software to the registers via the RS232 Universal Asynchronous Receiver/Transmitter (UART).

Some parameters are larger than 7 bits so are split over two control registers. Other parameters are boolean, where any value other than zero indicates `true`. The feedback logic can access parameters stored in the control registers during operation.

Table 4.1 lists the various parameters stored in the array clocked at 40 MHz along with the addresses of the registers used. The parameters stored in the 357 MHz array are listed in tables 4.2 and 4.3.

Operating parameter	Address	Size	Function
p1_scan_align_ch_sel	0	2 bits	Unused
p2_scan_align_ch_sel	1	2 bits	Unused
p3_scan_align_ch_sel	2	2 bits	Unused
p1_offset_delay	3	7 bits	Starting point for P1 ADC clock phase feedback
p2_offset_delay	4	7 bits	Starting point for P2 ADC clock phase feedback
p3_offset_delay	5	7 bits	Starting point for P3 ADC clock phase feedback
p1_scan_delay	6	6 bits	P1 ADC sampling delay (75 ps taps)
p2_scan_delay	7	6 bits	P2 ADC sampling delay (75 ps taps)
p3_scan_delay	8	6 bits	P3 ADC sampling delay (75 ps taps)
master357_delay	9	7 bits	Input delay on the 357 MHz clock
k1_b2_offset	10 & 11	13 bits	Constant added to bunch 2's kick at K1
k1_b3_offset	12 & 13	13 bits	Constant added to bunch 3's kick at K1
ring_clk_thresh	14	7 bits	Specify ring clock logic level (e.g. TTL)
trig_thresh	15	7 bits	Specify trigger logic level (e.g. TTL)
k1_fir_k1	16	7 bits	Coefficient of K1's first FIR filter tap
k2_b2_offset	17 & 18	13 bits	Constant added to bunch 2's kick at K2
k2_b3_offset	19 & 20	13 bits	Constant added to bunch 3's kick at K2
k2_fir_k1	21	7 bits	Coefficient of K2's first FIR filter tap
k1_bunch_strb_sel	22	Boolean	P2 or P3 data used in K1's low latency path

Table 4.1: Structure of the control register array clocked at 40 MHz.

Operating parameter	Address	Size	Function
trig_delay	0	7 bits	Unused
trig_out_delay	1	7 bits	Output delay of TTL amplifier trigger to K1
trig_out_en	2	Boolean	Enable both TTL amplifier triggers
p1_bunch1pos	3 & 4	8 bits	Sample # of bunch 1 peak in the P1 ADC data
p1_bunch2pos	5 & 6	8 bits	Sample # of bunch 2 peak in the P1 ADC data
p1_bunch3pos	7 & 8	8 bits	Sample # of bunch 3 peak in the P1 ADC data
p2_bunch1pos	9 & 10	8 bits	Sample # of bunch 1 peak in the P2 ADC data
p2_bunch2pos	11 & 12	8 bits	Sample # of bunch 2 peak in the P2 ADC data
p2_bunch3pos	13 & 14	8 bits	Sample # of bunch 3 peak in the P2 ADC data
p3_bunch1pos	15 & 16	8 bits	Sample # of bunch 1 peak in the P3 ADC data
p3_bunch2pos	17 & 18	8 bits	Sample # of bunch 2 peak in the P3 ADC data
p3_bunch3pos	19 & 20	8 bits	Sample # of bunch 3 peak in the P3 ADC data
feedback_flags	21	6 bits	Bitwise boolean flags to configure feedback mode
k1_const_dac_out	22 & 23	13 bits	Output on K1's DAC in constant output mode
k2_const_dac_out	24 & 25	13 bits	Output on K2's DAC in constant output mode
clk2_16_edge_sel	26	Boolean	Use rising or falling edge of the ring clock
sample_hold_off	27	7 bits	Delay of ADC sampling from the ring clock edge
big-trig_delay	28 & 29	10 bits	# ring clock cycles after trigger until ADC
trig_out_delay	30	7 bits	Output delay of TTL amplifier trigger to K2

Table 4.2: Structure of the control register array clocked at 357 MHz.

Bit #	Function
0	Enable K1 DAC output
1	Enable K2 DAC output
2	Enable K1 delay loop
3	Enable K2 delay loop
4	Overwrite K1 output with constant value
5	Overwrite K2 output with constant value

Table 4.3: Bit assignments in the 6 bit feedback flag control register.

4.2 ATF timing system

A number of timing signals from the ATF control system [97] are used in the FONT5 feedback system to synchronise the electronics with the beam. Firstly a 714 MHz signal, phase-locked to the damping ring cavity RF of the same frequency, is passed to the analogue electronics for use as a local oscillator (section 3.3.2). A frequency divider is used to generate a 357 MHz clock from the 714 MHz, ensuring the clock to the FPGA and ADCs is well locked to the beam.

Also provided by the ATF timing system is a 2.16 MHz NIM logic signal, referred to as the ring clock. Each period of the signal corresponds to an orbit of the damping ring by the beam. This ring clock is stable with respect to the beam at the few-ps level, so is suitable for use in defining the start of the ADC sampling period.

In order to select the ring clock cycle in which the beam is extracted from the damping ring, a further NIM trigger is obtained. It is routed from the damping ring extraction kicker and fires once per extraction. This trigger is stable only at the 10 ns level, so is not as well suited to specifying when ADC sampling should begin. It is however stable enough to select the ring clock cycle containing an extracted bunch train consistently.

4.2.1 357 MHz clock input electronics

Dedicated logic on the FONT5 board FPGA deals with the input of the 357 MHz system clock. A coaxial input is transformer-coupled to the FONT5 board, from where it is brought onto the FPGA via a differential input buffer. Virtex-5 I/O blocks' SelectIO™ [98] resources contain configurable input delay elements, or IDELAYS, with 75 ps resolution, which are used to adjust the system clock phase. Figure 4.1 shows a schematic of the input logic. The delay incrementor element is coded in Verilog and takes input from the control registers, driving the IDELAY control signals in response to user specified RS-232 commands.

After phase adjustment the 357 MHz is routed to a Virtex-5 Clock Management Tile [98].

A Digital Clock Manager element is used to synthesise a 200 MHz reference clock, required for the IDELAY control logic. The tile's phase-locked loop is configured as a jitter filter, the output of which is globally routed throughout the FPGA to provide the 357 MHz clock to the logic.

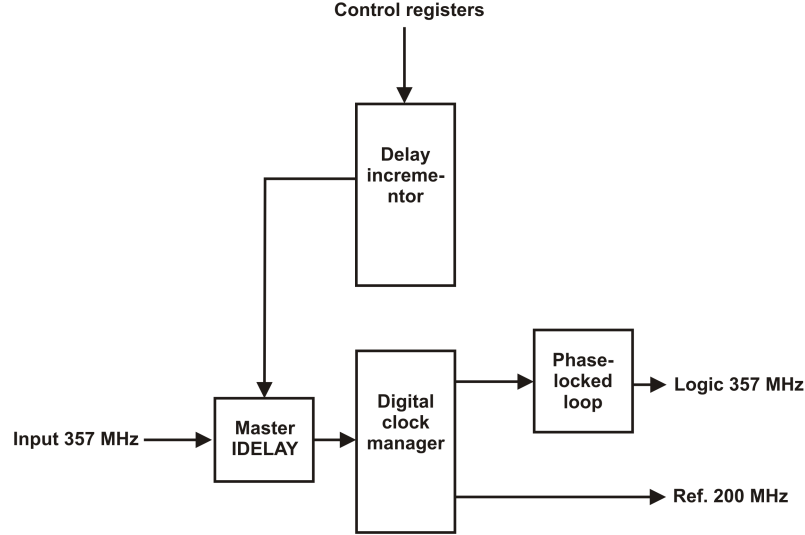


Figure 4.1: Schematic of the FPGA-based input electronics for the 357 MHz clock.

4.2.2 Digital input thresholds

In order to allow digital inputs of various signal levels to be recognised by the FONT5 board, the ring clock and trigger inputs are made via RC circuits on the PCB as shown in figure 4.2. Here, the MCX connector by which the input signal is brought onto the board is labelled **MCX STR JACK**. The wires labelled **DIGINPUT1A** and **DIGINPUT1B** are connected to tristate FPGA outputs, whose states are dependent on control register values. Each of the nine possible combinations of 0, 1 and Z gives rise to different threshold voltages for logic transitions in the input signal. At ATF2, both the ring clock and trigger are NIM signals, for which the control registers `ring_clk_thresh` and `trig_thresh` should be loaded with the value 5. This corresponds to 0 and Z on **DIGINPUT1A** and **DIGINPUT1B** respectively.

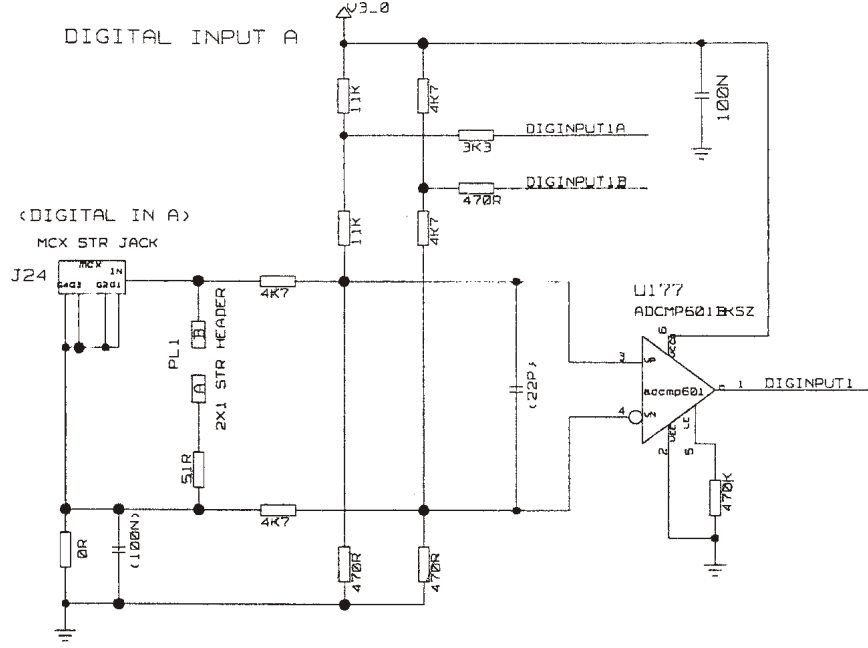


Figure 4.2: RC circuit used to vary detection thresholds for input logic signals [99].

4.2.3 Synchronisation with the ATF2 beam

Synchronising the feedback system with the ATF2 beam makes use of the ring clock signal. During three-bunch operation, each period of the ring clock corresponds to an orbit of the damping ring, with the three bunches being extracted during the final orbit. This brings the three bunches into the extraction line synchronous to a specific ring clock cycle, counted relative to the extraction kicker trigger.

Both the ring clock and trigger are sampled by registers in the FPGA fabric at 357 MHz. Edge detection logic waits for the rising edge of the trigger, then counts subsequent edges of the ring clock. Whether to use the rising or falling ring clock edge may be specified in the control registers. Also specified in the control registers are the two delays shown in figure 4.3, the amplifier trigger delay t_A and the main trigger delay t_D . They are given in terms of ring clock cycles, such that $t_D = \text{big_trig_delay} \times 463 \text{ ns}$ etc.




Figure 4.3: Timing diagram of the FONT5 triggering routine.

There are four stages to the logic's response to a trigger, as shown in figure 4.3. They are:

1. After t_A μs , a 1.4 μs trigger pulse is sent to both of the TMD kicker amplifiers driving K1 and K2, preparing them to accept the feedback signal.
2. For most of the time the ADCs are in a low power mode consuming ~ 50 mW per channel. The ADC power down signal is released t_D μs after the trigger is registered.
3. The ADCs are ready to return data 4.6 μs later, at which point the ADC store strobe goes high for a single cycle of the ring clock. This strobe instructs the FPGA logic to log those 463 ns of ADC samples, and to use those data to form a feedback signal. During installation, the delay t_D must be found such that the data returned correspond to the ring clock cycle containing the extracted beam.
4. Finally, a cycle after the ADC data have been logged, the ADC align enable signal is taken high for around 250 μs . During this time the ADCs remain powered and an ADC clock phase feedback routine is allowed to execute, as described in section 4.3.2. The routine ensures that the ADC data are well sampled despite any possible timing drifts due to temperature variation.

4.3 BPM signal digitisation

Once the FONT5 feedback board is synchronised to the correct ring clock cycle, further steps must be taken to ensure the analogue BPM signals are being well sampled. In particular, it

is important that the signals are sampled very accurately on peak to maximise the system resolution. The remaining timing parameters are stored in the control registers, and may vary depending on the specific machine configuration at ATF.

All timing settings related to the ADC sampling may be set independently for each ADC group. That is, the three ADCs assigned to digitise signals from a given BPM have their own timing settings.

4.3.1 ADC synchronisation

In the ATF damping ring the accelerating cavity operates at 714 MHz, with RF buckets spaced by $1/357 \text{ MHz} = 2.8 \text{ ns}$. The fill pattern of the three bunches in the ring determines the bunch spacing observed in the extraction line, with a symmetric fill resulting in three bunches all spaced by 154 ns. As described previously, a ring clock period's worth of continuous ADC samples are captured for each FONT5 channel on each pulse, resulting in 463 ns of data. This is sufficient to capture the BPM signal for the entire three-bunch train.

It is useful to have the ability to delay the start of the sampling period in 2.8 ns steps, allowing a particular 463 ns range to be sampled with the three individual bunch signals well placed. Figure 4.4 is an expanded view of the ADC store strobe from figure 4.3. The control register `sample_hold_off` is used to specify the delay t_S in units of 357 MHz cycles, relative to the ring clock edge.




Figure 4.4: Timing diagram showing the ADC sampling synchronisation.

All of the ADC samples during the period when the store strobe is high are logged and returned over RS232 as described in section 4.4. In generating the feedback signal, only

those three samples that fall on the bunch signal peaks are of interest. The position of the peak samples in the data stream are stored in control registers, and used to generate a bunch strobe signal consisting of three 2.8 ns pulses. Only the three samples coincident to the strobes are used to update the feedback signal.

Note that each group of three ADCs that serve a given BPM have independent clocks and bunch position registers. The exact sample number in which a bunch position lies depends on the beamline geometry and the lengths of cables carrying signals from the analogue processors.

4.3.2 ADC data phase feedback for sampling stability

The arrival phase of the ADC data with respect to the logic 357 MHz signal was not well known after construction of the FONT5 board, depending as it does on signal propagation times between the FPGA and the ADCs. Since the ADC data might only be held for less than 1 ns, the logic 357 MHz must sample quite accurately midway between ADC data transitions. By using the Virtex-5's configurable input delay, or IDELAY, elements, the arrival time of the ADC data may be shifted with respect to the logic 357 MHz in 75 ps taps. The configuration is similar to the IDELAY described in section 4.2.1, with each ADC data bit going through its own input delay element. Figure 4.5 shows well and poorly aligned ADC sampling.




Figure 4.5: Examples of badly (a) and well (b) aligned ADC data with respect to the logic 357 MHz clock.

Given the tight tolerances on the ADC clock alignment, phase variation due to, for example, temperature change is a concern. The ADCs provide a data ready signal consisting of a 357/2 MHz square wave synchronous to the data. This signal may be used in an ADC sampling phase diagnostic, ultimately allowing a phase feedback to compensate for timing drifts on a pulse-to-pulse basis.

The I/O blocks of the Virtex-5 contain a double data rate register, or IDDR. By clocking this register using the logic 357 MHz, three samples of the data ready signal may be obtained with just 1.4 ns spacing. Three such samples are shown in yellow in figure 4.5. In the FONT5

system, this set of three samples is taken 127 times. Each set of samples is passed onto the 40 MHz clock domain where they are added to three running totals. If the data ready signal is perfectly aligned with the logic 357 MHz then the central sample should fall on an unstable edge, meaning its total should be close to 64. The totals of the two adjacent samples would be 0 and 127.

If the total of the central samples is equal to the total of the first samples, the ADC data is arriving late. If equal to the total of the last samples, the ADC data is arriving too early. After every pulse is logged a logic module written in Verilog, the alignment monitor in figure 4.6, is used to perform a number of iterations in which the three sample totals are obtained. After each iteration the ADC data IDELAYs and the ADC data ready IDELAY are increased or decreased by a tap, or left unchanged if the central sample total is close to 64 ($20 \leq \text{total} \leq 107$). The largest phase adjustment that could be required is 1.4 ns or 19×75 ps taps, and with each iteration taking ~ 14 μs this phase feedback is allowed to run for 270 μs after every pulse acquisition, as shown in figure 4.3.

In addition, an extra delay may be added to the data IDELAY and not the data ready IDELAY to allow compensation for any phase offset between the data and data ready signals due to delays in the FPGA. This additional delay is stored in the `offset_delay` control registers.

4.3.3 Fine adjustment of the ADC sample time

As shown in section 3.3.2, the output of the analogue BPM processors is well modelled by a Gaussian curve, such that the best position resolution is obtained when the pulse is sampled exactly on peak. In this case the digital sample is least sensitive to sample time jitter. The Virtex-5 also provides output delay elements, or ODELAYs, again with 75 ps resolution. By delaying the ADC clock with an ODELAY element, the ADC sample time may be scanned with respect to the beam signal. If the delay added to the ADC clock is subtracted from the data IDELAYs, the optimised phase between the ADC data and the logic 357 MHz, as discussed in 4.3.2, is maintained.

As usual, this delay on the ADC sampling time is stored in a control register. The three `scan_delay` registers hold a delay for each ADC group. Figure 4.6 shows the logic block which combines all of the various factors contributing to the ADC clock output delay and the data input delay: the offset delay, the phase feedback and the sample time adjustment.

The total delays are calculated by:

$$\text{ADC clock output delay} = \text{scan_delay}$$

$$\text{ADC data ready input delay} = 32 + \text{feedback adjustment} - \text{scan_delay} \quad (4.3.1)$$

$$\text{ADC data input delay} = 32 + \text{feedback adjustment} + \text{offset_delay} - \text{scan_delay}$$

The `scan_delay` control registers and the phase feedback adjustment are both two's complement values that may be positive or negative. A default delay on the ADC data and data ready signals of 32 taps is specified, which is midway through the range of the IDELAY elements.

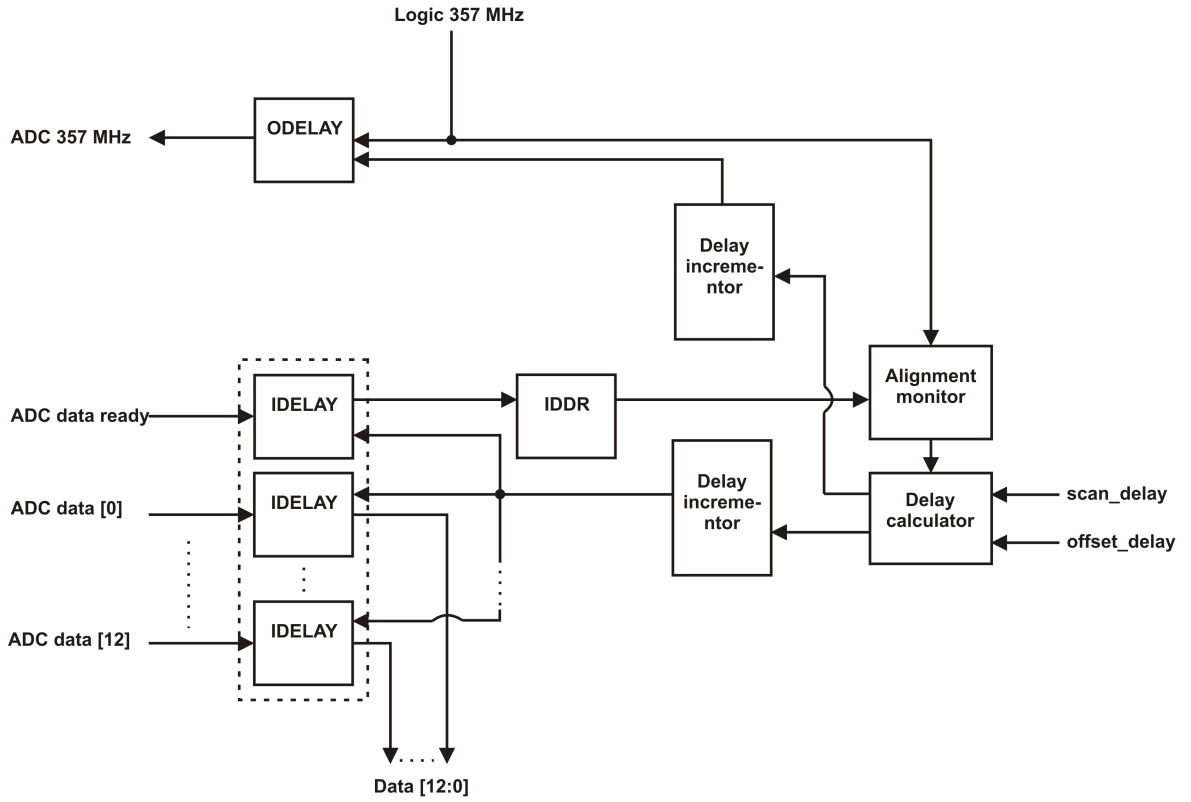


Figure 4.6: Schematic showing the logic of an `ADC_block` module that has been implemented in Verilog. Each block controls the timing of a group of three ADCs and is implemented on the Virtex-5 FPGA.

4.3.4 ADC inputs and pedestal subtraction with trim DACs

Sum and difference signals produced by the analogue front end processors have a non-zero baseline voltage, or pedestal, which varies from PCB to PCB. Should this baseline voltage not be subtracted before the charge normalised position is calculated, a charge-dependent error will be introduced. Subtraction of the baseline in software is trivial to perform by

determining the mean baseline voltage from samples which fall between the bunch signals. To do so in the FPGA firmware without increasing the latency of the digital signal processing logic is not so simple, however.

With this in mind, the ADC data inputs have trim DAC chips wired in parallel on the PCB. Each chip is capable of providing independent voltage offsets to three ADCs. Three chips are therefore used, one for each ADC group. An example of the wiring on an ADC input is shown in figure 4.7. The BPM signal to be digitised is brought in via the shown MCX connector, which is transformer-coupled to the trimming electronics. Wires labelled ADC1P and ADC1N carry a differential signal to the inputs on an ADC chip, while a voltage offset from a trim DAC is applied via the wires labelled TRIMV1P and TRIMV1N. Adjusting the trim DAC output allows the signal pedestal to be zeroed in an analogue manner, costing no additional latency.

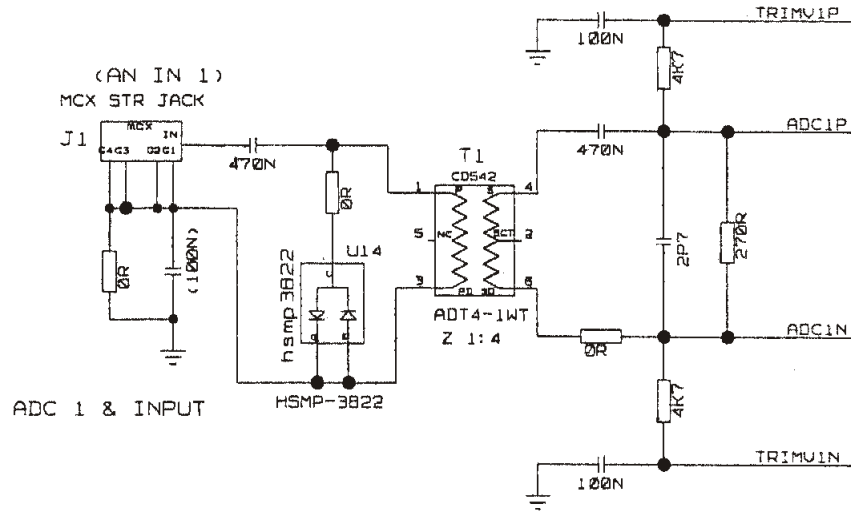


Figure 4.7: Input electronics and trimming circuit for the analogue BPM signals [99].

Each DAC output is a 12 bit value which is stored on board the FPGA in a Random Access Memory (RAM) structure. This memory must be loaded by the controlling software each time the FONT5 board is reset, as described in section 4.6.1. Figure 4.8 demonstrates the linear responses of the nine ADC channels' trim DACs. These data were obtained by leaving the FONT5 board data inputs disconnected and scanning the value output by the trim DACs. Each datapoint shows the resulting ADC baseline averaged over around 20 triggers.




Figure 4.8: Average recorded baseline for all ADC channels v.s. trim DAC output.

4.4 Pulse-by-pulse acquisition of BPM data

Virtex-5 FPGAs contain many 36 Kb block RAM (BRAM [90]) elements. Each of the nine FONT5 ADC channels makes use of one of these BRAMs for buffering ADC data prior to transmission over the RS232 UART to DAQ software running in the ATF control room. On each machine pulse 164 samples at 2.8 ns spacing are logged for each ADC channel. The data are transmitted at 115.2 kbaud, meaning that they can be processed and displayed in the DAQ software Graphical User Interface (GUI) before the next pulse arrives. Therefore the GUI is able to present the BPM data on a pulse-by-pulse basis.

4.4.1 Dual-port block RAM

Each ADC channel's BRAM is implemented as part of a Verilog module, `DAQ_RAM`, which contains all the associated transmission logic necessary to interface with the RS232 UART. The BRAMs are dual port, and data may be written to or read from either port A or port B. Dual port RAM is designed to allow simultaneous reads and writes to different memory locations using the two ports, which have independent clocks, addresses, enables, data inputs, data outputs and data widths. Using dual port BRAMs provides a safe means of bridging the 357 MHz and 40 MHz clock domains. Coregen [100] software from Xilinx is used to map dual port RAMs with port A of size 10×14 bits and port B of size 11×7 bits to the FPGA BRAMs.

It is the `store_strobe` signal that determines which 164 ADC samples are stored. Therefore it is used as the enable signal for port A of the BRAM. Port A is clocked by the logic

357 MHz and has a 14 bit wide data input. The 13 bit ADC samples are padded with a trailing zero. In this way 164 consecutive samples are written to the BRAM on each pulse. Each of the nine ADC channels has its own BRAM with data being written to them all in parallel.

Data are read from port B of the BRAMs for transmission over RS232 after the `store_strobe` signal goes low. Port B is clocked using the slower 40 MHz clock since the transmission logic is not time critical. The data output of port B is 7 bits wide, such that the zeroth address corresponds to the least significant 7 bits of the zeroth ADC sample, the first address to the most significant 7 bits of the zeroth ADC sample, the second address to the least significant 7 bits of the first ADC sample and so on (see figure 4.9). According to the communication protocol further described in section 4.4.2, these 7 bit data words have a leading 1 added to signify they are data rather than commands. The nine BRAM port B outputs are multiplexed before being passed to the UART. Dedicated logic enables each DAQ_RAM's transmission routine in turn, stepping through each of the nine channels and sequentially sending each stored byte out of the UART at 115.2 kbaud. Figure 4.10 is a schematic showing a single data channel and associated control and transmission logic.




Figure 4.9: Address structure of the Virtex-5 dual port block RAM used for data acquisition.

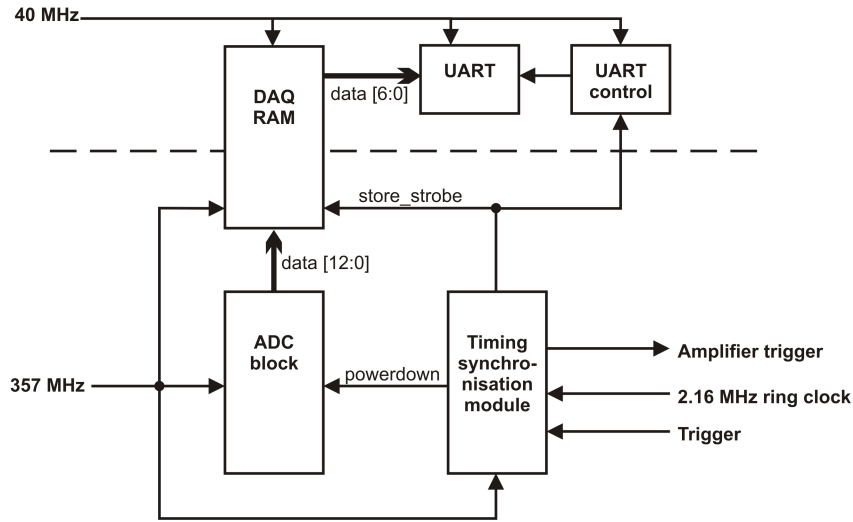


Figure 4.10: Schematic of a single data channel's DAQ logic as implemented on the Virtex-5 FPGA.

4.4.2 Data stream format

In addition to the $9 \times 164 \times 2 = 2,952$ bytes of ADC data returned by the FONT5 board on each pulse, additional diagnostic information is returned in the form of readback data. All current control register values are returned, along with the values of several internal firmware signals. The values output onto the two FONT5 DACs are returned, as is a simple timestamp. The format of the data stream is shown in figure 4.11. With the RS232 configured for transmission at 115.2 kbaud, the total of up to 3,051 bytes which constitute the data stream are transmitted in 212 ms.




Figure 4.11: Format of the RS232 data stream returned by the FONT5 board on each machine pulse. The blue line represents a single byte timestamp and the red lines represent single framing bytes.

The various sections of the data stream are identified by sending various framing bytes to the software interface within the stream. For example, the hexadecimal value 0x14 is sent before the 328 bytes which encode the digitised P1 sum signal. Table 4.4 lists the framing bytes used. Note that all of the framing bytes have values less than 0x80, i.e. their most significant bit is unset. This is common to all command bytes sent to or from the FONT5

board, and is used to distinguish commands from data bytes. As described in section 4.4.1, all 7 bit data words are transmitted over RS232 as 8 bit bytes with a leading 1.

Framing byte value	Section of data stream indicated
0x1F	Time stamp
0x10	P1 horizontal difference
0x12	P1 vertical difference
0x14	P1 sum
0x15	P2 horizontal difference
0x16	P2 vertical difference
0x17	P2 sum
0x18	P3 horizontal difference
0x19	P3 vertical difference
0x1A	P3 sum
0x1D	K1 DAC readback
0x1E	K2 DAC readback
0x1B	357 MHz control register readback
0x1C	40 MHz control register readback
0x0F	Internal signal monitor readback

Table 4.4: Values assigned as data stream framing bytes in hexadecimal.

In order to generate a simple timestamp, a 7 bit counter is implemented on the 40 MHz clock domain. This counter is incremented every 40 ms and allowed to wrap around repeatedly. On each machine pulse the current value of the counter is transmitted as part of the data stream, allowing the time difference between successive pulses to be monitored. An additional absolute timestamp is also logged by the DAQ software upon receipt of the data, but as the exact stamp obtained by the computer depends upon factors such as network and processing performance, this software stamp is less accurate.

All values which are put out onto one of the FONT5 DACs are returned in the data stream. The DAC readbacks again make use of `DAQ_RAM` modules as described in section 4.4.1. By duplicating a DAC's data and clock signals and using them to drive a `DAQ_RAM`, each time the 13 bit DAC output is updated the value is written into the RAM. The number of values logged on any given pulse depends on the configuration of the feedback system; if the DAC outputs never change then no values are written or returned in the data stream.

For the 357 MHz and 40 MHz control register readbacks, the values currently held in every single register listed in tables 4.2 and 4.1 are returned as part of the data stream on each pulse, contributing 31 and 23 bytes respectively. The final component of the data

stream, the monitor readbacks, are formed by registering a number of firmware signals in a 13×7 bit array implemented in the fabric of the FPGA before transmitting their current values. Table 4.5 lists the signals so monitored.

Address	Size (bits)	Description
0	7	Status flags 0: P1 ADC I/O delay saturated 1: P2 ADC I/O delay saturated 2: P3 ADC I/O delay saturated 3: 2.16 MHz ring clock present 4: IDELAYCTRL element ready 5: Digital clock manager locked 6: Phase locked loop locked
1	7	P1 ADC DRDY sample 1 total
2	7	P1 ADC DRDY sample 2 total
3	7	P1 ADC DRDY sample 3 total
4	6	P1 ADC input delay value
5	7	P2 ADC DRDY sample 1 total
6	7	P2 ADC DRDY sample 2 total
7	7	P2 ADC DRDY sample 3 total
8	6	P2 ADC input delay value
9	7	P3 ADC DRDY sample 1 total
10	7	P3 ADC DRDY sample 2 total
11	7	P3 ADC DRDY sample 3 total
12	6	P3 ADC input delay value

Table 4.5: Internal firmware signals monitored and returned as part of the FONT5 data stream on every pulse.

All but one of the monitor readbacks consist of signals related to the ADC data phase feedback described in section 4.3.2. The other byte returned contains seven single bit flags containing sampled firmware signals reflecting the FONT5 operation status. Bits 0-2 are high if, for whatever reason, any of the ADC data IDELAYS or ADC clock ODELAYS reach the end of their range. This provides a diagnostic for both the ADC data phase feedback and fine ADC sample time adjustment. The third bit is set whenever the 2.16 MHz ring clock is present. When bit 4 is high, the IDELAYCTRL element of the FPGA (required for proper operation of the Virtex-5 IDELAYS and ODELAYS) is correctly set up. The remaining two bits are set when the digital clock manager and the phase-locked loop in the 357 MHz input electronics (section 4.2.1) are locked and operating stably.

4.5 Feedback signal generation

4.5.1 Virtex-5 DSP48E slices

Dedicated digital signal processing electronics, known as DSP48E slices, are present on board the Virtex-5 FPGA. The XC5VLX50T devices contain 48 such slices, which are specifically designed to perform fast and efficient mathematical operations. Logic which is mapped onto DSP48E slices rather than being implemented in the FPGA fabric will run at higher clock frequency and lower latency. The Xilinx documentation [96] describes the wide array of applications to which the slices may be put.

Two main configurations of DSP48E slice are used in the FONT5 firmware:

- An adder with two 48 bit inputs and a two cycle latency.
- A multiplier with one 13 bit input, one 21 bit input, a 48 bit carry input and a three cycle latency.

4.5.2 Look-up tables for charge normalisation and gain application

As described in section 3.4.1, the first step in generating the FONT5 feedback signals is to convert the sum and difference ADC samples from the feedback BPM into a position measure which is independent of beam intensity. The position must then be multiplied by an appropriate gain constant. Because minimising system latency is a priority, all the signal processing logic is clocked at 357 MHz and the processing is performed in as few clock cycles as is possible.

A look-up table (LUT) approach is used to determine the reciprocal of the sum sample, and to multiply the reciprocal by a gain constant. This technique is simple, using extra FPGA resources to perform the necessary calculations in fewer clock cycles. Virtex-5 BRAMs are used to construct a LUT of 21 bit entries, and the 13 bit sum sample provides the LUT address. Each entry of the LUT is loaded with the required gain constant divided by the address as shown in section 4.6.1. Since all calculations are carried out using fixed point arithmetic, an additional factor of 2^{12} is applied to each entry of the LUT to ensure that rounding errors do not impair the feedback system performance. This factor is easily removed after the signal processing logic by discarding the least significant 12 bits with no latency cost.

As shown in figure 4.12, the LUT output is then multiplied by the difference sample using a DSP48E slice. The resulting 48 bit number is given by $4,096 \times G \times y$, where y is the measured bunch position. It requires five cycles of the 357 MHz clock, or 14 ns, to calculate the bunch position and apply the gain. In figure 4.12 an unconnected adder, or carry input,

is present in the DSP48E slice; this adder may be used to sum the calculated position with another value at no extra latency cost.

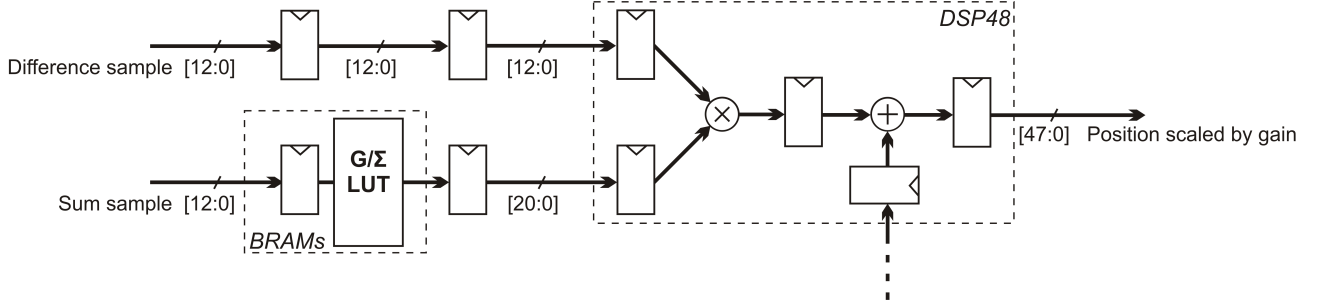


Figure 4.12: Schematic showing the firmware logic used to calculate bunch positions and apply feedback gain.

4.5.3 Digital signal processing for a single feedback loop

In this section the detailed digital signal processing logic required for a single feedback loop is described. Figure 4.13 is a detailed schematic of the processing logic, which includes the following mechanisms as described in section 3.5:

- Implement the feedback delay loop.
- Override the DAC output with a constant value.
- Add an additional constant kick on a bunch-to-bunch basis.
- Compensate for the kicker pulse droop with a FIR filter.
- Generate the clock required to drive the DAC.

Firstly, the bunch position in the feedback BPM must be calculated and scaled by the system gain, as described in section 4.5.2 using a LUT and DSP48E slice. During the cycles used to perform this calculation, the bunch strobe (which goes high for a cycle coincident with the ADC samples which contain the bunch signal peaks, section 4.3.1) is passed along a pipeline. Five cycles later the strobe is used to enable three registers, ensuring that the correct ADC samples are used in the calculation.

One of these registers implements the delay loop, which stores the current feedback signal (the gain-scaled position in the case of the pilot bunch). The value is registered via a multiplexer with its select line tied to a control register. This allows zeros to be stored instead, thus disabling the delay loop if desired. The value stored in the delay loop is presented, via two adders, to the carry input of the DSP48E slice used to calculate the position. In this way, when the delay loop is enabled, the gain-scaled position for the current bunch is summed with that of all previous bunches as required.

It is this sum of gain-scaled positions which forms the feedback signal to be output onto the DAC. The feedback signal is stored in an output register which again is enabled by the delayed bunch strobe. Firstly overflow control is performed, during which the 48 bit signal is reduced to a 13 bit DAC output saturated at -4,096 to +4,095 counts. During this overflow control the least significant 12 bits are discarded, removing the additional factor that was introduced in the LUT entries to improve the accuracy of the fixed point arithmetic.

A pair of multiplexers, controlled again by setting control registers, are used to pass the 13 bit feedback signal to the output register. One multiplexer allows zeros to be passed instead, disabling the DAC output. The other is used to overwrite the feedback signal with a specified constant value. The DAC clock is unchanged however, meaning that the constant DAC output shares the same time profile as the feedback signal.

Latency is the key concern when mapping the feedback logic to FPGA resources. When adding a constant to the feedback signal on a bunch-to-bunch basis, latency may be kept to a minimum by performing the extra arithmetic within the delay loop rather than in the feedback signal path. Another DSP48E slice, configured in this case as a two-input 48 bit adder, sums a constant value with the output of the delay loop register. These constant bunch kicks are specified in control registers and cycled through by external logic, being left shifted by 12 bits before being added into the delay loop.

In order to calculate the contribution from the FIR filter, another DSP48E slice is used to calculate the gain-scaled position in parallel with the main feedback path. The delay loop is not carried into this slice however, with the unadulterated gain-scaled position being stored in the FIR tap 1 register which is again enabled by the delayed bunch strobe. A final DSP48E slice is used to both scale the register's output by the tap weight (specified in a control register) and add the result into the delay loop. This single tap FIR filter would require extension for a train of more than three bunches.

Each time the DAC value is to be updated, a clock pulse must be generated. The bunch strobe is delayed by a further cycle in order to provide the pulse, and an additional register is used with an OR gate to extend the length of the pulse to 5.6 ns, ensuring the DAC clocking requirements are met [94]. The total latency of the feedback signal processing is 6 cycles of 357 MHz, or 16.8 ns. The ADC data are registered in the FPGA input blocks, and another pipeline stage between the input logic and the DSP modules combine to give a total firmware latency of 8 cycles of 357 MHz, or 22.4 ns.

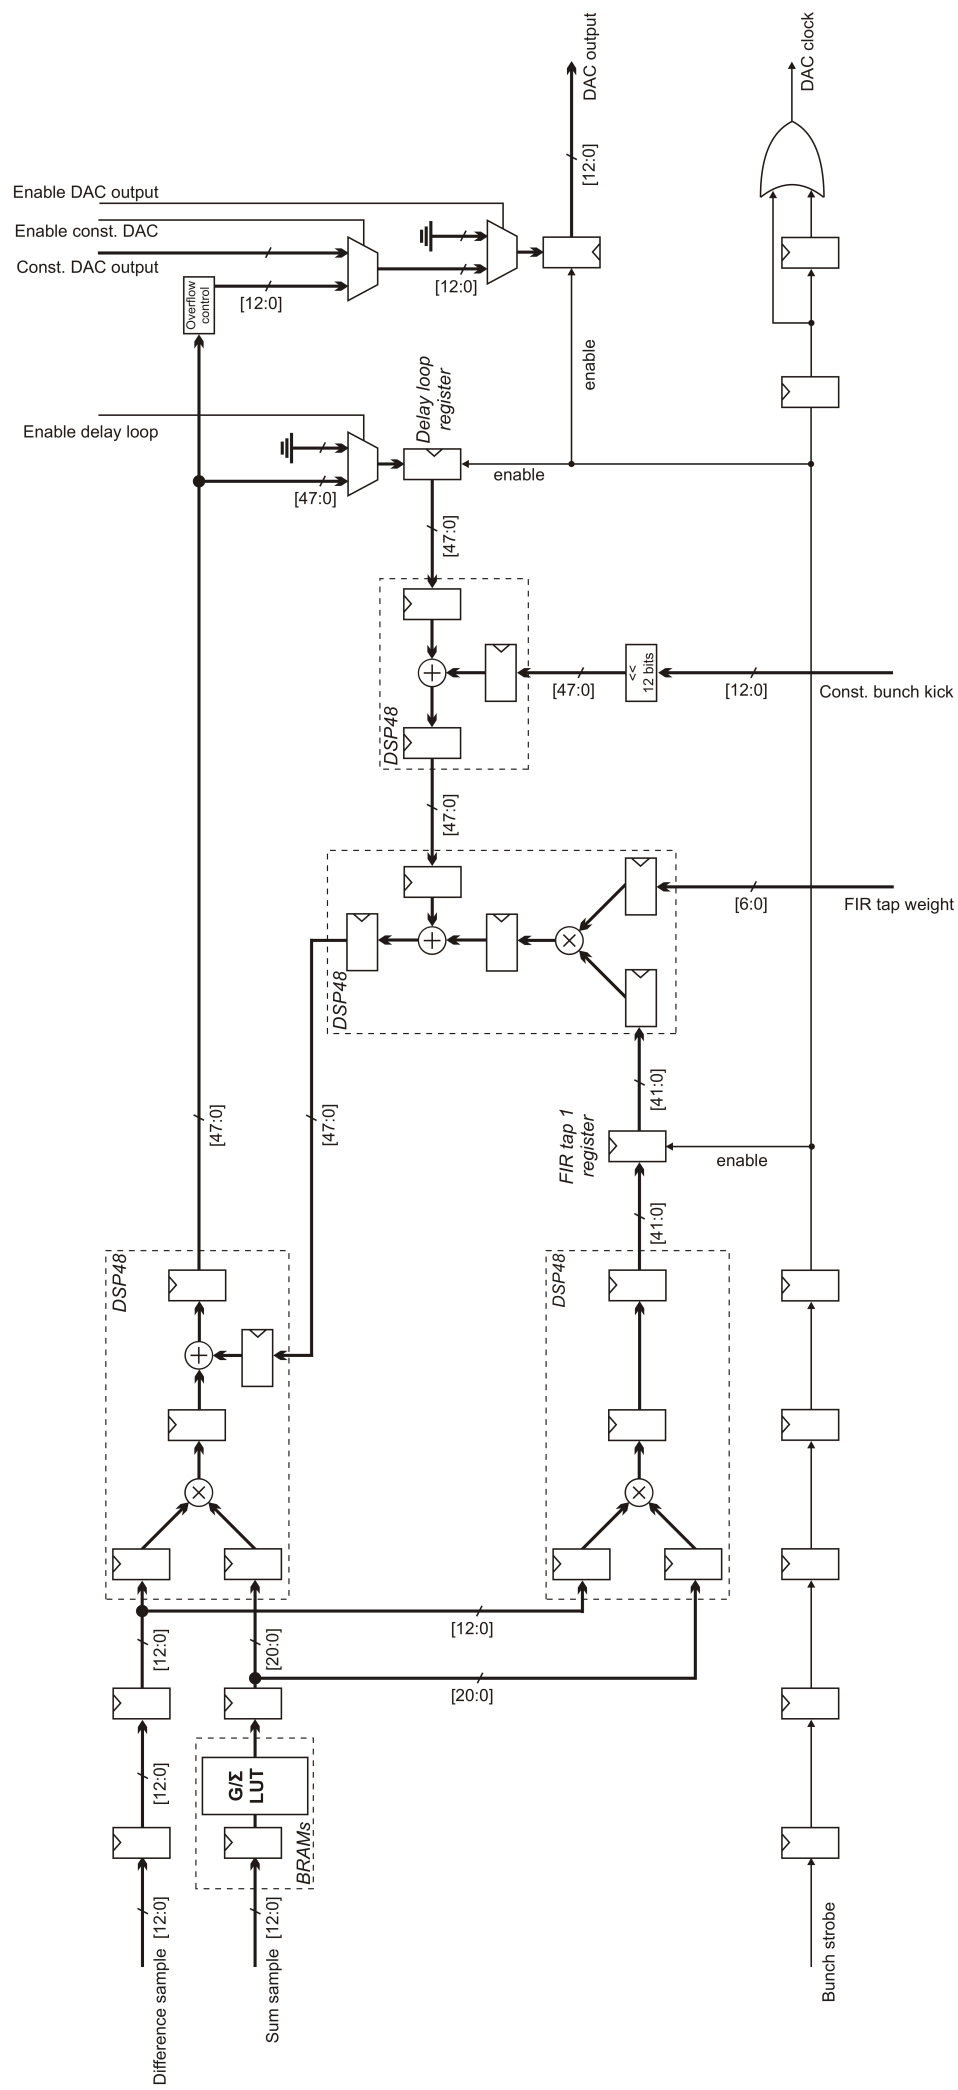


Figure 4.13: Schematic showing the firmware implementation of digital signal processing logic for single loop feedback.

4.5.4 Digital signal processing for coupled feedback loops

In the case of coupled feedback loops, a slightly more complex DSP module is required. It would be possible to process the BPM signals from both P2 and P3 in parallel using modules similar to that described in the previous section then combine their outputs linearly to form a feedback signal. Though this would generate an appropriate feedback signal for one of the DAC outputs in the coupled system, the additional combination stage would inevitably increase the system latency.

However, it is possible to exploit the geometry of the beamline and electronics installation to perform the combination in a more efficient manner. The FONT5 board itself is situated close to the BPM P2, with the other feedback BPM P3 located 2.9 m downstream. The time of flight of the beam added to the signal return path along coaxial cables means that the P3 signal will arrive at the board in excess of 20 ns later than P2's. By processing the P2 signal immediately, it may be added into the delay loop such that is ready to be combined with the P3 data when it arrives. No additional combination stage is required, and the total DSP latency remains 6 cycles.

Figure 4.14 shows a schematic of a complete module for driving one DAC in the coupled FONT5 system. Two such modules are implemented, one for the DAC driving each kicker. The module contains two data paths, each with its own gain value and LUT. P2's data are passed to the high latency path for preprocessing, where the P2 gain-scaled position is calculated and registered in response to the arrival of the P2 bunch strobe.

Generally, the data from P3 will be passed to the low latency path. The slight change in architecture for the low latency path in comparison to the single loop feedback described in section 4.5.3 ensures the value of the P2 position register is added into the delay loop as soon as possible. If the P2 data (and therefore bunch strobe) arrive at least 14 ns (5 cycles) before P3's, the preprocessed P2 gain-scaled position will propagate through the delay loop registers in time to be added to P3's gain-scaled position. The P2 gain-scaled position is also added into the FIR tap 1 register prior to it being weighted.

In order to allow flexible operation, the sum, difference and bunch strobe inputs to the low latency path driving K1's DAC are via multiplexers. With this addition, the logic can be configured to operate equivalently to the single loop DSP of section 4.5.3. Various modes of operation are possible as listed in table 4.6, where Gain 1 and Gain 2 represent the gains for K1's high and low latency paths respectively. Gain 3 and Gain 4 represent the same for K2's DSP paths.

Throughout each of the experiments described in this thesis two modules of this type, one for K1 and the other for K2, are implemented in the FONT5 FPGA firmware. Different modes of operation are entered using the settings in table 4.6, that is without changing the FPGA firmware.

Select	Gain 1	Gain 2	Gain 3	Gain 4	Mode of operation
1	G_1	G_2	G_3	G_4	Two coupled loops
1	G_1	0	0	G_4	Two independent loops
1	0	0	0	G_4	Single loop P3-K2
0	G_1	0	0	0	Single loop P2-K1

Table 4.6: Different FONT5 modes of operation.

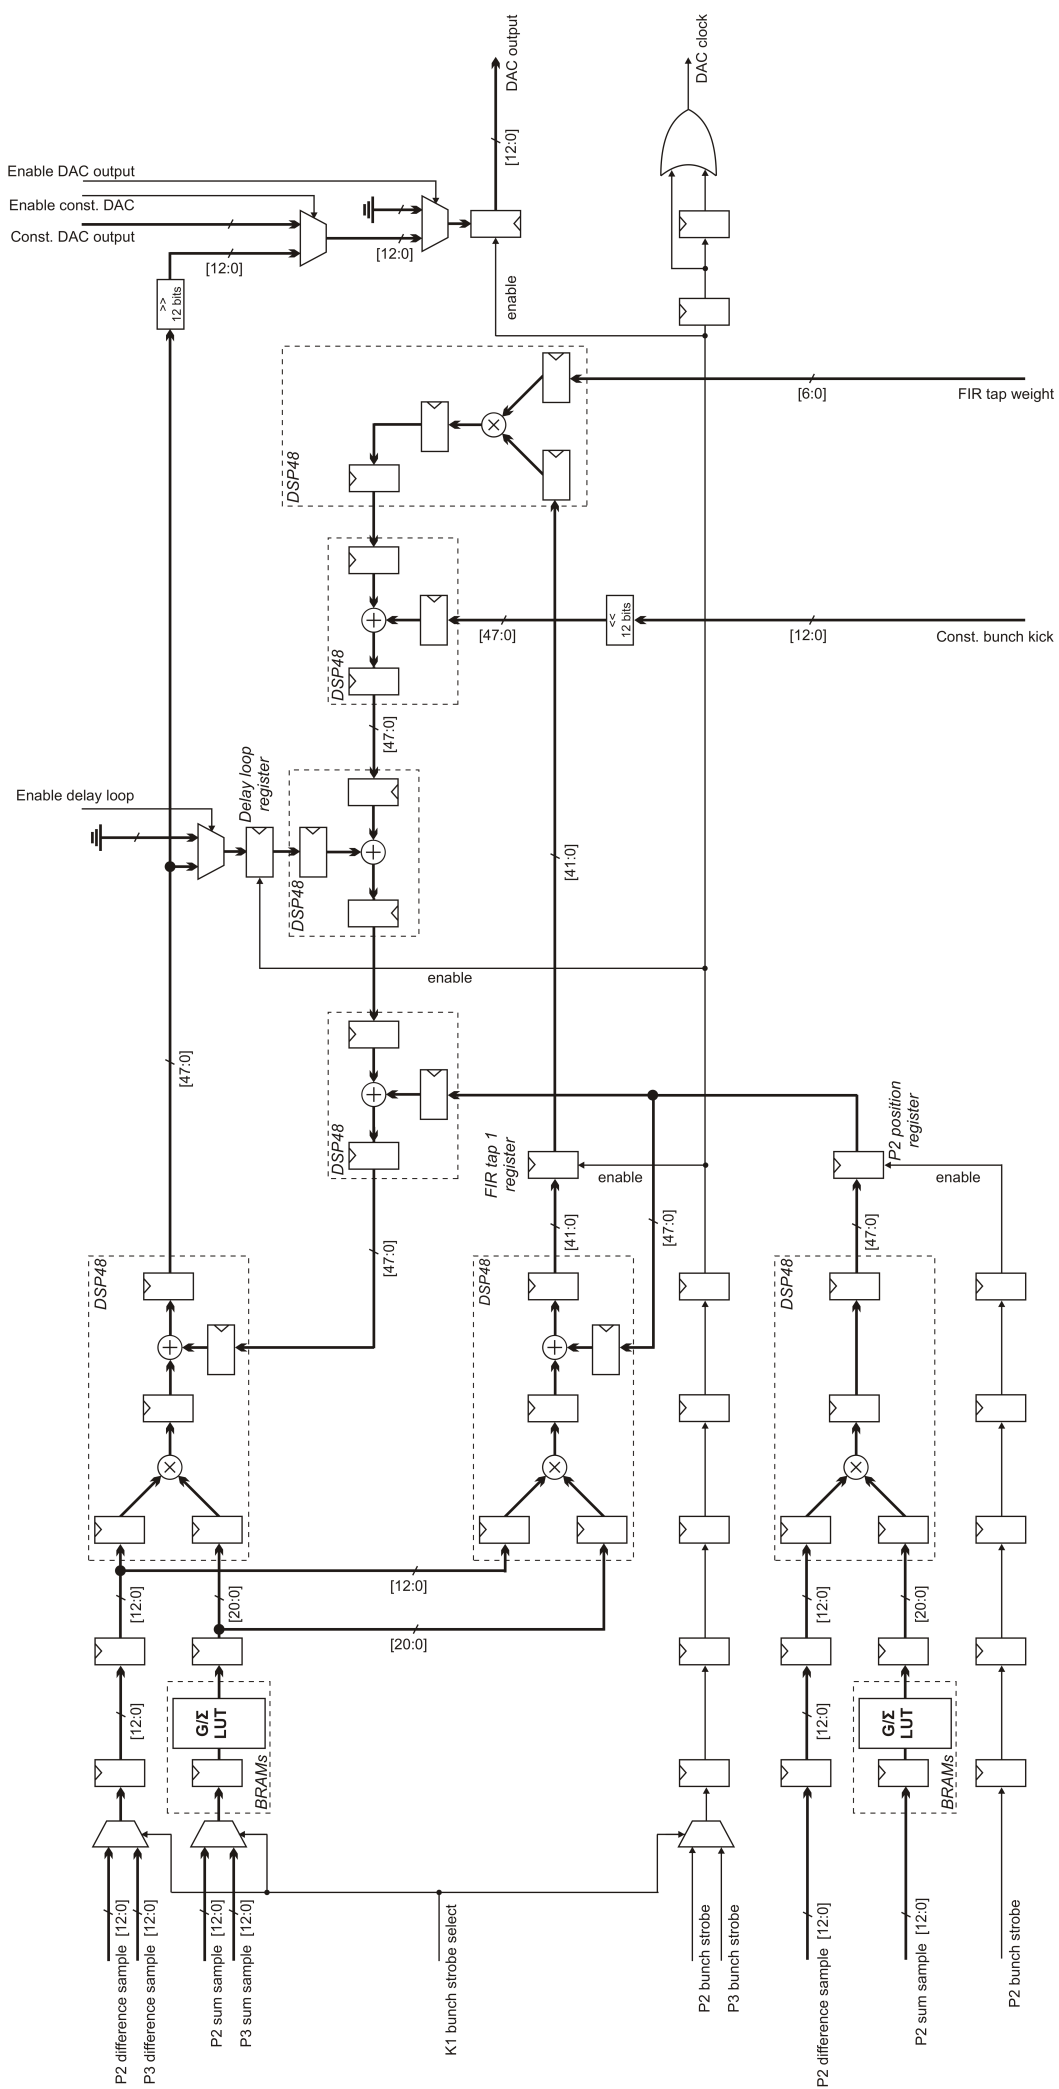


Figure 4.14: Schematic showing the firmware implementation of digital signal processing logic for coupled loop feedback.

4.6 RS232 interface for control of the FONT5 system

Logic running on the FPGA continuously monitors and decodes bytes received by the RS232 UART. As with the outgoing RS232 data stream, all bytes sent to the FONT5 board from control software are separated into commands and data. Data are sent as 7 bit words with the most significant bit of the RS232 byte set. Similarly, the command space is 7 bits and a command is recognised by an additional zero in the most significant position. All the FPGA based interface logic runs on the slower 40 MHz clock domain since it is not time critical, allowing place and route software to concentrate its efforts on achieving the best possible implementation of the fast feedback logic.

4.6.1 RS232 control protocol

There are three types of commands defined as part of the FONT5 board control protocol. Firstly there is the option to load a new value to a specified control register. The second type of command begins writing data to a LUT or a RAM in the firmware, either to specify a gain or to alter the trim DAC values. Finally, a subset of the commands, once received, drive internal firmware signals with particular values allowing for functions such as resets to be triggered.

Received RS232 byte	Command interpretation
011X XXXX	Address of a control register on the 40 MHz domain specified
010X XXXX	Address of a control register on the 357 MHz domain specified
001X XXXX	Address of a LUT to be loaded specified
000X XXXX	Miscellaneous commands

Table 4.7: Structure of RS232 command bytes as recognised by the FONT5 firmware.

Table 4.7 shows how the available command space is divided amongst the different types of commands. The addresses of the control registers are listed in tables 4.1 and 4.2, and the addresses of the LUTs below in table 4.8.

In order to set a particular control register's value, first the appropriate command must be sent over RS232 to make this the current register. Once active, any data byte sent to the FONT5 board will be written to this register. Take for example the `sample_hold_off` register on the 357 MHz domain with address $27 = 1\ 1011$. To make this register active, the byte $0101\ 1011 = 0x5B$ must be sent to the UART. Then to set the sample hold off to 18 cycles, the data byte $1001\ 0010 = 0x92$ must subsequently be sent.

A similar procedure is followed to load an FPGA RAM such as one of the gain LUTs in the DSP paths. Again a command must be sent to specify which RAM is to be loaded.

After a RAM is made current, it remains so until another command is received by the board. All data bytes subsequently received are written to the RAM, with an internal pointer starting at the zeroth address and incrementing after each byte arrives. Each gain LUT contains 8,192 entries of 21 bits each, such that each entry requires the transmission of three 7 bit data words. To load a LUT therefore requires the transmission of 24,576 bytes at 115.2 kbaud, a data stream lasting around 1.7 s. No error checking is currently implemented in the firmware, and it is software which ensures the correct number of bytes are transmitted.

Similarly, this procedure is also used to change the entries in the smaller RAM that stores the desired trim DAC outputs. This RAM has 18×7 bit entries, with each 12 bit trim DAC code stored across two consecutive entries. The lower 7 bits of the code are stored in the lower of the two addresses. Due to the FONT5 board layout, the order in which the trim DAC codes must be loaded is ADC channel 7, 4, 1, 8, 5, 2, 9, 6, then 3.

RAM address	Function
0	K1 high latency path gain
1	K1 low latency path gain
3	Trim DAC values
4	K2 high latency path gain
5	K2 low latency path gain

Table 4.8: Addresses and functions of the various FONT5 look-up tables.

A number of internal firmware signals are driven by the UART decoding logic, which takes those signals high in response to the receipt of certain command bytes. After such a byte arrives, the appropriate signal is taken high for one cycle of 40 MHz allowing further logic to be triggered as shown in table 4.9.

Command	Function
0x00	Full reset of all firmware registers and instantiated FPGA resources
0x01	Reset P1 ADC data I/O delays
0x02	Reset P2 ADC data I/O delays
0x03	Reset P3 ADC data I/O delays
0x04	Reset 357 MHz input delay
0x05	Refresh 357 MHz input delay
0x06	Refresh trim DAC outputs

Table 4.9: Command bytes for the FONT5 digital board in hexadecimal.

The command bytes of table 4.9 and the framing bytes of table 4.4, along with the values 0x11 and 0x13 which are reserved for future XON/XOFF RS232 flow control, make up 75% of the available protocol space allocated to miscellaneous commands. The remaining values are unused.

4.6.2 FONT5 graphical user interface

In parallel to the development of the firmware running on the FONT5 digital feedback board, a GUI was developed to interface with the system. It was written in C++/CLI [101] and is a standalone application which runs under Microsoft Windows. Primarily, the software was designed for operating the feedback system during laboratory testing and operations during the commissioning stage at ATF2. It is envisaged that a permanent installation would have a specifically designed interface which may integrate into the ATF control system, control the FONT BPM movers and automate certain operations which are currently performed manually. Figure 4.15 shows a screenshot of the C++/CLI GUI during a data acquisition run at ATF2.

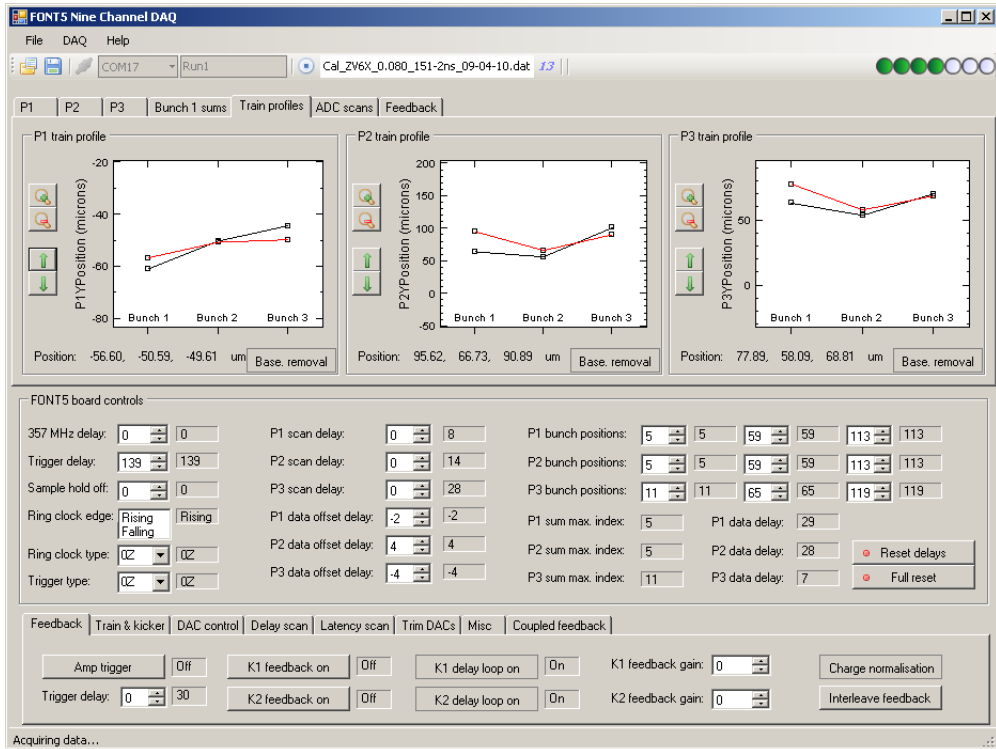


Figure 4.15: Screenshot of the FONT5 C++/CLI graphical user interface.

The GUI includes DAQ, control and real-time display functionality. Visible in the top panel of figure 4.15 are three plots showing the vertical position profile of the bunch train in BPMs P1, P2 and P3. The display updates on a pulse-to-pulse basis, providing valuable information to the operator. As well as the profile of the current machine pulse (black), a

rolling average over a variable number of previous pulses is overlayed (red). By using the tabs at the top of the panel, further plots may be displayed which show the raw digitised BPM signals from all nine FONT5 ADC channels, again updated after each machine pulse.

Each input control in the centre panel of the GUI corresponds to one or more firmware control registers. When the value of a control is modified, the appropriate command and data bytes are transmitted over RS232 to modify the value of the control registers to match. Beside each input control is a read-only text box which is updated with the control register's readback value after each machine pulse, providing a quick diagnostic. The virtual LEDs in the top-right of the form are updated based on the status flag readback (table 4.5).

Parameters controlling the operation of the feedback loops are located in the lower panel. A number of control registers may be set, and the feedback gains modified. When the value in a gain control is modified, the software calculates the appropriate LUT entries and loads them over RS232.

The DAQ may be started or stopped with a single click. When acquiring data, the raw ADC samples from all nine channels are logged on a pulse-by-pulse basis to a human readable ASCII file. The computer clock is used to add a timestamp to each pulse, which is recorded along with the timestamp transmitted by the FONT5 board. A separate ASCII file logs the value of all control registers on every pulse.

An additional mode of operation, interleaved mode, is implemented in the GUI software. When this mode is activated, the FONT5 DAC outputs are toggled on/off after each data stream is received. This means that the resulting data files will contain pulses with the feedback active on every other pulse. By appropriate analysis, the feedback off data may be used to monitor the unperturbed beam properties including any position shift or charge fluctuation, and thus the effects of the feedback system separated from external influence.

Chapter 5

Performance of the FONT5 feedback system at ATF

5.1 Feedback gain

The nominal gain of the feedback system G_0 is defined as the gain which, for a flat bunch train, results in a zero offset in the feedback BPM(s) for all bunches succeeding the pilot. For this purpose, any droop of the kicker pulse (see section 3.5.4) is ignored.

5.1.1 Nominal system gain for a single feedback loop

In a single loop feedback system, the nominal gain is easily defined. Consider a loop between the BPM P2 and kicker K1. Let M be the transfer matrix between the two elements. Then, if an offset of y_1 μm is measured for the pilot bunch at P2, an angular correction of $-y_1/m_{12}$ μrad must be applied to the second bunch at K1 in order to zero its position in the feedback BPM P2 (assuming again a flat bunch train). Hence we have:

$$G_0 = -\frac{1}{m_{12}} \quad (5.1.1)$$

In practice however, the nominal gain is determined by calibration of the kicker as described here. This removes any reliance on knowledge of the transfer matrix between the BPM and kicker. Firstly, note that the FONT5 board performs its digital signal processing on *uncalibrated* position measurements. Let Y_i denote the uncalibrated vertical position of the i^{th} bunch, as calculated from the digitised sum and difference voltages v_i^Σ and v_i^Δ :

$$y_i = cY_i = c\frac{v_i^\Delta}{v_i^\Sigma} \quad (5.1.2)$$

where c is the BPM calibration constant from section 3.4.2. It is then advantageous to rewrite 5.1.1 in terms of the gradients of kicker calibration curves of the type presented in section 3.5.5. If $v = GY_1$ represents the DAC output voltage for an arbitrary gain of G , then bunch two transforms at the feedback BPM P2 as:

$$\tilde{Y}_2 = Y_2 + \frac{dY_2}{dv}v = Y_2 + \frac{dY_2}{dv}GY_1 \quad (5.1.3)$$

Again the nominal gain is defined by assuming a flat train, $Y_1 = Y_2$, and requiring the corrected bunch two position to go to zero, $\tilde{Y}_2 = 0$. Then:

$$G_0 = - \left[\frac{dY_2}{dv} \right]^{-1} \quad (5.1.4)$$

By measuring dY_2/dv in terms of uncalibrated position per DAC count, the units of the gain are assured to be appropriate for use in the FONT5 feedback system. That is, the feedback system may be calibrated entirely with a single kicker scan with no dependence on knowledge of transfer matrices or on the absolute BPM calibration.

5.1.2 Nominal system gain in a coupled feedback system

To determine the set of nominal gains of the coupled feedback system by calibration, the effects of both K1 and K2 kicks must be considered at both feedback BPMs P2 and P3. If the outputs of the DACs which drive K1 and K2 are respectively v_1 and v_2 then:

$$\tilde{Y}_2^{P2} = Y_2^{P2} + \frac{\partial Y_2^{P2}}{\partial v_1}v_1 + \frac{\partial Y_2^{P2}}{\partial v_2}v_2, \quad \text{and} \quad \tilde{Y}_2^{P3} = Y_2^{P3} + \frac{\partial Y_2^{P3}}{\partial v_1}v_1 + \frac{\partial Y_2^{P3}}{\partial v_2}v_2 \quad (5.1.5)$$

The nominal gains for the coupled system are those which, for a flat train ($Y_1 = Y_2$), zero the bunch two offsets in both of the feedback BPMs P2 and P3. Solving the two equations (5.1.5) simultaneously and applying the constraint $\tilde{Y}_2^{P2} = \tilde{Y}_2^{P3} = 0$ gives:

$$v_1 = G_1Y_1^{P2} + G_2Y_1^{P3}, \quad \text{and} \quad v_2 = G_3Y_1^{P2} + G_4Y_1^{P3} \quad (5.1.6)$$

where the four nominal gains are given by:

$$\begin{aligned}
G_1 &= \frac{-1}{\Lambda} \frac{\partial Y_2^{P3}}{\partial v_2} \\
G_2 &= \frac{1}{\Lambda} \frac{\partial Y_2^{P2}}{\partial v_2} \\
G_3 &= \frac{1}{\Lambda} \frac{\partial Y_2^{P3}}{\partial v_1} \\
G_4 &= \frac{-1}{\Lambda} \frac{\partial Y_2^{P2}}{\partial v_1} \\
\text{where } \Lambda &= \frac{\partial Y_2^{P2}}{\partial v_1} \frac{\partial Y_2^{P3}}{\partial v_2} - \frac{\partial Y_2^{P2}}{\partial v_2} \frac{\partial Y_2^{P3}}{\partial v_1}
\end{aligned} \tag{5.1.7}$$

Examples of the kicker calibration curves used to obtain the various gradients are presented in section 3.5.5.

5.1.3 Betatron phase advance and feedback system location

When deciding on the location of the BPMs and kickers that constitute the FONT5 feedback system, attention must be paid to the expected betatron phase advance through the region's lattice. With a single degree-of-freedom system, for example a loop between P2 and K1, the phase advance between the elements determines the quality of the obtained angular correction. When operating at nominal gain, the RMS error in the angular deflection applied to bunch two is calculated from:

$$\text{Var}[\epsilon_i] = \frac{(R^{P2})^2}{a_{12}^2} \tag{5.1.8}$$

Rewriting in terms of the Twiss parameters using equation 2.1.22 and differentiating with respect to the phase advance:

$$\frac{d}{d(\Delta\psi)} \text{Var}[\epsilon_i] = \frac{(R^{P2})^2}{\beta(s_{K1})\beta(s_{P2})} \frac{d}{d(\Delta\psi)} \frac{1}{\sin^2 \Delta\psi} = \frac{8(R^{P2})^2}{\beta(s_{K1})\beta(s_{P2})} \frac{\cos \Delta\psi}{\sin(3\Delta\psi) - 3\sin \Delta\psi} \tag{5.1.9}$$

then the stationary point representing minimum error in the angular kick is found when the phase advance $\Delta\psi = \pi/2$. With a $\pi/2$ phase advance between the kicker and BPM, the angular error is then:

$$\text{Var}[\epsilon_i^{min.}] = \frac{(R^{P2})^2}{\beta(s_{K1})\beta(s_{P2})} \tag{5.1.10}$$

That is, angular error is at its lowest when the BPM is placed in a region of maximal β and the kicker is $\pi/2$ radians upstream. It is then possible to construct a two degree-of-freedom system using two such single loops. Keeping both the phase advances between K1 and P2

and between K2 and P3 at $\pi/2$ radians, the nominal feedback gains of section 5.1.2 may be rewritten in terms of the Twiss parameters as:

$$\begin{aligned}
G_1 &= \frac{-1}{\sin^2 \Delta\phi \sqrt{\beta(s_{K1})\beta(s_{P2})}} \\
G_2 &= \frac{\cos \Delta\phi}{\sin^2 \Delta\phi \sqrt{\beta(s_{K1})\beta(s_{P3})}} \\
G_3 &= \frac{\cos \Delta\phi}{\sin^2 \Delta\phi \sqrt{\beta(s_{K2})\beta(s_{P2})}} \\
G_4 &= \frac{-1}{\sin^2 \Delta\phi \sqrt{\beta(s_{K2})\beta(s_{P3})}}
\end{aligned} \tag{5.1.11}$$

where $\Delta\phi$ is the phase advance between the kickers K1 and K2. When the two loops are separated by $\Delta\phi = (2j + 1)\pi/2$ radians where ($j = 0, 1, 2, \dots$), the nominal gains G_2 and G_3 become zero. In this situation the two loops are operating independently yet are able to reduce both position and angle jitter. With both loops operating independently, the necessary signal path lengths are minimised as is the latency contribution from beam time of flight; this would represent a minimal latency installation. Reliance on a near-perfect phase advance between the beamline elements is undesirable in practice however, since lattice optics are tuned for a variety of operational goals. A feedback system must work effectively in the various lattice configurations that it may encounter. Furthermore, beamline real-estate is a valued commodity. As much as it may be beneficial to install BPMs and kickers at optimal locations, this desire must be balanced against the needs of the lattice's magnets, instrumentation and insertion devices.

The FONT5 installation is then a compromise. The feedback BPMs P2 and P3 are both located where the vertical beta function is expected to be large, with the phase advance between them approximately $\pi/2$ radians. The two kickers are located approximately $\pi/2$ radians upstream of their respective BPMs. Though the two loops need to be coupled, this compact arrangement leads to a low latency whilst remaining fully flexible.

5.1.4 Optimal gain for an imperfect beam

The nominal gain as discussed in section 5.1 is, for a bunch train with 100% position correlation (i.e. train-to-train jitter only) and BPMs with perfect resolution, the optimal system gain. Uncorrelated bunch-to-bunch jitter or noise can be amplified by the feedback system, altering the gain which achieves an optimal reduction in jitter. As an example, the effects of a single feedback loop at the feedback BPM are considered. The corrected jitter

of the second bunch may be calculated by noting that:

$$\text{Var}[\tilde{y}_2] = \text{Var}[y_2] + G_N^2 \text{Var}[y_1] - 2G_N \text{Cov}[y_1, y_2] \quad (5.1.12)$$

here the arbitrary gain G has been normalised by the nominal gain, that is $G_N = G/G_0$. Differentiating with respect to the normalised gain and finding the stationary point gives the optimal gain:

$$G_N^{opt.} = \frac{\text{Cov}[y_1, y_2]}{\text{Var}[y_1]} = \frac{\sigma_2}{\sigma_1} \rho(y_1, y_2) \quad (5.1.13)$$

The optimal gain then is linearly dependent on the bunch-to-bunch correlation $\rho(y_1, y_2)$. In the case that bunches one and two have jitters of equal magnitude, the optimal normalised gain is equal to the correlation; for a bunch-to-bunch correlation of 80% the jitter will be minimised when the gain is set to 80% of the nominal. The correlation of bunch one with the corrected bunch two may be written as:

$$\rho(y_1, \tilde{y}_2) = \frac{\text{Cov}[y_1, y_2] - G_N \text{Var}[y_1]}{\sqrt{\text{Var}[y_1] \text{Var}[\tilde{y}_2]}} = \frac{\sigma_2 \rho(y_1, y_2) - \sigma_1 G_N}{\sqrt{\sigma_2^2 + G_N^2 \sigma_1^2 - 2G_N \sigma_1 \sigma_2 \rho(y_1, y_2)}} \quad (5.1.14)$$

If the optimal gain from equation 5.1.13 is substituted for G_N , the correlation $\rho(y_1, \tilde{y}_2) = 0$. When the gain is set to its optimal value, the feedback system minimises the magnitude of bunch two's jitter by removing all of the component correlated with bunch one. These findings are valid whether the uncorrelated component of the observed jitter is a real beam effect or whether it is due to measurement error. Figure 5.1 shows the analytically expected change in bunch two's jitter magnitude and correlation with bunch one for a range of normalised feedback gains. Each curve represents a different initial bunch-to-bunch correlation, and shows the special case where $\sigma_1 = \sigma_2$.




Figure 5.1: Analytically expected variation of bunch two's jitter magnitude (a) and correlation with bunch one (b) for a range of normalised gains. Each curve represents a different initial bunch-to-bunch correlation.

5.2 Latency measurements

Measurement of the FONT5 feedback system latency is carried out by forcing a constant DAC output instead of the usual feedback signals (section 3.5.5). The constant kick has the same timing as the feedback signal, i.e. the kicker rises as soon as possible after the pilot bunch is measured at the feedback BPM. The latency is defined as the length of time from the pilot bunch entering the feedback system at the kicker K1 until the kicker pulse rises to 90% of its full output.

Additional delay may be added to (or subtracted from) the constant K1 kick's timing. While scanning this delay, the deflection of bunch two will decrease as the rising edge of the kicker pulse arrives increasingly late, as shown in figure 5.2. Since the bunch spacing is well known, the latency of the feedback system is easily calculated by subtracting the bunch spacing from the delay required to reduce the kick to 90%.




Figure 5.2: Schematic showing the procedure for measuring feedback system latency. The upper plot shows the kicker pulse, with the solid blue pulse at the nominal timing while the dashed yellow pulse has been delayed. The lower plot shows the resulting deflection to bunches 2 and 3 in each case.

The data in figure 5.3 were taken with a bunch spacing of 151.2 ns on the 7th of April 2010. The kick was applied in an interleaved mode (section 4.6.2) providing around 10 pulses each of kicked and unkicked measurements at every delay setting. The datapoints represent the mean difference between the kicked and unkicked pulses with the error bars showing the statistical error on the measurements. To calculate the 90% kick point, a four parameter sigmoid was fitted to the data. The sigmoid has the form:

$$f(t) = a_1 + \frac{a_2}{1 + e^{-a_3(t-a_4)}} \quad (5.2.1)$$

with the four parameters a_i . Figures 5.3 (a) and (b) are used to calculate the latency of the single feedback loop from P2-to-K1 and the coupled feedback system respectively. In the coupled case, the total latency of the system is set by the longest signal path, that from P3 to K1.

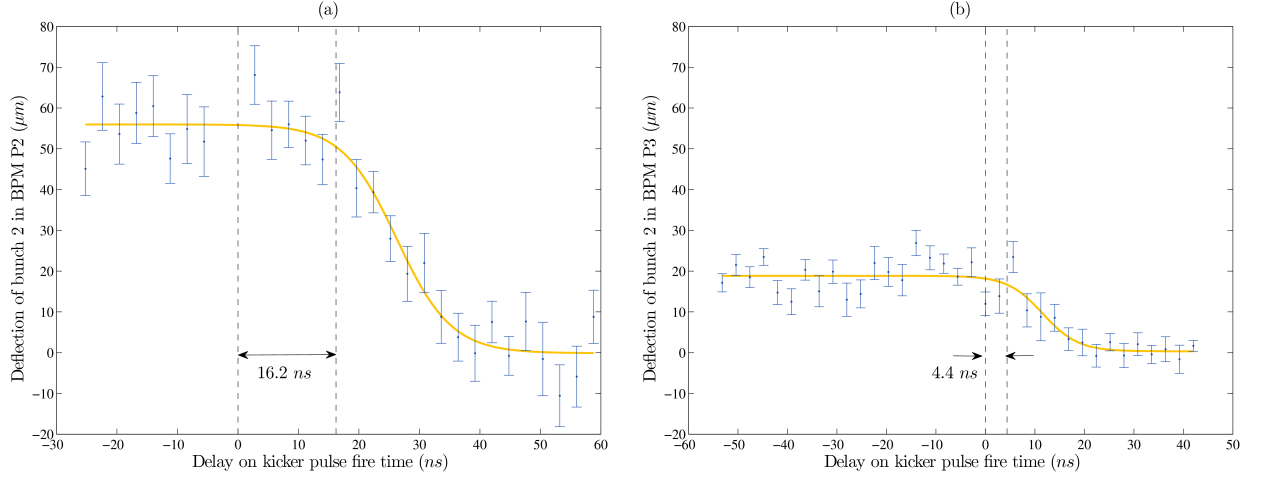


Figure 5.3: Mapping of K1 kicker pulse rising edge by delaying the nominal output time of the feedback signal. A constant kick is applied, with (a) representing feedback from P2-to-K1 while (b) represents P3-to-K1.

The latencies calculated from the above data are presented in table 5.1, along with the sigmoid fit parameters. Table 5.2 lists the expected latencies of the components comprising the FONT5 system, and further discussion of the latency measurement may be found elsewhere [102].

	P2-to-K1	P3-to-K1
a_1	56.0 ± 2.0	33.9 ± 1.5
a_2	-56.1 ± 3.5	-29.7 ± 3.3
a_3	0.22 ± 0.05	0.18 ± 0.07
a_4	-26.2 ± 1.3	-10.2 ± 2.5
χ^2/df	0.97	1.27
Latency	135.0 ± 2.3 ns	147.2 ± 3.1 ns

Table 5.1: Measured latencies of the P2-to-K1 single loop feedback system and the coupled loop feedback system. Parameters of the sigmoid fits performed are shown.

Source	P2-to-K1 latency (ns)	P3-to-K1 latency (ns)
Time of flight	12	21
Signal return	32	42
Analogue processor	10	10
ADC (3.5 cycles)	9.8	9.8
FPGA I/O	3	3
FPGA DSP (8 cycles)	22.4	22.4
DAC (1 cycle + 1 ns)	3.8	3.8
Amplifier	35	35
Kicker fill time	3	3
Estimated latency	131	150

Table 5.2: Breakdown of the sources of the anticipated FONT5 latencies.

5.3 Effects of LO phase jitter on the BPM resolution

By making the assumption that BPMs P1, P2 and P3 have the same magnitude measurement errors, the FONT5 BPM resolution was calculated in section 3.4.4. This assumption is shown here to be invalid however, and further information is used to provide estimates of the individual BPM resolutions.

5.3.1 Local oscillator phase jitter

In addition to the intrinsic BPM resolutions, another potential source of measurement error comes from the phase jitter of the 714 MHz LO signal. Any change in phase of the LO signal with respect to the beam affects the output level of each processor’s analogue mixer. This translates into a change in apparent beam position. Observations of the ATF2 LO phase jitter reveal its magnitude to be less than 1° , which in a well installed system is expected to lead to a position error negligible in comparison to the intrinsic BPM resolution. However, any phase difference between the signals from opposing striplines when they reach the analogue processor will exacerbate the induced position error. Such a phase offset might be introduced by using a pair of cables with poorly matched length, and might degrade BPM resolution without having an immediately obvious impact on system operation [103].

To test for the presence of LO phase jitter induced position error, three FONT5 channels were temporarily reassigned to record the amplified quadrature sum signals from the three analogue processors. The quadrature sum signal is maximally sensitive to changes in LO phase. If the quadrature component is normalised by the sum signal to remove the effects

of beam charge fluctuation, a good signal for monitoring the LO phase jitter is obtained. Although it is not trivial to calibrate accurately this signal, its correlation or lack thereof with observed position jitter is informative.




Figure 5.4: The quadrature sum signal, normalised against charge fluctuation, vs. the measured position of the pilot bunch in BPMs P1 (a), P2 (b) and P3 (c). Only P3 displays significant correlation.

A data set of 470 pulses is presented in figure 5.4. For BPMs P1, P2 and P3, the recorded position of the first bunch of each pulse is plotted against the normalised quadrature sum signal. It is clear that in P1 and P2 there is no significant correlation between the two observations. In P3 however, there is a significant link between the LO phase and the observed position. Therefore the resolution of the BPM P3 will be degraded. This has several consequences:

- Clearly, the assumption in section 3.4.4 is not appropriate when the resolution of P3 is poorer than its neighbours. The observed $\sim 1.5 \mu\text{m}$ can not be taken as correct, although it certainly sets the scale.
- The LO phase jitter may have a broad frequency spectrum. Jitter which is high frequency on the time-scale of the bunch spacing ($c. 100 \text{ ns}$) will lead to a worsening of P3's intrinsic resolution.
- Components of the LO phase jitter which are of lower frequency than the bunching frequency will introduce additional error into the position measurements which is *correlated* between bunches.

In order to understand the results of coupled feedback loop experiments in section 5.7, the additional correlated and uncorrelated noise in P3 stemming from LO phase jitter must be taken into consideration. This is necessarily done in a quantitative but approximate way. It is important to remember that this additional noise is due to a poorly behaved BPM in an experimental installation; the final installation of the FONT5 system is currently (November

2010) being undertaken with a view to eliminating the cable mismatch responsible. This work will be documented in [82].

5.3.2 Obtaining limits on the BPM resolutions

As discussed previously, the differing behaviour of the FONT5 BPMs makes a direct calculation of their individual resolutions impossible. Given additional information however, one can place constraints on those resolutions. The correlation between the positions of consecutive bunches in a single BPM can be used to gain insight into that BPM's resolution, for example. If the positions of two consecutive bunches are correlated over many pulses with a coefficient of ρ_{ij} , then the measured correlation coefficient ρ_{ij}^* will be reduced due to the measurement error. Given the variance of the position of the i^{th} bunch over many pulses, V_i , and of the BPM readings, V_i^* , the true and measured correlation coefficients for bunches one and two are written as:

$$\rho_{12} = \frac{\text{Cov}[y_1, y_2]}{\sqrt{V_1 V_2}}, \quad \text{and} \quad \rho_{12}^* = \frac{\text{Cov}[y_1^*, y_2^*]}{\sqrt{V_1^* V_2^*}} \quad (5.3.1)$$

where $V_1 = V_1^* - R^2$ for a BPM of resolution R . Then if r_i is again the measurement error of the i^{th} bunch on a given pulse, the covariance of the measured positions is, assuming no dependence on position for the resolution:

$$\text{Cov}[y_1^*, y_2^*] = \text{Cov}[y_1 + r_1, y_2 + r_2] = \text{Cov}[y_1, y_2] + \text{Cov}[r_1, r_2] \quad (5.3.2)$$

and where there is no correlation between the measurement errors of subsequent bunches, as is the case for P1 and P2, the true bunch position correlation can simply be deconvolved from the measured bunch position correlation by noting that:

$$\rho_{12} = \rho_{12}^* \sqrt{\frac{V_1^* V_2^*}{(V_1^* - R^2)(V_2^* - R^2)}} \quad (5.3.3)$$

In order to obtain the worst case limit on the resolution, R^{max} , the true bunch position correlation may be set to 100% such that all observed decorrelation is due to measurement error. The upper limit on the resolution is then obtained by finding the physical root of the following polynomial:

$$(R^{max})^4 - [V_1^* + V_2^*](R^{max})^2 + V_1^* V_2^* [1 - (\rho_{12}^*)^2] \quad (5.3.4)$$

Using this technique with data from the 16th April 2010, limits on the intrinsic resolutions (that is, the component of noise uncorrelated from bunch to bunch) of the three FONT5 BPMs are obtained as listed in table 5.3.

BPM	Limit (μm)
P1	0.45
P2	0.45
P3	1.50

Table 5.3: Worst-case resolution limits calculated from the observed bunch to bunch correlations in the FONT5 BPMs.

5.4 Measurement of the jitter covariance matrix

In order to model the operation of the feedback system it is necessary to reconstruct the position/angle covariance matrix. The BPM readings from at least two of the BPMs P1, P2 and P3 can be combined to determine the covariance matrix at K1, but the BPM resolution necessarily degrades the reconstruction. This degradation can be corrected if the resolutions are known or can be estimated.

5.4.1 Ideal case with perfect resolution

Each bunch has a coordinate vector at K1:

$$\mathbf{y}(s_{K1}) = \begin{pmatrix} y(s_{K1}) \\ y'(s_{K1}) \end{pmatrix} \quad (5.4.1)$$

The vector may be written in terms of the positions in the three BPMs as follows:

$$\mathbf{y}(s_{K1}) = U \begin{pmatrix} y(s_{P1}) \\ y(s_{P2}) \\ y(s_{P3}) \end{pmatrix} = U \mathbf{y}_{BPM} \quad (5.4.2)$$

where the 2×3 matrix U may be written in terms of the transfer matrices between K1 and the three BPMs. As usual, all statistical quantities are taken over an ensemble of many machine pulses. The covariance matrix Σ at K1 is then:

$$\begin{aligned} \Sigma(s_{K1}) &= \left\langle (\mathbf{y}(s_{K1}) - \langle \mathbf{y}(s_{K1}) \rangle)(\mathbf{y}(s_{K1}) - \langle \mathbf{y}(s_{K1}) \rangle)^\dagger \right\rangle \\ &= \left\langle \left(U \mathbf{y}_{BPM} - \langle U \mathbf{y}_{BPM} \rangle \right) \left(U \mathbf{y}_{BPM} - \langle U \mathbf{y}_{BPM} \rangle \right)^\dagger \right\rangle \\ &= U \Sigma_{BPM} U^\dagger \end{aligned} \quad (5.4.3)$$

where Σ_{BPM} is the 3×3 covariance matrix of the three sets of BPM positions.

5.4.2 Correcting for degradation by known resolutions

Let the measured BPM vector be written \mathbf{y}_{BPM}^* such that:

$$\mathbf{y}_{BPM}^* = \mathbf{y}_{BPM} + \mathbf{r}, \quad \mathbf{r} = \begin{pmatrix} r^{P1} \\ r^{P2} \\ r^{P3} \end{pmatrix} \quad (5.4.4)$$

The components r^{Pi} of the resolution vector are random variables drawn from Gaussian distributions with standard deviation R^{Pi} and zero mean. Assuming the resolution is independent of position (a good assumption if the bunches jitter around the BPMs' zero positions), the covariance matrix of the BPM readings is:

$$\Sigma_{BPM}^* = \Sigma(s_{K1}) + U \langle \mathbf{r} \mathbf{r}^\dagger \rangle U^\dagger \quad (5.4.5)$$

and if the three BPM resolutions are independent then the true covariance matrix at K1 can be calculated from:

$$\Sigma_{BPM} = \Sigma_{BPM}^* - U \begin{pmatrix} (R^{P1})^2 & 0 & 0 \\ 0 & (R^{P2})^2 & 0 \\ 0 & 0 & (R^{P3})^2 \end{pmatrix} U^\dagger \quad (5.4.6)$$

5.5 Modelling the FONT5 system

Summarised here are the mathematical models used to compare the FONT5 feedback system operation with expectations. They are each a specific implementation of the statistical formalism described fully in appendix A. In all cases, the three bunches of the ATF2 train are described by a 6-element state vector:

$$\mathbf{y} = \begin{pmatrix} y_1 \\ y'_1 \\ y_2 \\ y'_2 \\ y_3 \\ y'_3 \end{pmatrix} \quad (5.5.1)$$

At any point s in the lattice, an ensemble of machine pulses is characterised by the 6×6 jitter covariance matrix $\Sigma(s)$:

$$\Sigma(s) = \left\langle \left(\mathbf{y}(s) - \langle \mathbf{y}(s) \rangle \right) \left(\mathbf{y}(s) - \langle \mathbf{y}(s) \rangle \right)^\dagger \right\rangle \quad (5.5.2)$$

If M is the 2×2 vertical position transfer matrix between s_1 and s_2 , then a 6×6 block diagonal transfer matrix M may be used to transform the jitter covariance matrix:

$$M = \begin{pmatrix} M & 0 & 0 \\ 0 & M & 0 \\ 0 & 0 & M \end{pmatrix} \quad (5.5.3)$$

such that:

$$\Sigma(s_2) = M \Sigma(s_1) M^\dagger \quad (5.5.4)$$

5.5.1 Modelling of P2-to-K1 single loop feedback

Section 5.4 describes how the jitter covariance matrix at K1 may be reconstructed from data taken with the feedback turned off. As explained in appendix A, turning on the feedback will transform the covariance matrix. Firstly, the following matrices are defined:

$$F = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ Ga_{11} & Ga_{12} & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ Ga_{11}(1 + Ga_{12}) & Ga_{12}(1 + Ga_{12}) & Ga_{11} & Ga_{12} & 0 & 1 \end{pmatrix} \quad (5.5.5)$$

and

$$H = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ G & 0 & 0 \\ 0 & 0 & 0 \\ G(1 + Ga_{12}) & G & 0 \end{pmatrix} \quad (5.5.6)$$

where G is the feedback gain and a_{ij} the elements of the transfer matrix A between K1 and P2. Now the transformed jitter covariance matrix at K1 may be written:

$$\tilde{\Sigma}(s_{K1}) = F\Sigma(s_{K1})F^\dagger + H\Sigma_R H^\dagger \quad (5.5.7)$$

Σ_R , the resolution covariance matrix for the feedback BPM P2, is assumed to be diagonal. This is the case when the measurement errors are not correlated between different bunches, giving:

$$\Sigma_R = \begin{pmatrix} (R^{P2})^2 & 0 & 0 \\ 0 & (R^{P2})^2 & 0 \\ 0 & 0 & (R^{P2})^2 \end{pmatrix} \quad (5.5.8)$$

It is then possible to predict the feedback system's effects at any point downstream of K1 by transforming $\tilde{\Sigma}(s_{K1})$ with the appropriate block diagonal transfer matrix.

5.5.2 Modelling of P3-to-K2 single loop feedback

When modelling a single feedback loop closed between K2 and P3, the procedure is much the same as that for the P2-to-K1 loop. The jitter covariance matrix at K2 is determined before being transformed as in equation 5.5.7, using elements from the transfer matrix between P3 and K2 instead. A slight complication arises due to the bunch to bunch correlated noise stemming from LO phase jitter. In the model, this is accounted for by modifying the resolution covariance matrix to:

$$\Sigma_R = \begin{pmatrix} (R^{P3})^2 & P^2 & P^2 \\ P^2 & (R^{P3})^2 & P^2 \\ P^2 & P^2 & (R^{P3})^2 \end{pmatrix} \quad (5.5.9)$$

With this modification, an additional noise component, correlated between the bunches, is allowed. This model is somewhat simplified. Only two components of noise are allowed; high frequency in comparison to the bunch spacing, and in P3 a component slow in comparison to the bunch train length.

5.5.3 Modelling the coupled feedback system

A model of the full coupled system requires that transformations be applied at both kickers. As before, the jitter covariance matrix at K1 is determined and transformed. However, in the coupled case, the error due to resolution is not applied along with the K1 feedback transformation. The same measurement errors contribute to both kicks, introducing a

correlation between the angular errors in the kicks of K1 and K2. To account for this the combined angular error from both kickers is added at K2. At K1 then, the transformation is:

$$\tilde{\Sigma}(s_{K1}) = F_{K1} \Sigma(s_{K1}) F_{K1}^\dagger \quad (5.5.10)$$

where in the coupled case, the feedback transformation matrix takes the form (see again appendix A):

$$F = \begin{pmatrix} 1 & 0 & \dots \\ 0 & 1 & \dots \\ 0 & 0 & \dots \\ G_1 a_{11} + G_2 b_{11} & G_1 a_{12} + G_2 b_{12} & \dots \\ 0 & 0 & \dots \\ (G_1 a_{11} + G_2 b_{11})(1 + G_1 a_{12} + G_2 b_{12}) & (G_1 a_{12} + G_2 b_{12})(1 + G_1 a_{12} + G_2 b_{12}) & \dots \\ \dots & 0 & 0 & 0 & 0 \\ \dots & 0 & 0 & 0 & 0 \\ \dots & 1 & 0 & 0 & 0 \\ \dots & 0 & 1 & 0 & 0 \\ \dots & 0 & 0 & 1 & 0 \\ \dots & G_1 a_{11} + G_2 b_{11} & G_1 a_{12} + G_2 b_{12} & 0 & 1 \end{pmatrix} \quad (5.5.11)$$

For the coupled case the matrix H becomes a 6×6 matrix of the form:

$$H = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ G_1 & G_2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ G_1(1 + G_1 a_{12} + G_2 b_{12}) & G_2(1 + G_1 a_{12} + G_2 b_{12}) & G_1 & G_2 & 0 & 0 \end{pmatrix} \quad (5.5.12)$$

Let the transfer matrices from K1 to P2 and P3 be denoted A and B respectively. For K2 to P2 and P3, let them be C and D such that $\mathbf{y}(s_{K2}) = C^{-1} A \mathbf{y}(s_{K1})$. At K2, the effect on

the jitter covariance matrix including resolution errors from both kickers is then:

$$\begin{aligned} \tilde{\Sigma}(s_{K2}) &= [\mathbf{F}_{K2}\mathbf{C}^{-1}\mathbf{A}] \tilde{\Sigma}(s_{K1}) [\mathbf{F}_{K2}\mathbf{C}^{-1}\mathbf{A}]^\dagger \\ &+ [\mathbf{F}_{K2}\mathbf{C}^{-1}\mathbf{A}\mathbf{F}_{K1}\mathbf{H}_{K1} + \mathbf{H}_{K2}] \Sigma_R [\mathbf{F}_{K2}\mathbf{C}^{-1}\mathbf{A}\mathbf{F}_{K1}\mathbf{H}_{K1} + \mathbf{H}_{K2}]^\dagger \end{aligned} \quad (5.5.13)$$

where the resolution covariance matrix, taking into account the correlated noise in P3, is:

$$\Sigma_R = \begin{pmatrix} (R^{P2})^2 & 0 & 0 & 0 & 0 & 0 \\ 0 & (R^{P3})^2 & 0 & P^2 & 0 & P^2 \\ 0 & 0 & (R^{P2})^2 & 0 & 0 & 0 \\ 0 & P^2 & 0 & (R^{P3})^2 & 0 & P^2 \\ 0 & 0 & 0 & 0 & (R^{P2})^2 & 0 \\ 0 & P^2 & 0 & P^2 & 0 & (R^{P3})^2 \end{pmatrix} \quad (5.5.14)$$

5.6 Feedback performance in single-loop mode

As an initial test of the single degree-of-freedom vertical position feedback loop between P2 and K1, various currents were run through the dipole corrector magnet ZV7X upstream of the feedback system to introduce artificially a number of different beam position offsets at P2. The range of these offsets was an order of magnitude larger than the observed beam jitter and measured BPM resolutions. Given that the offsets are within the feedback system's linear range, the effects of feedback system resolution can be ignored and the position of the second and third bunches may be expected to approach zero in P2 regardless of introduced offset.




Figure 5.5: The effect of single degree-of-freedom vertical position feedback on a beam offset of $\sim 80 \mu\text{m}$, where the error bars represent the RMS beam jitter (a). The effect on a range of offsets covering approximately $\pm 100 \mu\text{m}$ (b).

Figure 5.5 shows the effect of the feedback when the beam is offset in this way. Each datapoint represents the average over ~ 50 beam pulses. The feedback system gain was set to the nominal, as discussed in section 5.1, which was determined to be 5070. In order to obtain a good correction for the third bunch, the first coefficient of the FIR filter k_1 was set to 0.254. In the feedback off data of figure 5.5 (a), a clear shape to the bunch train is visible with approximately a 20 μm sagitta. In order to flatten the bunch train, additional constant kicks of 100 and -200 DAC counts were applied to bunches 2 and 3 respectively (see section 3.5.3). Figure 5.5 (b) shows the effect of the feedback system for seven offsets approximately covering the range $\pm 100 \mu\text{m}$, and the position of bunches 2 and 3 are seen to approach zero for all cases.

5.6.1 Vertical position jitter reduction at P2

To observe the effect of the single loop feedback system on the vertical position jitter as measured at the feedback BPM, P2, the apparatus was allowed to run in interleaved mode (section 4.6.2). Presented in this section is a 1000-pulse dataset which was obtained on the 16th April 2010, when the nominal gain was measured to be 4625. The first coefficient of the FIR filter k_1 remained at 0.254.

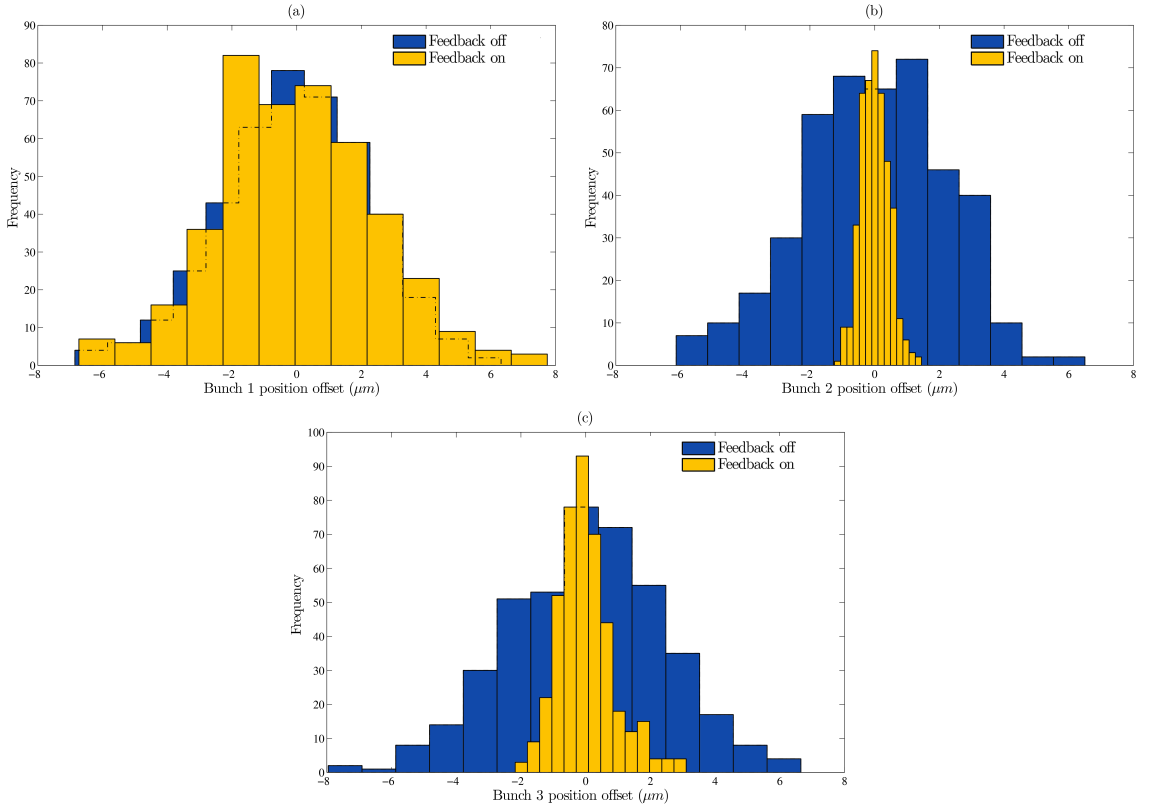


Figure 5.6: The effects of the P2-to-K1 single loop position feedback on the jitter magnitude measured at P2 for bunches one (a), two (b) and three (c).

The reduction in position jitter at P2 for bunches two and three is shown in figure 5.6, whilst the effect on the bunch-to-bunch position correlation is noted in figure 5.7. The feedback works to remove the correlated component of the jitter, reducing its overall magnitude. The magnitude of the residual jitter is consistent with the limit placed on the resolution of P2 in section 5.3.2.




Figure 5.7: The effects of the P2-to-K1 single loop position feedback on the bunch-to-bunch position correlations observed in measurements with P2.

Bunch	Feedback off	Feedback on	Bunches	Feedback off	Feedback on
2	2.2 μm	0.43 μm	1 to 2	98%	34%
3	2.4 μm	0.87 μm	2 to 3	92%	2.3%

Table 5.4: Summary of the effects of the P2-to-K1 single loop position feedback on the position jitter (left) and bunch-to-bunch correlations (right) as observed in measurements with P2.

The important results are summarised in table 5.4. Notice that bunch two has a better initial correlation with bunch one than does bunch three with bunch two. In this case a better correction is expected for the second bunch than the third, as is indeed observed. A more quantitative understanding of these figures requires the feedback model of section 5.5, as is used in section 5.8.1.

With a single feedback loop in operation, the reduction in jitter is focussed at the feedback BPM. Figures 5.8 and 5.9 show how the jitter has developed at the downstream BPM P3. The angular kicks used to minimise the position in P2 soon cause downstream position offsets as the betatron phase advances.

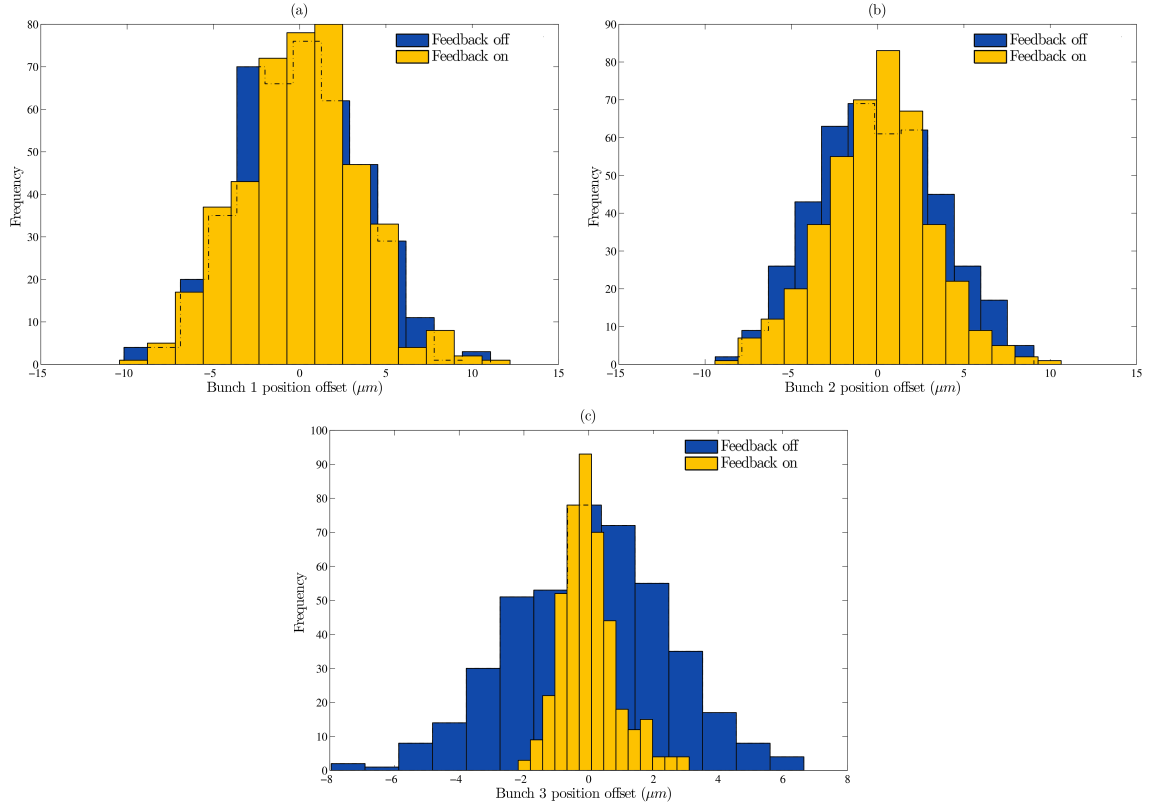


Figure 5.8: The effects of the P2-to-K1 single loop position feedback on the jitter magnitude measured at downstream BPM P3 for bunches one (a), two (b) and three (c).

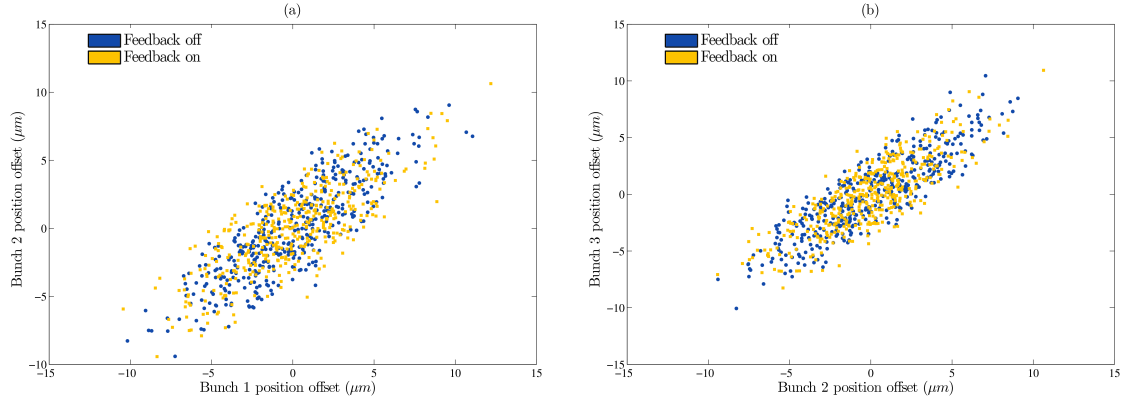


Figure 5.9: The effects of the P2-to-K1 single loop position feedback on the bunch-to-bunch position correlations observed in measurements with the downstream BPM P3.

It is important to note that these histograms are of BPM *measurements*. The width of the distributions are dependent on both the true beam jitter and the BPM resolution; when comparing the results with the predicted beam jitter reduction as in section 5.8.1, the resolution must be deconvolved from the measurements. Necessary deconvolutions are described in appendix B.

5.6.2 Vertical position jitter reduction at P3

A further test of single loop feedback was carried out by closing the loop between P3 and K2. The following data are from a 1000-pulse dataset taken on the 17th April 2010. For the P3 loop, the nominal feedback gain was measured at this time to be 4130 and the FIR tap weight k_1 to be again 0.254. As usual, data were taken in the interleaved mode of operation meaning 50% of pulses are with the feedback turned off.

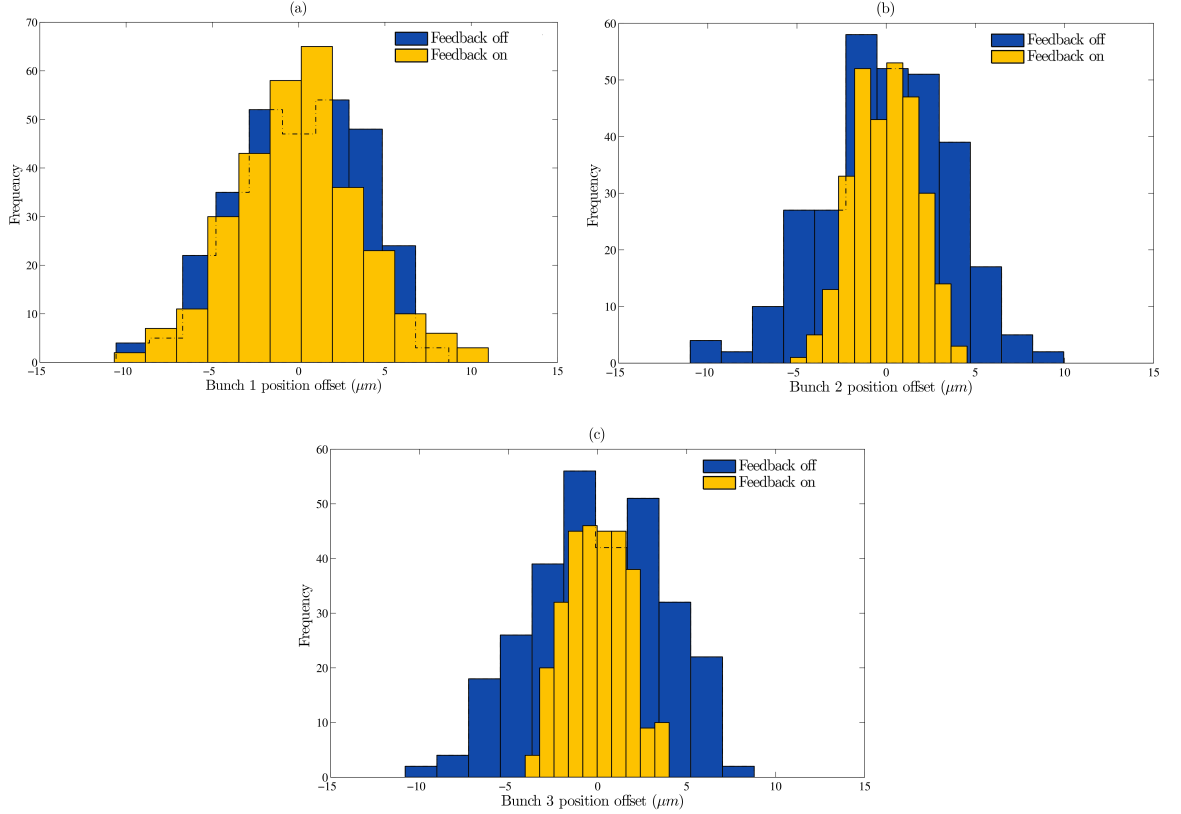


Figure 5.10: The effects of the P3-to-K2 single loop position feedback on the jitter magnitude measured at P3 for bunches one (a), two (b) and three (c).

Again, the distributions are of BPM measurements and include contributions both from beam jitter and measurement error. In the P3 case however, the fact that these measurement errors are correlated between bunches means that although the width of the measured distributions may be reduced by the feedback, the magnitude of the true beam jitter is not necessarily reduced. In section 5.7, where the performance of the coupled loops are presented, this effect is discussed in more detail.

Figure 5.11 shows how the feedback has affected the bunch-to-bunch correlations of the position measurements, and the important results are listed in table 5.5. For this data set the initial jitters of bunches two and three are comparable, as are the bunch one-to-two and two-to-three correlations. As a result, the width of the distributions of bunches two and three when the feedback is turned on are also comparable.

Also note that the initial bunch-to-bunch correlations in this dataset are significantly lower than those observed in the P2-to-K1 feedback of section 5.6.1. For this reason the P2-to-K1 loop is seen to provide better jitter reduction. A number of factors affect the true position correlations between the bunches: betatron phase, extraction errors and x-y coupling all contribute. Here however, the most significant cause is the degradation of P3's intrinsic resolution. This will be seen in section 5.8, where the feedback performance is modelled and compared with the data.




Figure 5.11: The effects of the P3-to-K2 single loop position feedback on the bunch-to-bunch position correlations observed in measurements with P3.

Bunch	Feedback off	Feedback on	Bunches	Feedback off	Feedback on
2	3.6 μm	1.8 μm	1 to 2	87%	-26%
3	3.7 μm	1.7 μm	2 to 3	88%	-58%

Table 5.5: Summary of the effects of the P3-to-K2 single loop position feedback on the position jitter (left) and bunch-to-bunch correlations (right) as observed in measurements with P3.

5.7 Performance of two coupled feedback loops

For these data, the two feedback loops P2-to-K1 and P3-to-K2 ran simultaneously in the coupled mode of operation. In addition, to understand the effect on angle jitter downstream of the feedback system, the FONT5 channels dedicated to BPM P1 were temporarily re-assigned to record data from the downstream witness BPM MQF15X.

5.7.1 Effect on vertical position and angle jitter

On the 22nd of April 2010, the dataset presented in this section was taken. Table 5.6 shows the parameters of the feedback system as set using the gain calibration procedure outlined in section 5.1.

Parameter	K1 loop	K2 loop
P2 gain	-5878	1997
P3 gain	1458	-5131
FIR k_1	0.111	0.238

Table 5.6: Operating parameters of the FONT5 coupled loop feedback system for the 22nd of April 2010.

A similar quality of position jitter reduction in P2 is obtained with the coupled system as was demonstrated in section 5.6 for the single loop case (figures 5.12 and 5.13). Because the feedback signals are functions of measurements at both P2 and P3 however, the error in the kick will be larger meaning that the coupled loops will never be able to achieve as good a correction as a single loop.

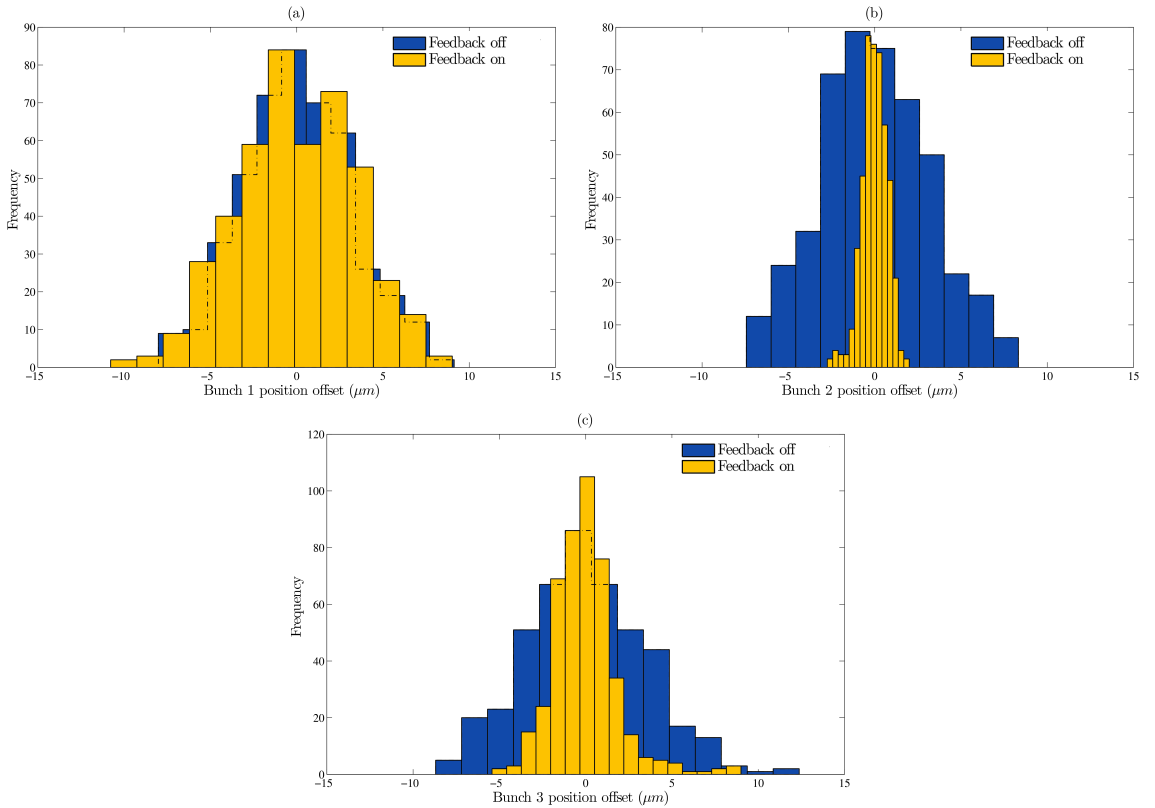


Figure 5.12: The effects of the coupled loop position and angle feedback on the jitter magnitude measured at P2 for bunches one (a), two (b) and three (c).

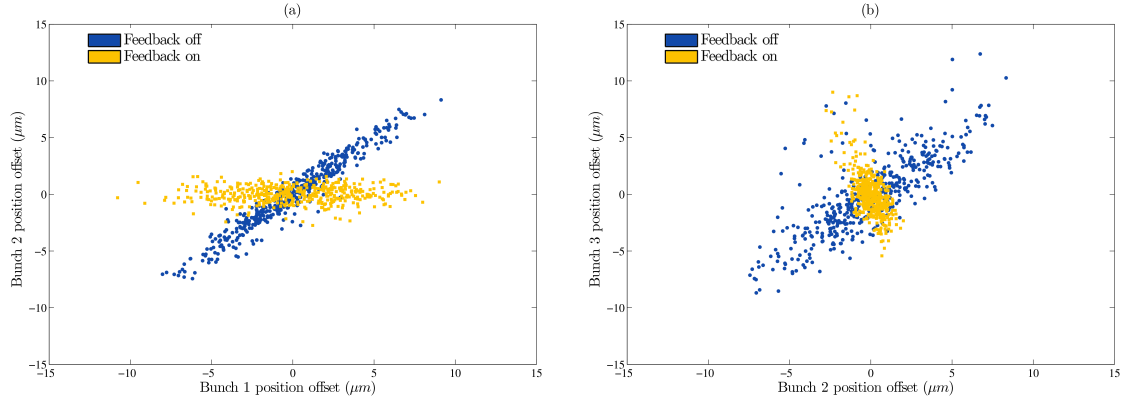


Figure 5.13: The effects of the coupled loop position and angle feedback on the bunch-to-bunch position correlations observed in measurements with P2.

With both coupled loops operating, an apparent correction is maintained in the second feedback BPM P3 as presented in figures 5.14 and 5.15. There is an important caveat to these results however. The plots are of the BPM measurements, not the true beam position. In P3, the noise on the measurement is correlated bunch-to-bunch due to installation errors as discussed in section 5.3.1. The feedback system performs admirably, applying corrections which work to reduce the component of the measured jitter which is correlated between the bunches. However, some of this correlated jitter as measured by P3 is spurious; it stems not from true beam position correlation but from measurement error induced by LO phase variation. When the feedback attempts to correct for this measurement error, it does so by introducing a real beam position deflection in opposition to the error. In effect then, the feedback system couples the bunch-to-bunch correlated noise of P3 into the true beam jitter. That this occurs cannot be deduced from measurements using P3 itself. Figures 5.16 and 5.17 which describe the beam jitter measured at the downstream witness BPM MQF15X make the degradation so introduced very clear.

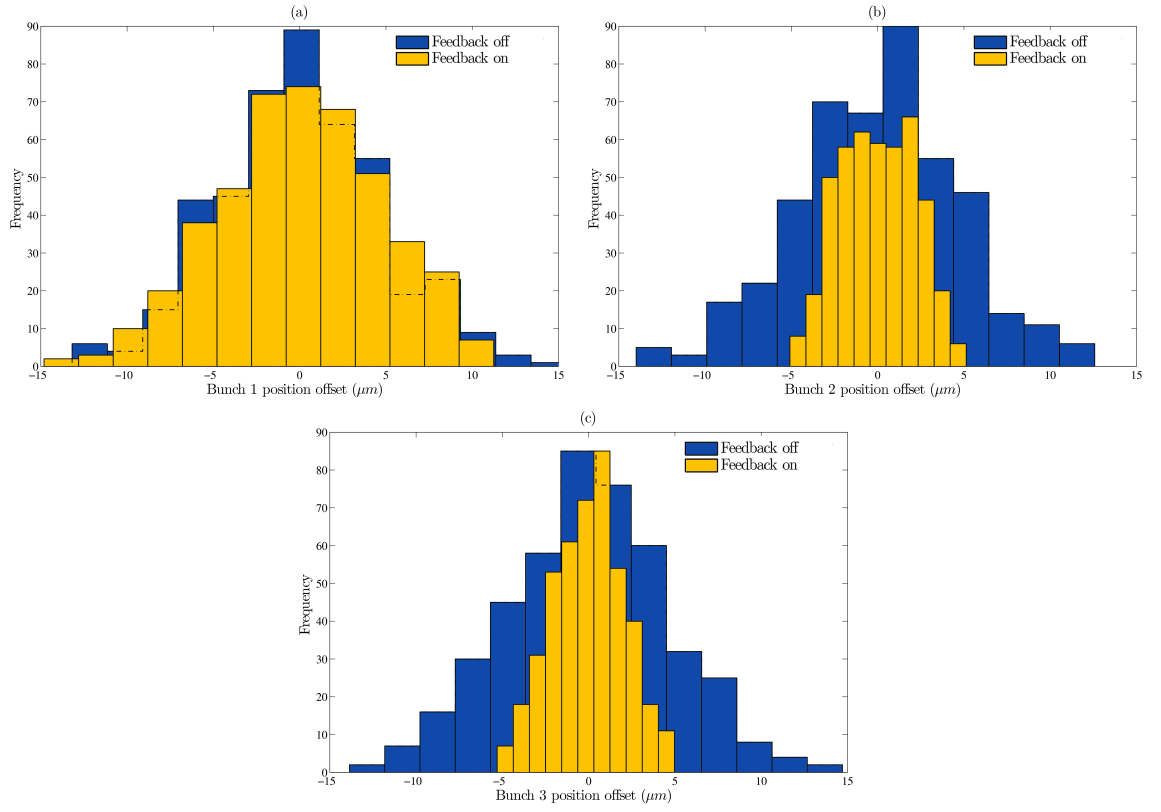


Figure 5.14: The effects of the coupled loop position and angle feedback on the jitter magnitude measured at P3 for bunches one (a), two (b) and three (c).




Figure 5.15: The effects of the coupled loop position and angle feedback on the bunch-to-bunch position correlations observed in measurements with P3.

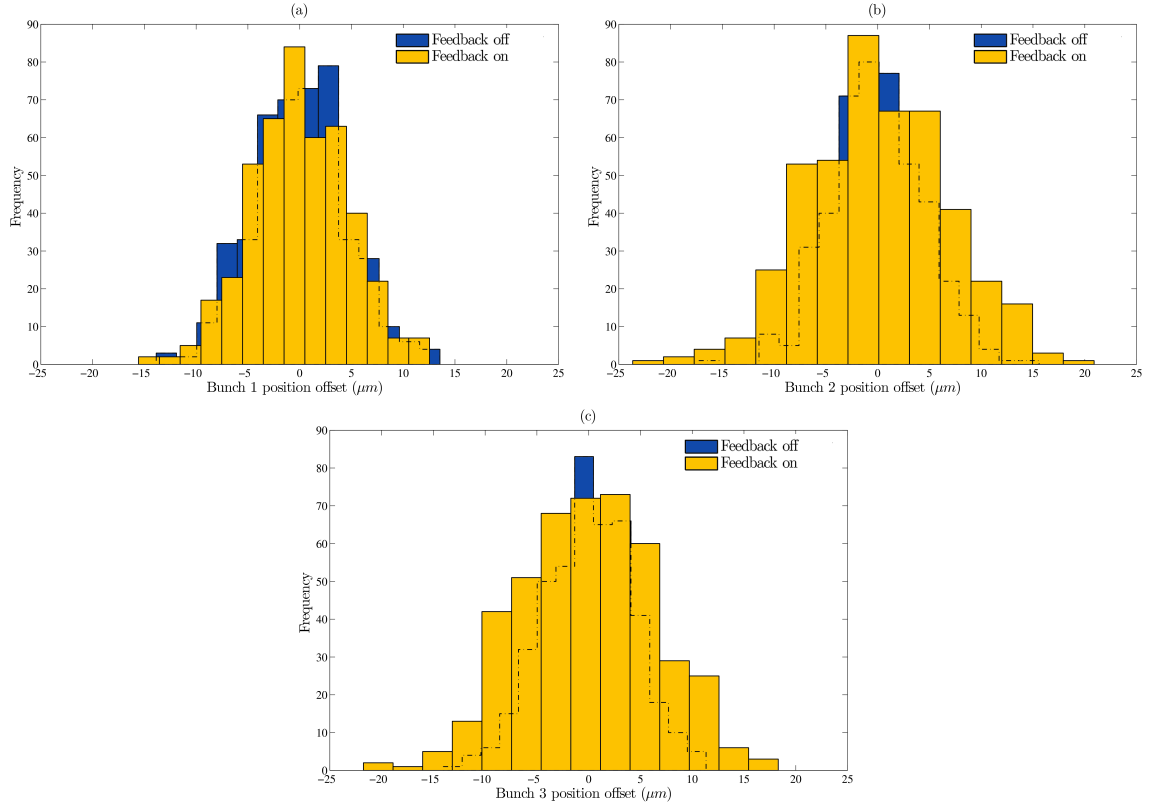


Figure 5.16: The effects of the coupled loop position and angle feedback on the jitter magnitude measured at the downstream witness BPM MQF15X for bunches one (a), two (b) and three (c).

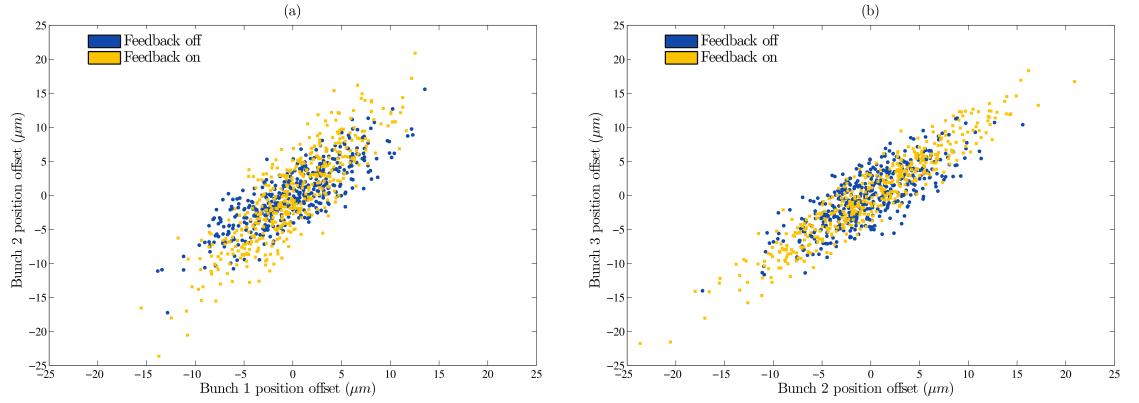


Figure 5.17: The effects of the coupled loop position and angle feedback on the bunch-to-bunch position correlations observed in measurements with the downstream witness BPM MQF15X.

Tables 5.7 - 5.9 summarise the properties of the measured distributions from P2, P3 and MQF15X.

Bunch	Feedback off	Feedback on	Bunches	Feedback off	Feedback on
2	3.2 μm	0.73 μm	1 to 2	97%	3.7%
3	3.5 μm	1.9 μm	2 to 3	80%	-62%

Table 5.7: Summary of the effects of two coupled feedback loops on the position jitter (left) and bunch-to-bunch correlations (right) as observed in measurements with P2.

Bunch	Feedback off	Feedback on	Bunches	Feedback off	Feedback on
2	4.7 μm	2.1 μm	1 to 2	89%	-19%
3	4.7 μm	2.1 μm	2 to 3	90%	-63%

Table 5.8: Summary of the effects of two coupled feedback loops on the position jitter (left) and bunch-to-bunch correlations (right) as observed in measurements with P3.

Bunch	Feedback off	Feedback on	Bunches	Feedback off	Feedback on
2	4.5 μm	6.7 μm	1 to 2	82%	85%
3	4.3 μm	6.6 μm	2 to 3	79%	95%

Table 5.9: Summary of the effects of two coupled feedback loops on the position jitter (left) and bunch-to-bunch correlations (right) as observed in measurements with MQF15X.

5.8 Effects of gain variation on feedback performance

Varying the operational gain of the feedback system provides both a valuable insight into the system’s operation, along with a powerful means of verifying the mathematical model outlined in section 5.5. The plots in this section show a comparison between measured beam jitter for various feedback gains and the response as predicted by the model. The steps in producing the model predictions are outlined below. Using a dataset of interleaved feedback on and off data at each gain setting:

- The three-bunch jitter covariance matrix at K1 is reconstructed using feedback off data from P2 and P3 (section 5.4).

- As outlined in section 5.4.2, the jitter covariance matrix is corrected for errors introduced by P2 and P3 resolutions. The intrinsic resolutions of P2 and P3 are taken to be the worst case limits from section 5.3.2, that is 0.45 μm and 1.5 μm respectively. The additional correlated component of P3's resolution is determined approximately by fitting the model to coupled loop feedback data. The result of this fit gives 3 μm .
- The average jitter covariance matrix at K1 over all datasets is transformed used to predict the effects of the feedback system according to the model in section 5.5. Expected jitter magnitudes and correlations are extracted from the predicted covariance matrix at BPMs of interest. This is done for a range of feedback gain settings.
- In order to compare the predicted statistical quantities with actual measurements, the effect of BPM resolution must be deconvolved from the measurements such that they reflect the true beam properties. The necessary deconvolutions are described in appendix B.

In the following plots the gain scale is normalised by the nominal gain. That is, the normalised gain $G_N = G/G_0$. The datapoints represent the BPM measurements of the beam jitter after deconvolution of the assumed resolution, and the error bars the standard errors in these measurements. Predictions of the model are shown by solid lines.

5.8.1 Scan of the P2-to-K1 loop gain

In the following scans of the gain of the single feedback loop between P2 and K1, good agreement is seen between the model predictions and observed performance. Due to the high initial bunch-to-bunch correlations and excellent P2 resolution, the optimal gain is very close to the nominal gain. Better agreement is observed for the second bunch than the third. This might be expected given that the model deals only approximately with the degradation of P3's resolution due to LO phase jitter, meaning any discrepancies will accumulate over successive bunches.




Figure 5.18: The change in bunch two position jitter magnitude and correlation with bunch one at the feedback BPM P2 vs. the single loop feedback gain. Solid lines give the model predictions.




Figure 5.19: The change in bunch three position jitter magnitude and correlation with bunch two at the feedback BPM P2 vs. the single loop feedback gain. Solid lines give the model predictions.

5.8.2 Scan of the P3-to-K2 loop gain

In this section, the gain of the single loop operating between P3 and K2 is varied. The agreement of the model and data is not as good as was the case for the P2-to-K1 loop, and again bunch two agrees better than bunch three. This is unsurprising given that the treatment of P3's correlated noise in the model is approximate. Due to the degraded performance of P3, the optimal gain is lower than the nominal. Even with the presence of correlated noise, a significant reduction in jitter is observed with lower gains.




Figure 5.20: The change in bunch two position jitter magnitude and correlation with bunch one at the feedback BPM P3 vs. the single loop feedback gain. Solid lines give the model predictions.




Figure 5.21: The change in bunch three position jitter magnitude and correlation with bunch two at the feedback BPM P3 vs. the single loop feedback gain. Solid lines give the model predictions.

5.8.3 Scans of the coupled loop gains

When scanning the gains of the coupled feedback system, one loop is varied independently whilst the other loop remains set nominally. Both gains of the varied loop are scanned in proportion. In addition, the downstream witness BPM MQF15X is monitored in preference to P1.




Figure 5.22: The change in position jitter magnitude of bunches two and three at P2 as the gains G_1 and G_2 of the coupled system are scanned together. Solid lines give the model predictions.

Figure 5.22 shows how varying the gains G_1 and G_2 of the K1 loop affects the observed and predicted corrections at P2. There is good agreement with the model, and jitter is reduced to around 1 μm with the gain at 90% of nominal. The gains of the K2 loop remain constant at their nominal.




Figure 5.23: The change in position jitter magnitude of bunches two and three at P3 as the gains G_3 and G_4 of the coupled system are scanned together. Solid lines give the model predictions.




Figure 5.24: The change in position jitter magnitude of bunches two and three at the witness BPM MQF15X as the gains G_3 and G_4 of the coupled system are scanned together. Solid lines give the model predictions.

The effect of scanning the gains G_3 and G_4 of the K2 loop as part of the coupled system are shown in figures 5.23 and 5.24, where in this case the K1 loop gains remain nominally set. As expected, the performance of this loop is impaired relative to K1 due to P3's additional correlated noise. However, a degree of jitter reduction is again observed in P3 for lower gains. Importantly, this jitter reduction is maintained in the downstream witness BPM MQF15X, albeit marginally. At nominal and higher gains the system is unable to reduce the jitter due to P3's correlated noise.

5.8.4 Expected future performance and extrapolation to the final focus

Given that the model has agreed well with the data of the previous sections, particularly in the cases involving the P2-to-K1 loop where there was no bunch-to-bunch correlated noise, it is possible to predict how well the FONT5 system should perform once P3's installation issues have been addressed. Using the data of section 5.8.3, the initial jitter covariance matrix is calculated again and used as input to the model. Figure 5.25 shows the results of the model for the well correlated bunch two (99%). The resolutions of both P2 and P3 are taken to be $0.45 \mu\text{m}$ with no correlated noise present. As before, the gains of the K1 loop are held at nominal and those of the K2 loop varied.




Figure 5.25: Predicted change in position jitter magnitude of bunch two at the witness BPM MQF15X (left) and the interaction point (right) as the gains G_3 and G_4 of the K2 loop of the coupled system are scanned together. Solid yellow and blue lines represent expectations with and without feedback respectively.

With both P2 and P3 operating with $0.45\text{ }\mu\text{m}$ resolution, the jitter in the extraction line will be reduced to sub-micron level as required for the ATF2 goal 1. Using transfer matrices from the ATF2 MAD model the jitter has been propagated to the interaction point after the final focus, giving the results shown in table 5.10. This is a linear extrapolation, and a full investigation would require a treatment of non-linear effects and misalignments. It shows however that, as expected, the FONT5 installation in the extraction line alone is not sufficient to maintain nanometre stability at the interaction point as required for the ATF goal 2. With the beam at the IP stable to below 10 nm however, it will be well within the operating range of an additional high-resolution cavity BPM based position feedback.

	MQF15X	IP
Feedback off	$2.6\text{ }\mu\text{m}$	20 nm
Feedback on	$0.7\text{ }\mu\text{m}$	4 nm

Table 5.10: Predicted optimal jitter reduction in the extraction line (MQF15X) and at the interaction point with both BPMs operating at $0.45\text{ }\mu\text{m}$ resolution on a 99% correlated beam.

Chapter 6

Conclusions

6.1 Summary

Making precise measurements of the new physical processes discovered at the Terascale will require a machine which, despite its size and complexity, is able to operate with appropriate, unprecedented precision and stability. As a leading contender to provide the necessary high energy lepton collisions, the ILC is a well studied and understood design, and many of the challenges inherent to its operation have been discussed in this thesis. In particular, the required beam stability in the presence of dynamic disturbances such as ground motion makes fast feedback and feed-forward systems essential. One specific technical challenge is found at the interaction point, where the two beams must be steered into collision with nanometre accuracy.

In order to demonstrate the feasibility of a feedback system capable of providing this level of stability at the ILC, the FONT group designed the digital, intra-train FONT5 system. By making use of the beam-beam deflection experienced by beams as they exit the interaction point, a large amplification of the beam offset is obtained. This allows nanometre-level separations to be measured with micron-resolution stripline BPMs, which coupled with low latency processing electronics ensures the system is fast enough to correct individual bunches within an ILC train.

After initial design work was complete, the three stripline BPMs and two stripline kickers of the prototype FONT5 system were installed at ATF2, along with their associated electronics. Low latency, flexible firmware for the FONT5 digital feedback board was designed and implemented using Verilog. During a number of ATF2 user shifts, consistent calibration constants were obtained for the BPMs using ATF2 corrector magnets, and for the kickers by driving them with constant currents. BPM resolution was shown to be below $2\text{ }\mu\text{m}$ by assuming three identical BPMs. However, analysis of the apparent position jitter due to local oscillator phase variation revealed that P3 was not as well behaved as the other

two BPMs. Observed bunch-to-bunch correlations in the BPMs allowed upper limits of $0.45\text{ }\mu\text{m}$, $0.45\text{ }\mu\text{m}$ and $1.50\text{ }\mu\text{m}$ to be placed on the intrinsic resolutions of P1, P2 and P3 respectively. An additional component to the resolution of P3 was found to be correlated bunch-to-bunch, and fitting witness BPM measurements to data taken with the feedback in operation later showed the magnitude of this component to be $3\text{ }\mu\text{m}$.

Measurements of the system latency were made for the two operating modes with lowest and highest latencies. In the fastest mode, a single loop between P2 and K1, the latency was measured to be 135 ns . In the worst-case mode, where two loops are coupled and the K1 kick depends also on the position in P3, a latency of 147 ns was determined.

The system was then allowed to feed back on the ATF beam over a number user shifts and in a number of modes of operation, with results as summarised over the following sections.

6.2 Demonstration of an ILC interaction point feedback system

As a technology demonstration of a digital, intra-train beam based feedback for the ILC interaction point, experiments with the ATF2 FONT5 installation have been a complete success. Position feedback at the ILC interaction point requires single degree-of-freedom correction and, as discussed in section 3.2.1, must achieve micron-level stability for a beam with an energy of $\sim 1\text{ GeV}$ such as ATF's. Using a single loop between P2 and K1, the measured RMS jitter at the feedback BPM P2 was stabilised for bunches two and three at $0.43\text{ }\mu\text{m}$ and $0.87\text{ }\mu\text{m}$ respectively.

This stability must be provided with a latency smaller than the ILC bunch spacing. Results with the FONT5 loop between P2 and K1 demonstrated that, for a well correlated bunch train, the RMS vertical jitter may be stabilised below 500 nm with a latency of just 135 ns . Not only is this fast enough for the nominal ILC bunch spacing of 369 ns , but FONT5 would still work in a true bunch-to-bunch sense were the ILC to operate with the proposed Low-N parameter set with a bunch spacing of just 184.5 ns .

Scaling a 500 nm correction from 1.6 GeV to the ILC's 250 GeV leads to an anticipated stability of 2.6 nm or $0.5\sigma_y$ at the ILC interaction point. As such, FONT5 is expected to be capable of keeping the luminosity loss at less than 10% , as per the requirements discussed in section 2.4. Indeed, this is a worst-case limit on the eventual performance at ILC. At ATF, the prototype FONT5 installation is limited by the resolution of its BPMs, whereas at the ILC interaction point the system will benefit from the large amplification provided by the beam-beam deflection.

6.3 Jitter reduction in two degrees of freedom

When a single loop operated between P3 and K2, the measured jitter appeared to stabilise at $1.8\text{ }\mu\text{m}$ and $1.7\text{ }\mu\text{m}$ for bunches two and three respectively, yet this apparent level was lower than in reality due to P3's correlated noise. Later investigation of the coupled loop system made use of the downstream witness BPM MQF15X to allow unbiased measurement of the jitter magnitude.

Initial measurements with the coupled loops, set to their nominal gains, showed that the feedback system increased the magnitude of the jitter at the downstream witness BPM MQF15X, despite reducing it at P2. Variation of the feedback gains allowed an investigation into the nature of this effect, and to aid this analysis a statistical model of the system was developed. The model is very pragmatic, in that the measured multi-bunch jitter covariance matrix characterises the beam sufficiently to predict the feedback system's response, regardless of whether any decorrelation arises from the extraction kicker, x-y coupling, or other factors.

This model was seen to describe well the response of the feedback to gain variation. As expected, lower gains performed better than the nominal due to P3's resolution. A marginal reduction in jitter was then obtained in MQF15X using the two coupled loops. The operation of the system is well understood, allowing predictions of the achievable jitter reduction in the ATF2 extraction line and interaction point to be made.

With appropriate modification to the FONT5 installation at ATF2, the two coupled loops P2-to-K1 and P3-to-K2 should achieve sub-micron jitter stability downstream of the feedback system. After demagnification by the ATF2 final focus, the beam is expected to be stable at around 4 nm RMS.

6.4 Outlook

Clearly the important next step for the FONT5 feedback system at ATF2 is to transition from an experimental to permanent installation, with care taken to avoid cable length mismatches. This work is currently being undertaken with a view to demonstrating the sub-micron stability at the final focus entrance of which the system is capable.

For the digital signal processing stage of the system, current development focusses on extending the FPGA firmware to deal with trains of more than three bunches, such as the 20 or 60 bunch trains expected at ATF2. This extension requires moving to LUT or pattern-based approaches to specifying the following:

- Position of bunches within a train
- Constant kicks for specific bunches enabling train flattening

- FIR coefficients for more bunches in the kicker pulse shape compensation

The FIR filter must be extended to cover several bunches, and an IIR filter implemented to deal with the tendency of the kicker pulse to decay exponentially. Also, the amount of data logged per pulse will necessarily increase, and a corresponding increase in data rate from the FONT5 board is required to retain pulse-by-pulse acquisition.

Finally, the software interface will require upgrading to deal with the longer data streams. As part of this upgrade the mover controls may be integrated with the DAQ. A fair amount of automation may also be implemented, including a slow software feedback for bunch train flattening, automatic BPM calibration using the movers and automatic kicker calibration. With these modifications the permanent FONT5 installation will be a turn-key system.

Forthcoming FONT group theses [82, 102] will document the next stage of the feedback system development.

Appendix A

Statistical modelling of intra-train feedback systems

A.1 Statistical transport of beam jitter

Where each beam pulse consists of a single bunch, it is possible to prepare a single ensemble of bunches over many pulses. The ensemble may be characterised by the covariance matrix, where it is assumed that both bunch position and angle are normally distributed. A given bunch has a coordinate vector \mathbf{y} :

$$\mathbf{y} = \begin{pmatrix} y \\ y' \end{pmatrix} \quad (\text{A.1.1})$$

which is transformed during transport from s_1 to s_2 by a standard lattice transfer matrix M :

$$\mathbf{y}(s_2) = M\mathbf{y}(s_1), \quad \text{where} \quad M = \begin{pmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{pmatrix} \quad (\text{A.1.2})$$

The covariance matrix is a function of s and given by:

$$\Sigma = \langle (\mathbf{y} - \langle \mathbf{y} \rangle)(\mathbf{y} - \langle \mathbf{y} \rangle)^\dagger \rangle = \begin{pmatrix} \Sigma_{11} & \Sigma_{11'} \\ \Sigma_{11'} & \Sigma_{1'1'} \end{pmatrix} \quad (\text{A.1.3})$$

Using the variances Σ_{11} and $\Sigma_{1'1'}$, the RMS position and angle jitter of the bunch are given by:

$$\sigma_y = \sqrt{\Sigma_{11}} \quad \text{and} \quad \sigma_{y'} = \sqrt{\Sigma_{1'1'}} \quad (\text{A.1.4})$$

In transporting the coordinate vector from s_1 to s_2 , the covariance matrix is transformed as follows:

$$\begin{aligned}\Sigma(s_2) &= \left\langle \left(M\mathbf{y}(s_1) - \langle M\mathbf{y}(s_1) \rangle \right) \left(M\mathbf{y}(s_1) - \langle M\mathbf{y}(s_1) \rangle \right)^\dagger \right\rangle \\ &= M\Sigma(s_1)M^\dagger\end{aligned}\tag{A.1.5}$$

Given that the transfer matrix \mathbf{M} is symplectic with $\det[\mathbf{M}] = \det[\mathbf{M}^\dagger] = 1$, then the determinant of the covariance matrix forms a conserved quantity under transport. By analogy with the beam emittance the following *jitter emittance* may be defined:

$$\varepsilon^{jitter} = \sqrt{\det(\Sigma)} = \sqrt{\Sigma_{11}\Sigma_{1'1'} - \Sigma_{11'}^2}\tag{A.1.6}$$

A.2 Treatment of multi-bunch trains

Where each beam pulse consists of multiple bunches, multiple ensembles of bunches may be prepared. The covariance matrix Σ_n characterises the ensemble of the n^{th} bunch prepared over many beam pulses. Given differences in extraction and instabilities experienced, it is not expected that each bunch will have identical distributions.

Whilst the covariance matrix of each bunch can be transported independently, the cross-correlation between bunches is of interest for feedback operation. In the case of two bunches, a vector \mathbf{y} can be written that describes the state of both bunches at s_1 :

$$\mathbf{y} = \begin{pmatrix} y_1 \\ y'_1 \\ y_2 \\ y'_2 \end{pmatrix}\tag{A.2.1}$$

The transport of the two bunches is described by the following 4×4 block matrix equation:

$$\mathbf{y}(s_2) = \mathbf{M}\mathbf{y}(s_1), \quad \text{where} \quad \mathbf{M} = \begin{pmatrix} M & 0 \\ 0 & M \end{pmatrix}\tag{A.2.2}$$

The two-bunch covariance matrix Σ is calculated from \mathbf{y} :

$$\Sigma = \langle (\mathbf{y} - \langle \mathbf{y} \rangle)(\mathbf{y} - \langle \mathbf{y} \rangle)^\dagger \rangle = \begin{pmatrix} \Sigma_1 & \bar{\Sigma} \\ \bar{\Sigma}^\dagger & \Sigma_2 \end{pmatrix}\tag{A.2.3}$$

where $\bar{\Sigma}$ is used to denote the bunch-to-bunch covariances:

$$\bar{\Sigma} = \begin{pmatrix} \Sigma_{12} & \Sigma_{12'} \\ \Sigma_{1'2} & \Sigma_{1'2'} \end{pmatrix} \quad (\text{A.2.4})$$

As in the single bunch case, the two-bunch covariance matrix transforms under linear transport as:

$$\Sigma(s_2) = M\Sigma(s_1)M^\dagger \quad (\text{A.2.5})$$

The independent jitter emittances of the two bunches, $\sqrt{\det(\Sigma_1)}$ and $\sqrt{\det(\Sigma_2)}$, are conserved, and since:

$$\det(M) = \det(M) \times \det(M) = 1 \quad (\text{A.2.6})$$

it is possible to identify $E = \det(\Sigma)$ as another conserved quantity.

Note that extending to an arbitrary number of bunches n is trivial: the state vector \mathbf{y} becomes a $2n$ -dimensional column vector, the transport matrix M becomes a $2n \times 2n$ block-diagonal matrix and the covariance matrix Σ becomes a $2n \times 2n$ symmetric matrix with the individual bunch covariance matrices on the main diagonal.

A.3 Feedback transformations - the two-bunch case

In the case of a fast position feedback system where a kicker imparts an angular deflection to the second bunch in a train, the deflection being proportional to the signal from the first bunch in a downstream BPM, the second bunch angle at the kicker location s_K transforms as:

$$y'(s_K) \longrightarrow \tilde{y}'(s_K) = y'(s_K) + Gy(s_{BPM}) \quad (\text{A.3.1})$$

for a gain factor G . If A is the 2×2 transfer matrix between kicker and BPM with elements a_{ij} , then the two-bunch state vector leaving the kicker is transformed by F_A :

$$\tilde{\mathbf{y}}(s_K) = F_A \mathbf{y}(s_K), \quad \text{where} \quad F_A = \begin{pmatrix} 1 & 0 \\ F_A & 1 \end{pmatrix} \quad \text{and} \quad F_A = \begin{pmatrix} 0 & 0 \\ Ga_{11} & Ga_{12} \end{pmatrix} \quad (\text{A.3.2})$$

As before, it is straightforward to note that the effect of such a transformation on the covariance matrix is:

$$\tilde{\Sigma}(s_K) = F_A \Sigma(s_K) F_A^\dagger \quad (\text{A.3.3})$$

The unitriangular matrix F_A has determinant 1, and thus $E = \det(\Sigma)$ is conserved under a feedback transformation.

A.4 Generalisation of feedback transformations to n bunches

A general n -bunch state vector transforms as:

$$\tilde{\mathbf{y}}(s_K) = \mathbf{F}_A \mathbf{y}(s_K) \quad \text{where} \quad \mathbf{y}(s_K) = \begin{pmatrix} y_1 \\ y'_1 \\ y_2 \\ y'_2 \\ y_3 \\ y'_3 \\ \vdots \\ y_{n-1} \\ y'_{n-1} \\ y_n \\ y'_n \end{pmatrix} \quad (\text{A.4.1})$$

Taking into account the effect of the feedback delay loop, the n^{th} bunch is transformed as:

$$y'_n(s_K) \longrightarrow \tilde{y}'_n(s_K) = y'_n(s_K) + G\tilde{y}_{n-1}(s_{BPM}) + \delta y_{n-1} \quad (\text{A.4.2})$$

where δy_{n-1} is the correction applied to the $(n-1)^{th}$ bunch. This may be written in matrix form as:

$$\mathbf{F}_A = \begin{pmatrix} 1 & 0 & 0 & 0 & \dots & 0 & 0 \\ F_A^1 & 1 & 0 & 0 & \dots & 0 & 0 \\ F_A^2 & F_A^1 & 1 & 0 & \dots & 0 & 0 \\ F_A^3 & F_A^2 & F_A^1 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ F_A^{n-2} & F_A^{n-3} & F_A^{n-4} & F_A^{n-5} & \dots & 1 & 0 \\ F_A^{n-1} & F_A^{n-2} & F_A^{n-3} & F_A^{n-4} & \dots & F_A^1 & 1 \end{pmatrix} \quad (\text{A.4.3})$$

where:

$$F_A^p \equiv \begin{pmatrix} 0 & 0 \\ Ga_{11}(1 + Ga_{12})^{p-1} & Ga_{12}(1 + Ga_{12})^{p-1} \end{pmatrix} \quad \text{with} \quad p \geq 1 \quad (\text{A.4.4})$$

The feedback transformation matrix F_A stays unitriangular for any number of bunches, conserving $E = \det(\Sigma)$ for arbitrarily long bunch trains.

A.5 Modelling of a single intra-train feedback loop

Let a single loop feedback system consist of a kicker at s_K upstream of a feedback BPM at s_{BPM} . Let A be the standard 2×2 lattice transfer matrix between them. For an n -bunch beam let the $2n$ -dimensional state vector be denoted $\mathbf{y}(s)$, the $2n \times 2n$ block diagonal transfer matrix as A and the $2n \times 2n$ feedback transformation vector at s_K as F_A . Then, noting that A is invertible, the corrected state vector becomes:

$$\tilde{\mathbf{y}}(s_{BPM}) = A\tilde{\mathbf{y}}(s_K) = AF_A\mathbf{y}(s_K) = AF_AA^{-1}\mathbf{y}(s_{BPM}) = T_{BPM}\mathbf{y}(s_{BPM}) \quad (\text{A.5.1})$$

where $T_{BPM} = AF_AA^{-1}$ is the transformation matrix representing the feedback system at the feedback BPM. The covariance matrix transforms as:

$$\tilde{\Sigma}(s_{BPM}) = T_{BPM}\Sigma(s_{BPM})T_{BPM}^\dagger \quad (\text{A.5.2})$$

It is then possible to extract the 2×2 covariance matrices from $\tilde{\Sigma}(s_{BPM})$. The corrected covariance matrix of the n^{th} bunch, $\tilde{\Sigma}_n(s_{BPM})$ is simply:

$$\tilde{\Sigma}_n(s_{BPM}) = \begin{pmatrix} \left[\tilde{\Sigma}(s_{BPM}) \right]_{2n-1, 2n-1} & \left[\tilde{\Sigma}(s_{BPM}) \right]_{2n-1, 2n} \\ \left[\tilde{\Sigma}(s_{BPM}) \right]_{2n, 2n-1} & \left[\tilde{\Sigma}(s_{BPM}) \right]_{2n, 2n} \end{pmatrix} \quad (\text{A.5.3})$$

and the jitter emittance of the n^{th} bunch is reduced by a factor κ_n :

$$\kappa_n = \sqrt{\frac{\det(\tilde{\Sigma}_n)}{\det(\Sigma_n)}}, \quad \tilde{\varepsilon}_n^{jitter} = \kappa_n \varepsilon_n^{jitter} \quad (\text{A.5.4})$$

Indeed, given that:

$$\sigma_y(s) = \sqrt{\beta(s) \times \varepsilon^{jitter}}, \quad \text{and} \quad \sigma_{y'}(s) = \sqrt{\gamma(s) \times \varepsilon^{jitter}} \quad (\text{A.5.5})$$

where $\gamma(s)$ and $\beta(s)$ are Twiss parameters, then the conserved jitter emittance of the n^{th} bunch may be used to predict the quality of the position and angle correction at an arbitrary position s downstream of the feedback system. It immediately follows that for all downstream s :

$$\frac{\tilde{\sigma}_{y,n}(s)}{\sigma_{y,n}(s)} = \frac{\tilde{\sigma}_{y',n}(s)}{\sigma_{y',n}(s)} = \sqrt{\kappa_n} \quad (\text{A.5.6})$$

A.6 Modelling of coupled intra-train feedback loops

It is now fairly straightforward to extend the analysis to coupled feedback systems. Firstly, where the kick from a kicker is a linear combination of two BPM signals and the BPMs are separated from the kicker by transfer matrices A and B , equation (A.4.4) becomes:

$$F_{AB}^p \equiv \begin{pmatrix} 0 & 0 \\ (G_1 a_{11} + G_2 b_{11})(G_1 a_{12} + G_2 b_{12} + 1)^{p-1} & (G_1 a_{12} + G_2 b_{12})(G_1 a_{12} + G_2 b_{12} + 1)^{p-1} \end{pmatrix}$$

with $p \geq 1$

(A.6.1)

for the gain factors G_1 and G_2 . Equation (A.4.3) for the feedback transformation then remains unchanged. In the coupled system we have two kickers, K1 and K2 at s_{K1} and s_{K2} . Let A and B be the $2n \times 2n$ block diagonal transfer matrices from K1 to the two downstream feedback BPMs at s_{P2} and s_{P3} . Similarly let C and D be the transfer matrices from K2 to the same BPMs.

The $2n \times 2n$ feedback transformation matrices for K1 and K2, based on linear combinations of the signals from the two BPMs, are denoted F_{AB} and F_{CD} respectively. For the $2n$ -dimensional state vector $\mathbf{y}(s)$ therefore:

$$\mathbf{y}(s_{K2}) = C^{-1} A \mathbf{y}(s_{K1})$$
(A.6.2)

and the corrected vector at the downstream BPM P3 is:

$$\begin{aligned} \tilde{\mathbf{y}}(s_{P3}) &= D \tilde{\mathbf{y}}(s_{K2}) = D F_{CD} C^{-1} A \tilde{\mathbf{y}}(s_{K1}) = D F_{CD} C^{-1} A F_{AB} \mathbf{y}(s_{K1}) \\ &= D F_{CD} C^{-1} A F_{AB} B^{-1} \mathbf{y}(s_{P3}) \end{aligned}$$
(A.6.3)

This then is T_{P3} , the transformation matrix for the coupled feedback system at the downstream feedback BPM, $T_{P3} = D F_{CD} C^{-1} A F_{AB} B^{-1}$. This matrix may be used in equation (A.5.2) to obtain the corrected covariance matrix and hence the factor κ .

A.7 Treatment of BPM resolution in a single feedback loop

As previously noted, by far the dominant contribution to the overall resolution of the feedback system is the position resolution of the feedback BPM measurements. Equation A.4.2

for single loop feedback becomes, in the presence of imperfect BPM readings:

$$y'_n(s_K) \longrightarrow \tilde{y}'_n(s_K) = y'_n(s_K) + G\tilde{y}_{n-1}^\star(s_{BPM}) + \Delta y_{n-1} \quad (\text{A.7.1})$$

where the measured position, denoted with a star, is:

$$\tilde{y}_{n-1}^\star(s_{BPM}) = \tilde{y}_{n-1}(s_{BPM}) + r_{n-1}; \quad (\text{A.7.2})$$

The random variable r_i is drawn from a population that over many pulses is normally distributed with standard deviation equal to the feedback BPM's resolution. It is assumed that the BPM resolution does not vary from bunch to bunch, such that:

$$\text{Std}[r_i] = R_i = R \quad (\text{A.7.3})$$

By defining the following n -element resolution vector:

$$\mathbf{r} = \begin{pmatrix} r_1 \\ r_2 \\ r_3 \\ \vdots \\ r_{n-1} \\ r_n \end{pmatrix} \quad (\text{A.7.4})$$

along with a $2n \times n$ matrix closely related to the feedback transformation matrix:

$$\mathbf{H}_A = \begin{pmatrix} 0 & 0 & 0 & 0 & \dots & 0 & 0 \\ H_A^1 & 0 & 0 & 0 & \dots & 0 & 0 \\ H_A^2 & H_A^1 & 0 & 0 & \dots & 0 & 0 \\ H_A^3 & H_A^2 & H_A^1 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ H_A^{n-2} & H_A^{n-3} & H_A^{n-4} & H_A^{n-5} & \dots & 0 & 0 \\ H_A^{n-1} & H_A^{n-2} & H_A^{n-3} & H_A^{n-4} & \dots & H_A^1 & 0 \end{pmatrix} \quad (\text{A.7.5})$$

where:

$$H_A^p \equiv \begin{pmatrix} 0 \\ G(1 + Ga_{12})^{p-1} \end{pmatrix} \quad \text{with } p \geq 1 \quad (\text{A.7.6})$$

it is then possible to write the effects of feedback, including resolution, in matrix form as:

$$\tilde{\mathbf{y}}(s_K) = \mathbf{F}_A \mathbf{y}(s_K) + \mathbf{H}_A \mathbf{r} \quad (\text{A.7.7})$$

Using the feedback transformation matrix for a single loop (equation A.5.1):

$$\tilde{\mathbf{y}}(s_{BPM}) = \mathbf{T}_{BPM} \mathbf{y}(s_{BPM}) + \mathbf{A} \mathbf{H}_A \mathbf{r} \quad (\text{A.7.8})$$

It is worth noting here that for the i^{th} bunch, the angular error ϵ_i in the correction due to resolution may be written as:

$$\epsilon_i = [\mathbf{H}_A \mathbf{r}]_{2i} = \sum_{l=0}^{i-2} G(1 + G a_{12})^l r_l \quad (\text{A.7.9})$$

and that if the gain of the feedback system is set to its nominal value $G_0 = -1/a_{12}$, then the magnitude of the correction error is the same for all bunches. That is, the feedback system does not compound the correction error for subsequent bunches. This may be seen by noting that for uncorrelated bunch-to-bunch errors:

$$\begin{aligned} \text{Var}[\epsilon_i] &= G_0^2 \sum_{l=0}^{i-2} (1 + G_0 a_{12})^{2l} \text{Var}[r_l] \\ &= G_0^2 R^2 \left[\frac{(1 + G_0 a_{12})^{2(i-1)} - 1}{2G_0 a_{12} + G_0^2 a_{12}^2} \right] \\ &= \frac{R^2}{a_{12}^2} \quad \forall i > 1 \end{aligned} \quad (\text{A.7.10})$$

Now, assuming that resolution is independent of position, the feedback system causes the jitter covariance matrix at the feedback BPM to transform as:

$$\tilde{\Sigma}(s_{BPM}) = \mathbf{T}_{BPM} \Sigma(s_{BPM}) \mathbf{T}_{BPM}^\dagger + \mathbf{A} \mathbf{H}_A \Sigma_R (\mathbf{A} \mathbf{H}_A)^\dagger \quad (\text{A.7.11})$$

In the case that the resolution contributions to measurements of different bunches are uncorrelated, the $n \times n$ resolution covariance matrix is diagonal:

$$\Sigma_R = \langle \mathbf{r} \mathbf{r}^\dagger \rangle = \begin{pmatrix} R^2 & 0 & \dots & 0 \\ 0 & R^2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & R^2 \end{pmatrix} \quad (\text{A.7.12})$$

A.8 Treatment of BPM resolution in coupled feedback loops

Coupling two feedback loops complicates the inclusion of resolution effects. In particular, both K1 and K2 kicks are based upon the same readings from the two feedback BPMs, with identical measurement errors affecting each kick. To account for this correlation between the errors in each kick, the combined error from both kickers is applied at exit of the downstream K2. Using the FONT5 convention, the two BPMs are labelled P2 and P3. Since the coupled loops act on both degrees of freedom, the resolution vector of equation A.7.4 is extended to $2n$ elements:

$$\mathbf{r} = \begin{pmatrix} r_1^{P2} \\ r_1^{P3} \\ r_2^{P2} \\ r_2^{P3} \\ \vdots \\ r_n^{P2} \\ r_n^{P3} \end{pmatrix} \quad (\text{A.8.1})$$

while the resolution transformation matrix of equation A.7.5 becomes a $2n \times 2n$ matrix by setting:

$$H_{AB}^p = \begin{pmatrix} 0 & 0 \\ G_1(1 + G_1 a_{12} + G_2 b_{12})^{p-1} & G_2(1 + G_1 a_{12} + G_2 b_{12})^{p-1} \end{pmatrix} \quad \text{with } p \geq 1 \quad (\text{A.8.2})$$

Keeping with the notation of section A.6, the matrix equation describing two coupled feedback loops and including BPM resolution effects is:

$$\begin{aligned} \tilde{\mathbf{y}}(s_{K2}) &= \mathbf{F}_{CD} \mathbf{C}^{-1} \mathbf{A} \tilde{\mathbf{y}}(s_{K1}) + \mathbf{H}_{CD} \mathbf{r} \\ &= \mathbf{F}_{CD} \mathbf{C}^{-1} \mathbf{A} [\mathbf{F}_{AB} \mathbf{y}(s_{K1}) + \mathbf{H}_{AB} \mathbf{r}] + \mathbf{H}_{CD} \mathbf{r} \\ &= \mathbf{F}_{CD} \mathbf{C}^{-1} \mathbf{A} \mathbf{F}_{AB} \mathbf{y}(s_{K1}) + [\mathbf{F}_{CD} \mathbf{C}^{-1} \mathbf{A} \mathbf{F}_{AB} \mathbf{H}_{AB} + \mathbf{H}_{CD}] \mathbf{r} \end{aligned} \quad (\text{A.8.3})$$

and again, it is possible to rewrite the transformation at the downstream feedback BPM in terms of the feedback transformation matrix \mathbf{T}_{P3} :

$$\begin{aligned} \tilde{\mathbf{y}}(s_{P3}) &= \mathbf{D} \mathbf{F}_{CD} \mathbf{C}^{-1} \mathbf{A} \mathbf{F}_{AB} \mathbf{B}^{-1} \mathbf{y}(s_{P3}) + \mathbf{D} [\mathbf{F}_{CD} \mathbf{C}^{-1} \mathbf{A} \mathbf{F}_{AB} \mathbf{H}_{AB} + \mathbf{H}_{CD}] \mathbf{r} \\ &= \mathbf{T}_{P3} \mathbf{y}(s_{P3}) + \mathbf{D} [\mathbf{F}_{CD} \mathbf{C}^{-1} \mathbf{A} \mathbf{F}_{AB} \mathbf{H}_{AB} + \mathbf{H}_{CD}] \mathbf{r} \end{aligned} \quad (\text{A.8.4})$$

Finally, as before, the jitter covariance matrix at the downstream feedback BPM transforms as:

$$\tilde{\Sigma}(s_{P3}) = T_{P3}\Sigma(s_{P3})T_{P3}^\dagger + D \left[F_{CD}C^{-1}AF_{AB}H_{AB} + H_{CD} \right] \Sigma_R \left[F_{CD}C^{-1}AF_{AB}H_{AB} + H_{CD} \right]^\dagger D^\dagger \quad (\text{A.8.5})$$

and in the case that the resolution contributions to measurements of different bunches are uncorrelated, the now $2n \times 2n$ resolution covariance matrix is again diagonal:

$$\Sigma_R = \langle \mathbf{r} \mathbf{r}^\dagger \rangle = \begin{pmatrix} (R^{P2})^2 & 0 & \dots & 0 & 0 \\ 0 & (R^{P3})^2 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & (R^{P2})^2 & 0 \\ 0 & 0 & \dots & 0 & (R^{P3})^2 \end{pmatrix} \quad (\text{A.8.6})$$

Appendix B

Deconvolution of resolution effects from BPM measurements

For BPM measurements to be consistent with the true beam position jitters as predicted by the feedback models of appendix A, the intrinsic resolution of the BPM must be deconvolved. This is the case whenever an RMS jitter or a bunch-to-bunch jitter correlation coefficient is measured. The simplest case is for a witness BPM such as P1 or MQF15X, that is a BPM which does not form part of the feedback system. In this case the deconvolution of the resolution from the measured RMS jitter is trivial. Using V_i for the variance of the i^{th} bunch's jitter:

$$V_i = V_i^* - R^2 \quad (\text{B.0.1})$$

and as shown in section 5.3.2, the resolution may be deconvolved from a measured correlation coefficient:

$$\rho_{i-1,i} = \rho_{i-1,i}^* \sqrt{\frac{V_{i-1}^* V_i^*}{(V_{i-1}^* - R^2)(V_i^* - R^2)}} \quad (\text{B.0.2})$$

The situation becomes a little more complicated in the case of measurements at a feedback BPM. Equation B.0.1 is still valid for the RMS jitter, and when no feedback is in operation equation B.0.2 is valid for the jitter correlation between two adjacent bunches. However, when feedback is operating it introduces a correlation between the measurement error of the earlier bunch and the position of the subsequent bunch which must be accounted for. Taking the case of a single feedback loop, the covariance between the corrected positions of two adjacent bunches as measured by the feedback BPM is:

$$\begin{aligned} \text{Cov}[\tilde{y}_{i-1}^*, \tilde{y}_i^*] &= \text{Cov}[\tilde{y}_{i-1} + r_{i-1}, \tilde{y}_i + r_i] \\ &= \text{Cov}[\tilde{y}_{i-1}, \tilde{y}_i] + \text{Cov}[r_{i-1}, \tilde{y}_i] \\ &= \text{Cov}[\tilde{y}_{i-1}, \tilde{y}_i] + G_N \text{Cov}[r_{i-1}, r_{i-1}] \end{aligned} \quad (\text{B.0.3})$$

for normalised feedback gain G_N . Thus the deconvolution becomes:

$$\rho_{i-1,i} = \rho_{i-1,i}^* \left[\frac{\text{Cov}[\tilde{y}_{i-1}^*, \tilde{y}_i^*] - G_N R^2}{\text{Cov}[\tilde{y}_{i-1}^*, \tilde{y}_i^*]} \right] \sqrt{\frac{V_{i-1}^* V_i^*}{(V_{i-1}^* - R^2)(V_i^* - R^2)}} \quad (\text{B.0.4})$$

Further complication arises when the feedback BPM is plagued by noise which is correlated bunch-to-bunch, as is the case for P3 (see section 5.3.1). Let the measurement error for a given pulse be denoted by the random variable p , which over many pulses is normally distributed with standard deviation P . The measurement error p in the feedback BPM is the same for all bunches in a given pulse. Then the measured variance of the corrected position of the i^{th} bunch is given by:

$$\begin{aligned} \tilde{V}_i^* &= \text{Var}[\tilde{y}_i + r_i + p] \\ &= \tilde{V}_i + R^2 + P^2 + 2\text{Cov}[p, \tilde{y}_i] \\ &= \tilde{V}_i + R^2 + P^2 + 2\text{Cov} \left[p, \sum_{l=0}^{i-2} G_N (1 + G_N)^l p \right] \end{aligned} \quad (\text{B.0.5})$$

and the required deconvolution is:

$$\tilde{V}_i = \tilde{V}_i^* - R^2 - \left[1 + 2G_N \sum_{l=0}^{i-2} (1 + G_N)^l \right] P^2 \quad (\text{B.0.6})$$

Bibliography

- [1] M. Peskin and D. Schroeder, *An introduction to quantum field theory* (Westview press, 1995).
- [2] F. Halzen and A. Martin, *Quarks and leptons: an introductory course in modern particle physics* (Wiley, 1984).
- [3] The LEP collaborations and the LEP electroweak working group, “A combination of preliminary electroweak measurements and constraints on the standard model,” hep-ex/0612034 (2006).
- [4] “Fermilab experiments narrow allowed mass range for higgs boson,” http://www.fnal.gov/pub/presspass/press_releases/Higgs-mass-constraints-20100726.html (2011).
- [5] CERN, “LHC machine outreach,” <http://lhc-machine-outreach.web.cern.ch/lhc-machine-outreach/> (2011).
- [6] ILC collaboration, “International linear collider reference design report volume 2: Physics at the ILC,” <http://www.linearcollider.org/about/Publications/Reference-Design-Report/> (2007).
- [7] SNO Collaboration, “Low-energy-threshold analysis of the phase I and phase II data sets of the Sudbury Neutrino Observatory,” Phys. Rev. C **81**, 055504 (2010).
- [8] The Super-Kamiokande Collaboration, “Atmospheric neutrino oscillation analysis with subleading effects in Super-Kamiokande I, II, and III,” Phys. Rev. D **81**, 092004 (2010).
- [9] A. Habig, “MINOS neutrino oscillation results,” Phys. Lett. A **25**, 1219 (2010).
- [10] G. Jungman, M. Kamionkowski, and K. Griest, “Supersymmetric dark matter,” Phys. Rep. **267**, 195 (1996).
- [11] S. Martin, “A supersymmetry primer,” hep-ph/9709356 (2008).

- [12] L. Langacker and M. Luo, “Implications of precision electroweak experiments for m_t , ρ_0 , $\sin \theta_w$, and grand unification,” *Phys. Rev. D.* **44**, 817 (1991).
- [13] CMS collaboration, “CMS physics (TDR vol. II),” CERN/LHCC 2006-021 (2006).
- [14] ATLAS collaboration, “ATLAS detector and physics performance (TDR vol. II),” CERN/LHCC 99-15 (1999).
- [15] R. -D. Heuer, “Roll of accelerators in ‘dark world’,” presented at the ICFA seminar, Korea (2005).
- [16] D. G. Charlton, “LEP, SLC and the Standard Model,” in *XXXth SLAC summer institute topical conference* (2002).
- [17] Fermilab, “Internal Notes of the E288 Upsilon-Discovery experiment,” http://history.fnal.gov/jyoh_docs/e288_internal_notes.html (2011).
- [18] H. Tajima and the Belle collaboration, “Belle B physics results,” *Int. J. Mod. Phys. A* **17**, 2967 (2002).
- [19] BaBar collaboration, “BaBar physics book,” SLAC-R-504 (1998).
- [20] “ICFA statement on linear colliders,” <http://www.fnal.gov/directorate/icfa/>. (2004).
- [21] S. Y. Lee, *Accelerator Physics*, 2nd ed. (World Scientific Publishing Co., 2004).
- [22] E. J. N. Wilson, “A review of accelerators beyond 2000,” *Physica Scripta* **T30**, 69 (1990).
- [23] V. Shiltsev, “When will we know a muon collider is feasible? status and directions of muon accelerator R&D,” *Mod. Phys. Lett. A* **25**, 567 (2010).
- [24] “International muon ionization cooling experiment,” <http://www.mice.iit.edu/>. (2011).
- [25] “International Technology Review Panel executive summary,” <http://www.interactions.org/pdf/ITRPexec.pdf>. (2004).
- [26] ILC collaboration, “International linear collider reference design report volume 3: Accelerator,” <http://www.linearcollider.org/about/Publications/Reference-Design-Report/> (2007).
- [27] “ILC Newslines 10th February 2011,” <http://newsline.linearcollider.org/2011/02/10/>. (2011).

- [28] The CLIC study team, “A 3 TeV e^+e^- collider based on CLIC technology,” CERN 2008-008 (2000).
- [29] J. P. Delahaye, “CLIC-ILC collaboration,” presented at LCWS08 Chicago. http://clic-study.web.cern.ch/CLIC-Study/CLIC_ILC_Collab_Mtg/JPD_ILC08_Nov08.pdf (2008).
- [30] A. Wolski, “The accelerator hamiltonian in a curved coordinate system,” Lecture for the Cockcroft Institute, Daresbury, UK. <http://pcwww.liv.ac.uk/~awolski/Teaching/Cockcroft/LinearDynamics/LinearDynamics-Lecture3.pdf> (2006).
- [31] E. Courant and H. Snyder, “Theory of the alternating gradient synchrotron,” *Annals of Physics* **281**, 360 (1957).
- [32] E. J. N. Wilson, *An Introduction to Particle Accelerators* (Oxford University Press, 2001).
- [33] J. Buon, “Beam phase space and emittance,” CERN Accelerator School: 5th General Accelerator Physics Course (1992).
- [34] K. Brown and R. Servranckx, “First- and second-order charge particle optics,” SLAC-PUB-3381 (1984).
- [35] M. Sands, “The physics of electron storage rings: an introduction,” SLAC-R-121 (1970).
- [36] L. C. Teng, “Minimizing the emittance in designing the lattice of an electron storage ring,” SLAC-R-121 (1984).
- [37] T. Naito *et al.*, “Beam extraction using strip-line kicker at KEK-ATF,” in *Proceedings of PAC95* (1995) pp. 1–3.
- [38] D. Proch, “The TESLA cavity: Design considerations and RF properties,” in *Proceedings of the Sixth Workshop on RF Superconductivity* (October 1993) pp. 382–397.
- [39] K. Sennyu *et al.*, “Status of superconducting cavity development for ILC at MHI,” in *Proceedings of IPAC10* (2010) pp. 463–465.
- [40] L. Lilje, “Superconducting RF R&D for the ILC,” presented at PAC2007. http://accelconf.web.cern.ch/accelconf/p07/TALKS/THXKI01_TALK.PDF (2007).
- [41] P. Raimondi and A. Seryi, “A novel final focus design for future linear colliders,” *Phys. Rev. Lett.* **86**, 3779 (2001).

- [42] B. Parker *et al.*, “Challenges and concepts for design of an interaction region with push-pull arrangement of detectors - an interface document,” in *Proceedings of EPAC08* (2008) pp. 616–618.
- [43] International Detector Advisory Group, “IDAG report on the validation of letters of intent for ILC detectors,” http://ilcdoc.linearcollider.org/record/23970/files/IDAG_report_090816.pdf (2009).
- [44] H. Aihara, P. N. Burrows and M. Oreglia, “SiD Letter of Intent,” <https://www.silicondetector.org/download/attachments/46170132/SiliconDetectorLetterOfIntent.pdf> (2009).
- [45] ILD Concept Group, “International Large Detector Letter of Intent,” http://arxiv.org/PS_cache/arxiv/pdf/1006/1006.3396v1.pdf (2010).
- [46] Y. Nosochkov *et al.*, “ILC extraction line for 14 mrad crossing angle,” in *Proceedings of the 36th ICFA Advanced Beam Dynamics Workshop* (2005).
- [47] B. Muratori, “Luminosity in the presence of offsets and a crossing angle,” CERN AB-Note-2003-026 (2003).
- [48] R. Assman, “LEP luminosity revisited: design and reality,” in *Proceedings of the Second Asian Particle Accelerator Conference* (2001) pp. 74–78.
- [49] M. Furman and M. Zisman, “Luminosity,” in *Handbook of Accelerator Physics and Engineering*, edited by A. Chao and M. Tigner (World Scientific Publishing Co. Pte. Ltd., 1999) pp. 277–281.
- [50] P. Chen, “Beam-beam effects in linear colliders,” in *Handbook of Accelerator Physics and Engineering*, edited by A. Chao and M. Tigner (World Scientific Publishing Co. Pte. Ltd., 1999) pp. 159–162.
- [51] K. Yokoya and P. Chen, “Beam-beam phenomena in linear colliders,” in *Frontiers of Particle Beams: Intensity Limitations* (Springer Verlag, 1990) pp. 415–445.
- [52] F. Zimmerman, “Accelerator physics and technologies for linear colliders,” Lecture notes for the University of Chicago. http://hep.uchicago.edu/~kwangje/LectureNotes_Zimmermann.pdf (2002).
- [53] J. Brossard *et al.*, “Evaluation of luminosity reduction in the ILC head-on scheme from parasitic collisions,” in *EUROTeV-Report-2007-052* (2007).
- [54] J. Smith *et al.*, “Comparison of beam-based alignment algorithms for the ILC,” in *Proceedings of PAC05* (2005) pp. 1847–1849.

- [55] A. Seryi *et al.*, “Issues of stability and ground motion in ILC,” SLAC-PUB-11661 (2006).
- [56] J. Resta-Lopez, “Beam dynamics simulations for linear colliders: CLIC & ILC,” seminar at the John Adams Institute. <http://www.physics.ox.ac.uk/pp/seminars/> (2007).
- [57] J. Resta-Lopez, “Luminosity performance studies of linear colliders with intra-train feedback systems,” arXiv:0902.2915v1 (2008).
- [58] R. Amirkas *et al.*, “Ground motion & comparison of various sites,” EUROTeV-Report-2005-023-1 (2005).
- [59] S. Shin *et al.*, “Design of a low-Q S-band cavity beam position monitor,” J. Korean. Phys. Soc. **52**, 992 (2008).
- [60] T. Naito *et al.*, “Emittance measurement at KEK-ATF damping ring,” in *Proceedings of the 1999 Particle Accelerator Conference* (1999) pp. 2143–2145.
- [61] ATF collaboration, “Introduction of ATF accelerator,” <http://www-atf.kek.jp/atf/introduction.html> (2011).
- [62] C. Swinson, *Development of Beam Position Monitors for Final Focus Systems at the International Linear Collider*, Ph.D. thesis, University of Oxford (2010).
- [63] Y. Honda, *Experimental studies of a low emittance electron beam in the KEK-ATF damping ring with a laserwire beam profile monitor*, Ph.D. thesis, Kyoto University (2004).
- [64] Z. Farkas, “Passive pulse compression,” in *Handbook of Accelerator Physics and Engineering*, edited by A. Chao and M. Tigner (World Scientific Publishing Co. Pte. Ltd., 1999) pp. 407–410.
- [65] ATF collaboration, “Aiming for nanobeams: Accelerator test facility,” <http://www-atf.kek.jp/atf/files/ATF-E1227.pdf> (2011).
- [66] J. Urakawa *et al.*, “Experimental results and technical research and development at ATF (KEK),” in *Proceedings of EPAC2000* (2000) pp. 63–67.
- [67] M. Woodley *et al.*, “Beam Based Alignment at the KEK-ATF damping ring,” in *Proceedings of EPAC 2004* (2004) pp. 36–38.
- [68] S. Sakanaka *et al.*, “Design of an RF system for the ATF damping ring,” in *Proceedings of PAC95* (1995) pp. 1788–1790.

- [69] T. Imai *et al.*, “Double kicker system at ATF,” in *XX International Linac Conference* (2000) pp. 77–79.
- [70] P. Bambade *et al.*, “Present status and first test results of the final focus beamline at the KEK accelerator test facility,” *Phys. Rev. Spec. Top. Accel. Beams* **13**, 042801 (2010).
- [71] N. Terunuma *et al.*, “Instrumentation for the ATF2 facility,” in *Proceedings of IPAC10* (2010) pp. 2397–2401.
- [72] ATF2 collaboration, “ATF2 proposal,” SLAC-R-771 (2005).
- [73] K. Oide, “Limit of focusing of electron beam due to synchrotron radiation,” in *Handbook of Accelerator Physics and Engineering*, edited by A. Chao and M. Tigner (World Scientific Publishing Co. Pte. Ltd., 1999) pp. 252–253.
- [74] T. Yamanaka *et al.*, “Shintake monitor in ATF2: Present status,” arXiv:1006.3626v1 (2010).
- [75] T. Nakamura *et al.*, “High resolution cavity BPM for ILC final focal system (IP-BPM),” arXiv:0709.2254v1 (2007).
- [76] CERN, “Methodical accelerator design homepage,” <http://mad.web.cern.ch/mad/> (2011).
- [77] MathWorks, “MATLAB - The Language of Technical Computing,” <http://www.mathworks.com/products/matlab/> (2011).
- [78] S. Palnitkar, *Verilog HDL* (Prentice Hall, 1996).
- [79] R. Shafer, “Beam position monitoring,” in *AIP conf. proc*, Volume 212 (Los Alamos National Laboratory, 1990) pp. 26–58.
- [80] C. Clarke, *The Interaction Point Collision Feedback System at the International Linear Collider and its Sensitivity to Expected Electromagnetic Backgrounds*, Ph.D. thesis, University of Oxford (2008).
- [81] S. Molloy, *A Fast Feedback System Designed to Maintain Luminosity at a Linear Collider*, Ph.D. thesis, Queen Mary College, University of London (2006).
- [82] R. Apsimon, *The Development and Implementation of a Beam Position Monitoring System for use in the FONT Feedback System at ATF2*, Ph.D. thesis, University of Oxford (2011).

- [83] R. Apsimon, (2010), private communication.
- [84] A. Faus-Golfe *et al.*, “Accelerator R&D for future large colliders at IFIC,” wwwae.ciemat.es/ilcspain/talks/AcceleratorsIFIC_ciemat09.ppt (2009).
- [85] J. Rogers, “Feedback systems for coupled bunch instabilities,” in *Handbook of Accelerator Physics and Engineering*, edited by A. Chao and M. Tigner (World Scientific Publishing Co. Pte. Ltd., 1999) pp. 534–539.
- [86] “John adams institute for accelerator science,” <http://www.adams-institute.ac.uk/> (2011).
- [87] “TMD technologies,” <http://www.tmd.co.uk/> (2011).
- [88] P. N. Burrows *et al.*, “Beam test results with the FONT4 ILC prototype intra-train beam feedback system,” in *Proceedings of PAC09* (2009) pp. 2676–2678.
- [89] G. Christian, (2007), internal FONT note.
- [90] Xilinx, “Virtex-5,” <http://www.xilinx.com/support/documentation/virtex-5.htm> (2011).
- [91] Texas Instruments, “ADS5474,” <http://www.focus.ti.com/lit/ds/symlink/ads5474.pdf> (2011).
- [92] Linear Technology, “LTC2624,” <http://www.linear.com/product/LTC2624/> (2011).
- [93] “Mini-circuits,” <http://www.minicircuits.com/> (2011).
- [94] Analog Devices, “AD9744,” http://www.analog.com/static/imported-files/data_sheets/ad9744.html. (2011).
- [95] IEEE Standard 1149.1-2001, “Test access port and boundary scan architecture,” (2001).
- [96] Xilinx, “Virtex-5 FPGA XtremeDSP design considerations,” http://www.xilinx.com/support/documentation/user_guides/ug193.pdf (2011).
- [97] T. Naito *et al.*, “Timing system of the ATF,” KEK preprint 97-230 (1997).
- [98] Xilinx, “Virtex-5 FPGA user guide,” http://www.xilinx.com/support/documentation/user_guides/ug190.pdf (2011).
- [99] C. Perry, “FONT5 feedback board schematics,” (2010), internal FONT note.

- [100] “Xilinx CORE Generator System,” <http://www.xilinx.com/tools/coregen.htm> (2011).
- [101] Microsoft, “C++: The Most Powerful Language for .NET Framework Programming,” <http://msdn.microsoft.com/en-us/library/ms379617.aspx> (2011).
- [102] D. Bett, (unpublished Ph.D. thesis).
- [103] C. Perry, (2010), private communication.