Development of a Prototype Cavity Beam Position Monitor for the Compact Linear Collider

Francis Cullinan

Department of Physics

Royal Holloway, University of London

A thesis submitted for the degree of

Doctor of Philosophy

October, 2014

Declaration of Authorship

I, Francis Cullinan, hereby declare that this thesis and the work presented in it is entirely my own. Where I have consulted the work of others, this is always clearly stated.

Signed: _____

Date:

Acknowledgements

I am extremely grateful for the many opportunities that allowed this thesis to be written. These opportunities are in large part, down to the following people: my one and only PhD supervisor, Dr. Stewart Boogert, who guided me excellently from start to finish, Alexey Lyapin who I could, and many times did, rely on for his careful and considered advice and Manfred Wendt, Thibaut Lefevre and Lars Søby at CERN who worked hard to ensure that the experiment at CTF3 went ahead and was a success. Additionally, I would like to thank all the members of CTF3, especially CALIFES user and main operator, Wilfrid Farabolini. I am grateful to my fellow students of cavity BPMs for the casual yet fruitful discussions: Jack Towler who I worked very closely with whilst at CERN, and Nirav Joshi and Young-Im Kim who I worked alongside before that. From other research areas, I would like to thank William Shields for his outsider insight and Robert Ainsworth for his private ACE3P tutorials.

I acknowledge the work of Andrei Lunin, who did the RF design of the prototype cavity BPM, and of Nicolas Chritin, who made the mechanical design. The mechanics of the experimental installation were designed by David Bastard. Phillippe Levanchy and Laurent Payraud constructed the electrical installations and Gary Boorman and Ola Popoola are responsible for the attenuator control. Without these people, as well as the many other staff at CERN whose work ensured that the full installation was complete before the start of beam operation, the experiment documented in this thesis would not have happened.

I am proud to have been a member of the John Adams Institute, the QP section at CERN, the global CLIC collaboration and CTF3, not just professional entities but collections of bright and welcoming individuals. I am grateful to the numerous friends who I met at Royal Holloway, CERN and the other various places along the way. Finally, I would like to thank my parents and my whole extended family for being so big and strong. At the risk of spreading it a bit thin, this thesis is dedicated to all of them.

Abstract

The Compact Linear Collider (CLIC) is a proposed electron positron collider with a maximum centre of mass energy of 3 TeV. In order for it to operate at its maximum luminosity, precise alignment of the accelerator elements in the main linac and beam delivery system is required. The beam based alignment algorithm designed to do this requires measurements of the transverse beam position with a resolution of 50 nm and multiple position measurements within a single 156 ns long bunch train. The proposed solution is a cavity beam position monitor (BPM) with a low quality factor. A prototype cavity BPM pick-up has been designed and manufactured for tests on the probe beamline of the third CLIC test facility (CTF3) at the European Organisation for Nuclear Research (CERN). This thesis presents the measurements performed on this prototype in the laboratory, the development of the associated processing electronics and tests with beam.

The resonant frequencies and quality factors of the modes of interest in the two cavities of the prototype BPM pick-up were measured in the laboratory using a network analyser. Problems with the design of the pick-up geometry are identified and solutions proposed. Analytical expressions for the amplitude and phase and total energy of the multiple bunch signals are derived and a processing algorithm for the deconvolution of the single bunch waveform from the multiple bunch signal is presented. A set of receiver electronics for analogue processing of the microwave frequency signals has been designed and built from connectorised components. Beam-based measurements of the sensitivity of the position cavity signal to beam position and the reference cavity signal to charge have been made. The cavity BPM has been calibrated and used to measure the beam position jitter at the BPM location. Finally, the performance of the whole system is discussed and possible tests that could be used to determine the measurement bandwidth of the cavity BPM are proposed.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 The Compact Linear Collider

The Compact Linear Collider (CLIC) is a proposed electron-positron collider with a maximum centre of mass energy of 3 TeV. Since the electron and its anti-particle, the positron, are fundamental particles, the collision energy is known to greater accuracy than in a proton collider so more precise measurements of physical phenomena can be made. The physics experiments at CLIC will be focused on measuring particle interactions to further test the Standard Model of particle physics and new physics discovered at the current high energy accelerator, the Large Hadron Collider (LHC) at the European Organisation for Nuclear Research (CERN).

In order to minimise the statistical uncertainties in its measurements, CLIC is designed to bring as many particles as possible into collision at the design energy. The number of potential particle interactions is quantified by the luminosity $\mathcal{L}$ given by

$$\mathcal{L} = \frac{f_{\text{rep}} N_1 N_2}{4\pi\sigma_x\sigma_y} \tag{1.1}$$

where f_{rep} is the collider repetition rate, N_1 and N_2 are the number of particles in one bunch of each beam and σ_x and σ_y are the transverse beam sizes in the horizontal and vertical directions respectively at the interaction point (IP) where the bunches collide. The luminosity has units of inverse area so that when multiplied by the probability of a particle interaction expressed as a cross-section σ,

the number of interactions expected per unit time $N_{\text{exp}} = \mathcal{L}\sigma$ is obtained.

In a circular collider, a high energy electron beam loses significant energy to synchrotron radiation as its trajectory is curved around the ring. A 1 TeV electron beam in an accelerator with the bending radius of the LHC would lose 48 TeV per turn. A linear collider does not suffer from this so it can be more energy efficient for the same accelerator length but typically has a lower repetition rate since each bunch can only be collided once. The smallest possible transverse beam size is therefore targeted to achieve maximum luminosity in a single pass. In the CLIC design for 3 TeV centre of mass energy, the horizontal beam size at the interaction point is 45 nm while the vertical beam size is just 1 nm [5].

The transverse beam size in the horizontal direction σ_x is given by

$$\sigma_x = \sqrt{\beta_x^* \epsilon_x} \tag{1.2}$$

where β_x^* is the horizontal beta function at the interaction point and ϵ_x is the horizontal transverse emittance. The expression for the vertical beam size σ_y is the same except the vertical beta function and emittance are used. More details about these parameters can be found in Section A.1 of the Appendix. The beta function changes throughout an accelerator depending on the beam optics. The emittance, on the other hand, is conserved in linear beam transport. It may, however, be degraded by several mechanisms. It is inversely proportional to the beam momentum so the normalised emittance, which is independent of the beam momentum, is often used. The normalised emittance is calculated as $\epsilon_{\text{N}} = \gamma_0 \beta_0 \epsilon$ where γ_0 and β_0 are the Lorentz beta and gamma factors respectively.

A diagram of the CLIC design for 3 TeV centre of mass energy is shown in Figure 1.1. Its distinctive feature is the two-beam acceleration scheme. The drive beam is a high current electron beam at low energy. It is used to distribute radio frequency (RF) power efficiently to the 12 GHz travelling-wave structures along the main linac that accelerate the main colliding beams to their maximum energy. These high power, high frequency accelerating structures have a nominal accelerating gradient of 100 MV m^{-1} which keeps the length of the linac short considering its maximum energy.

The drive beam is generated in a unique way. First, it is accelerated to

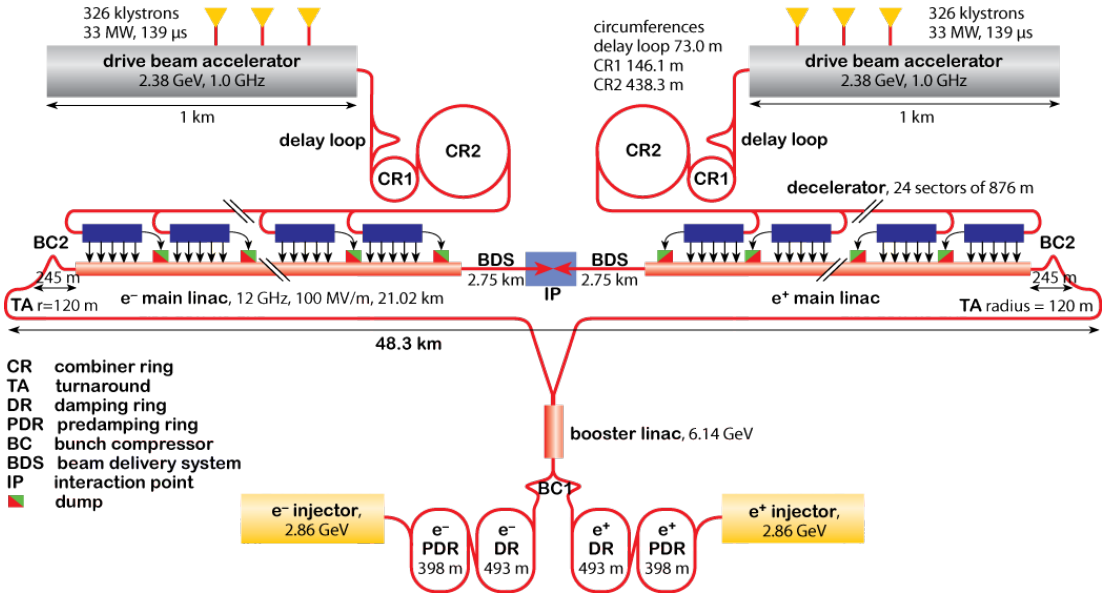

Figure 1.1: Diagram of CLIC for 3 TeV centre of mass energy where both the drive beam and the main beam and the different sub-systems are shown [1].

just above 2 GeV with a high efficiency from RF power to beam power with a bunching frequency of 0.5 GHz. The first half of the bunch train then goes round a delay loop. This lengthens its path distance so that, after the delay loop, it can be merged with the second half of the bunch train such that their bunches are interleaved and the bunching frequency is doubled to 1 GHz. Next, the train enters a combiner ring. The 2 following bunch trains are injected into the combiner ring with their bunches interleaved with the bunches in the first so that the bunching frequency is tripled to 3 GHz. A second combiner ring stores 4 bunch trains extracted from the first, giving the final bunching frequency of 12 GHz. After this recombination, the average beam current during a train has increased to 100 A which, at 2.38 GeV, corresponds to a beam power of 238 GW. Power extraction and transfer structures (PETS) throughout the tunnel of the main linac then decelerate the drive beam to 10 % of its initial energy and transfer the RF power extracted to the main beam accelerating cavities [5].

Both the electron main beam and the positron main beam are accelerated in trains consisting of 312 bunches separated in time by 0.5 ns. The trains are initially accelerated to an energy of about 3 GeV. They are then injected into

damping rings which reduce the normalised transverse beam emittances to 500 nm in the horizontal direction and 5 nm in the vertical direction through radiation damping. The beams must then be transported to the main linac where the two beam acceleration scheme is used to accelerate each beam to its final energy of 1.5 TeV. The collider is described as compact because, with the main linac, main beam and drive beam transport lines and drive beam decelerator, it has over 200 km of beamline in a single 50 km tunnel [5]. Throughout the main linac, the normalised beam emittances must be preserved within certain tolerances to achieve the target luminosity. For this purpose, the demands on the beam position monitors (BPMs) in the main linac, which measure the transverse position of the beam, and the beam diagnostics in general, are high.

Two ways that emittance can be degraded are dispersion and wakefields. Dispersion in a particle accelerator is the difference between the trajectories of a particle at the design momentum p_0 and a particle whose momentum differs by 100 %. It is brought about in locations where the perpendicular magnetic field along the design beam trajectory is non-zero, such that lower momentum particles are deflected more than higher momentum particles. The position and angular coordinates of an off-momentum particle, $x(s)$ and $x'(s)$ respectively, at beamline location s are given by the coordinates of a particle with the design momentum $(x_d(s), x'_d(s))$ plus the dispersion scaled by the fractional momentum difference $\Delta p/p_0$:

$$\begin{pmatrix} x(s) \\ x'(s) \end{pmatrix} = \begin{pmatrix} x_d(s) + D(s)\frac{\Delta p}{p_0} \\ x'_d(s) + D'(s)\frac{\Delta p}{p_0} \end{pmatrix} \tag{1.3}$$

where $D(s)$ and $D'(s)$ are the position and angular dispersion respectively [7]. Once introduced, dispersion evolves down the beamline in the same way as the particle coordinates and, in sections where there are no dispersive elements, such as bending magnets, the quantity $\mathcal{F}$, which is similar to the emittance, is conserved and is given by

$$\mathcal{F} = \gamma D^2 + 2\alpha D D' + \beta D'^2 \tag{1.4}$$

where β, γ and α are the Twiss parameters. These are described in Section A.1 of the Appendix along with the use of position-angle coordinates. The increased

emittance ϵ is then given by

$$\epsilon = \epsilon_0 \sqrt{1 + \frac{\mathcal{F}}{\epsilon_0} \left(\frac{\sigma_p}{p_0} \right)^2}. \tag{1.5}$$

where σ_p is the root mean square momentum deviation and ϵ_0 is the emittance with no dispersion [8]. In linear colliders, the design dispersion is, for the most part, zero and this is achieved by ensuring that any introduced dispersion is cancelled out in another part of the beamline. The high order magnets that are used are designed to have no magnetic field along the beam axis and to introduce no dispersion. The most common of these is the quadrupole magnet which is used for beam focusing. Its transverse magnetic field increases linearly with offset from its centre. If the beam does not pass through the centre, it will be deflected and additional dispersion will be introduced, increasing the emittance and therefore, the beam size at the interaction point. The beam must therefore be precisely aligned to the magnetic centres of the high order magnets to preserve the emittance and maximise the luminosity.

A travelling electromagnetic monopole mode excited in the accelerating cavities by an external source of RF power (in this case, the drive beam) is used to give energy to the main beam. Higher order modes can also be excited in these cavities by the beam itself. These can build up along the bunch train and give a transverse kick to each bunch and degrade the emittance. The kick can be described by a single function, the transverse wakefield. Wakefields exist in many structures but are most damaging in the accelerating cavities because there are so many. The angular kick x'_k due to the transverse wakefield $W_\perp$ of a single bunch is given by

$$x'_k(s) = \frac{NeW_\perp(s)x_0}{p_0} \tag{1.6}$$

where N is the number of particles in the bunch, x_0 is its position offset and e is the unit electronic charge. Here, the coordinate s describes the longitudinal position behind the wakefield-exciting bunch. Because there are many bunches with different offsets and many overlapping wakefields in several locations, the resulting increase in the emittance is commonly determined in simulation. More detail on the form of the beam-excited higher order modes is given in Section 1.2.1.1.

Their electric fields are zero at the centre of the cavities where there is no coupling to the beam. The accelerating cavities must therefore, also be precisely aligned to minimise the effects of the transverse wakefields.

The most precise method for aligning accelerating structures and quadrupole magnets is a beam based alignment (BBA). There are several ways of doing this that make use of the beam position monitors. The first is one to one steering where the beam is positioned at the centre of each BPM. Because all the elements cannot be perfectly aligned to the centre of the BPMs, the second technique, dispersion free steering is also used. Here, the phase of the field in the accelerating structures is altered for one beam pulse that consequently, sees less acceleration. The difference between the orbits of the low and high energy pulses is minimised by moving the quadrupole magnets so that they are better aligned. This makes the orbit independent of particle energy and therefore, dispersion free. More precisely, the procedure is the minimisation of a quantity χ^2 given by

$$\chi^2 = \sum_{i=1}^{n} w_{0,i} x_{0,i}^2 + \sum_{j=1}^{m} \sum_{i=1}^{n} w_{j,i} (x_{j,i} - x_{0,i})^2 \tag{1.7}$$

where $x_{i,j}$ is the position of beam pulse j as measured at BPM n, $j = 0$ corresponds to the nominal energy beam and $w_{j,i}$ are weighting parameters to be optimised [9]. Wakefield monitors, which directly measure the wakefields excited in the accelerating cavities, can be used to further align the accelerating structures [10]. One option that may be accommodated is the use of a single bunch train with an energy chirp along it rather than multiple beam pulses of different energy. This would reduce the time required for the alignment procedure and eliminate uncertainties arising from variations between beam pulses beyond the intentional energy shift. For this option to be possible, the BPMs must be able to make multiple position measurements within a single bunch train [9].

Another system which makes use of the beam position monitors is the orbit correction. This maintains the alignment of the beamline elements throughout running. The beam positions are measured and fed back into a controller that calculates actuator settings for the positioning of quadrupole magnets off-axis so that they steer the beam. This system deals with small changes in the alignment

due to factors such as ground motion and temperature induced mechanical motion. It requires 50 nm resolution from the BPMs but only one measurement per beam pulse. It is also sensitive to errors in the BPM calibration scales, requiring them to be smaller than 1 % [11]. The scales must therefore first be determined with this accuracy and then remain stable at this level for as long as possible to maximise the time between calibrations. Certainly, having to recalibrate more than once a month would be too disruptive to the experimental Physics program. A further requirement for the BPMs in the main linac is that they must be able to operate over a position range of ± 100 μm [5].

The greatest demands placed on the BPM system are in the beam delivery system (BDS), which comes after the main linac and prepares the beams for collision. The vertical size of each beam when they collide is 1 nm. The final doublet quadrupole magnets that focus the beams to these small sizes must therefore be stabilised to within a few nanometers to keep the beams colliding. In order to achieve such a precise alignment, the beam position in the final doublet must be measured with 3 nm resolution [5].

1.1.1 Beam Position Monitoring in CLIC

The majority of beam position monitors in the compact linear collider can be divided into three main types. One type is a high resolution cavity beam position monitor described in detail in Section 1.2 and whose development specifically for CLIC is documented in this thesis. Cavity BPMs will be installed throughout the main linac and BDS and will be used for the beam based alignment algorithm described in the previous section. It has been shown that they are able to achieve the required 50 nm resolution [12]. The other two types of BPM are known as stripline and button beam position monitors. The pick-ups for both of these types are made up of electrodes that take up a fraction of the beam pipe circumference. As the beam passes down the centre of the beam pipe, its electric field gives rise to an image current on the inside surface that is uniformly distributed around the circumference. This is a consequence of Maxwell's equations and the behaviour of electromagnetic fields on conducting boundaries, discussed briefly in Section 1.2.1.1. Having electrodes on the beam pipe surface allows the image

current to be detected. This is the principle on which button and stripline BPMs rely [13]. Table 1.1 lists the major CLIC sub-systems that are instrumented with each of these types of BPM as outlined in the CLIC conceptual design report [5]. Stripline BPMs in the drive beam decelerator account for 75 % of all the BPMs in CLIC.

BPM type	Sub-system	Number	Total
Button	Damping rings	600	3992
	Predamping rings	600	
	Drive beam transfer	872	
	Drive beam turnarounds	1920	
Stripline	Drive beam decelerator	41,484	41,484
Cavity	Main linac	4196	4396
	Beam delivery system	200	

Table 1.1: List of the major sub-systems where the three main types of BPM in CLIC are located and their number in each [5].

A button BPM pick-up has small, typically circular electrodes that take up a fraction of the beam pipe circumference. A diagram of such a pick-up is shown in Figure 1.2. Assuming that the beam is ultrarelativistic and that the beam pipe material is a perfect conductor, the image current has the same magnitude and longitudinal profile as the beam current. Using this approximation (as for all the equations in this section), the time dependent beam image charge on a rectangular electrode that extends over azimuthal angle ϕ and length l is

$$Q_{\mathrm{e}}(t) = \frac{-\phi l}{2\pi}\left[\frac{I_{\mathrm{b}}(t)}{\beta_0 c}\right] \tag{1.8}$$

where $I_{\mathrm{b}}(t)$ is the beam current, β_0 is the Lorentz beta of the beam and c is the speed of light. The term in brackets in Equation 1.8 is the charge density of the beam along its direction of motion. The equivalent circuit of a button BPM is shown on the right of Figure 1.2. The resistance in the circuit is equal to the impedance of the transmission line and the capacitance is small. It can therefore be assumed that all the current, the time derivative of $Q_{\mathrm{e}}(t)$, flows across the

load resistance R such that the signal voltage can be approximated by

$$V_{\mathrm{R}}(t) = R\frac{dQ_{\mathrm{e}}(t)}{dt} = \frac{R\phi l}{2\pi\beta_0 c}\frac{dI_{\mathrm{b}}(t)}{dt}\,. \tag{1.9}$$

The signal therefore follows the derivative of the bunch current. A button BPM pick-up does not detect the direct current (DC) component and is at its maximum amplitude at the rising and falling edges of the bunch [13]. In CLIC, the greatest demands on the performance of the button BPMs are in the damping rings where turn by turn position measurements must be made [5]. Typically, button BPMs are able to make single pass beam position measurements at resolutions in the order of 10 μm and when they are used in circular machines, their resolution can be reduced to less than 10 μm by averaging over many turns.

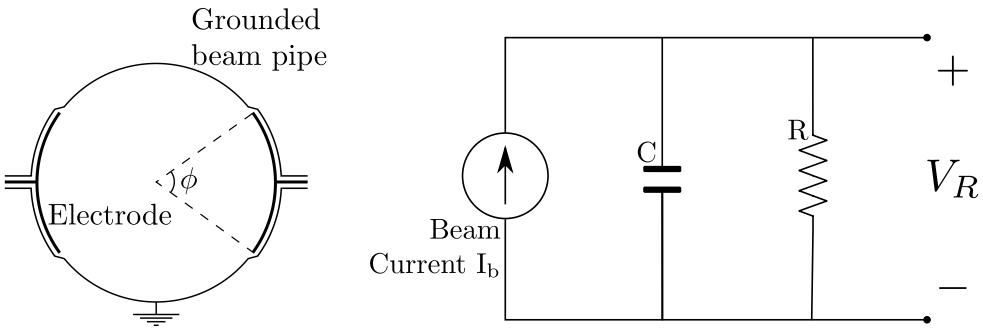

Figure 1.2: Cross section of a button BPM pick-up (left) and equivalent circuit of a single electrode (right).

The pick-up of a stripline BPM, a diagram of which is shown in Figure 1.3, is the same in cross section as the button-type pick-up. The length of the electrodes is typically much longer than the bunch length and their impedance is matched to the impedance of the output transmission line, which is connected at their upstream end. A fraction of the beam image current flows along each electrode. At the start of the electrode, the current must cross the impedance Z of the transition line in parallel with the equal impedance of the electrode giving a total impedance of $Z/2$. With the additional approximation that the bandwidth of the pick-up is large enough for its effect on the signal shape in time to be neglected,

the induced voltage V is then given by

$$V(t) = \frac{\phi}{2\pi}\frac{Z}{2} I_b(t) \tag{1.10}$$

where ϕ is the azimuthal angle covered by the electrode as before. This voltage signal propagates in two directions: down the transmission line and along the electrode. The downstream end of the electrode is either short circuited to the beam pipe, open circuited or connected to a matched transmission line or load. The total output signal from the transmission line at the start of the electrode $V_{\mathrm{S}}(t)$ is given by the general expression

$$V_{\mathrm{S}}(t) = \frac{\phi Z}{4\pi}\left[I_{\mathrm{b}}(t) + AI_{\mathrm{b}}\left(t - \frac{2l}{\beta_{\mathrm{s}}c}\right) + BI_b\left(t - \frac{l}{\beta_0 c} - \frac{l}{\beta_{\mathrm{s}}c}\right)\right] \tag{1.11}$$

where A and B are constants that depend on the electrode's end termination and β_{s} is the Lorentz factor of the signal propagating down the electrode. In the case where the end of the electrode is short-circuited, the signal excited at the start of the electrode is reflected back with opposite polarity so that $A = -1$ and $B = 0$. In the other two cases, the beam excites another signal at the end of the electrode but with opposite polarity because the beam current this time crosses from the electrode end to the beam pipe. If the end is open circuited, the reflected signal maintains its polarity and the beam excited signal has the opposite polarity but is twice as large ($A = 1, B = -2$). If the electrode ends in a matched impedance, the beam excited signal is of the same magnitude as the signal excited at the beginning of the electrode and there is no reflection so that $A = 0$ and $B = -1$. In all three cases, $A + B = -1$, so that if the speed of the electrode signal is equal to the speed of the beam $\beta_{\mathrm{s}} = \beta_0$, the total signal $V_S(t)$ has a positive peak followed, after a time interval of $2l/(\beta_0 c)$, by a negative peak of the same amplitude. In the aforementioned approximations, both peaks have the same shape as the bunch current profile. If the end of the electrode is also connected to a matched transmission line, the signal extracted there $V_{\mathrm{E}}(t)$ is

$$V_{\mathrm{E}}(t) = \frac{\phi Z}{4\pi}\left[I_{\mathrm{b}}\left(t - \frac{l}{\beta_0 c}\right) - I_{\mathrm{b}}\left(t - \frac{l}{\beta_{\mathrm{s}}c}\right)\right]. \tag{1.12}$$

In the case where $\beta_s = \beta_0$, Equation 1.12 evaluates to zero. Stripline pick-ups of this design are therefore directional and the ratio of the powers from the start and end ports can be as high as 40 dB [13].

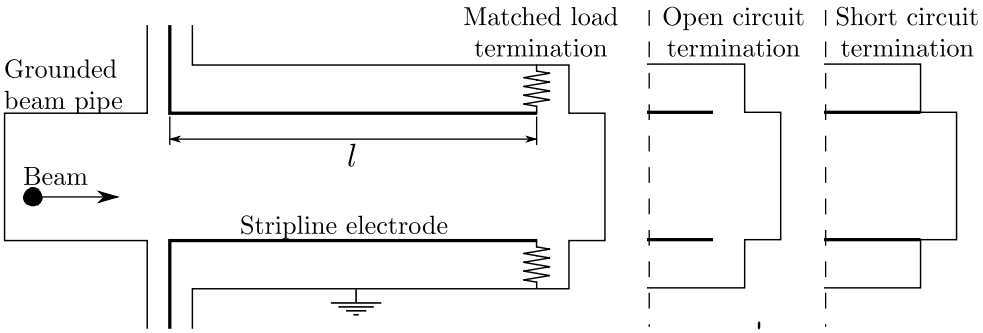

Figure 1.3: Side on view of a beam pipe with a stripline BPM pick-up with the three different types of electrode termination.

The stripline BPM designed for the CLIC drive beam has an electrode that ends with a matched load. Its design length is equal to the bunch spacing so that the negative signal excited by a bunch at the end of the electrode cancels the positive signal excited at the start by the next bunch but one. Therefore, if the bunch charge and position offset is constant along the whole bunch train, the pick-up output signal will have two positive peaks at the beginning of the bunch train followed by two negative peaks at the end and will be zero in between. The bunching frequency of the drive beam in the decelerator is 12 GHz which corresponds to an electrode electrical length of 25 mm. The required resolution is 2 μm, which stripline BPM systems are typically able to achieve [14]. The main challenge in the development of the stripline BPM system for CLIC is the quantity of BPMs required.

The electrode signals for both the button and stripline type pick-ups that are given by Equations 1.9 and 1.11 respectively are from a beam that is centred in a circular beam pipe. The signals are sensitive to the beam position offset from the centre and so position measurements can be made. Usually, a pair of electrodes is used for each transverse direction so that the difference in their signals is position dependent and the sum of the two signals can be used to normalise for charge. Electrode pick-ups have weaker wakefields than cavity pick-ups and

their theoretical sensitivity is much easier to estimate. However, they are, in general, less sensitive and are limited in performance by their nonlinear response to large beam position offsets. In a circular beam pipe, an electrode that covers a larger azimuthal angle is more sensitive but has a position response that is more nonlinear [13]. In certain locations in the CLIC beamline, such as the damping ring wigglers, the beam pipe is not circular, which makes the BPM position response more complicated [5].

1.1.2 The CLIC Test Facility

The 3rd CLIC Test Facility (CTF3) at CERN is the latest experiment built to demonstrate the feasibility of the two beam acceleration scheme. The facility layout is shown in Figure 1.4. The first step in demonstrating feasibility was to

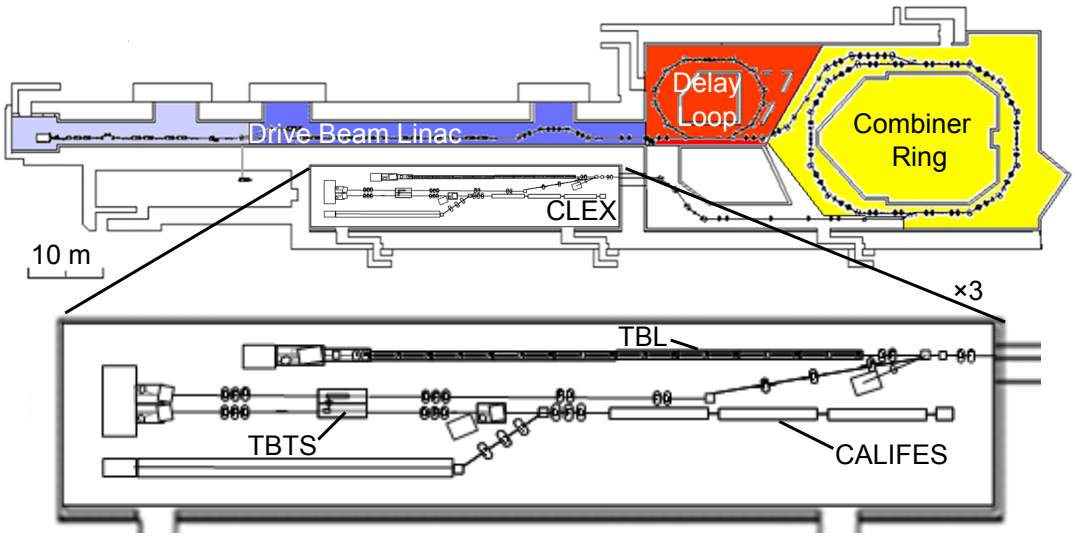

Figure 1.4: Diagram of CTF3 with a zoom in view of CLEX.

show that the drive beam can be accelerated with an efficiency from RF power to beam power that is greater than 90 % [15]. At the final energy of 600 MeV, a factor eight recombination has been achieved which increases the bunching frequency from 1.5 GHz to 12 GHz [16]. The prototype drive beam generator has a delay loop for a multiplication factor of 2 and then a combiner ring with a maximum multiplication factor of 4. The design maximum beam current after

the combiner ring is 32 A. The beam then enters the CLIC experimental area (CLEX). Here, it may continue straight down the test beamline (TBL) where there are five power extraction and transfer structures (PETS) connected to high power loads. These are used to test the stable deceleration of the drive beam. Alternatively, it may be steered towards the two beam test stand (TBTS) where there are two PETS that are connected to accelerating structures in the probe beamline.

The probe beamline CALIFES (Concept dAccélérateur Linéar pour Faisceau d'Electron Sonde) [6] is entirely based in CLEX. Its injection linac uses 3 GHz accelerating structures to generate a beam of up to 226 bunches at a bunching frequency of 1.5 GHz and energy of 200 MeV. The probe beamline includes a table for testing accelerating structures which receive power from the PETS in the drive beamline and have successfully achieved an accelerating gradient of 150 MV m^{-1} [17]. At the end of the probe beamline, there is a spectrometer which is used to measure the achieved acceleration [16].

1.2 The High Resolution Cavity Beam Position Monitor

The most basic cavity beam position monitor consists of a cavity pick-up, some receiver electronics and digital readout. The position of the beam is determined from the resonant electromagnetic modes excited in the cavity by the beam. These modes are necessarily accompanied by surface currents and so the method of detection is physically no different from the button and stripline BPMs described in Section 1.1.1. A variety of designs exist for cavity BPMs but the majority, especially those designed for applications where a high resolution is required, include many of the same features.. The next two sections give a detailed description of these features and the underlying physics.

1.2.1 Cavity Pick-up

The cavity pick-up for a high resolution cavity BPM consists of one or more resonant cavities, each commonly a short cylinder with metallic walls and a hole in

each end for the beam pipe. When a bunch of charged particles passes through the cavity, it excites resonant electromagnetic cavity modes. The pick-up has output couplers so that these modes can be measured. Two modes are usually used to make beam position measurements: the first resonant dipole mode whose amplitude is proportional to the beam position offset and the first resonant monopole mode. The latter is used to remove bunch charge dependence and provide a reference phase to remove components of the dipole mode signal arising from the beam trajectory angle and bunch tilt. The two modes are usually measured in two separate cavities, the position cavity for the dipole mode and the reference cavity for the monopole mode. Further details of these modes, their characteristics and the beam excitation are given in the rest of this section.

1.2.1.1 Resonant Modes in Cylindrical Cavities

The dominant interaction between macroscopic particles carrying electric charge is electromagnetic and can be described in terms of electric and magnetic vector fields, $\vec{E}$ and $\vec{B}$ respectively. They are described by Maxwell's fundamental equations for electromagnetic fields. These are shown in derivative form in Equations 1.13 to 1.16 where $\vec{j}$ is the charged current per unit area, ρ is the volume charge density, $\vec{\nabla}$ is the Del operator of vector calculus, t is time and ϵ and μ are known as the electric permittivity and magnetic permeability respectively and are material dependent.

$$\vec{\nabla} \cdot \vec{B} = 0 \tag{1.13}$$

$$\vec{\nabla} \times \vec{E} = -\frac{\partial \vec{B}}{\partial t} \tag{1.14}$$

$$\vec{\nabla} \times \vec{B} = \mu\vec{j} + \mu\epsilon\frac{\partial \vec{E}}{\partial t} \tag{1.15}$$

$$\vec{\nabla} \cdot \vec{E} = \frac{\rho}{\epsilon} \tag{1.16}$$

From Equations 1.14 and 1.15 it can be shown that at the boundary with a perfectly conducting surface, any electric field must be perpendicular to the surface while any magnetic field must be tangential. From Maxwell's equations, the

Helmholtz wave equations can also be derived for electric and magnetic fields:

$$\nabla^2 \vec{E} - \mu\epsilon \frac{\partial^2 \vec{E}}{\partial t^2} = 0 \tag{1.17}$$

$$\nabla^2 \vec{B} - \mu\epsilon \frac{\partial^2 \vec{B}}{\partial t^2} = 0\,. \tag{1.18}$$

Solving these differential equations with the boundary conditions given by the properties of electric and magnetic fields on a conducting surface, it is possible to obtain the electromagnetic fields that are able to exist within a cylindrical metallic cavity. These can be expressed as the linear combination of an infinite number of electromagnetic resonant modes. It is conventional to categorise these modes as either transverse magnetic TM, where there is no magnetic field in the longitudinal direction ($B_z = 0$), or transverse electric TE, where there is no electric field in the longitudinal direction ($E_z = 0$).

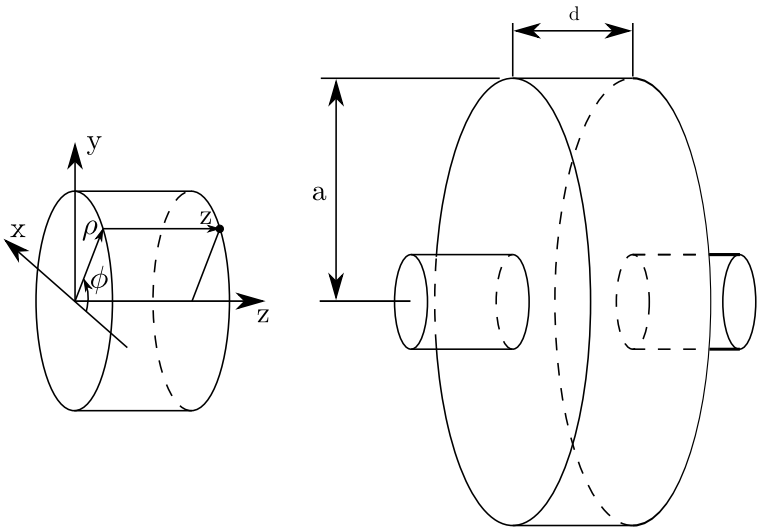

Figure 1.5: Diagram of a cylindrical cavity with beam pipes and cylindrical coordinates defined.

Since only modes with longitudinal electric fields are able to extract or give energy to a particle beam, TM mode are of most interest in accelerator physics. An illustration of a cylindrical cavity and the local coordinate system is shown in Figure 1.5. If it is of length d and radius a, the field components of a TM mode

are given in cylindrical coordinates ϕ, ρ and z as

$$B_\rho = -i\frac{2\pi f_{nml}\mu\epsilon n a^2}{p_{nm}^2\rho} J_n\left(\frac{p_{nm}\rho}{a}\right) E_0 \sin(n\phi)\cos\left(\frac{l\pi z}{d}\right) e^{i2\pi f_{nml}t} \tag{1.19}$$

$$B_\phi = -i\frac{2\pi f_{nml}\mu\epsilon a}{p_{nm}} J_n'\left(\frac{p_{nm}\rho}{a}\right) E_0 \cos(n\phi)\cos\left(\frac{l\pi z}{d}\right) e^{i2\pi f_{nml}t} \tag{1.20}$$

$$B_z = 0 \tag{1.21}$$

$$E_\rho = -\frac{l\pi a}{d p_{nm}} J_n'\left(\frac{p_{nm}\rho}{a}\right) E_0 \cos(n\phi)\sin\left(\frac{l\pi z}{d}\right) e^{i2\pi f_{nml}t} \tag{1.22}$$

$$E_\phi = \frac{l\pi n a^2}{d p_{nm}^2\rho} J_n\left(\frac{p_{nm}\rho}{a}\right) E_0 \sin(n\phi)\sin\left(\frac{l\pi z}{d}\right) e^{i2\pi f_{nml}t} \tag{1.23}$$

$$E_z = J_n\left(\frac{p_{nm}\rho}{a}\right) E_0 \cos(n\phi)\cos\left(\frac{l\pi z}{d}\right) e^{i2\pi f_{nml}t} \tag{1.24}$$

where E_0 is the complex amplitude of the electric field. J_n is a Bessel function of the first kind and order n and p_{nm} is its m^{th} root and J_n' its derivative. Bessel functions are the solution to a particular kind of differential equation and are a key feature in the solutions of electromagnetic fields in cylindrical structures. The indices n, m and l give the order of the mode (the number of oscillations) along the cylindrical coordinates ϕ, ρ and z respectively. The axis where $\phi = 0$ and thus the mode orientation, are arbitrary. However, any mode orientation can be obtained from the sum of several orthogonal polarisations, the number of which depends on the mode order in the ϕ coordinate. TM modes with $n = 0$ have only one possible polarisation and are referred to as monopole modes. Dipole modes have two possible polarisations and have $n = 1$. Beyond this, there are quadrupole modes $n = 2$ with four polarisations, sextupole modes $n = 3$ with six, octupole modes $n = 4$ with eight, etc. [18]. Once excited, the amplitude of a cavity resonant mode oscillates in time t at its characteristic resonant frequency f_{nml} given by

$$f_{nml} = \frac{1}{2\pi\sqrt{\mu\epsilon}}\sqrt{\left(\frac{p_{nm}}{a}\right)^2 + \left(\frac{l\pi}{d}\right)^2}. \tag{1.25}$$

Figure 1.6 shows illustrations of the electric and magnetic fields for the modes that are measured in order to extract a beam position. They are the first transverse magnetic monopole mode TM_{010} and first dipole mode TM_{110} where the

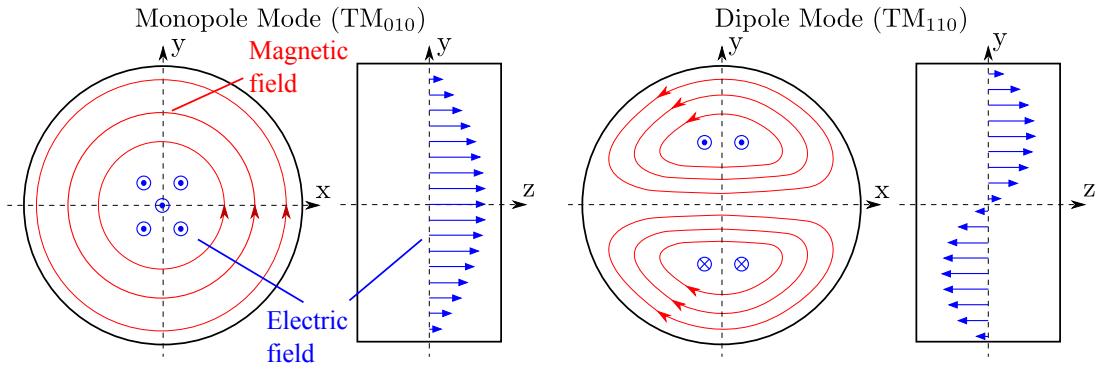

Figure 1.6: Diagram of the electromagnetic fields of the first monopole (left) and dipole (right) modes in a cylindrical cavity [2].

three subscripts denote the values of n, m and l respectively. The modes are measured using two separate cavities so that they may have the same resonant frequency. The dipole mode cavity is referred to as the position cavity while the monopole mode cavity is the reference cavity. Further mode properties are discussed in the following sections.

1.2.1.2 Quality Factor

The quality factor Q of a resonator is 2π times the energy stored divided by the energy lost per oscillation period. It is a useful number because it describes the response of an oscillator in both the frequency domain and the time domain and can also be separated in terms of resistive energy losses and the energy extracted from the system.

For many resonances, including electromagnetic cavity modes, good approximations can be made for high quality factors. In this approximation, the frequency response in terms of amplitude $A(f)$ squared can be described by a Lorentzian curve given by

$$\left[\frac{A(f)}{A(f_0)}\right]^2 = \frac{1}{1+\left(\frac{2(f-f_0)}{\Delta f}\right)^2} \tag{1.26}$$

where f_0 is the resonant frequency and Δf is a bandwidth such that the resonance amplitudes at frequencies $f = f_0 \pm \frac{\Delta f}{2}$ are a factor of $\sqrt{2}$ times smaller than the

amplitude at resonance. The quality factor is then given by

$$Q = \frac{f_0}{\Delta f}\,. \tag{1.27}$$

Applying the same approximation to the time domain, if a resonance is excited and left to ring down, its amplitude will decay exponentially in time t

$$A(t) = A(0)e^{-\frac{t}{\tau}} \tag{1.28}$$

with a characteristic decay time τ given by

$$\tau = \frac{Q}{\pi f} \tag{1.29}$$

In the context of electromagnetic modes in resonant cavities, the overall quality factor, referred to as the loaded quality factor Q_{L}, is divided into its two contributions, the internal quality factor Q_0 and external quality factor Q_{ext}, such that

$$\frac{1}{Q_{\mathrm{L}}} = \frac{1}{Q_0} + \frac{1}{Q_{\mathrm{ext}}}\,. \tag{1.30}$$

The different quality factors are defined in terms of the mechanism of the energy loss. The three terms in Equation 1.30 can be expressed as

$$Q_{\mathrm{L}} = \frac{2\pi f_0 U}{P_{\mathrm{d}}} \tag{1.31}$$

$$Q_0 = \frac{2\pi f_0 U}{P_{\mathrm{wall}}} \tag{1.32}$$

$$Q_{\mathrm{ext}} = \frac{2\pi f_0 U}{P_{\mathrm{out}}} \tag{1.33}$$

where U is the energy stored in the mode. P_{d} is the total dissipated power and is the sum of the power lost in the finite conducting walls of the cavity P_{wall} and the power extracted from the cavity P_{out} [19].

Figure 1.7 shows the dipole mode position cavity of a cavity BPM pick-up. The power is extracted from the cavity via waveguides that magnetically couple to the mode via slots. It is this coupling that determines the output power P_{out}. The slots in the diagram are aligned to selectively extract energy from the horizontal

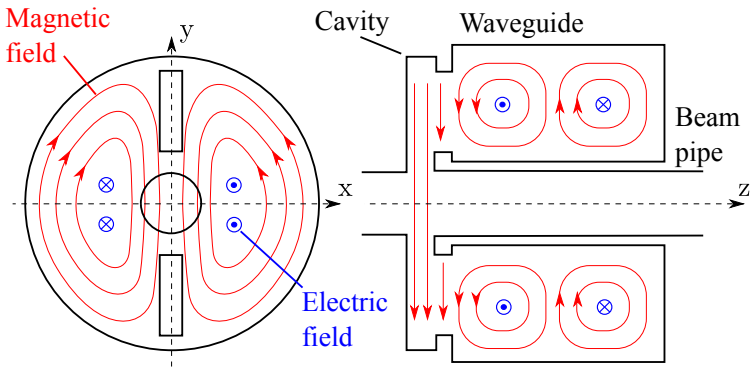

Figure 1.7: Diagram of waveguide couplers used to couple out the first dipole mode from the position cavity [2].

dipole mode polarisation whose magnetic field runs parallel to them. Another pair of slots aligned horizontally, not shown in the diagram, selects the vertical polarisation [20].

1.2.1.3 Signal Excitation

Excitation of a cavity mode is analogous to a charged particle moving across a potential. The voltage V seen by a test particle moving at the speed of light c parallel to the longitudinal axis z can be found by integration along its path:

$$V(f) = \int_{-\infty}^{\infty} E_z(x, y, z) e^{i\frac{2\pi f z}{c}} dz \tag{1.34}$$

where f is the mode frequency and E_z is its longitudinal electric field. The impedance of a given mode is defined in terms of frequency as

$$R(f) = \frac{|V(f)|^2}{P_\mathrm{d}} \tag{1.35}$$

where P_d is the dissipated power. The energy left behind by a bunch is given by the power integrated over time

$$U = \int_{-\infty}^{\infty} I(t)V(t)dt = \int_{-\infty}^{\infty} V(f)I(f)df \tag{1.36}$$

where the power theorem has been used for the second equality. Since $V(f) = R(f)I(f)$,

$$U = \int_{-\infty}^{\infty} R(f)I^2(f)df \tag{1.37}$$

Assuming a point charge excitation, the beam current frequency distribution $I(f)$ may just be replaced by the charge q. If the cavity mode has a reasonably high quality factor, its frequency response can be approximated to a Lorentzian around the resonant frequency f_0 with bandwidth $\Delta f = \frac{f_0}{Q}$. The integral then evaluates to

$$U = q^2 \int_{-\infty}^{\infty} R(f)df = q^2 \int_{-\infty}^{\infty} \frac{R_\mathrm{s}}{1 + \left(\frac{2Q}{f_0}(f - f_0)\right)^2} df = \frac{q^2 \pi f_0 R_\mathrm{s}}{2Q} \tag{1.38}$$

where R_s is known as the shunt impedance and is given by the impedance $R(f)$ at the mode resonant frequency f_0. The total energy is then

$$U = \frac{\pi f_0}{2} \left(\frac{R_\mathrm{s}}{Q}\right) q^2 \,. \tag{1.39}$$

The term in brackets is called the normalised shunt impedance or R/Q. It can be found by using Equation 1.31 to eliminate P_out from Equation 1.35. As far as the wall losses and power extraction do not influence the electromagnetic field shape, it is independent of the cavity material and only dependent on the geometry [21].

From Equation 1.39 and Equation 1.33, the power coupled out of the cavity is

$$P_\mathrm{out} = \frac{\pi^2 f_0^2}{Q_\mathrm{ext}} \left(\frac{R_\mathrm{s}}{Q}\right) q^2 \,. \tag{1.40}$$

Assuming constant power over one oscillation of the cavity mode, the peak output voltage is then

$$V_\mathrm{out} = \sqrt{2P_\mathrm{out}Z} = q\pi f_0 \sqrt{\frac{2Z}{Q_\mathrm{ext}} \left(\frac{R_\mathrm{s}}{Q}\right)} \,. \tag{1.41}$$

After excitation, the output voltage $V(t)$ oscillates at the mode resonant frequency f_0 and decays with decay time τ given by Equation 1.29:

$$V(t) = V_\mathrm{out} e^{-\frac{t}{\tau}} \sin(2\pi f_0 t) \,. \tag{1.42}$$

The total energy in the signal coupled out is obtained by integrating the output power over time [12]:

$$E_{\text{out}} = \int_0^\infty P_{\text{out}} e^{-\frac{2t}{\tau}} dt = P_{\text{out}} \frac{\tau}{2} . \tag{1.43}$$

As illustrated in Figure 1.6, the maximum longitudinal electric field in the monopole mode is at the centre of the cavity. Since the beam trajectory is close to the centre of the cavity by design, the result of Equation 1.34 is larger for the monopole mode than for higher order modes. The waveguide couplers illustrated in Figure 1.7 not only select the dipole mode polarisation, they also couple weakly to the monopole mode whose magnetic field runs perpendicular to the slots. The tails of the monopole mode in the frequency domain extend to the dipole mode frequency. These tails are still large enough in amplitude to degrade the resolution if there is no dipole mode selection. The monopole mode reference cavity requires no mode selection since the monopole mode is most strongly excited by the beam and is not as critical to the final resolution as the dipole mode. The geometry of the reference cavity is therefore, typically much simpler.

1.2.1.4 Sensitivity

The R/Q is dependent on the offset from the longitudinal axis at which it is evaluated using Equations 1.34 and 1.35. It is this dependence that is exploited to make position measurements and determines the sensitivity of the pick-up to changes in position. As an approximation for small position offsets in a cylindrical cavity, the R/Q of a monopole mode is independent of the offset and the R/Q of a dipole mode is proportional to the square of the offset along its axis of polarisation. Explicitly, for a horizontally polarised dipole mode,

$$\frac{R_{\text{s}}}{Q} = \left(\frac{R_{\text{s}}}{Q}\right)_0 \frac{x^2}{x_0^2} \tag{1.44}$$

where $(R_{\text{s}}/Q)_0$ is the normalised shunt impedance evaluated at a horizontal offset x_0. In this case, the output voltage given by Equation 1.41 is therefore proportional to the offset of the exciting charge. For dipole modes, it is convenient to define a sensitivity S_x that gives the peak output voltage for a unit offset and

unit charge such that

$$V_{\text{out}} = S_x q x \tag{1.45}$$

where q is the charge of the exciting particle and x is its offset in position [12].

If instead of a point charge, the mode is excited by a bunch with length σ_z and a longitudinal Gaussian charge distribution given by

$$dq = \frac{q_{\text{t}}}{\sqrt{2\pi}\sigma_z} e^{-\frac{z^2}{2\sigma_z^2}} dz \tag{1.46}$$

where q_{t} is the total charge, the charge distribution must be convolved with Equation 1.41 along the longitudinal axis z. This results in a factor that reduces the sensitivity so that

$$\frac{S_x(\sigma_z)}{S_x(0)} = e^{-\frac{2\pi^2 f_0^2 \sigma_z^2}{c^2}} \tag{1.47}$$

where c is the speed of light [21].

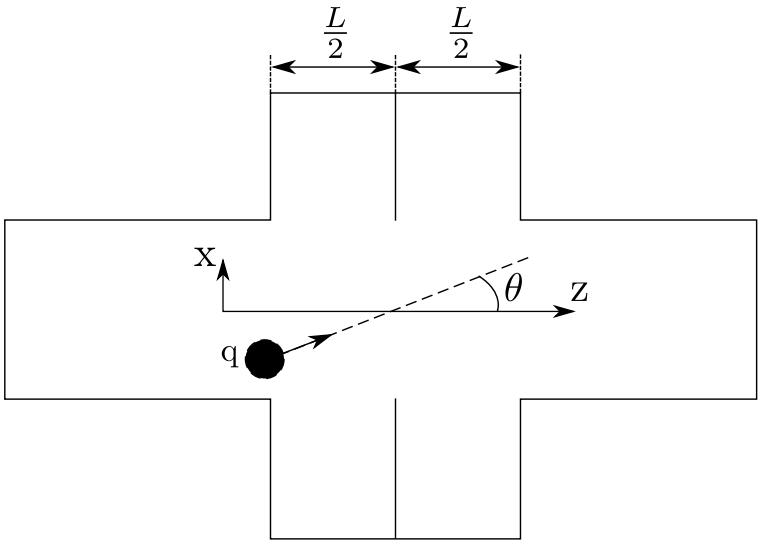

Figure 1.8: Diagram of a single charged particle crossing two successive cavities with a trajectory angle θ.

The amplitude of the dipole mode is also sensitive to tilted bunches or bunches with an angled trajectory. An approximation of the latter effect can be obtained by dividing a cavity of length L into two cavities of length $L/2$ joined front to back as illustrated in Figure 1.8. The exciting bunch is chosen to be moving

along a trajectory with a small angle θ that crosses the central longitudinal axis at the point at which the two cavities meet. In the small angle approximation, the bunch will have an average offset of $-\frac{1}{4}L\theta$ in the first cavity and $+\frac{1}{4}L\theta$ in the second. Each of the two cavities, being half as long as the original cavity, is also set to be half as sensitive. Summing the signals in the two cavities, the total output signal $V_\theta(t)$ is then given by

$$V_\theta(t) = -\frac{S_x}{2} q \frac{L\theta}{4} \sin\left(2\pi f_0 \left[t + \frac{L}{4c}\right]\right) + \frac{S_x}{2} q \frac{L\theta}{4} \sin\left(2\pi f_0 \left[t - \frac{L}{4c}\right]\right) \tag{1.48}$$

$$V_\theta(t) = -S_x q \theta \frac{\pi f_0 L^2}{8c} \cos(2\pi f_0 t) \tag{1.49}$$

where the cavity length is assumed to be much shorter than one mode oscillation multiplied by the speed of light so that $\sin(\frac{2\pi f_0 L}{2c}) \approx \frac{2\pi f_0 L}{2c}$. Equation 1.49 suggests that the voltage from a particle moving along a trajectory with a small angle θ is proportional to the angle and oscillates in quadrature phase relative to the signal excited by a particle at a constant offset, given by Equation 1.42. The amplitude of the signal V_a from an angled beam can also be compared to the signal V_p expected with a constant offset of $\frac{L\theta}{2}$:

$$\frac{V_\mathrm{a}}{V_\mathrm{p}} = \frac{\pi f_0 L}{4c} \,. \tag{1.50}$$

Equation 1.50 only gives an order of magnitude approximation but shows that the signal is expected to be small if the cavity length is short. A similar treatment has previously been applied to a tilted bunch of non-zero length [22]. The signal amplitude in this case is similar to Equation 1.50 but with opposite sign and the length scale L replaced by the bunch length σ_z. A more accurate analysis of both cases has also been carried out with similar conclusions [12]. Introducing constants for sensitivity to bunch tilt α and trajectory angle θ, S_α and S_θ respectively, the voltage signal extracted from a cavity dipole mode excited by a bunch of charge q, tilt α, offset x and trajectory angle θ is given by

$$V(t) = [S_x q x \sin(2\pi f_0 t) + q(S_\alpha \alpha - S_\theta \theta)\cos(2\pi f_0 t)] e^{-\frac{t}{\tau}} \,. \tag{1.51}$$

From Equation 1.41, it is clear that the amplitude of every mode is proportional to the charge q of the particle that excites it. Since a monopole mode signal is essentially independent of the offset of the exciting charge, it's amplitude is only dependent on bunch charge and its phase is fixed relative to the time of the bunch crossing. A monopole mode signal can therefore be used to remove the charge dependence and provide a reference phase. The phase is required to remove the contributions to the dipole mode signal that come from the bunch tilt and trajectory angle. It is also used to determine whether the bunch offset is positive or negative.

1.2.1.5 Simulation

The inclusion of the waveguide couplers and beam pipe make it very difficult, if not impossible, to find analytical solutions for the electromagnetic fields in cavity BPM pick-ups. Computer based field solvers are therefore used for pick-up design and the interpretation of experimental results. These solvers are either based on finite difference or finite element methods. For both methods, the path to the solution is similar: first the defined geometry is filled with a mesh of small volumes called cells, the corners of which are called nodes. This discretises the computational volume and allows the differential equations for the electromagnetic fields and the boundary conditions to be turned into matrix equations with unknowns corresponding to the fields at the nodes. Matrix inversion is then performed to obtain the fields. Two simulation packages that are popular in the field of accelerator physics are GdfidL [23], which uses the finite difference method, and ACE3P [24], which uses the finite element method.

The main difference between finite difference solvers and finite element solvers is the way in which the matrix equations are constructed. The differential of a function $f(x)$ with respect to variable x is defined as

$$\frac{df(x)}{dx} = \lim_{\Delta x \to 0} \frac{f(x + \Delta x) - f(x)}{\Delta x} \tag{1.52}$$

where Δx is a change in x. Finite difference methods express differential equations in the form given by Equation 1.52 with Δx as the distance between neighbouring mesh nodes. Imposing the boundary conditions sets the values of some of the

fields at certain nodes. The fields at each remaining node are then expressed as a linear combination of the fields at its neighbouring nodes to give one equation for each unknown. These equations are combined into one matrix equation and matrix inversion is used to find the problem solution.

In general, the nodes in a finite difference solver are uniformly distributed throughout the volume (ie. Δx is constant) so a cubic mesh, such as the one used in GdfidL, is well suited. More complex shaped mesh cells are possible by excluding the nodes at some of the cube corners to produce triangular prisms, tetrahedrons, etc. Because the solver performs best with a uniform mesh, the whole mesh must be dense enough to accurately represent the smallest details of the geometry. Since, for both finite element and finite difference methods, the computer memory required is proportional to the number of mesh cells, this is computationally expensive. However, because the finite difference method is less complex than the finite element method, a larger mesh density can be tolerated with the same computing resources.

In the finite element method, the fields within each mesh cell are expressed in terms of basis functions such that

$$\phi^e(x) = \sum_{i=0}^{N} c_i \psi_i^e \tag{1.53}$$

where ϕ^e is a field distribution inside element e, c_i is a coefficient and ψ_i^e is one of the element's basis functions. The basis functions can have high order components to improve the accuracy but this slows down the computation. The fields in all the elements are then combined into a single matrix equation as with the finite difference method. However, with the finite element method, the method of determining the matrix coefficients is more complicated. Variational methods are employed which work by finding solutions that drive a governing function to an extremum, a minimum or maximum value. At extrema, the function changes only slightly with a small change to the solution so the best accuracy can be obtained. Several variational methods exist and the one used is dependent on the solver.

The finite element method is well suited to a mesh of varying cell size. Because of this, the mesh can be less dense on average than for the finite difference method which, in terms of computing time, makes up for the increased complexity of the

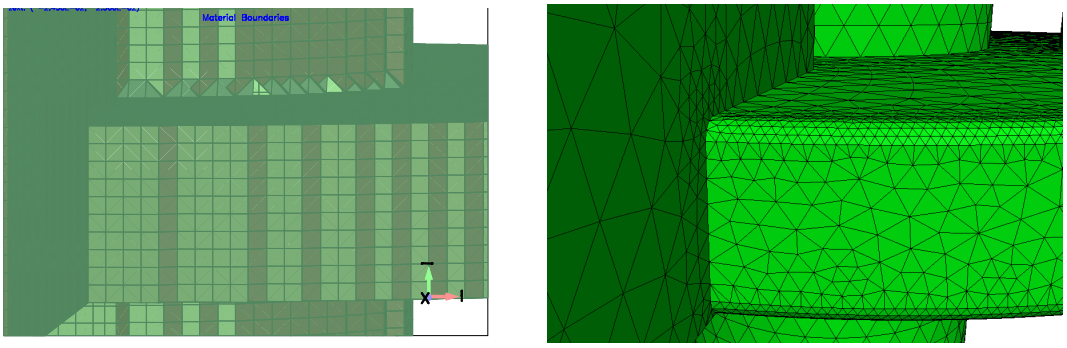

Figure 1.9: Cubic mesh generated by GdfidL (left) and tetrahedral mesh used by ACE3P (right).

calculation. ACE3P uses tetrahedral mesh cells and, because of the way the basis functions are defined, can have mesh cells with curved edges. This improves the representation of curved geometries. An example of the mesh used by each solver is shown in Figure 1.9.

Both GdfidL and ACE3P can be used to solve Equations 1.17 and 1.18 for the resonant cavity modes. They do this by expressing the wave equation as a matrix eigenvalue problem where each eigenvalue corresponds to a different mode wavenumber. Solving the problem for a set number of eigenvalues returns the field distributions of the same number of resonant modes, sometimes referred to as eigenmodes, each with a different frequency. Knowledge of the exact field distributions and characteristics from electromagnetic solvers is useful for estimating the mode shunt impedance using Equations 1.34 and 1.35. ACE3P and GdfidL are also able to simulate finite conducting cavity walls and waveguide ports so can give estimates for the resonant mode quality factors.

The two simulation packages can also be used to predict the reflections at waveguide ports and the transmission between them if power enters the port with a given frequency. Similar to a resonant cavity, electromagnetic fields propagate down a waveguide in the form of waveguide modes. Waveguide modes are much more broadband than cavity modes and so are able to transfer power over a wider range of frequencies. The simulation must therefore, first solve for the waveguide modes of the different ports before determining the transfer of power between them. Because the modes are broadband, knowledge of only one or two

waveguide modes is sufficient [25].

1.2.1.6 Equivalent Circuit

It can be useful to represent the pick-up of a cavity beam position monitor as an electrical circuit that is mathematically equivalent. While features of the actual problem are not easily identifiable within the complex geometry, the equivalent circuit separates them into individual electrical components with associated parameters that can be easier to refer to. The appropriate circuit is shown in Figure 1.10. The particle beam is replaced by a current source and the mode resonance, by a parallel RLC circuit. The waveguide coupling is replaced by a mutual inductance [26].

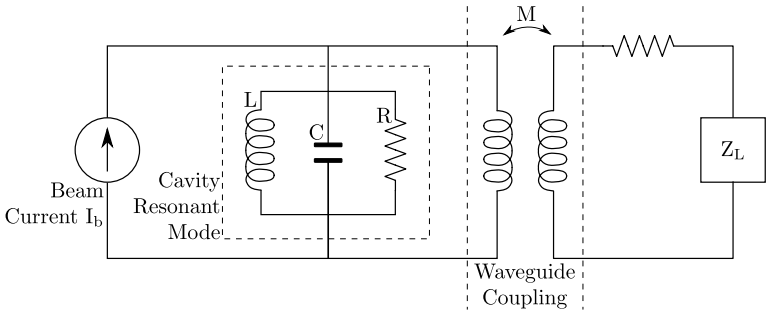

Figure 1.10: Equivalent circuit of a cavity beam position monitor pick-up.

An RLC circuit contains a resistance R, capacitance C and inductance L. The inductor and capacitor in an RLC circuit are able to store energy and will exchange energy at a resonant frequency f_0. Energy is dissipated in the resistance so the internal quality factor Q_0 can be expressed in terms of the three values R, L and C. An RLC circuit also has an impedance Z defined as the ratio between the voltage and current in the circuit. The impedance is complex and so includes the phase difference between oscillating voltages and currents. Its real part is equal to the resistance and its imaginary part, called the reactance, is defined by the capacitance and inductance. The impedance is given by

$$Z = R + i\left(2\pi f L - \frac{1}{2\pi f C}\right) . \tag{1.54}$$

The response of an RLC circuit in time and frequency is the same as the oscillator described in Section 1.2.1.2 and is similar to the response of a single resonant cavity mode.

Current flowing through an inductor creates a magnetic field. If there is a change in current over time, the magnetic field then acts back on the inductor and causes a potential difference across it. Two circuits may share a mutual inductance where the magnetic field in one inductor not only effects the potential across itself, but also across an inductor in the other circuit. This is how the waveguide coupling is represented in the equivalent circuit. A large mutual inductance M lowers the external quality factor and also has a small influence on the resonant frequency of the resonant RLC circuit [27].

1.2.2 Electronics

From Equation 1.25, the resonant frequency of a cavity mode is approximately inversely proportional to the size of the cavity. For the size of the pick-up to be practical (< 30 cm diameter), the resonant frequency must be larger than 3 GHz. Such a high frequency is far higher than the first Nyquist band of many digitisers. At high frequencies, the amount of power dissipated in signal transfer cables is also large and electrical components are more difficult to manufacture, more expensive and less reproducible. Furthermore, as discussed in Section 1.2.1.4, the relative phase of the position and reference cavity signals must be measured so a diode detector cannot be used to simply rectify the output signal.

One common solution is to use receiver electronics to mix down to an intermediate frequency (IF) of a few tens of megahertz. These are installed close to the pick-up in order to limit cable losses. The down-mixing is performed by a mixer which essentially multiplies the signal by a continuous-wave (CW) local oscillator (LO) of frequency f_{LO}:

$$\begin{aligned} e(t)\sin(2\pi f_0 t) \times \sin(2\pi f_{\mathrm{LO}} t + \phi) = \\ \frac{e(t)}{2}[\cos(2\pi(f_0 - f_{\mathrm{LO}})t - \phi) - \cos(2\pi(f_0 + f_{\mathrm{LO}})t + \phi)] \end{aligned} \tag{1.55}$$

Here, $e(t)$ is an envelope function describing the change in amplitude of a signal

oscillating at frequency f_0. The second sine function describes the local oscillator that comes in with an arbitrary phase ϕ. It can be seen that the mixer output includes two components at frequencies $f = f_0 \pm f_{\mathrm{LO}}$. The higher frequency component can be removed by filtering. The lower frequency is the IF. The signals can then be transported out of the tunnel with lower losses and then digitised at IF or at baseband after another down-conversion stage.

Amplification in the electronics also improves the spatial resolution by conserving the signal to noise ratio better (see Section 1.2.2.1) and matching the range of signal levels to the range that the digitiser is able to measure. The smallest change in voltage that the digitiser is able to measure corresponds to the smallest change in beam position and therefore, limits the spatial resolution. The maximum voltage range of the digitiser corresponds to the maximum measurable beam position and therefore, limits the range over which the BPM is able to operate. A compromise must be reached that delivers the spatial resolution required for the precise alignment of accelerator components across a position range that covers the expected beam trajectory perturbations and related alignment tolerances. Saturation of the electronics components also limits the amplitude of signals that are able to be processed and so to use the full range of the digitiser most effectively, this should occur outside the digitiser voltage range.

1.2.2.1 Noise

One of the limits to the resolution of a cavity BPM is the noise that is produced in the electronics since changes in the signal level that are smaller than the electronics noise cannot be measured. Thermal noise is present in an electronic circuit merely because of the random motion of charge carrying particles at nonzero temperature. Random noise has a flat frequency response and so the total noise power is proportional to the bandwidth. The thermal noise power N_{therm} is given by

$$N_{\mathrm{therm}} = 4k_{\mathrm{b}}T\Delta f \tag{1.56}$$

where k_{b} is Boltzmann's constant, T is the ambient temperature and Δf is the bandwidth. The power in a signal transferred through an electronics component will change by a factor that is equal to the power gain G of the component.

The power in the thermal noise will also change by this factor but there is also additional noise N_{a} generated inside the component given by

$$N_{\mathrm{a}} = (F-1)GN_{\mathrm{therm}} \tag{1.57}$$

where F is the noise factor of the component defined at the ambient temperature T. Therefore the signal power changes by a factor G and the total noise changes by a factor $(F-1)G+G = FG$. In this simple scenario, the signal to noise ratio at the output $SNR_{\mathrm{out}} = \frac{S_{\mathrm{out}}}{N_{\mathrm{out}}}$ is the signal to noise ratio at the input SNR_{in} divided by the noise factor:

$$SNR_{\mathrm{out}} = \frac{SNR_{\mathrm{in}}}{F}\,. \tag{1.58}$$

When a set of electronics is made up of several components, the noise after N components is the noise generated in each component multiplied by the gain of the following components plus the thermal noise background multiplied by the total gain:

$$N_{\mathrm{out}} = \left(\prod_{i=1}^{N} G_i\right) N_{\mathrm{therm}} + \sum_{i=1}^{N}\left[(F_i-1)\left(\prod_{j=i}^{N} G_j\right) N_{\mathrm{therm}}\right] \tag{1.59}$$

In terms of the signal to noise ratio, this gives

$$SNR_{\mathrm{out}} = \frac{SNR_{\mathrm{in}}}{F_{\mathrm{tot}}} = \frac{\left(\prod_{i=1}^{N} G_i\right) S_{\mathrm{in}}}{N_{\mathrm{out}}} \tag{1.60}$$

where S_{in} is the input signal power level and F_{tot} is the combined noise factor of all the electronics components. Substituting Equation 1.59 into the right hand side of Equation 1.60 gives the following expression for F_{tot}:

$$F_{\mathrm{tot}} = F_0 + \sum_{i=1}^{N} \frac{F_i - 1}{\prod_{j=0}^{i-1} G_j}\,. \tag{1.61}$$

Equation 1.61 is known as Frii's formula for noise and gives the noise power factor for a set of electronics from the noise factors of the individual constituent components. When expressed in dB, the noise factor is referred to as the noise figure *NF* [28].

In general, the noise factor of a component with only resistive losses is equal to the inverse of the gain. The noise factor of active components will be greater than this, although their gain may be greater than 1. No component has a noise factor that is less than 1 which would mean it absorbs noise and improves the signal to noise ratio. Some conclusions can therefore be drawn from Equation 1.61. A set of electronics made up of only resistive components will attenuate the signal while maintaining the noise at the thermal level. Having an amplifier (with $G > 1$) at the front end of the electronics will increase the denominator in the last term and mitigate the effects that the following components have on the noise. In other words, amplifying the noise along with the signal and raising it above the thermal level allows it to be attenuated by some amount in the following components and there will be less degradation of the signal to noise ratio.

1.2.2.2 Linearity

Ideally, the response of the electronics will be linear. That is, an increase in the signal level at the input will increase the output signal by the same factor. Since, as discussed in Section 1.2.1.4, the dipole mode signal amplitude is linearly proportional to beam offset, this would make the overall system linear and minimise the calibration constants needed to extract a position. In reality, many components such as mixers and amplifiers behave nonlinearly and their transfer functions, that give the output signal in terms of the input signal, have terms of order greater than 1 in V_{in}:

$$V_{\text{out}} = a_0 + a_1 V_{\text{in}} + a_2 V_{\text{in}}^2 + a_3 V_{\text{in}}^3 + \dots \tag{1.62}$$

where V_{in} and V_{out} are the signals at the electronics input and output respectively and a_n is a coefficient corresponding to the term of order n. If the input voltage

is a harmonic oscillator given by

$$V_{\text{in}} = V_0 \sin(2\pi f t) \tag{1.63}$$

where f is the signal frequency and V_0 is a constant amplitude, the third order term will be given by

$$a_3 V_{\text{in}}^3 = a_3 V_0^3 \sin^3(2\pi f t) = a_3 V_0^3 \left[\frac{3}{4}\sin(2\pi f t) - \frac{1}{4}\sin(2\pi(3f)t)\right] \tag{1.64}$$

The input power at the third order intercept point $IIP3$ is defined as the signal power for which the third order term in Equation 1.62 is equal to the first order term at the signal frequency and is given, from inspection of Equation 1.64, as

$$IIP3 = k\left|\frac{4a_1}{3a_3}\right| \tag{1.65}$$

where k is a conversion factor from voltage amplitude squared to power.

A common method used to determine the $IIP3$ of a set of electronics is known as a two-tone measurement. Two signals from separate sources that are close but not equal in frequency are combined resistively so that their voltages are added and are then fed to the electronics input. Taking only the first and third order terms in Equation 1.62, the input and output signals, V_{in} and V_{out}, will be equal to

$$V_{\text{in}} = V_1 \sin(2\pi f_1 t) + V_2 \sin(2\pi f_2 t)) \tag{1.66}$$

$$\begin{aligned} V_{\text{out}} = &\left[a_1 V_1 - \frac{3}{4}a_3 V_1^3\right]\sin(2\pi f_1 t) + \left[a_1 V_2 - \frac{3}{4}a_3 V_2^3\right]\sin(2\pi f_2 t) \\ &+ \frac{3}{4}a_3 V_1^2 V_2\left[\sin(2\pi(2f_1 + f_2)t) + \sin(2\pi(2f_1 - f_2)t)\right] \\ &+ \frac{3}{4}a_3 V_2^2 V_1\left[\sin(2\pi(2f_2 + f_1)t) + \sin(2\pi(2f_2 - f_1)t)\right] \\ &+ \frac{1}{4}a_3\left[V_1^3\sin(2\pi(3f_1)t) + V_2^3\sin(2\pi(3f_2)t)\right] \end{aligned} \tag{1.67}$$

where V_1 and V_2 are the amplitudes of the signals at f_1 and f_2 respectively. The output signal has components at more frequencies than the input signal. Of these

$2f_1 - f_2$ and $2f_2 - f_1$ are close to the input frequencies and show up as spurs either side of the main signal peaks in the output frequency spectrum. Ignoring terms in V_1^3 or V_2^3, the output power at the third order intercept point $OIP3$ is given in terms of the power P at these frequencies and the main signal frequency as

$$OIP3 = P(f_1)\sqrt{\frac{P(f_2)}{P(2f_1 - f_2)}} = P(f_2)\sqrt{\frac{P(f_1)}{P(2f_2 - f_1)}} \tag{1.68}$$

If the power is measured in a logarithmic scale such as dBm, Equation 1.68 becomes

$$OIP3 = P(f_1) + \frac{P(f_2) - P(2f_1 - f_2)}{2} = P(f_2) + \frac{P(f_1) - P(2f_2 - f_1)}{2} \tag{1.69}$$

It is important that the input powers during the measurement are low so that ignoring the cubic terms in the first line of Equation 1.67 is justified. The relationship between $IIP3$ and $OIP3$ is simply

$$OIP3 = a_1^2 IIP3\,, \tag{1.70}$$

or the input power multiplied by the power gain of the electronics [29].

Another related characteristic is the input or output power at 1 dB compression. This refers to the power for which the output power at the signal frequency is 1 dB lower than if the response was perfectly linear. Using the first and third order terms of Equation 1.62 and assuming that a cubic curve is an accurate representation of the power response of the electronics, the input power for 1 dB compression $P_{1\mathrm{dB}}$ is given by

$$P_{1\mathrm{dB}} = k(1 - 10^{-\frac{1}{20}})\left|\frac{4a_1}{3a_3}\right| = (1 - 10^{-\frac{1}{20}})IIP3\,. \tag{1.71}$$

The input power at 1 dB compression is therefore $10\log_{10}(1 - 10^{-\frac{1}{20}}) = 9.6$ dB lower than the input third order intercept point. It is a good number for the upper limit on the dynamic range of the electronics. The output power at 1 dB compression is, like the intercept, equal to the input power multiplied by the power gain.

Chapter 2

Microwave Cavity Pick-Up for CLIC

2.1 Pick-Up Design

A cavity BPM pick-up for CLIC has been designed at Fermilab [30]. The design of the position cavity follows the basic description given in Section 1.2.1 with waveguide couplers for dipole mode selection. The first dipole mode of the position cavity is at 14 GHz and the waveguides are coupled to the cavity via rectangular slots. The 14 GHz resonant frequency is a multiple of the 2 GHz CLIC bunching frequency such that signals from consecutive bunches add constructively while contamination from other modes at different frequencies adds destructively [31].

In order to be able to make several position measurements within a single CLIC bunch train without any signal subtraction, the cavity is designed such that the signal excited by any one bunch decays by a factor of 1000 within 50 ns. From Section 1.2.1.2, the signal decay time τ is given by

$$\tau = \frac{Q_{\mathrm{L}}}{\pi f_0} \tag{2.1}$$

where Q_{L} is the loaded quality factor of the cavity and f_0 is the mode resonant frequency. The loaded quality factor required for the desired decay time is therefore

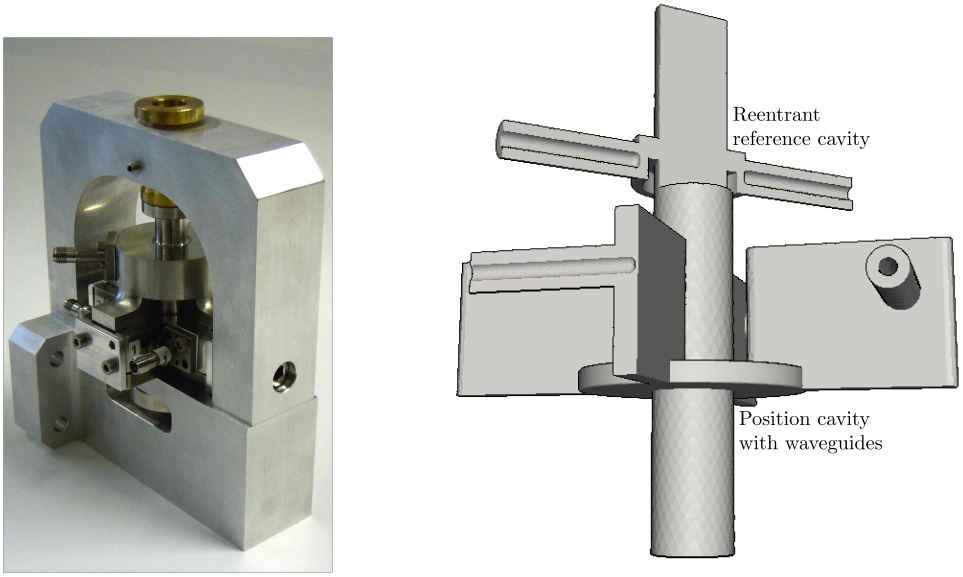

Figure 2.1: The prototype CTF3 pick-up assembled in a clamp specifically designed for the bench measurements (left) and the vacuum geometry (right) where part has been cut away so that the antenna coupling in the two cavities is visible.

given by

$$Q_{\mathrm{L}} = \frac{50 \times 10^{-9} \pi 14 \times 10^{9}}{\ln(1000)} = 320\,. \tag{2.2}$$

For this reason, stainless steel, which has low conductivity in comparison to other metals, has been chosen as the material from which the pick-up is made [30].

In order to test the design and ensure that the desired performance can be achieved, a prototype cavity BPM pick-up, shown in Figure 2.1 and also made from stainless steel, has been built to be tested on the probe beamline at CTF3. It is made from four separate parts and includes two cavities, a position cavity that is similar to the CLIC design and a reference cavity with a trench around its diameter that makes it reentrant. The design frequency of the first resonant dipole mode in the position cavity is 15 GHz since the bunch arrival frequency from the probe beam linac is 1.5 GHz and not the 2 GHz of CLIC. Instead of being separated from the cavity by slots, the waveguides join the cavity at their ends. This increases the power coupled out but reduces the isolation of each output port. A feedthrough antenna at the end of each waveguide forms the transition to a coaxial cable through a resonant antenna coupling. There are tuning bolts

that push against the waveguide wall opposite each antenna for fine adjustment of the coupling to minimise reflections [31].

There are two output ports for each of the dipole mode polarisations. As well as decreasing the quality factor by extracting more energy, this offers two further possible advantages. The first is that signals from two opposite position cavity ports can be combined in 180° hybrids via phase-matched cables to increase the sensitivity. This also improves the quadrupole mode rejection since the dipole mode signal from the two ports are 180° out of phase while the quadrupole mode contributions are in phase. The second is that by adding adjustable short-circuit terminations to one of each pair of ports, the electrical length can be adjusted to tune out the cross-coupling between adjacent output ports. The reentrant reference cavity also has two feedthrough antennas, one on each side. Having an extra output port for each mode to be measured also allows the possibility of having a spare set of electronics for each that can be used in the event of an electronics failure.

Mode parameter	Value	
	Position cavity TM_{110}	Reference cavity TM_{010}
Resonant frequency/GHz	15	15
Loaded quality factor Q_L	261	120
External quality factor Q_{ext}	727	301

Table 2.1: Parameters of the important resonant modes in the two cavities of the prototype pick-up as determined from simulations of the design geometry.

The dimensions of the prototype cavity BPM pick-up are illustrated in Figure 2.2 and the important mode parameters and summarised in Table 2.1. Each cavity dimension affects each particular mode parameter to a different extent. The cavity diameter (along with the trench depth in the case of the reference cavity) has a large influence on the mode resonant frequencies while the cavity length greatly affects their R/Q. As seen for the position cavity in Section 2.2.2, the spacing between the antenna and the opposite surface, the waveguide wall in the case of the position cavity and the trench wall in the case of the reference cavity, has a large effect on the mode quality factors. The waveguide height for the position cavity couplers is chosen so that the waveguide cut-off frequency is above the resonant frequency of the position cavity monopole mode TM_{010}. The

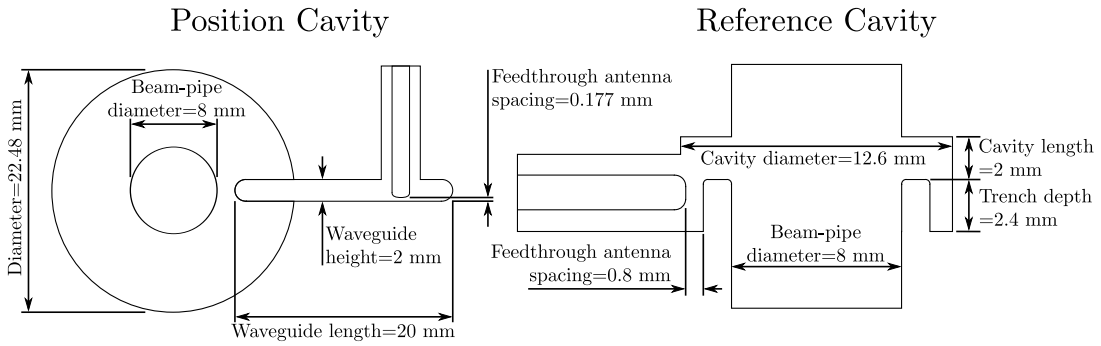

Dimension	Value/mm	
	Position cavity	Reference cavity
Diameter	22.48	12.6
Cavity length	2	2
Trench depth	0	2.4
Spacing between feedthrough antenna and opposite internal surface	0.177	0.8
Waveguide height	14	N/A
Waveguide length	20	N/A
Waveguide width	2	N/A
Beam pipe diameter	8	8

Figure 2.2: Diagram of the cavity BPM prototype pick-up and a table of the important dimensions. Only one output coupler for each cavity is shown in the diagram. In the case of the position cavity, the internal surface opposite the feedthrough antenna is the waveguide wall while for the reference cavity, it is the trench wall.

length of the waveguide is chosen to avoid resonances near the modes of interest. The beam-pipe diameter throughout the pick-up is the same as the beam pipe of the CLIC main linac. At certain points in the CLIC beam delivery system, the beam pipe diameter increases to 25 mm. This is larger than the cavity diameter and so at these points, a different pick-up design with a lower resonant frequency will have to be used [5].

2.2 Radio Frequency Characteristics

The characteristics of the modes of interest in the two cavities, specifically the resonant frequencies and quality factors, were measured using a network analyser. Initially, the assembly was not brazed and so was held together during the

measurements with a specifically designed clamp that is visible in the photo in Figure 2.1. The feedthrough antennas were to be mounted to the rest of the assembly with metal gaskets for vacuum tightness. Aluminium spacers were used in place of the gaskets so that each antenna was positioned with the correct spacing between it and the opposite conducting surface.

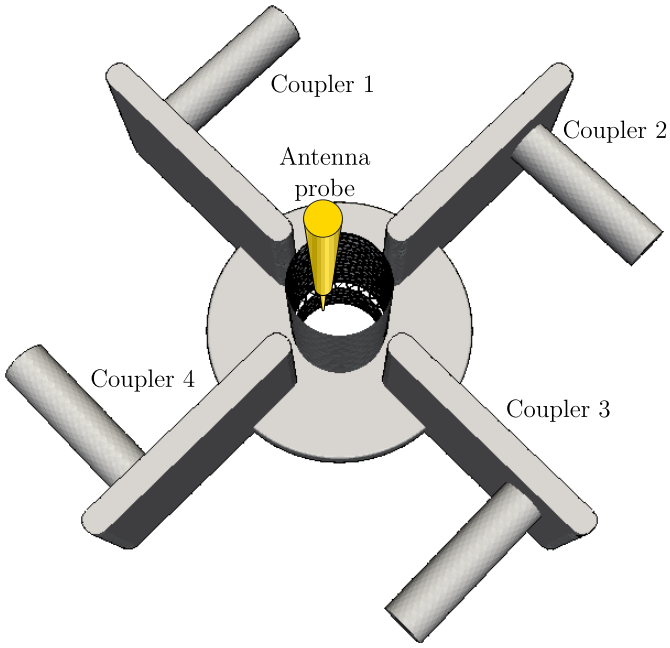

Figure 2.3: Diagram of coupler labels and approximate position of the antenna probe used for quality factor measurements.

The BPM cavity modes were characterised using the scattering S_{ij} parameters defined as

$$S_{ij} = \frac{V_i^-}{V_j^+} \tag{2.3}$$

where V_i^- is the incoming wave amplitude from port i and V_j^+ is the outgoing wave amplitude of the source connected to port j. The scattering parameters were measured using a four port network analyser. The waveguide couplers of the position cavity were numbered 1 to 4 clockwise as shown in Figure 2.3. As discussed in Section 1.2.1.1, the electromagnetic fields that are able to exist in a cavity can be expressed as the linear combination of the resonant modes and so these have a large effect on the scattering parameters. The couplers that are opposite

each other couple strongly the same dipole mode polarisation. The transmission between them is expected to follow the Lorentzian curve that describes the frequency response of the resonant mode as described in Section 1.2.1.2. A peak in transmission between two ports will be accompanied by a minimum in the power reflected at each since the total energy is conserved. When two port measurements are made, the scattering parameter is referred to as a transmission if the two output ports are opposite each other and cross-coupling if they are adjacent. In a single port measurement ($i = j$), it is the reflection that is measured.

2.2.1 Position Cavity

The four ports of the network analyser were each connected to a different output port of the position cavity. Figure 2.4 shows example network analyser traces measuring the reflections and cross-couplings of each ports. Each trace shows a frequency spectrum of measured S-parameters.

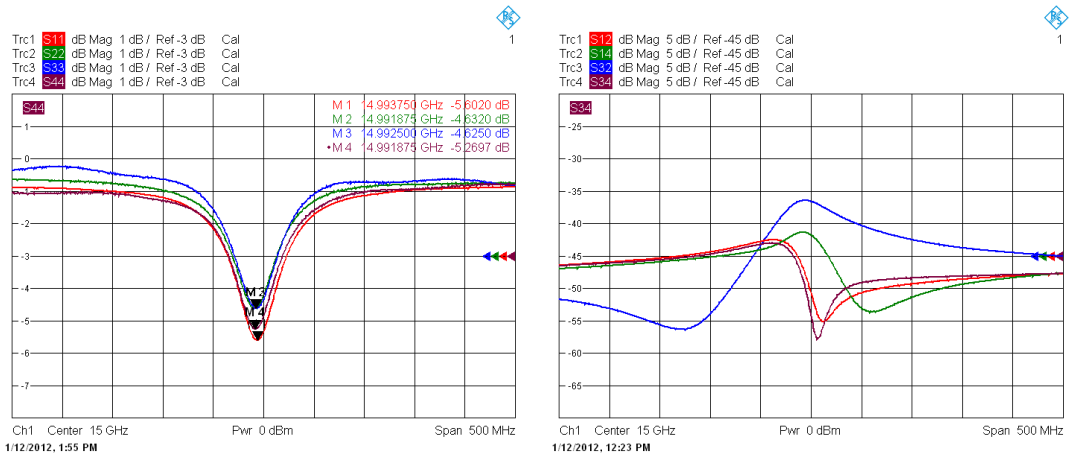

Figure 2.4: Example network analyser traces showing the reflection at the different ports of the position cavity (left) and cross-coupling between them (right).

The resonant frequency of the first dipole mode was measured from the minima in the port reflection spectra. This was 14.993 GHz which is within 10 MHz of the design resonant frequency of 15 GHz. The cross-couplings between adjacent ports (which should couple to orthogonal mode polarisations) were also measured within 200 MHz of the dipole mode resonant frequency. The maximum values in

this bandwidth were -43, -37, -43 and -42 dB so the ports are well isolated.

The maximum specified position offset of 100 μm and resolution of 50 nm [5] requires an operational dynamic range of 66 dB. If the beam has a position offset in one transverse direction, a fraction of the signal coupled out of the cavity will be reflected at the output port. The reflected power will then cross-couple to an adjacent output port and appear as a signal in the other transverse direction. Therefore, providing the reflection at the correct output port is no greater than -25 dB, all but one of the measured cross-couplings are small enough to allow for the required dynamic range.

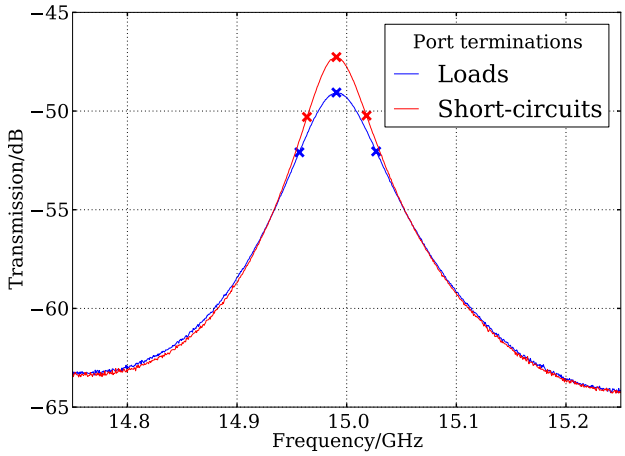

Figure 2.5: Transmission around the frequency of the first dipole mode of the position cavity measured using an antenna probe and one of the couplers. The points mark out the peak transmitted frequency and the measured 3 dB points.

The waveguide to coaxial transition at the end of each waveguide coupler has a finite bandwidth so does not couple out all frequencies equally. Because of the low quality factor and corresponding large bandwidth of the resonant modes in the position cavity, this finite bandwidth would introduce a nonnegligible systematic uncertainty to any quality factor measurement made using the output ports alone. Therefore, an external, weakly coupled coaxial antenna probe, as illustrated in Figure 2.3, was used. Due to the small size of the pick-up aperture, it is impossible to use 2 external probes and so the transmission spectrum from the single probe to each output port was measured.

As illustrated in Figure 2.5, the 3 dB bandwidth of the transmission between the external probe and each output port was measured twice: once with the other output ports terminated with matched loads, and once with them connected to short-circuit terminations. In this manner, an estimate of all of the quality factors (loaded, internal and external) could be made using Equation 1.27. With the other ports loaded, the measured quality factor is the loaded quality factor Q_{L} given, as explained in Section 1.2.1.2, by

$$\frac{1}{Q_{\mathrm{L}}} = \frac{1}{Q_0} + \frac{1}{Q_{\mathrm{ext}}} \tag{2.4}$$

where Q_0 and Q_{ext} are the internal and external quality factors respectively. With the other ports shorted so that they extract no power, the measured loaded quality factor $Q_{\mathrm{L,shorts}}$ will lie between two values. If the mode excited by the external probe is aligned to the waveguide coupler of the output port being used, assuming zero cross coupling, the minimum value given by

$$\frac{1}{Q_{\mathrm{L,shorts}}} = \frac{1}{Q_0} + \frac{1}{2Q_{\mathrm{ext}}} \tag{2.5}$$

will be measured. If there is no coupling between the probe and the output port, the measured quality factor will be at its maximum value and equivalent to the internal quality factor,

$$\frac{1}{Q_{\mathrm{L,shorts}}} = \frac{1}{Q_0}\,. \tag{2.6}$$

Eliminating the internal quality factor in Equation 2.4 using Equation 2.5 and Equation 2.6, maximum and minimum values for the external quality factor can be obtained.

$$Q_{\mathrm{ext,min}} = \frac{Q_{\mathrm{L}}Q_{\mathrm{L,shorts}}}{2(Q_{\mathrm{L,shorts}} - Q_{\mathrm{L}})}\,; \tag{2.7}$$

$$Q_{\mathrm{ext,max}} = \frac{Q_{\mathrm{L}}Q_{\mathrm{L,shorts}}}{Q_{\mathrm{L,shorts}} - Q_{\mathrm{L}}}\,. \tag{2.8}$$

The median of these two cases was calculated and averaged over all the ports to give the measured value for Q_{ext}. The results are summarised in Table 2.2. The averaged median value for Q_{ext} is 830, which is higher than the predicted value

of 727 in Table 2.1.

Coupler	Q_{L}	$Q_{\mathrm{L,shorts}}$	$Q_{\mathrm{ext,min}}$	$Q_{\mathrm{ext,max}}$	$Q_{\mathrm{ext,median}}$
1	235	302	530	1059	794
3	217	270	553	1105	829
2	215	277	480	961	720
4	227	275	650	1300	975
Average	224 ± 9				830

Table 2.2: Results of the quality factor measurement for the dipole mode of the position cavity. The error on the average loaded quality factor is the standard deviation.

Using the average values for the loaded and external quality factors, the internal quality factor calculated using Equation 2.4 is Q_0 =306. This is quite a bit lower than the predicted value of 446, most probably due to surface roughness and the nonideal electrical contacts of the pick-up assembly. There is large variation between the results for the different output ports, even for the loaded quality factor where the choice of port should not make a difference to the power coupled out. This suggests that the bandwidth of the output coupler used for the measurement still has an influence on the result.

2.2.2 Feedthrough Antenna Effects

2.2.2.1 Measurement

The variation in the quality factor measurements of the position cavity suggests that that the output couplers have a large effect on the overall mode parameters. This is because the cavity has a wide bandwidth comparable to that of the couplers. Their influence was therefore investigated.

Figure 2.6 shows the geometry of a waveguide to coaxial transition. The spacing between the end of each feedthrough antenna and the wall of the corresponding waveguide, labelled s, was varied using a number of shims of varying thickness along with the aluminium spacers that were used instead of the feedthrough gaskets. The shims were made from 0.10 mm thick copper and 0.23 mm thick aluminium. At each spacing, the spectra of the transmission between opposite

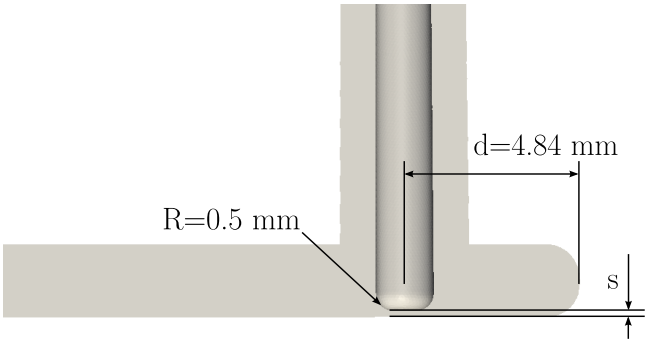

Figure 2.6: Diagram of the waveguide to coaxial transition with the dimensions labelled.

ports were recorded. A Lorentzian curve given by

$$S_{ij}(f) = \frac{A}{1 + \frac{(f-B)^2}{C^2}}\,, \tag{2.9}$$

where f is the frequency and A, B and C are fit parameters, was fitted to the spectra peaks around the resonant frequency. The fits were made in a linear scale. A selection of the fits made at different spacings is shown in Figure 2.7. From the fit parameters, the resonant frequency f_0, bandwidth Δf, peak transmitted amplitude $\hat{S}_{ij}$ and loaded quality factor Q_{L} can be determined as

$$f_0 = B\,, \tag{2.10}$$

$$\Delta f = 2C\,, \tag{2.11}$$

$$\hat{S}_{ij} = A\,, \tag{2.12}$$

$$Q_{\mathrm{L}} = \frac{B}{2C} \tag{2.13}$$

respectively.

Figure 2.8 shows the loaded quality factors measured using the fits for different antenna positions. The internal quality factor can be estimated from the results since, as the separation is increased, less power is coupled out and the external quality factor increases. From Equation 2.4, when the external quality factor is infinite so that there is no coupling to the output ports, the loaded quality factor

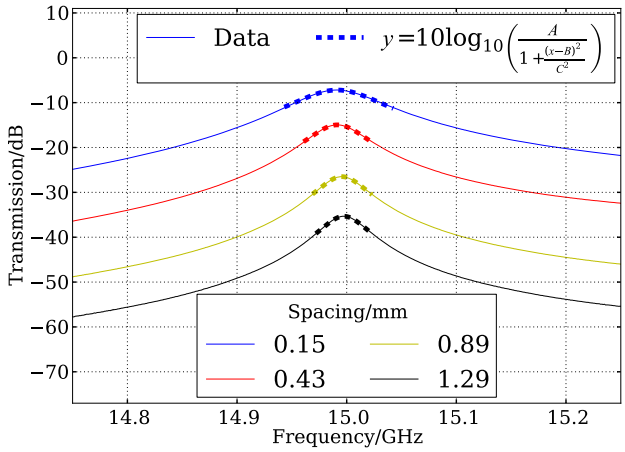

Figure 2.7: Transmission spectra for different feedthrough positions with fits shown by the dashed lines.

is equal to the internal quality factor. Therefore, the loaded quality factor Q_{L} was plotted against the spacing s, and a rising exponential function of the form

$$Q_{\mathrm{L}} = A + B\left(1 - e^{-\frac{s-C}{D}}\right) \tag{2.14}$$

with fit parameters A, B, C and D was fitted. The fit is also shown in Figure 2.8. The internal quality factor was calculated as the sum of A and B, the limit of the loaded quality factor with infinite spacing. For the two pairs of opposite ports, this gave $Q_0 = 357$ and $Q_0 = 352$. These are significantly higher than the value measured in Section 2.2.1, and may be better estimates of the internal quality factor of the dipole mode since the influence of the output coupler is minimised. The point for the smallest spacing was ignored in both fits because it falls in a region where the cavity behaviour is different as is confirmed below.

In order to investigate the effect with higher resolution, shims made from aluminium foil of 0.01 mm thickness were used within certain ranges. With all results combined, the variations in peak transmitted frequency and the measured loaded quality factor are shown in Figure 2.9.

The variation in the loaded quality factor is indeed different with small values for the spacing. Its measured value reaches a minimum with a spacing of about

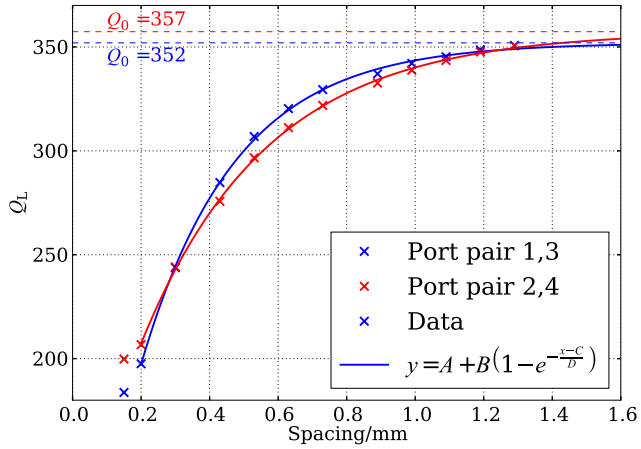

Figure 2.8: Loaded quality factor measured with different spacings between the ends of the feedthrough antennas from the wall of the corresponding waveguide.

0.10 mm, which differs from the design spacing of 0.18 mm. The offset of the peak transmitted frequency from 15 GHz varies by more than 15 MHz between the different values for the spacing and the variation is largest for small values. The peak transmitted frequency also reaches a minimum with a spacing of just over 0.2 mm. Further work in simulation was performed in order to interpret these results more fully and is described in the next section.

2.2.2.2 Simulation

Simulations were performed in order to better understand the influence of the output couplers on the resonant modes of the position cavity. This was done using the eigenmode solver Omega3P [24], part of the ACE3P package discussed in Section 1.2.1.5. The full geometry of the position cavity was simulated with different values for the spacing between each feedthrough antenna and the wall of the corresponding waveguide.

The effect on the first resonant dipole mode, as determined from the simulation results, is shown in Figure 2.10. As in the measurement, the loaded quality factor generally increases with the spacing. The same fit using Equation 2.14 was therefore made and as expected, the predicted loaded quality factor in the limit of infinite spacing is equal to the internal quality factor as determined by

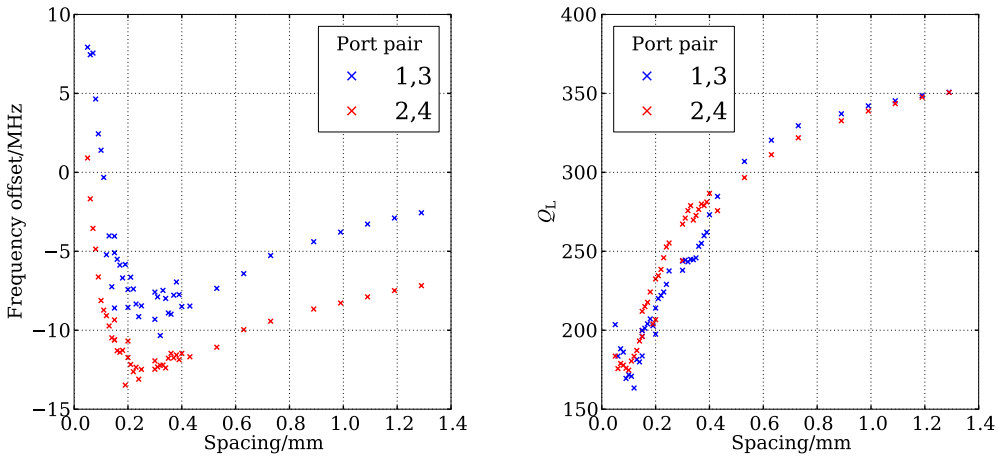

Figure 2.9: Offset of the peak transmitted frequency from 15 GHz for different feedthrough positions (left) and loaded quality factors with more data points (right).

the simulation ($Q_0 = 512$). There is also some deviation from the exponential fit at small values for the spacing as seen in the measurement. Less variation in the resonant frequency of the dipole mode is seen in simulation than in measurement but the overall trend is the same. The simulation model is symmetric and so the frequency difference between the two mode polarisations indicates an asymmetry in the mesh.

More results were obtained for smaller separations in order to investigate the behaviour in this region. The minimum internal quality factor is seen at a separation of 0.05 mm. This is close to the measured result of around 0.10 mm considering the mesh size (minimum edge length=0.04 mm).

The resonances that are part of the resonance antenna couplings are seen in the simulation along with the cavity modes. A visualisation of a feedthrough antenna location where one such resonance is located, is shown in Figure 2.11. These resonances couple magnetically to the cavity dipole mode at the point where the waveguide joins the cavity. As discussed in Section 1.2.1.6, this is analogous to two magnetically coupled resonant electric circuits made from passive components.

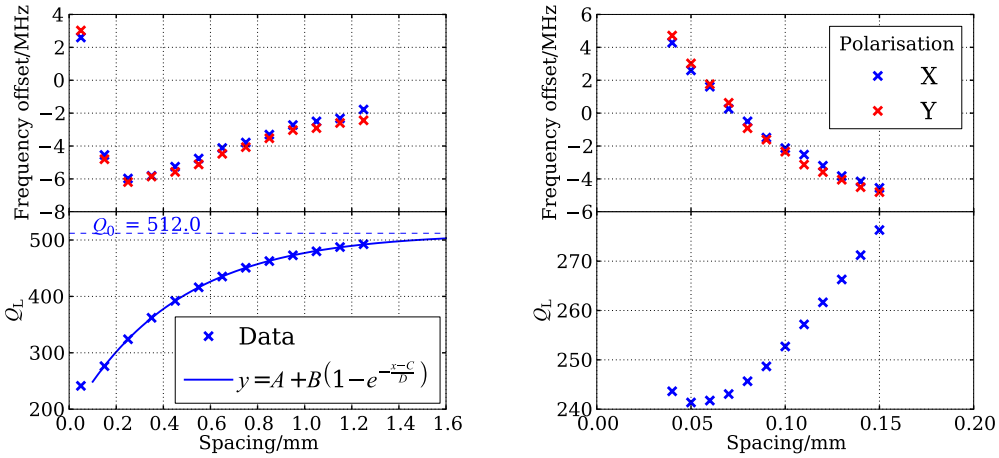

Figure 2.10: Dipole mode loaded quality factor and offset in resonant frequency from 15 GHz over a wide (left) and fine (right) range of simulated antenna positions. Only the results for the horizontal polarisation are shown in the plots of loaded quality factor because the two polarisations are indistinguishable.

The effective impedance Z_{eff} of a circuit representing the cavity is given by

$$Z_{\text{eff}} = Z_{\text{c}} + \frac{(2\pi f)^2 M^2}{Z_{\text{a}}} = Z_{\text{c}} + \frac{(2\pi f)^2 M^2}{|Z_{\text{a}}|^2} Z_{\text{a}}^* \tag{2.15}$$

where f is the driving frequency, M is the mutual inductance and Z_{c} and Z_{a} are the impedances of the uncoupled cavity and antenna circuits respectively [27]. The second term on the right hand side of Equation 2.15 is known as the reflected impedance. It is proportional to the complex conjugate of the antenna circuit impedance. This means that if the driving frequency f is above the resonant frequency of the antenna circuit, the reflected impedance is capacitive and increases the resonant frequency of the cavity circuit. Conversely, if the frequency is lower than the antenna circuit resonant frequency, the reflected impedance is inductive and this decreases the resonant frequency of the cavity circuit. If the frequency is sufficiently far from the antenna resonance, the imaginary part of the antenna circuit impedance, the reactance, becomes larger than the real part, the resistance, and the imaginary part of the reflected impedance is reduced.

The results can therefore be understood in the following way. The measured

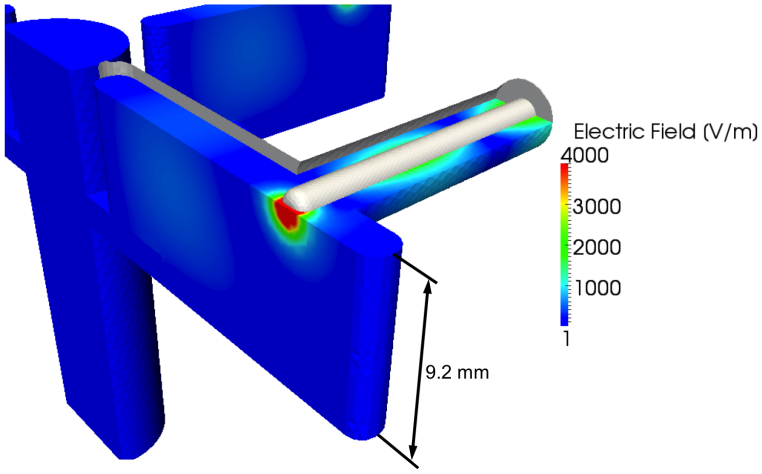

Figure 2.11: Electric field magnitude at the location of a feedthrough antenna resonance that appears in simulation as an eigenmode solution for the full position cavity geometry where there is 4.4×10^{-12} J of energy stored in the mode.

resonant frequency is the resonant frequency of the coupled system. As the spacing between the end of the feedthrough antenna and the wall of the waveguide is enlarged, the frequency of the antenna resonance is increased. With small values for the spacing, the frequency of the antenna resonance is below the cavity dipole mode frequency and so the measured resonant frequency is pushed up. As the separation is increased, the antenna resonance crosses the resonant frequency of the cavity dipole mode and begins to push the measured resonant frequency down. As the frequency of the antenna resonance is increased further, its reactance at the frequency of the cavity dipole mode becomes high enough that the reflected impedance is reduced. At this point, the measured resonant frequency reaches a minimum and then begins to return to the resonant frequency of the uncoupled cavity dipole mode. In the actual system, the antenna also has the function of matching the impedance of the waveguide to the 50 Ω impedance of the coaxial output. Therefore, the maximum output is not necessarily seen when the antenna is at resonance with the cavity dipole mode because power may be reflected back into the waveguide. For the best performance, the waveguide coupler should be designed so that the antenna resonance and the minimum reflection should both occur at the frequency of the cavity dipole mode.

2.2.3 Redesign of Waveguide to Coaxial Transition

Decreasing the sensitivity of the system to the antenna position in the waveguide to coaxial transition would loosen the mechanical tolerances and make the pick-up assembly easier to manufacture and more reproducible. Further simulations were performed in order to find ways to decrease the sensitivity to this dimension.

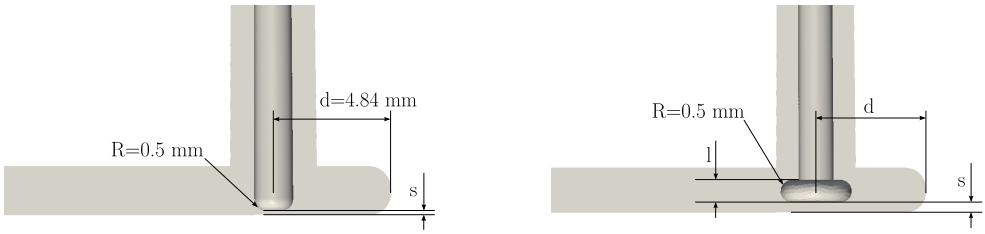

Figure 2.12: Diagram of a waveguide to coaxial transition with the original design (left) and an alternate antenna design with a bead at the end (right).

The original design of the transition between the waveguide and the coaxial using the resonance antenna coupling was based on simulations of the end of the waveguide including the transition, the geometry shown in Figure 2.12. Two ports are defined, the open end of the waveguide and the cross section of the coax, and the simulation predicts the port S-parameters. An improvement on the design of the antenna was attempted using the S-parameter simulation tool S3P, which is part of the ACE3P package [24].

First the old geometry of the transition was simulated and tuned to minimise the reflection from the antenna. The results are shown in Figure 2.13. It can be seen that, for a change in the separation of just 0.02 mm, the spectrum of reflected power changes significantly. The minimum reflection at the matching point does not occur at 15 GHz because the ends of the simulated antennas have a curvature of 0.5 mm, same as the manufactured ones, instead of the design value of 0.4 mm. The separation at the matching point, where the reflection is minimised, is however, still close to the design value of 0.177 mm.

The capacitance between two flat surfaces is proportional to their area and inversely proportional to their separation. Increasing the area of the antenna end should therefore allow for a larger spacing between it and the wall of the waveguide. Any absolute change in the spacing will then be smaller as a percentage

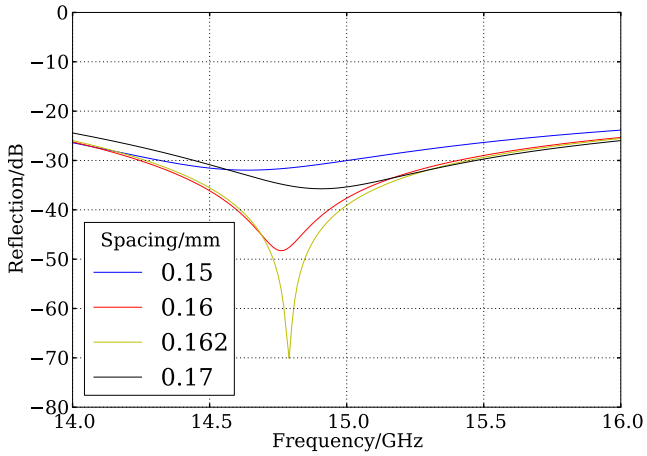

Figure 2.13: Return loss of the actual waveguide to coaxial transition determined from simulations with different waveguide-antenna separations.

Dimension	Value/mm		
	Original	Alternative 1	Alternative 2
Antenna radius	0.8		
Bead radius	N/A	1.6	
Curvature R	0.5		
Bead length l	N/A	1.0	1.2
Distance from waveguide end d	4.84	4.93	4.44
Optimum separation s	0.162	0.45	0.35

Table 2.3: Dimensions for the original and modified design of the waveguide to coaxial transition. Some of the dimensions are illustrated in Figure 2.12.

of the total separation and so the fractional change in capacitance and therefore, resonant frequency, will be reduced. An attempt to loosen the tolerance on the separation in this way was made by adding a flat bead to the end of the antenna design as illustrated in Figure 2.12. Similar features appear in other BPM designs [32]. The antenna radius was kept at 0.4 mm and a bead radius of 0.8 mm was chosen to fit easily inside the 0.9 mm coaxial outer radius. The curvature of 0.5 mm was also taken from the manufactured feedthrough antennas. Two lengths of bead were simulated, the minimum as limited by the curvature, 1 mm, labelled alternative 1, and a slightly larger value of 1.2 mm, labelled alternative 2. These were matched to the waveguide at 15 GHz by varying the distance to the

waveguide end, the dimension labelled d. The final values for all the dimensions are summarised in Table 2.3. The antenna was also shifted to the centre of the waveguide, 7 mm from each edge instead of 9 mm from one and 5 mm from the other as this is where the field of the waveguide mode is strongest.

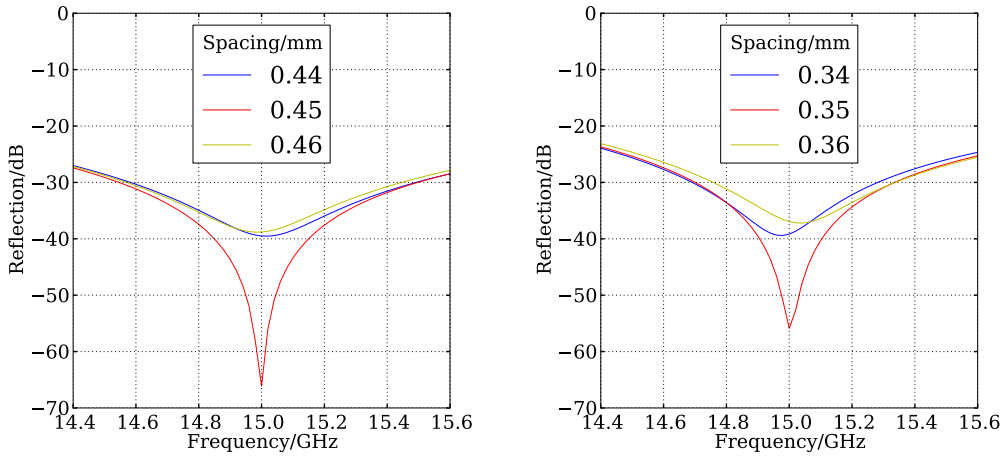

Figure 2.14: Reflection spectra of the waveguide to coaxial transitions determined from simulations with two alternative, beaded feedthrough antennas: 1. 1 mm bead length (left), 2. 1.2 mm bead length (right).

The results of the simulations of the two antennas with beads are shown in Figure 2.14. It can be seen that over a change in spacing of 0.02 mm, the change in the reflection spectrum is much less than in the case where there is no bead. The bead length of 1 mm performs better, with almost no change in the frequency range for which the reflected power is less than -30 dB. The design for the second iteration of the prototype pick-up includes a modified feedthrough antenna with a bead, the details of which are in [3].

The same simulations that were performed in Section 2.2.2.2 were repeated with the new design of feedthrough antenna with a bead length of 1 mm. The results are shown in Figure 2.15 where they have been compared with the original design. For small values of the spacing, there is only a slight reduction in the sensitivity in terms of both the loaded quality factor and the resonant frequency. One further advantage of the new design is apparent, however. The minimum in the resonant frequency is shallower and so, by operating at this point, the

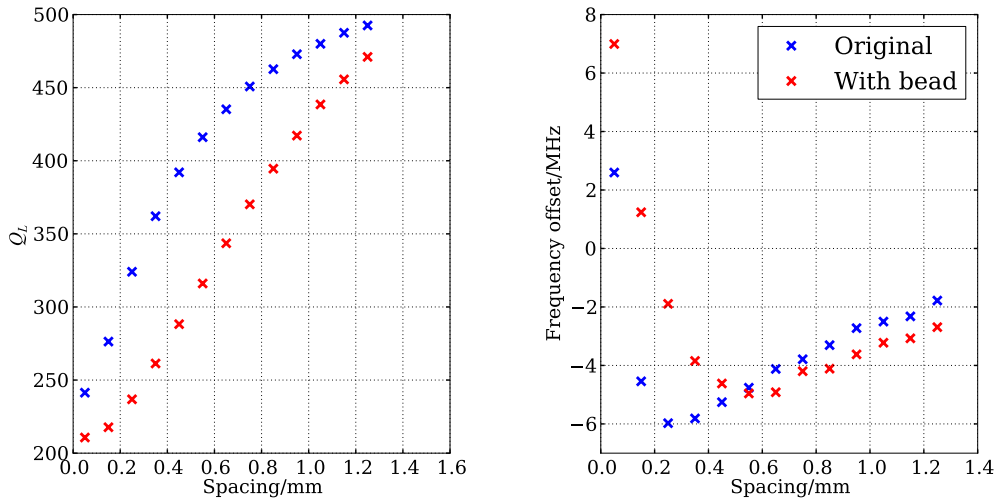

Figure 2.15: Position cavity dipole mode loaded quality factor (left) and resonant frequency (right) for different simulated antenna positions with the original antenna design and a design with a 1 mm long bead added at the end.

tolerances in terms of frequency may be reduced at the expense of increased antenna reflections.

2.2.4 Reference Cavity

2.2.4.1 Measurements

The internal quality factor of the reference cavity first monopole mode was measured using the same setup as for the position cavity with one port of the network analyser connected to an external coaxial antenna positioned inside the pick-up and another connected to one output port of the reference cavity. The one remaining output port was either loaded or short-circuited. Since there are only two output ports and the mode is axially symmetric, the contribution from each $Q_{\mathrm{ext,n}}$ to the total quality factor can be determined when the other port is short-circuited since, in this case

$$\frac{1}{Q_{\mathrm{L,shorts}}} = \frac{1}{Q_0} + \frac{1}{Q_{\mathrm{ext,n}}}\,. \tag{2.16}$$

Figure 2.16 shows the transmission spectra for the two cases. In the instance shown, the transmission between the coaxial antenna and the connected port was too high and so the antenna may have been causing significant changes to the field inside the cavity. Furthermore, there is a clear difference in the frequency of peak transmission between the two cases and a secondary peak at 17.2 GHz that is visible when the disconnected port is loaded and greatly enhanced when it is short-circuited. This suggests that there is a resonant mode close in frequency to the mode being measured, probably caused by the presence of the strongly coupled antenna.

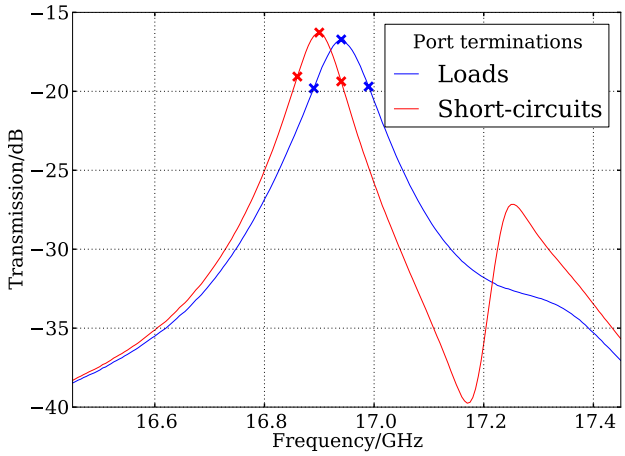

Figure 2.16: Transmission spectra for the monopole mode of the reference cavity excited using an antenna and measured using one coupler with the other coupler terminated in a matched load (left) and shorted (right).

The results of the measurements are summarised in Table 2.4. The measured value for the external and loaded quality factors can then be used to calculate the internal quality factor Q_0 from Equation 2.4 which gives Q_0=343. Unlike in the position cavity measurements, the loaded quality factors measured using each port were identical. This supports the previous assertion that the output coupler is influencing the results of the position cavity measurements.

The resonant frequency of the loaded reference cavity was measured to be 16.940 GHz. This is too far from the 15 GHz design frequency and so the cavity geometry had to be modified mechanically.

Coupler	Q_{L}	$Q_{\mathrm{L,shorts}}$	$Q_{\mathrm{ext,n}}$
1	153	209	578
2	153	215	529
		Q_{ext}	276

Table 2.4: Results of the quality factor measurement for the monopole mode of the reference cavity.

2.2.4.2 Frequency Modification

The resonant frequency of the reference cavity monopole mode was measured to be 16.940 GHz, far higher than the 15 GHz of the position cavity. This was because the length of the reference cavity in the mechanical design was too long by 0.5 mm. Simulations of the incorrect geometry in the eigenmode solver, Omega3P, underestimate the measured resonant frequency by 46 MHz. Since the exact geometry of the original manufactured pick-up, including manufacturing errors, was not known, the simulation model was altered to produce a resonant frequency of 14.954 GHz. Similar alterations to the actual cavity would then result in a resonant frequency close to 15 GHz.

It was possible to modify three different dimensions by extra milling: the

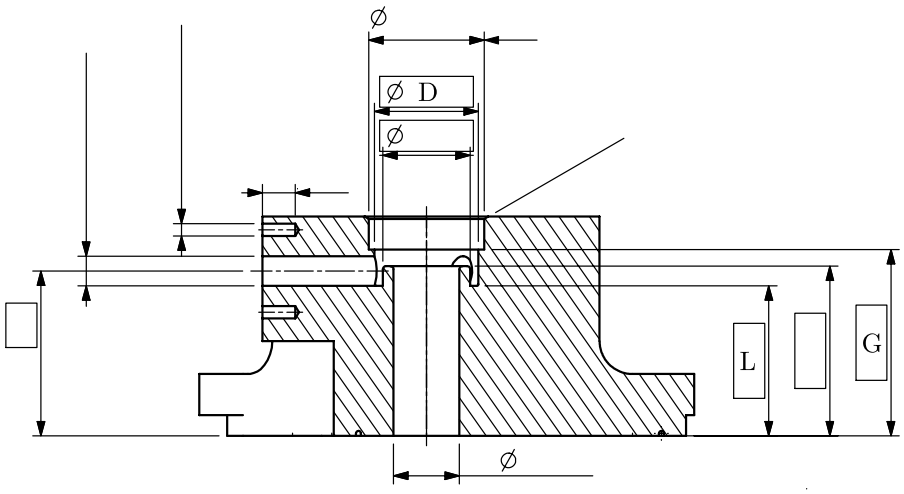

Figure 2.17: The CTF3 prototype BPM intermediate body including the reference cavity with the dimensions that could be modified labelled.

length of the reference cavity, the depth of the trench that makes the cavity reentrant and the cavity diameter. Figure 2.17 shows a mechanical drawing in which three dimensions, D, L and G, have been highlighted. D is the cavity diameter, decreasing G increases the cavity length and decreasing L increases the trench depth.

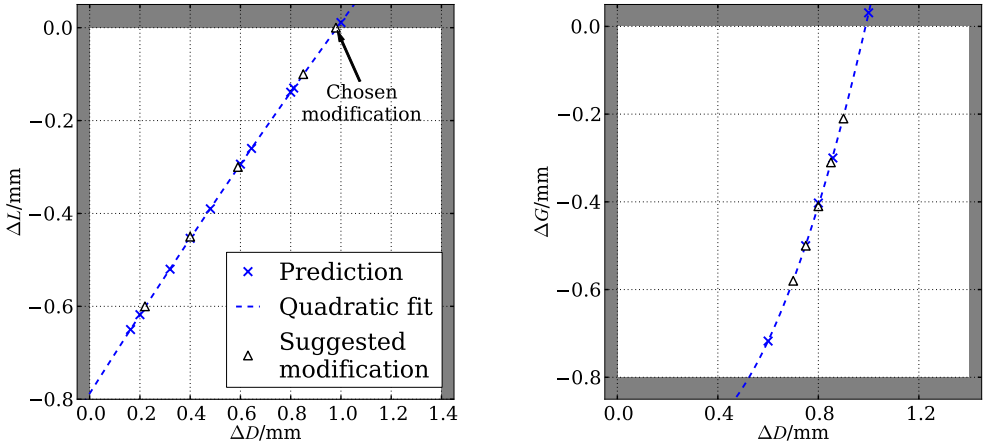

Figure 2.18: Modifications involving a change in D. The blue crosses represent predicted geometries that give the correct resonant frequency. It is predicted that all geometries on the fitted polynomial will also give the correct resonant frequency. Of these, several have been chosen and simulated and these are represented with triangles. The limits on the change in each dimension are indicated by the grey shaded areas.

There were practical limits on the amount by which each dimension could be modified. The modification was to be carried out by milling only, no extra material was to be added. Therefore, D could only be increased while L and G could only be decreased. The top of the cavity is made from a separate cap that is 1.4 mm larger in diameter than the cavity itself. D could not, therefore, be made larger than the cap. G could only be decreased by up to 0.8 mm before it met the holes for the feedthrough antennas. Having these holes partially open would complicate the geometry unnecessarily. The amount by which L could be decreased was essentially unlimited in terms of material that could be cut away.

Possible modifications were determined, each of which involved changing at most, two of the three dimensions. These are displayed by the lines in Figure 2.18

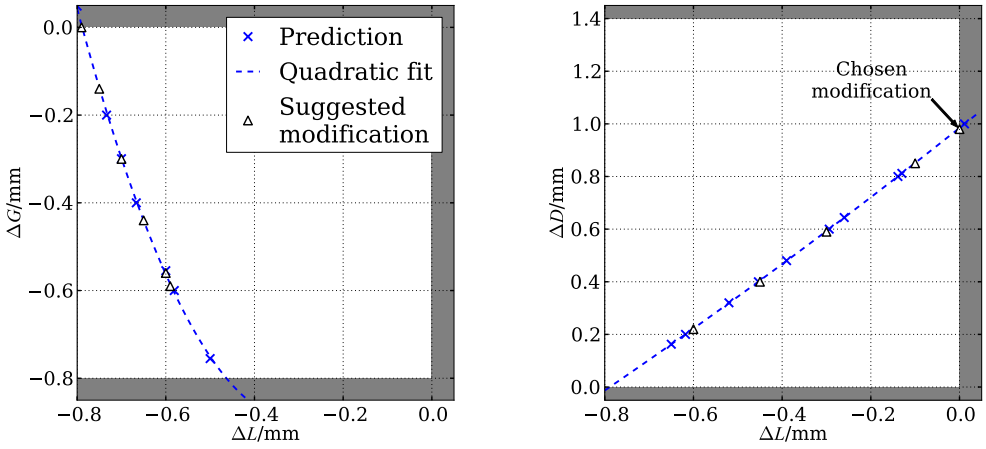

Figure 2.19: Modifications involving a change in L. The plot on the right is equivalent to the left-hand plot in Figure 2.18 but with the axes swapped. It is shown here for clarity.

and Figure 2.19 where the grey regions bound the areas where the modifications are possible. Since the resonant frequency is most sensitive to L and D, these were swept over the possible range that could give the desired outcome. At each point in the sweep, two values for one of the other two dimensions were chosen that would give resonant frequencies above and below the target of 14.954 GHz. The results of these simulations were then used to predict the value for the second dimension that would be required to meet the target. When at least 5 possible combinations had been determined, a 2nd order polynomial was fitted to the predictions. A selection of points on this line where the changes were exact multiples of 10 μm was put forward as suggested modifications. These were all simulated individually and the results are listed in Table 2.5.

The monopole mode resonant frequency is sufficiently sensitive to both L and D that the target could be reached by altering one of these dimensions only. Decreasing L by 0.79 mm would be enough. Increasing D by 0.98 mm would also be sufficient but would leave a rim of only 0.21 mm to support the cavity cap. Because the resonant frequency is not so sensitive to G, if required, greater precision could be achieved with a looser tolerance if it were to be altered in combination with one of the other dimensions. Ultimately, the modification was

ΔL/mm	ΔG/mm	Resonant frequency/GHz	Q_{L}
-0.59	-0.59	14.944	131.3
-0.60	-0.56	14.955	130.6
-0.65	-0.44	14.958	128.2
-0.70	-0.30	14.960	125.9
-0.75	-0.14	14.958	123.8
-0.79	0.00	14.948	122.3

ΔD/mm	ΔL/mm	Resonant frequency/GHz	Q_{L}
0.22	-0.60	14.958	126.3
0.40	-0.45	14.962	130.3
0.59	-0.30	14.959	135.2
0.85	-0.10	14.957	144.1
0.98	0.00	14.963	149.5

ΔD/mm	ΔG/mm	Resonant frequency/GHz	Q_{L}
0.70	-0.58	14.950	154.5
0.75	-0.50	14.950	153.5
0.80	-0.41	14.953	152.5
0.85	-0.31	14.956	151.5
0.90	-0.21	14.953	150.7

Table 2.5: Results from simulation of suggested modifications.

made by simply increasing the cavity diameter by 0.98 mm. A final simulation was performed using the measured dimensions of the manufactured reference cavity rather than the design values. The increase in diameter of 0.98 mm was then applied and the resonant frequency predicted by the simulation was 14.999 GHz. Therefore, the initial difference between the simulation and the measured resonant frequency could be attributed to errors in the mechanical production and the chosen modification was still suitable.

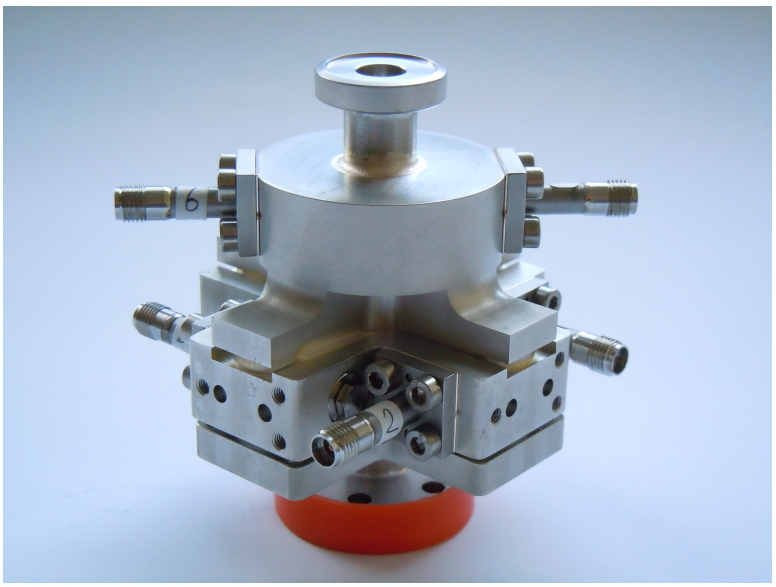

Figure 2.20: Image of the brazed cavity assembly.

After the reference cavity had been modified as described, the resonant frequency and loaded quality factor were measured before the assembly was brazed. The measured resonant frequency was 14.960 MHz and the loaded quality factor was 74. Both of these values are lower than the predicted values but when the measurement was made, only one feedthrough spacer was available and so one of the feedthroughs was inserted further into the cavity body than the other. A simulation was performed with the feedthroughs in these positions. This lowered the resonant frequency by 33 MHz and the loaded quality factor by a factor of two compared with the simulation with the feedthrough in the correct position.

2.2.5 Brazing and Tuning

Figure 2.20 shows an image of the pick-up assembly after the four parts had been brazed together. The measurements of the quality factors and resonant frequencies were repeated after brazing, the results of which are shown in Table 2.6. The

Coupler	Q_{L}	$Q_{\mathrm{L,shorts}}$	$Q_{\mathrm{ext,min}}$	$Q_{\mathrm{ext,max}}$	$Q_{\mathrm{ext,median}}$
1	198	267	383	766	575
3	191	236	501	1002	751
2	209	269	469	937	703
4	193	293	282	565	424
Average	198				613

Coupler	Q_{L}	$Q_{\mathrm{L,shorts}}$	$Q_{\mathrm{ext,n}}$
1	128	191	388
2	131	189	427
Average	130	Q_{ext}	203

Table 2.6: Results of the quality factor measurement for the first resonant dipole mode of the position cavity (top) and for the first resonant monopole mode of the reference cavity (bottom).

results are also summarised in Table 2.7 along with those taken before brazing for comparison. The resonant frequencies have remained close to their design values while the loaded quality factor is clearly reduced in both cavities. The internal quality factors have also changed but within the variation seen between the different ports. This suggests that the coupling has increased while the losses in the cavity walls have not changed by a significant amount.

The tuning bolts that push against the waveguide walls opposite the feedthrough antennas were used to tune the position cavity output ports. The largest effect

Cavity	Position		Reference		
Parameter	Before	Brazed	Before	Modified	Brazed
f_0	14.993	15.012	16.940	14.960	14.997
Q_{L}	224	198	153	74	130
Q_0	306	292	343	-	357

Table 2.7: Results of the mode measurements of both the reference and position cavities before and after brazing.

is seen in the signal reflection. By tuning a pair of opposite ports, it is possible to decrease the reflected power at one while increasing it at the other until they are equal. The deformation of the metal by the tuning bolt is permanent so adjustment is only possible in one direction which may not be the right one. Figure 2.21 shows the final spectra of the port reflections and cross-couplings. Despite the tuning, there is more variation in the reflections of the different ports than before brazing (Figure 2.4). The cross coupling is also worse between all the ports, although it is still lower than -33 dB for all port pairs and lower than -35 dB for three of the four. Ports 1 and 4 show the highest reflections and the poorest cross-coupling and so they were terminated with matched loads during the beam tests while ports 2 and 3 were used to measure the dipole mode signals.

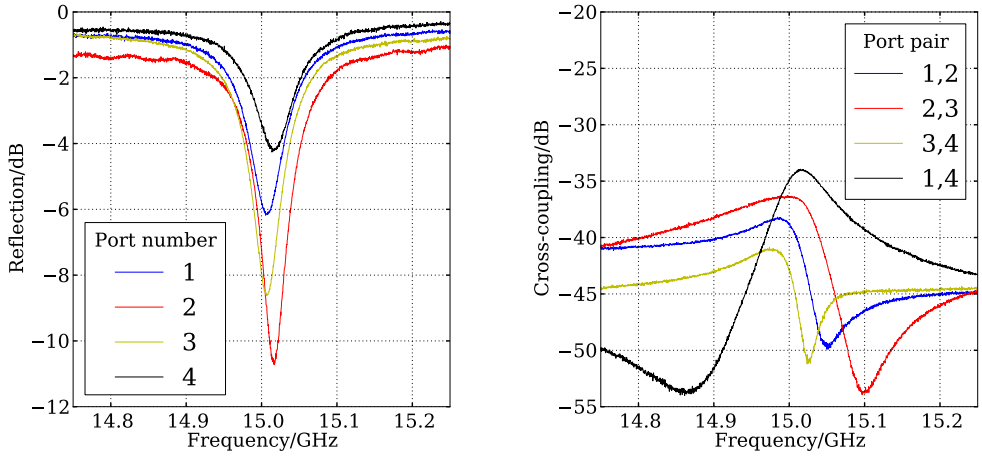

Figure 2.21: Final reflected power spectra of all the ports (left) and cross coupling between them (right).

2.2.5.1 Influence of Short-circuited Ports

The measurements of the internal quality factor made in Sections 2.2.1, 2.2.4.1 and 2.2.5 assume that when a port is terminated with a short-circuit, it is removed so that is has no impact on the power loss in the pick-up. In reality, reflections off of a short-circuit termination and resonances excited inside a short-circuited port can alter the measured resonant frequency and quality factor of a cavity mode. Furthermore, in the coaxial section just before each of the pick-up's output ports,

Cavity	Q_0		Change /%
	without dielectric	with dielectric	
Position	512	479	-6.4
Reference	379	353	-6.9

Table 2.8: Internal quality factors of the modes of interest in the position and reference cavities as determined from simulations with and without the transition to the dielectric signal medium included.

there is a transition in the signal medium from vacuum to Boron Nitride dielectric where further reflections may occur. In order to estimate the influence of this transition and the short-circuiting of the ports, simulations of both the position and reference cavities were performed using ACE3P. First, the length of the coaxial sections was varied to simulate short-circuiting of the ports at different electrical lengths. This effect was found to be insignificant in terms of the change in the internal quality factor (less than 2 % for a change in length of over 30 mm for the position cavity). Next, each simulation was run twice, once with the simulated geometry ending at the transitions from vacuum to dielectric and once with the dielectric in the feedthroughs included. The transition to the dielectric (permittivity $\epsilon = 4.2\epsilon_0$ where ϵ_0 is the permittivity of free space) is accompanied by a step change in the radius of the central conductor of the coax to 0.33 mm so that the impedance after the transition remains at 50 Ω.

Table 2.8 shows the results of the simulations with and without dielectric for the two cavities. The effect is predicted to be similar for both cavities, a decrease of around 6.5 % in the internal quality factor when the effects of the dielectric are included. For the reference cavity, where the power loss from the monopole mode is dominated by the external quality factor, this has very little effect on the external quality factor that would be measured, an overestimate by just over 1 %. In the position cavity, however, the predicted overestimate in the external quality factor is around 12 %. This would lead to an underestimate of the predicted sensitivity by 6 %.

2.2.6 Frequency Variation with Temperature

Because of the high beam power and power dissipation expected throughout the CLIC tunnel, large temperature changes of up to 20 °C are expected during beam operation. The variation in the resonant frequencies with temperature was therefore investigated, both before and after brazing of the pick-up assembly.

Before brazing, since the reference cavity had not yet been modified, only the resonant dipole mode frequency in the position cavity was measured. The setup was left over one weekend, over which, two clear fluctuations in the ambient temperature and a net change of -6 °C were observed. The electronic temperature probe used was a HumidiProbe from Pico Technology [33], which can measure the ambient temperature with a resolution of 0.1°C. The Lorentzian fit of Equation 2.9 was applied to the port to port transmission spectra so that the resonant frequency could be determined with greater precision than the resolution bandwidth of the network analyser. A temperature dependence of -243 kHz $°C^{-1}$ was observed as shown in Figure 2.22. The maximum 20 °C temperature shift therefore corresponds to a change in frequency of about 5 MHz. How this affects the phase stability of the BPM depends on how the frequency of this mode changes relative to the monopole mode in the reference cavity, whose temperature stability was not measured before brazing. The hysteresis from the heating and cooling of the cavity material during the two temperature fluctuations is clearly visible in the fit since the temperature probe measured the ambient temperature and was not in contact with the cavity body.

After the assembly was brazed, the temperature stability was measured again. This time, both cavities were included in the measurement. The results are shown in Figure 2.23. The resonant frequency of the position cavity dipole mode is now significantly more sensitive to temperature with a dependence of -359 kHz $°C^{-1}$. The reference cavity monopole mode frequency is less sensitive at -308 kHz $°C^{-1}$, giving the stability of the difference in mode frequencies as -51 kHz $°C^{-1}$. Over the 20 °C maximum temperature shift, this corresponds to a relative frequency shift of 1 MHz. This is significant for measurements of one or a few bunches but as discussed later in Section 2.3, should not be a problem for multiple bunch measurements. This is because as the number of bunches is increased, the harmonic

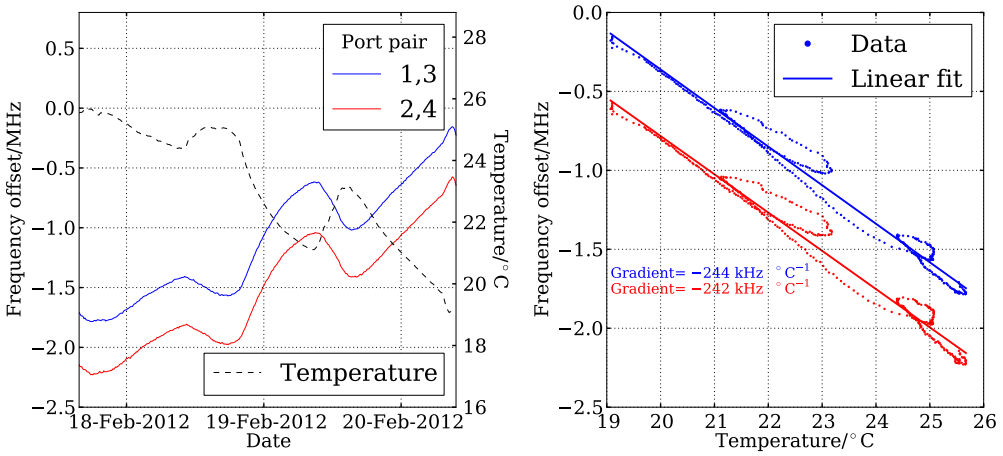

Figure 2.22: Variation in the ambient temperature and resonant frequency of the position cavity first dipole mode over three days (left) and correlation of the two (right).

of the bunch arrival frequency that is closest to the cavity resonant frequency becomes dominant and so shifts in temperature will not affect the signal frequency. A change in the resonant frequency of 1 MHz will have an effect on the signal amplitude but this will not be larger than 2 %.

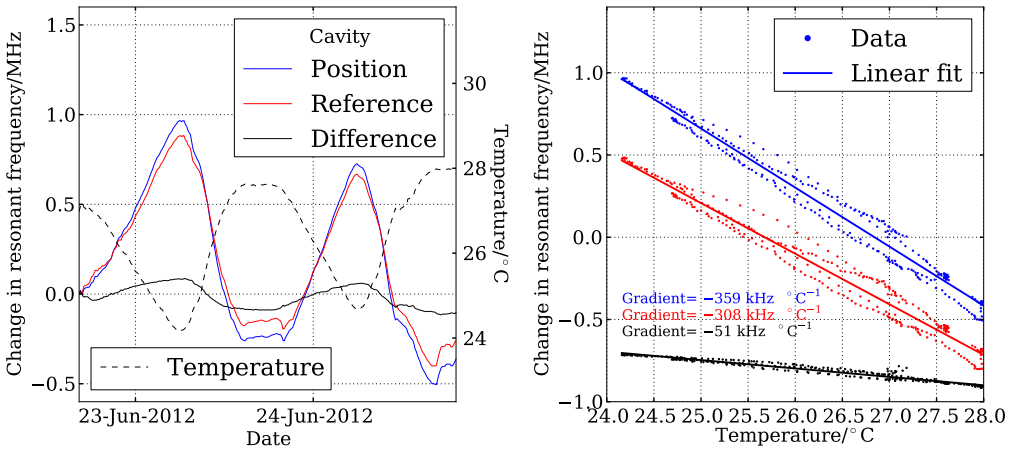

Figure 2.23: Variation in the ambient temperature and resonant frequencies of the modes of interest in both cavities (left) and correlation (right) where the data for the different fits has been offset for clarity.

The effect of the cavity temperature was also simulated. The mesh was scaled according to the thermal expansivity of stainless steel (16.5×10^{-6} °C^{-1}) and solved again for the resonant frequencies. The results for the reference and position cavities were both exactly -247 kHz °C^{-1} which is close to the measured value for the position cavity before brazing and equal to the expansivity multiplied by the resonant frequency as might be assumed. The large increase in the sensitivity of the resonant frequency to temperature after brazing suggests that the assumption that all dimensions of the cavity are scaled equally by a change in temperature is not valid. There may be an additional effect on the vacuum sealing of the feedthroughs for example.

2.2.7 Sensitivity

The sensitivity of the monopole mode in the reference cavity to charge and the sensitivity of the dipole mode in the position cavity to charge and position have a large effect on the ultimate performance of the cavity BPM. A high sensitivity will make it possible to resolve a smaller change in beam position above the noise. If the sensitivity is too high, the electronics will be saturated by large beam charges or large offsets in beam position and attenuation will have to be added to maintain the range over which the BPM performs effectively.

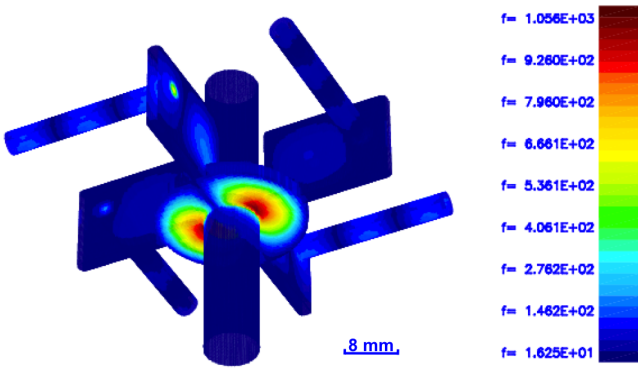

Figure 2.24: Distribution of the electric field magnitude of the eigenmode corresponding to the first cavity dipole mode in the position cavity with 1.3×10^{-12} J of energy stored in the cavity as determined using GdfidL.

Eigenmode simulations were performed using both Omega3p and GdfidL in

order to determine the normalised shunt impedances R/Qs of the modes of interest in the pick-up's two cavities. A visualisation of the position cavity dipole mode is shown in Figures 2.24. In the case of the dipole mode, the normalised shunt impedance was calculated by applying Equation 1.34 to the field distributions of the simulated eigenmodes in both polarisations. The integration was performed along the beam axis at an offset of 1 mm in the direction of the mode polarisation. For the reference cavity monopole mode, the integration was performed along the beam axis with no offset. In both cases, the integration was performed along the whole length of the model to account for the evanescent fields in the beam-pipe. The sensitivity of the pick-up in terms of beam charge and offset can then be estimated by applying Equation 1.41 and using the measured external quality factors for the two cavity modes taken from Table 2.6. The results for the sensitivity estimations are summarised in Table 2.9. They assume the use of

Cavity	Code	$(R/Q)/\Omega$	Measured Q_{ext}	Sensitivity/V nC^{-1}
Position	Omega3P	3.27	617	24.2 mm^{-1}
	GdfidL	3.26		24.2 mm^{-1}
Reference	Omega3P	50.7	203	167
	GdfidL	50.3		165

Table 2.9: Normalised shunt impedances calculated from electromagnetic field distributions determined from simulation and the resulting cavity sensitivities.

only one of the two ports for each cavity mode, which leads to a reduction in the sensitivity by a factor of $1/\sqrt{2}$.

2.2.7.1 Non-Signal Modes

Electromagnetic modes that are not used to determine the beam position are also excited by the beam and can contaminate the BPM signals. Table 2.10 lists those that contribute most to the pollution alongside the signal modes themselves for comparison. The parameters of each mode, including the modes of interest, have been determined using ACE3P simulations of the final, modified pick-up geometry. The table includes the first and second monopole modes and the first quadrupole mode TM_{210} in the position cavity (TM_{010}, TM_{020} and TM_{210} respectively) as well as the first dipole TM_{110} and second monopole TM_{020} modes

	Mode	Frequency/GHz	Q_L	Q_{ext}	$(R/Q)/\Omega$
Position	TM_{010}	11.15	419	1.63×10^{10}	46.4
	TM_{110}	14.99	301	727	3.27
	TM_{210}	20.45	100	122	0.0657
	TM_{020}	25.89	811	5.05×10^{7}	7.98
Reference	TM_{010}	15.00	150	247	50.7
	TM_{110}	19.17	66	78	4.40
	TM_{020}	29.27	207	250	1.47

Table 2.10: Dominant modes in the position and reference cavities of the prototype pick-up as determined using ACE3P. The signal modes are highlighted in yellow.

in the reference cavity. For the dipole and quadrupole modes, the normalised shunt impedances R/Q have been evaluated at a 1 mm offset in the direction of maximum increase in the longitudinal electric field E_z. For the quadrupole mode, this is at 45 degrees to the offset used for the dipole modes as can be seen by comparing the visualisation in Figure 2.25 to the on e in Figure 2.24. This is

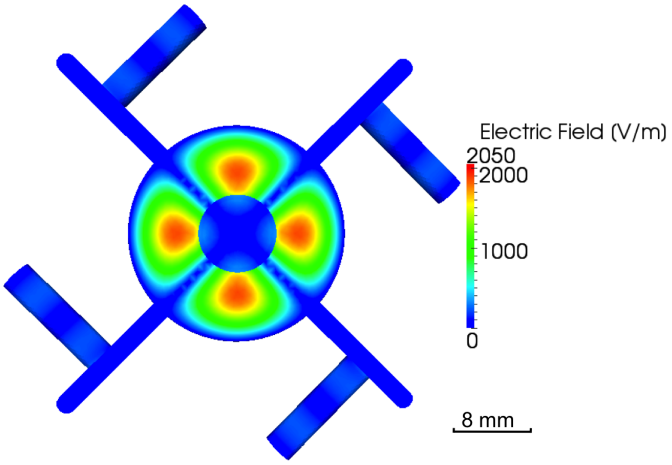

Figure 2.25: Distribution of the electric field magnitude in the first quadrupole mode of the position cavity with 4.4×10^{-12} J of stored energy as determined using ACE3P.

the mode orientation that couples most strongly to the waveguide couplers and is therefore, the most significant. For small offsets, the R/Q of the quadrupole mode is approximately proportional to the fourth power of the offset at which it is calculated and the R/Qs of all monopole modes are approximately offset independent.

Table 2.11 predicts the sensitivity of a single pick-up output channel to each

	Mode	Sensitivity/V nC^{-1}	Voltage rejection at 15 GHz	Long pulse gain factor
Position	TM_{010}	0.00933	2.2×10^{-5}	0.52
	TM_{110}	22.3 mm^{-1}	1	9.4
	TM_{210}	7.42 mm^{-2}	1.9×10^{-5}	0.66
	TM_{020}	0.160	7.2×10^{-7}	0.71
Reference	TM_{010}	150	1	5.3
	TM_{110}	101 mm^{-1}	7.4×10^{-4}	0.96
	TM_{020}	49.4	6.4×10^{-6}	0.57

Table 2.11: Output sensitivities of the dominant modes excited in the position and reference cavities of the prototype pick-up along with their rejection at the signal frequency and the gain factor from their interference when excited by a long train of bunches at a 1.5 GHz bunch arrival frequency. The signal modes are highlighted in yellow.

of the modes listed in Table 2.10 at their resonant frequency using Equation 1.41. Their contribution at the signal frequency, where they are not at resonance is significantly lower. The rejection of each mode at the signal frequency is therefore, also listed in Table 2.11, where a Lorentzian frequency response has been assumed. Additionally, excitations of a given mode by consecutive bunches interfere either destructively or constructively according to the signal frequency. The multiple bunch gain factor of each mode, which gives the output voltage when excited by a long train of bunches as a multiple of the single bunch output voltage and which is defined later by Equation 2.23, is also listed in the table for a bunching frequency of 1.5 GHz. It can be seen that for every non-signal mode, the interference is destructive while for the modes of interest, it is constructive. This type of analysis has previously been carried out for the design of the pick-up for CLIC itself [30]. Modes at frequencies other than the signal frequency can be further rejected using filters in the analogue processing, as discussed later in Section 3.1.1.

2.3 Multiple Bunch Excitation

The CLIC bunch train is unique because of the high 2 GHz bunching frequency. Any hypothetical beam diagnostic designed to measure each bunch with no interference from preceding bunches would require a very large bandwidth that would

be impractical for a cavity BPM. The overlap of BPM signals from consecutive bunches is therefore unavoidable. It is important, for the hardware design, to understand how these signals are summed inside the cavity.

The cavity modes excited by a charged particle bunch passing through a cavity oscillate at their resonant frequency f_0 and decay exponentially with a decay time τ. The output signal is the voltage across an output port where a fraction of the mode energy is coupled out. When multiple bunches pass through the cavity separated by a time interval t_b, their signals overlap and the resulting output voltage V_out is a sum of exponentially decaying oscillators. If the signal from each consecutive bunch is of the same initial amplitude A_0 and phase ϕ_0, the summed signal after the Nth bunch is given by

$$V_\mathrm{out}(N) = \sum_{n=0}^{N-1} A_0 e^{-\frac{nt_\mathrm{b}}{\tau}} e^{i(n\omega t_\mathrm{b}+\phi_0)}. \tag{2.17}$$

where ω is the cavity mode resonant angular frequency and is given by $\omega = 2\pi f_0$. After the addition of each bunch, the signal will evolve in time in the same way as a single bunch with initial amplitude $A(N)$ and phase $\phi(N)$:

$$V(t) = V_\mathrm{out}(N) e^{-\frac{t}{\tau}} e^{i\omega t} = A(N) e^{i\phi(N)} e^{-\frac{t}{\tau}} e^{i\omega t}. \tag{2.18}$$

Equation 2.17 is the summation of a geometric series [26]. The summation can be evaluated to give an analytic expression for the amplitude and phase of the multiple bunch signal after N bunches. With reference to the single bunch amplitude A_0 and phase ϕ_0, these are

$$\frac{A(N)}{A_0} = \sqrt{\frac{1 - 2e^{-\frac{Nt_\mathrm{b}}{\tau}} \cos(N\omega t_\mathrm{b}) + e^{-\frac{2Nt_\mathrm{b}}{\tau}}}{1 - 2e^{-\frac{t_\mathrm{b}}{\tau}} \cos(\omega t_\mathrm{b}) + e^{-\frac{2t_\mathrm{b}}{\tau}}}} \text{ and} \tag{2.19}$$

$$\phi(N) - \phi_0 = \tan^{-1}\left(\frac{e^{-\frac{t_b}{\tau}}\sin(\omega t_b) + e^{-\frac{(N+1)t_b}{\tau}}\sin([N-1]\omega t_b) - e^{-\frac{Nt_b}{\tau}}\sin(N\omega t_b)}{1 - e^{-\frac{t_b}{\tau}}\cos(\omega t_b) + e^{-\frac{(N+1)t_b}{\tau}}\cos([N-1]\omega t_b) - e^{-\frac{Nt_b}{\tau}}\cos(N\omega t_b)}\right). \tag{2.20}$$

In the case where the cavity mode resonant frequency $\omega/(2\pi)$ is a harmonic of the bunching frequency $1/t_b$, the signals from consecutive bunches add constructively. The phase is the same after the arrival of each bunch and the amplitude is an exponential rise given by

$$\frac{A(N)}{A_0} = \frac{1 - e^{-\frac{Nt_b}{\tau}}}{1 - e^{-\frac{t_b}{\tau}}} \text{ for } \frac{\omega}{2\pi} = \frac{k}{t_b} \tag{2.21}$$

where k is a positive integer. Conversely, when the cavity resonant frequency is a half integer multiple of the bunch arrival frequency, the signals add destructively and the amplitude after N bunches is given by

$$\frac{A(N)}{A_0} = \frac{1 + e^{-\frac{Nt_b}{\tau}}}{1 + e^{-\frac{t_b}{\tau}}} \text{ for } \frac{\omega}{2\pi} = \frac{k + \frac{1}{2}}{t_b}. \tag{2.22}$$

Equation 2.19 and Equation 2.20 were evaluated for up to 150 bunches for a decay time τ of 10 ns and a bunching frequency $1/t_b$ of 1.5 GHz. These values were chosen to be similar to the CTF3 prototype cavity BPM but with a longer decay time so that the convergence is not too fast. The results are shown in Figure 2.26. It can be seen that for a resonant frequency of 15 GHz, a harmonic of the bunching frequency, the rise in amplitude is purely exponential; the amplitude goes as Equation 2.21 and there is no change in the relative phase with the addition of more bunches. If there is some frequency offset, there is beating between the bunch arrival harmonic and the cavity resonant frequency, which shows up in the amplitude and phase of the multiple bunch signal.

The common ratio of the geometric series in Equation 2.17, $e^{-\frac{t_b}{\tau}+i\omega t_b}$ has a magnitude less than one. The summation therefore converges for an infinite

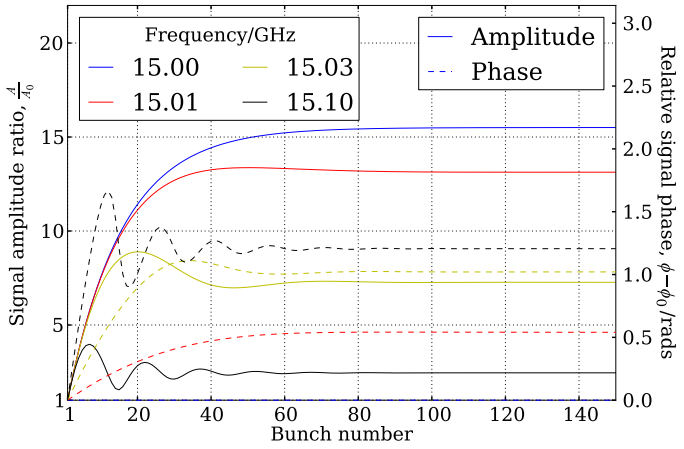

Figure 2.26: Signal amplitude and phase at the arrival of each bunch compared with the single bunch case.

number of bunches. The limits in amplitude and phase are given by

$$\lim_{N\to\infty}\frac{A(N)}{A_0}=\frac{1}{\sqrt{1-2e^{-\frac{t_\mathrm{b}}{\tau}}\cos(\omega t_\mathrm{b})+e^{-\frac{2t_\mathrm{b}}{\tau}}}}\text{ and} \tag{2.23}$$

$$\lim_{N\to\infty}\phi(N)-\phi_0=\tan^{-1}\left(\frac{e^{-\frac{t_\mathrm{b}}{\tau}}\sin(\omega t_\mathrm{b})}{1-e^{-\frac{t_\mathrm{b}}{\tau}}\cos(\omega t_\mathrm{b})}\right). \tag{2.24}$$

The result of Equation 2.25 is referred to here as the multiple bunch gain factor for decay time τ and angular frequency ω. In the case where the cavity resonant frequency is a harmonic or a half integer multiple of the bunching frequency, the limits are given by

$$\lim_{N\to\infty}\frac{A(N)}{A_0}=\frac{1}{1-e^{-\frac{t_\mathrm{b}}{\tau}}}\text{ for }\frac{\omega}{2\pi}=\frac{k}{t_\mathrm{b}} \tag{2.25}$$

$$\lim_{N\to\infty}\frac{A(N)}{A_0}=\frac{1}{1+e^{-\frac{t_\mathrm{b}}{\tau}}}\text{ for }\frac{\omega}{2\pi}=\frac{k+\frac{1}{2}}{t_\mathrm{b}} \tag{2.26}$$

where k is again a positive integer. The speed of the convergence, is only dependent on the signal decay time. Figure 2.26 shows that for all values of the cavity mode resonant frequency, the signal amplitude and phase approach their limits

after roughly the same number of bunches.

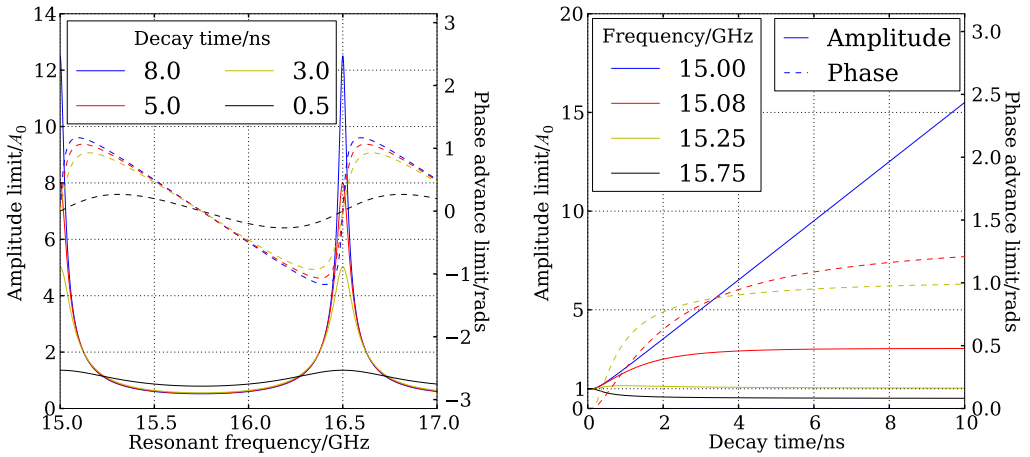

Figure 2.27: Convergence limits of the multiple bunch signal amplitude and phase against cavity mode resonant frequency for different decay times (left) and against signal decay time for different resonant frequencies (right).

The convergence at different resonant frequencies and decay times is plotted in Figure 2.27 by evaluating Equation 2.23 and Equation 2.24. The peaks where the resonant frequency is a harmonic of the bunch arrival frequency are clearly visible. Shortening the decay time means there is less variation with frequency since there is less overlap between consecutive bunches. The variation with decay time τ shows that at all resonant frequencies, if the decay time is 0, the signal amplitude limit becomes the same as the single bunch amplitude. At the resonant frequency, the relationship between the signal amplitude limit and the signal decay time approaches proportionality for long decay times with the bunching frequency as the gradient. This is expected since applying a first order Taylor expansion to the exponential in Equation 2.25 gives

$$\frac{A(N)}{A_0} \approx \frac{1}{1-(1-\frac{t_\mathrm{b}}{\tau})} = \frac{\tau}{t_\mathrm{b}}. \tag{2.27}$$

and this is valid when $\tau \gg t_\mathrm{b}$. For resonant frequencies that are not a harmonic of the bunch arrival frequency, the amplitude limit will level off for infinite τ. When the cavity mode resonant frequency is a sixth-integer multiple of the bunch arrival

frequency (15.025 GHz here), the amplitude limit tends to the single bunch amplitude for large decay times. For frequencies that are further than this from the nearest harmonic of the bunching frequency, there will be destructive interference if the decay time is long enough. When the cavity mode resonant frequency is a half-integer multiple of the bunch arrival frequency, the most destructive interference is seen. As expected from Equation 2.26, the multiple bunch signal amplitude for long decay times will then tend to half the single bunch amplitude.

Another consequence of the convergence of Equation 2.17 is that the signal phase at the arrival of each bunch is the same. This means that the signal becomes periodic at the bunch arrival frequency and this will show up in the frequency spectrum.

2.3.1 Simulation

The equations in the previous section give amplitudes and phases at the arrival of each bunch but do not describe the signal between bunches. This was investigated separately by adding up single bunch signals numerically. This shows how much the periodicity of the signal at the bunch arrival frequency influences the signal compared to the cavity mode resonant frequency.

To simulate the excitation of a cavity mode with a resonant frequency of 15.1 GHz and a decay time of 4.8 ns, decaying complex waveforms were generated with these parameters. They were added together with the start of each waveform coming 2/3 ns after the previous one to simulate the 1.5 GHz bunch arrival frequency of the CTF3 probe beamline. Because the waveforms are complex, the amplitude is easy to determine and the phase can be extracted by simply multiplying by a complex local oscillator. The results are shown in Figure 2.28. The frequency used for the local oscillator was 15 GHz. Although this is not the frequency of the individual waveforms, the overall phase becomes flat. This is the convergence of Equation 2.20 which means the phase is the same at the start of each added waveform. Between the beginning of each waveform and the addition of the next, the phase processes linearly at the angular frequency offset $(2\pi) \times 100$ MHz, which accounts for the oscillations seen in the phase. The overall trend in the signal amplitude and phase is the same as in Figure 2.26.

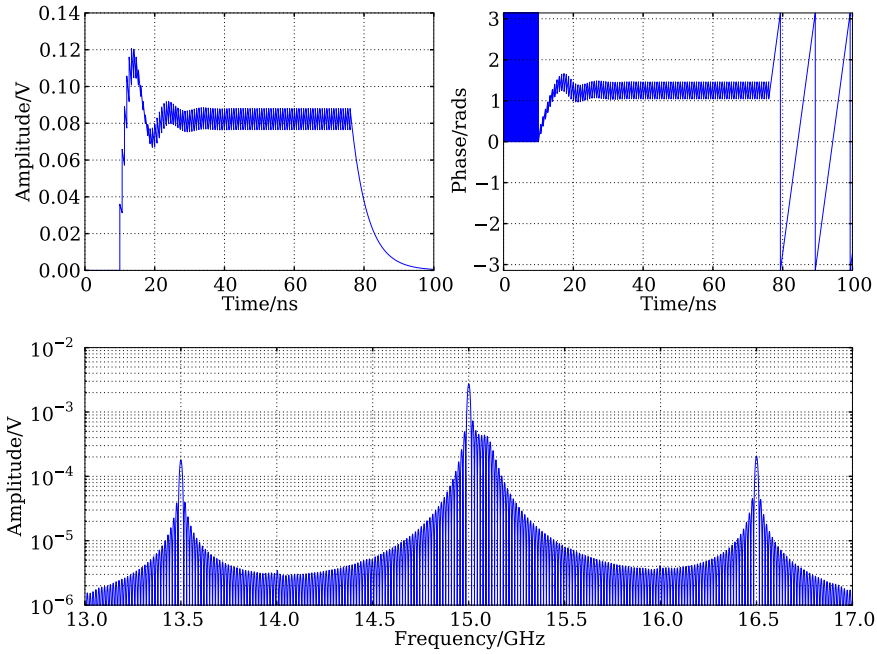

Figure 2.28: Addition of decaying complex oscillators offset by the bunch separation time: amplitude (top left), phase measured relative to a 15 GHz complex local oscillator (top right) and multiple bunch signal spectrum (bottom).

It is clear from the fast Fourier transform (FFT) of the summed signal that the dominant signal frequency is the harmonic of the simulated bunch arrival frequency that is closest to the resonant frequency of the cavity mode. Despite the 100 MHz frequency offset, the individual waveform frequency appears only as a slight shoulder on the central peak. The other harmonics of the bunch arrival frequency are also present but at lower amplitude.

2.3.2 Signal Energy

It can be useful to measure the total energy extracted from a cavity BPM pick-up instead of the maximum signal amplitude because this is a conserved quantity that is not affected greatly by finite processing bandwidths. The total energy can be determined by integrating the output power P_{out} which is given in terms of

the peak output voltage amplitude V_{out} as

$$P_{\text{out}} = \frac{V_{\text{out}}^2}{2Z} \tag{2.28}$$

where Z is the load impedance. In the case of a single bunch excitation, the energy is evaluated by integrating over the signal envelope [12] as

$$E_0 = \int_0^\infty P_{\text{out}} e^{-\frac{2t}{\tau}} dt = P_{\text{out}} \frac{\tau}{2} . \tag{2.29}$$

In the case of multiple bunches, the output power at the arrival of bunch n goes as the squared amplitude ratio $A(n)/A_0$ and decays exponentially over the time t_{b} between bunches. After the arrival of the final bunch, the signal decays to indefinitely. The total energy for N bunches is then given by

$$E(N) = \sum_{n=0}^{N-1} \left[\left(\frac{A(n)}{A_0} \right)^2 \int_0^{t_{\text{b}}} P_{\text{out}} e^{-\frac{2t}{\tau}} dt \right] + \left(\frac{A(N)}{A_0} \right)^2 \int_0^\infty P_{\text{out}} e^{-\frac{2t}{\tau}} dt . \tag{2.30}$$

In the case where the cavity mode frequency is a harmonic of the bunch arrival frequency, Equation 2.21 is used for $A(n)/A_0$, which means that the summation before the first integral includes the sum of two geometric series. Evaluating the summation and the integral gives

$$\begin{aligned} E(N) = & \left[N - 2 \left(\frac{1 - e^{-\frac{Nt_{\text{b}}}{\tau}}}{1 - e^{-\frac{t_{\text{b}}}{\tau}}} \right) + \left(\frac{1 - e^{-\frac{2Nt_{\text{b}}}{\tau}}}{1 - e^{-\frac{2t_{\text{b}}}{\tau}}} \right) \right] \frac{P_{\text{out}} \tau \left(1 - e^{-\frac{2t_{\text{b}}}{\tau}} \right)}{2 \left(1 - e^{-\frac{t_{\text{b}}}{\tau}} \right)^2} \\ & + P_{\text{out}} \left(\frac{1 - e^{-\frac{Nt_{\text{b}}}{\tau}}}{1 - e^{-\frac{t_{\text{b}}}{\tau}}} \right)^2 \frac{\tau}{2} \end{aligned} \tag{2.31}$$

for the total signal energy. When the cavity mode frequency is not an exact multiple of the bunch arrival frequency, Equation 2.19 is used for $A(n)/A_0$ and so the summation includes the sum of three geometric series and Equation 2.30

evaluates to

$$
\begin{aligned}
E(N) = \Bigg[N + \left(\frac{1 - e^{-\frac{2Nt_\mathrm{b}}{\tau}}}{1 - e^{-\frac{2t_\mathrm{b}}{\tau}}} \right) \\
-2 \left(\frac{1 + e^{-\frac{(N+1)t_\mathrm{b}}{\tau}} \cos((N-1)\omega t_\mathrm{b}) - 2e^{-\frac{Nt_\mathrm{b}}{\tau}} \cos(N\omega t_\mathrm{b}) - 2e^{-\frac{t_\mathrm{b}}{\tau}} \cos(\omega t_\mathrm{b})}{1 - 2e^{-\frac{t_\mathrm{b}}{\tau}} \cos(\omega t_\mathrm{b}) + e^{-\frac{2t_\mathrm{b}}{\tau}}} \right) \Bigg] \\
\times \frac{P_\mathrm{out}\tau \left(1 - e^{-\frac{2t_\mathrm{b}}{\tau}}\right)}{2\left(1 - 2e^{-\frac{t_\mathrm{b}}{\tau}} \cos(\omega t_\mathrm{b}) + e^{-\frac{2t_\mathrm{b}}{\tau}}\right)} \\
+ P_\mathrm{out} \left(\frac{1 - 2e^{-\frac{Nt_\mathrm{b}}{\tau}} \cos(N\omega t_\mathrm{b}) + e^{-\frac{2Nt_\mathrm{b}}{\tau}}}{1 - 2e^{-\frac{t_\mathrm{b}}{\tau}} \cos(\omega t_\mathrm{b}) + e^{-\frac{2t_\mathrm{b}}{\tau}}} \right) \frac{\tau}{2} .
\end{aligned}
\tag{2.32}
$$

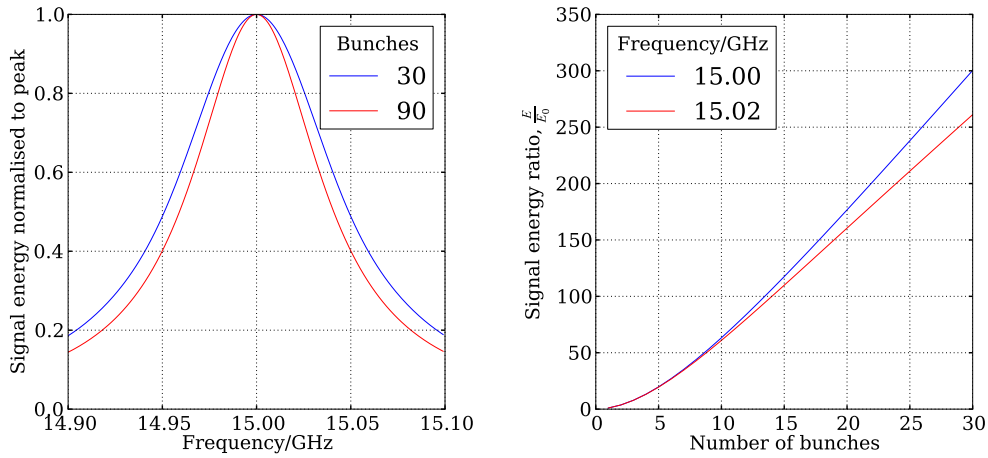

Figure 2.29: Total signal energy from 30 and 90 bunches against cavity resonant frequency as a fraction of the energy at 15 GHz (left) and against the number of bunches (right).

Figure 2.29 shows how the total energy output from a cavity mode changes with the cavity mode resonant frequency. The change depends on the number of bunches but for a frequency that is 20 MHz from the bunch arrival frequency harmonic, the reduction is about 20 % after 90 bunches. Figure 2.29 also shows the signal energy against number of bunches in the train. With a few bunches, the energy increases as the signal amplitude gets larger. When the signal nears

convergence and the amplitude reaches a steady state, the extracted energy increases linearly with the number of bunches as the signal grows longer in time. As the number of bunches increases, the energy in the signal with the 20 MHz frequency offset reaches a near constant fraction of the energy when there is no frequency offset.

Chapter 3

Down-Converter Electronics

3.1 Electronics Design

The analogue processing electronics for the prototype CTF3 cavity BPM perform three primary functions. The first is the rejection of noise from unwanted modes at frequencies outside the bandwidth of the mode of interest. The second is amplification of the signal and the third is down-conversion to an intermediate frequency for signal transport and digitisation. Aliasing frequencies must also be removed. The CLIC conceptual design report includes a proposed electronics outline [5] for the cavity BPMs of the CLIC main linac and BDS. It is based on the principle described briefly in Section 1.2.2 and was used as a template for the prototype cavity BPM electronics.

3.1.1 Outline

Figure 3.1 shows the layout of a single electronics channel used for the first beam tests of the prototype cavity BPM at CTF3. The bandwidth of the electronics is designed to be wider than the 70 MHz bandwidth of the cavity pick-up. This minimises the response time and reduces the influence of the electronics on the shape of the signal in the time domain. To maintain a large bandwidth while ensuring good rejection of the image frequency from the mixer, a high intermediate frequency (IF) of 200 MHz was chosen. The fast 2 GS s^{-1} digitiser allows for a large bandwidth and for the signal to be digitised in the first Nyquist band. It

has 10 bit resolution and a pre-amp that allows selection of the voltage range between $\pm 0.025\ V$ and $\pm 2.5\ V$ [34]. Two distinct sections are labelled in Figure 3.1. These are the 'microwave front end', which comes before the mixer and processes the signals at the resonant mode frequency of 15 GHz, and the 'intermediate frequency section', which comes after the mixer.

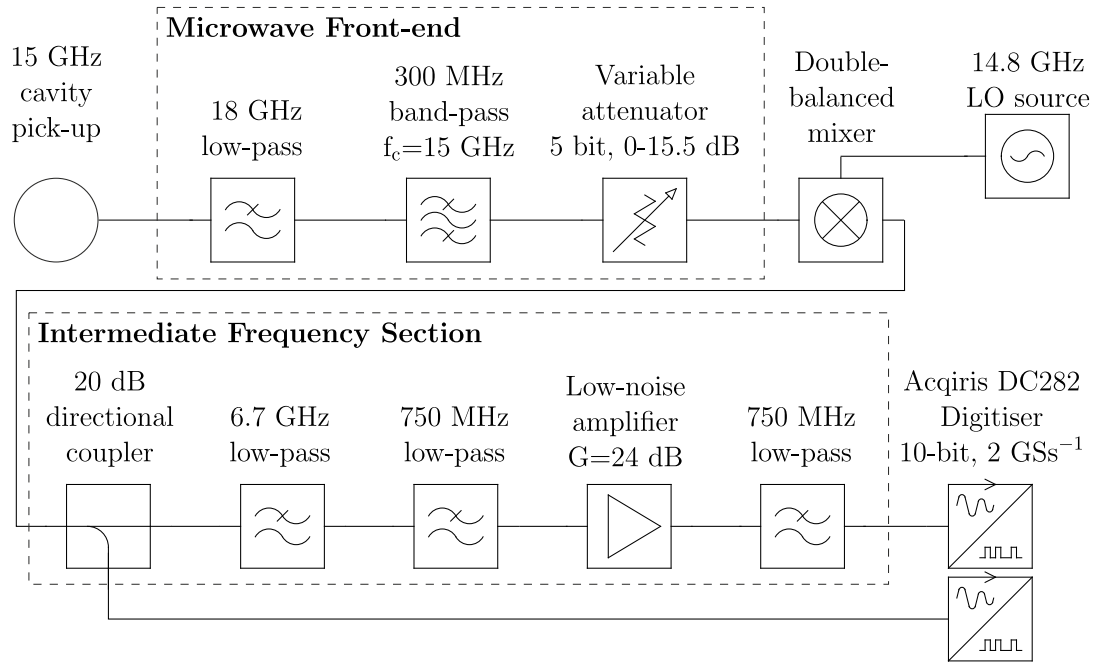

Figure 3.1: Diagram of electronics layout for the first tests of the prototype cavity BPM.

The microwave front end prepares the signal for down-mixing to a lower frequency. The 18 GHz lowpass filter at the front end removes noise from higher order modes, such as those listed in Table 2.10, and random noise at high frequencies. The 300 MHz bandwidth bandpass filter, centred at 15 GHz, is for image rejection purposes: signal content at the 14.6 MHz image frequency must be attenuated below the digitiser resolution in order to achieve the best possible phase noise [35]. The variable attenuator allows the input power range of the electronics to be adjusted so that signals from multiple and single bunches and bunches of different charge may be measured.

The IF Section is to amplify the signal and prepare it for digitisation. The

20 dB directional coupler allows the output of the mixer to be measured so that the influence of the IF section can be inspected during beam operation. The 6.7 GHz lowpass filter is to suppress cross-talk from the local oscillator and high frequency inputs of the mixer, which are at around 15 GHz.

Since the amplifier has some nonlinearity in its response, there will be interference between different frequencies. This may cause some features outside the first Nyquist band to show up inside at the amplifier output. Conversely, some features inside the first Nyquist band at the input may appear outside at the output and show up in the digitised signal through aliasing. 750 MHz lowpass filters either side of the amplifier mitigate these effects.

3.1.2 Gain

Several schemes for signal amplification were investigated based on amplifiers from Hittite Microwave [36] or Mini-Circuits [37]. For each, the output at 1 dB compression, the total gain and the noise figure of a complete down-converter channel were calculated. These are shown in Table 3.1. The values have been calculated from the specifications of all the components in a channel. The gain is simply the sum of the individual gains and insertion losses; the noise figure has been calculated using Frii's formula for noise (Equation 1.61).

The compression point of each nonlinear component has been extrapolated back to the channel input. The minimum value obtained then defines the input power at 1 dB compression for the full channel, to which the gain minus one is added to arrive at the output power at 1 dB compression. An estimated signal level of 1.0 V for a single 0.6 nC bunch with a 100 μm position offset is taken from the predicted pick-up sensitivity of 24.2 V nC^{-1} mm^{-1} and the finite rise time of the electronics, which prevents the signal from reaching the predicted peak value. This has been used to estimate the resolution from the thermal noise, noise figure and gain of the electronics and to estimate the beam offset for 1 dB compression. The root-mean-square (RMS) thermal noise was estimated using Equation 1.56 to be 16 μV or -83 dBm for an electronics bandwidth of 300 MHz.

The baseline layout, which is shown in Figure 3.1, consists of a single amplifier (HMC616LP3 from Hittite [36]) at IF with a gain of 24 dB. A cascaded scheme

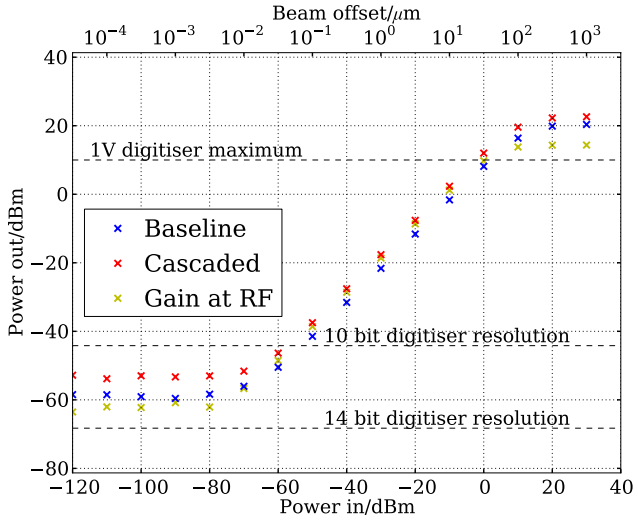

Figure 3.2: Simulation results of the BPM electronics proposed for the first test with several different amplification schemes.

Scheme	Gain/ dB	Noise figure/dB	Resolution/ nm	Output at 1 dB compression/dBm	Offset at 1 dB compression/μm
Baseline	8.4	16	9.7	19	91
Cascaded	12.4	17	12	21	72
Gain at front-end	11.4	10	5.1	16	52

Table 3.1: Simulation results for the three different schemes of amplification where the resolution and offset at 1 dB compression are for a single bunch of 0.6 nC and a pick-up sensitivity of 24.2 V nC^{-1} mm^{-1}.

has also been investigated where this is replaced by two amplifiers (2×ZX60-P103LN from Mini-Circuits [37]), each with a gain of 14 dB so that the overall gain is increased. It turns out that this offers no advantages over the baseline in terms of dynamic range or resolution. The third scheme has one of these amplifiers for 14 dB of gain at IF but also includes an amplifier with a gain of 13 dB (HMC962LC4) just before the mixer. Having a gain stage at the front end improves the resolution by moving away from the thermal noise floor but decreases the maximum measurable beam offset, which is now limited by the compression of the mixer. The python package wfProc [38] was used to simulate a full down-converter channel and determine the output power for a given input power and beam offset. The simulation was applied to continuous-wave (CW) signals and

the results are shown in Figure 3.2. The simulation behaves as expected from the calculations.

The signal range covered by digitisers with 10 and 14 bit resolutions and a range of ±1 V are also displayed. All three schemes are linear inside the range of the 10 bit digitiser but appear to be too noisy to exploit the 14 bit resolution.

Neither the simulation nor the calculations account for cable losses between the pick up and the down-converter. These will add to the noise figure and increase the beam offset for 1 dB compression. The lines indicating the digitiser ranges on Figure 3.1 will also be shifted up by the losses in the transfer cables running from the electronics output to the input of the digitiser.

The baseline scheme was chosen since it provides the best compromise between resolution and maximum beam offset, particularly for the first beam tests, where the setup is not suitable for measurements of the position resolution and the noise will be limited by the 10 bit digitiser.

3.1.3 Construction

Three electronics channels, one for each dipole mode polarisation in the position cavity and one for the monopole mode reference cavity, were constructed from connectorised components. An image of the complete down-converter box with the three channels is shown in Figure 3.3. It includes input connections for the three local oscillator signals, the attenuator control logic and a direct current power supply.

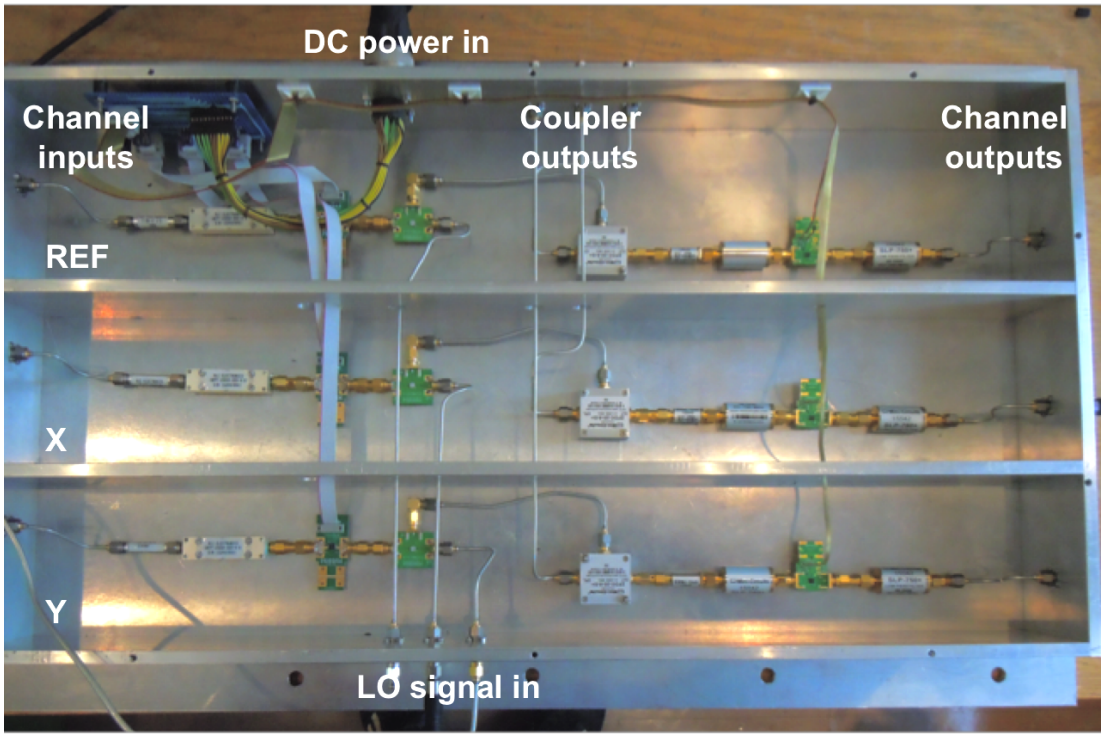

Figure 3.3: Top-down view of the three channel down-converter in its aluminium housing with the different ports labelled.

3.2 Laboratory Measurements

The down-converter was characterised in the laboratory, mostly at CERN, before being installed in the accelerator hall. A network analyser was used to measure the individual components and the separate sections either side of the mixer while a signal generator and spectrum analyser were used for the full channel measurements. An oscilloscope was also used to perform time domain measurements. For all the full channel measurements apart from those in Section 3.2.6, a second signal generator was used to provide the local oscillator signal.

3.2.1 Measurement of Individual Components

Since the electronics are made from connectorised components, it is easy to measure each one individually before construction of a full channel. The results for all of the filters are grouped together in Section 3.2.1.1 since similar measurements

were performed on each and the data was analysed in the same way. Next, the results for the components showing nonlinear behaviour, the mixer, amplifier and attenuator, are presented. The amplifier is compared with a second amplifier whose contrasting responses in both power and frequency make it a viable and distinct alternative. Finally, measurements of the coupler are briefly discussed. With the exception of the first amplifier and mixer, which were measured at Royal Holloway, all the components were measured at CERN. Knowledge of the responses of each individual component can be useful for identifying the sources of unexpected features that appear in the signal and determining how the design can be altered to mitigate them. The results could also be used along with a simulation package like wfProc [38] to make a realistic model of the electronics for development purposes.

3.2.1.1 Filters

There are four filters per channel in the down-converter box and each was measured individually using two ports of a network analyser. The two filters in the microwave front end, F-30-18.0 and CBPF-15000-300-4-R, are from RLC Electronics [39]. The first is a lowpass filter with a cut-off frequency of 18 GHz. It is not possible to observe the stop band because of the 20 GHz maximum frequency of the network analyser but the transmission and return loss in the passband were measured and are shown in Figure 3.4. The transmission is reasonably flat across the passband except for in a couple of frequency regions where there is a dip of a few dB. Fortunately, neither of these regions are close to the signal frequency of 15 GHz. There is also a fortunate minimum in return loss observed close to the signal frequency.

The 300 MHz bandpass filter centred at 15 GHz was measured across a 2 GHz bandwidth. The rejection at the signal image frequency of 14.6 GHz is -34 dB. This is not by itself, better than the 54 dB dynamic range of an ideal 10 bit digitiser. However, since the image frequency is several bandwidths away from the signal, it is already more than 60 dB lower in amplitude and so this level of rejection is more than sufficient provided that the image signal is above the thermal noise level. Having more rejection at the image frequency would therefore

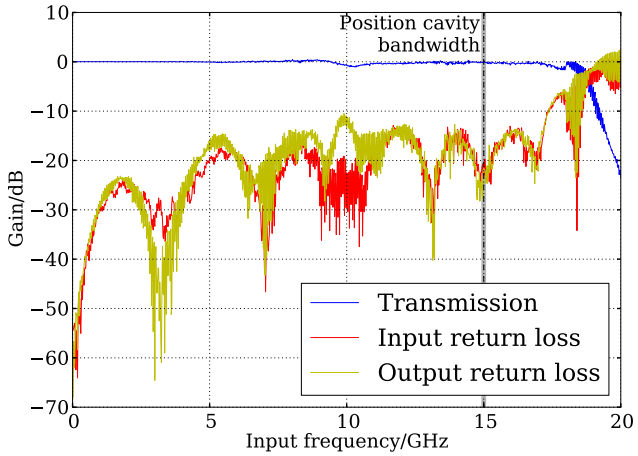

Figure 3.4: Transmission and return loss of the front end lowpass filter.

reduce the noise in low amplitude signals.

Since the complex transmission spectrum is measured, it is possible to perform an inverse Discrete Fourier Transform (DFT) to obtain the impulse response. The effect of the filter is equivalent to the convolution of the impulse response with the input signal in the time domain. This measurement is therefore useful for understanding the time domain response of the electronics. In order to determine the impulse response with good time resolution and no under-sampling, the transmission spectrum of the filter was measured from 10 MHz to 16 GHz in 10 MHz steps. A direct-current response of zero was added to the data to give 1601 equally spaced frequency bins that form a valid input to an inverse DFT. The results are shown in Figure 3.5. The impulse response shows that the filter acts as a resonator that rings with an initial pulse length of about 3 ns followed by bursts of smaller amplitude. This pattern would be instantly recognisable if observed in the BPM output signals.

The lowpass filters in the intermediate frequency section are SLP-750 and VLF-6700 from Mini-Circuits [37]. Their transfer functions and impulse responses are shown in Figure 3.6. It can be seen that neither filter has a minimum transmitted amplitude in the region around 15 GHz which would be optimal for preventing interference from RF or from the local oscillator. However, together they provide greater than 35 dB attenuation in the band from 8 GHz to 17 GHz. The two

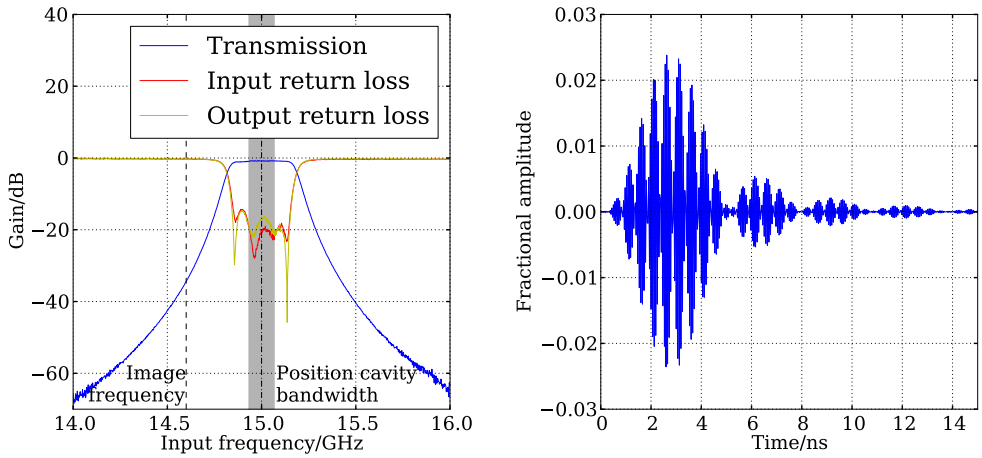

Figure 3.5: Transmission and return loss of the bandpass filter centred at 15 GHz (left) and its impulse response (right).

impulse responses are both simple but non-exponential decaying waveforms and the scaling in time with the inverse passband width is clear.

3.2.1.2 Mixer

The double-balanced mixer, HMC554LC3B from Hittite Microwave [36], was measured using the mixer-measurement option of a network analyser. The input and output of the mixer were connected to two ports of the network analyser with a third port used to provide a local oscillator signal. This allowed both the input power and the local oscillator power to be swept as shown in Figure 3.7. The loss of the mixer with 13 dBm LO power was measured to be 9.3 dB, with a flatness of ± 0.7 dB within a 200 MHz bandwidth of 15 GHz. The input power for 1 dB compression was measured to be 11.2 dBm and is almost independent of local oscillator power. Decreasing the local oscillator power below the nominal level of 13 dBm increases the loss of the mixer and the linearity below the compression point suffers. Increasing the LO power above 13 dBm has little effect as the mixer saturates.

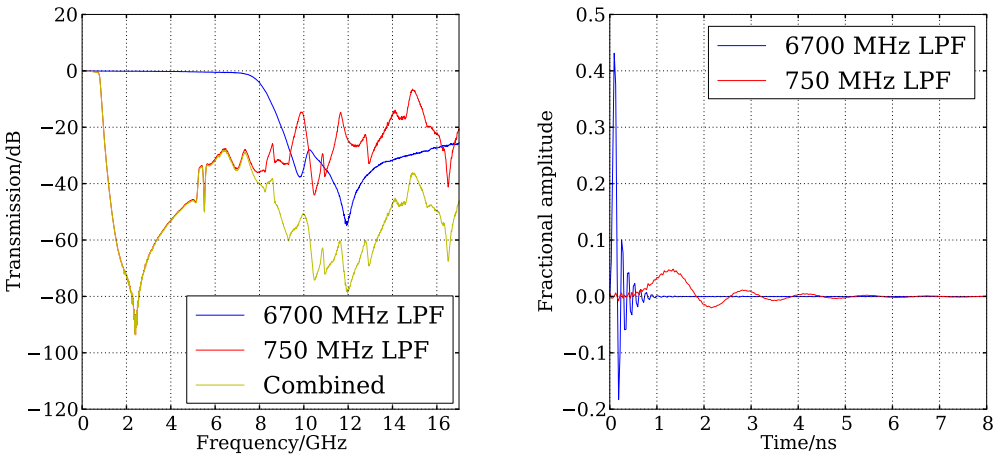

Figure 3.6: Transmission of the lowpass filters in the intermediate frequency section (left) and their impulse responses (right).

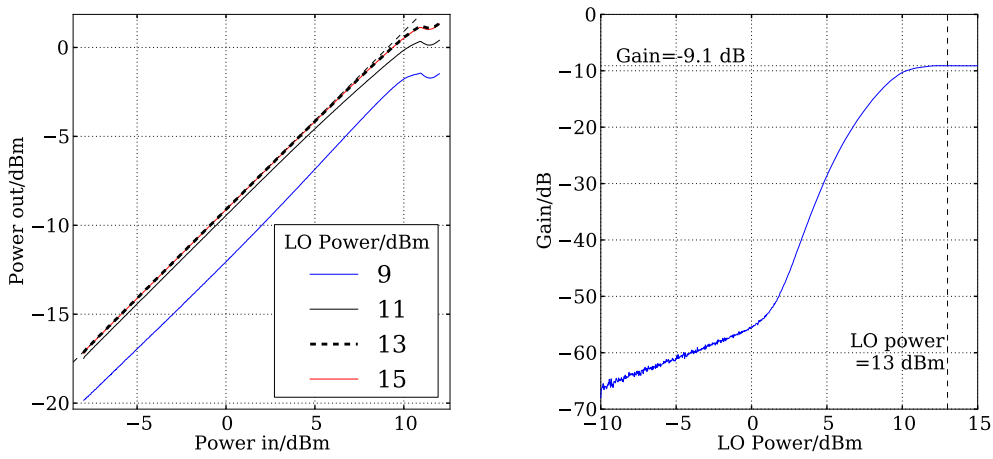

Figure 3.7: Power sweeps of the mixer input (left) and local oscillator (right).

3.2.1.3 Amplifiers

Two amplifiers for the IF section were tested. The first choice was HMC616LP3 from Hittite Microwave [36], which has an advertised maximum noise figure of 0.8 dB but a gain that is strongly dependent on frequency. A second amplifier from Mini-Circuits [37], ZX60-33LN, that shows less frequency dependence but has a higher advertised noise figure of 1.9 dB was also tested. This can be used

in addition in case more gain is needed or as an alternative if the frequency dependence of the first is too great.

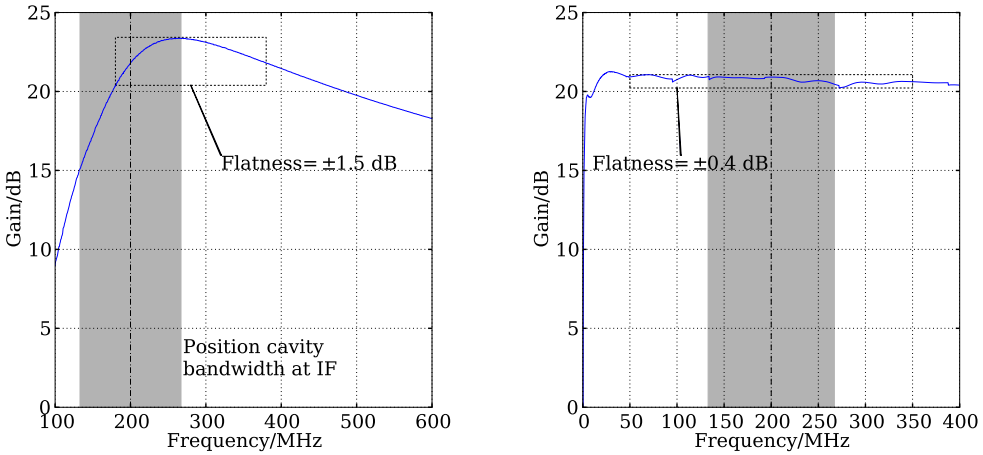

Figure 3.8: Transfer function of the first choice IF amplifier (left) and of the second amplifier (right).

The results for the responses to frequency and power are shown in Figure 3.8 and Figure 3.9 respectively. The first amplifier has a gain of 22 dB at the intermediate frequency of 200 MHz with a flatness of ±1.5 dB between 180 MHz and 380 MHz. The gain of the second amplifier is 21 dB at 200 MHz and it is clearly much more broadband with a flatness of ±0.4 dB between 50 MHz and 350 MHz. The second amplifier also has a higher output power at 1 dB compression, 16.9 dBm compared to 13.9 dBm for the first and its phase advance changes less as the input power is increased from -15 dBm to -5 dBm, 0.6° compared with 4° for the first. The first amplifier was originally preferred because of the advertised maximum noise figure and concerns that there would not be enough gain. Unless otherwise stated, the partial and full channel measurements of Sections 3.2.2 and 3.2.3 are with the first amplifier.

3.2.1.4 Attenuator

The large dynamic range of the beam induced signals from different numbers of bunches and different bunch charges means that it is necessary to include some variable attenuation that can be controlled from outside the accelerator hall.

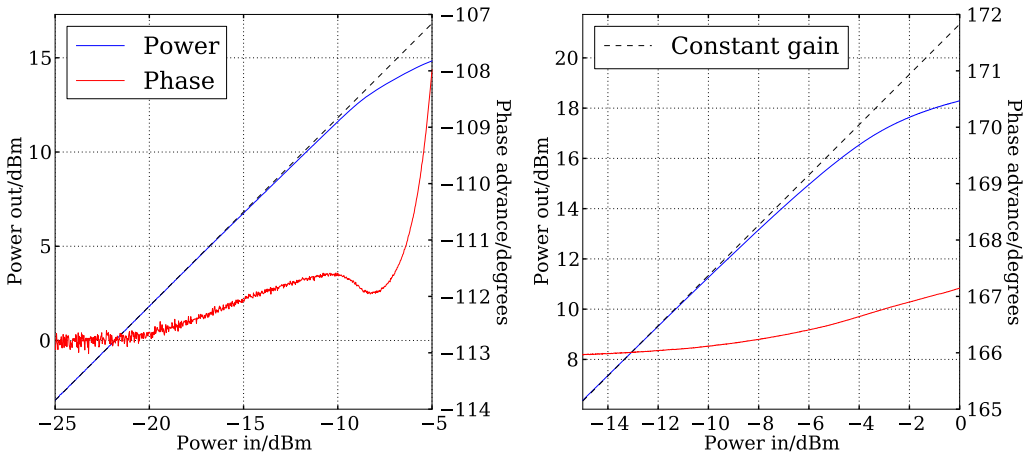

Figure 3.9: Output power and phase advance as a function of input power for the first amplifier (left) and second amplifier (right).

The chosen solution is the digitally controlled attenuator, HMC941LP4E from Hittite Microwave [36]. An ideal variable attenuator should have a flat power and frequency response and a perfect match between the nominal and observed attenuation.

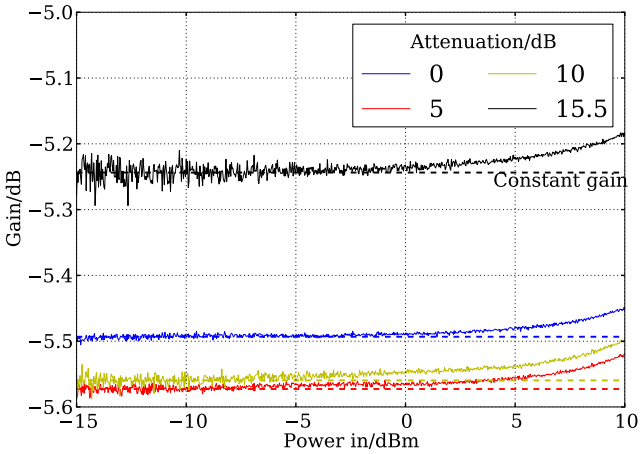

Figure 3.10: Output power as a function of input power for the digitally controlled attenuator where each curve has been offset by the nominal attenuation.

A power sweep was performed on the digitally controlled attenuator at four

different attenuation settings. The results are shown in Figure 3.10 where each curve has been offset by the nominal attenuation setting. The nonlinearity of the attenuator can be seen at high power but it does not reach a 1 dB point. The nonlinearity also appears to be independent of the attenuation setting.

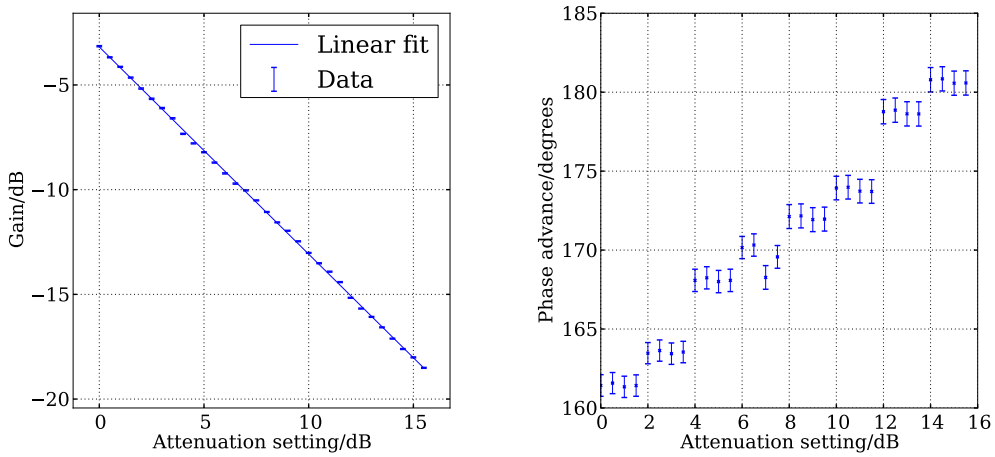

Figure 3.11: Gain (left) and phase advance (right) of the digitally controlled attenuator at 15 GHz for every attenuation setting. The error bars represent the standard deviation over a 200 MHz bandwidth.

The attenuation and phase advance was also measured for each setting and averaged over a 200 MHz bandwidth around a central frequency of 15 GHz. The results are shown in Figure 3.11 where the standard deviation over the bandwidth is indicated by the error bars. The gradient of the measured attenuation against the nominal value is 0.989 which is close to but slightly less than the ideal value of 1.0. The phase advance varies by almost 20° across the full range of attenuation settings but in clear steps of 4 attenuation settings each. Certain fine adjustments to the signal level could therefore be made with small changes in phase.

3.2.1.5 Coupler

Three ports of the network analyser were used to measure the characteristics of the 20 dB coupler in its 3 GHz bandwidth. The results are shown in Figure 3.12. The coupling remains between -17 dB and -20 dB across the whole bandwidth and the directionality is better than 20 dB under 1.5 GHz. It is therefore effective

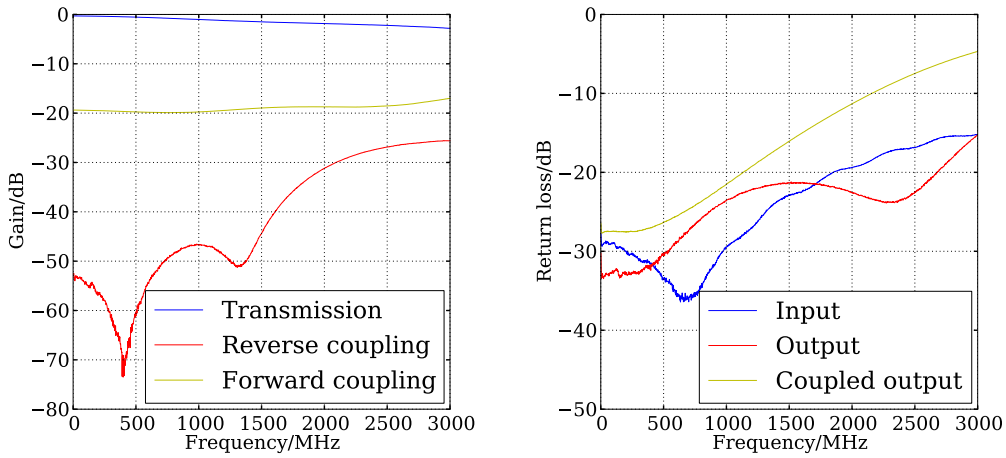

Figure 3.12: Transmission and coupling measurements of the 20 dB coupler (left) and return losses (right).

as a means of measuring the mixer output as it shouldn't have too much of an influence itself.

3.2.2 Partial Channel Measurements

Measurements of a full channel including the change in frequency could not be made with the available network analyser. One channel was therefore split into its two parts, the microwave front end before the mixer and the Intermediate Frequency Section after the mixer. A two port measurement was made on each section using a network analyser. This allowed more data points to be taken than if a separate spectrum analyser and source had been used.

3.2.2.1 Microwave Front End

The frequency response of the microwave front end from the electronics input to the input of the mixer, was measured for a 1 GHz bandwidth about the signal frequency of 15 GHz and is shown in Figure 3.13. It is defined in amplitude by the bandpass filter. The return losses in the passband are smaller than for any one component which suggests that they are dominated by the connections. The input return loss is independent of the attenuation setting while the output

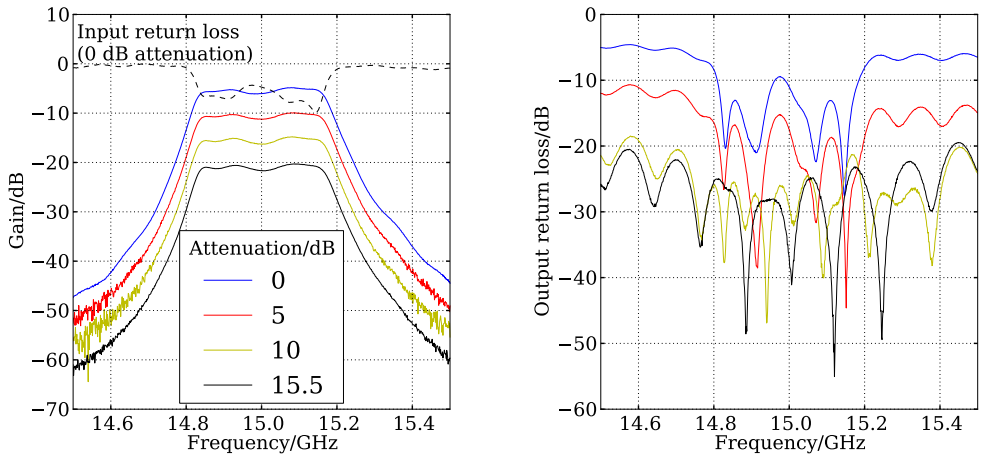

Figure 3.13: The transfer function of the microwave front end (left) and the output return loss (right), both for different attenuation settings.

return loss approaches the output return loss of the attenuator for attenuation settings higher than 10 dB. This suggests that the bulk of the reflections are occurring upstream of the attenuator. The power response of the microwave front end, shown in Figure 3.14, is also defined by the variable attenuator in both amplitude and phase. The phase advance remains stable to within 1° up to almost 5 dBm of input power.

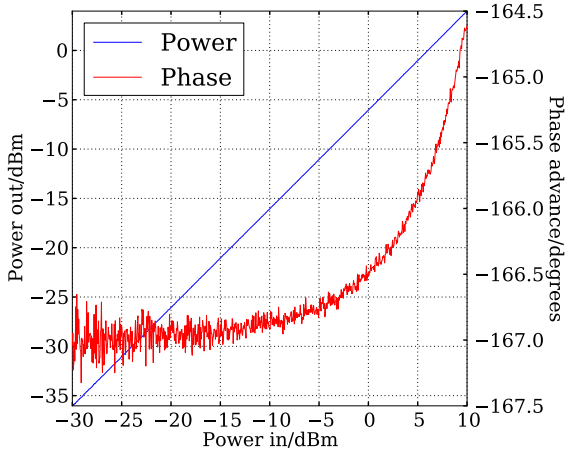

Figure 3.14: Output power as a function of input power for the microwave front end, measured at 15 GHz.

3.2.2.2 Intermediate Frequency Section

The frequency response of the intermediate frequency (IF) section, from the input of the 20 dB coupler to the electronics output is shown in Figure 3.15. It was also measured for 10 MHz to 16 GHz input to determine the impulse response. The lower 3 dB point is defined by the amplifier which was chosen for low noise and high gain rather than a large bandwidth. The upper 3 dB point is defined by the 750 MHz lowpass filter. Below this, the frequency dependence of the amplifier is clearly visible. The frequency response above 1.6 GHz cannot be resolved above the measurement noise and so is not shown. This means there is excellent rejection of high frequency interference around 15 GHz. The minimum attenuation outside the first Nyquist band of the digitiser is 31 dB and occurs at the Nyquist frequency of 1 GHz. Relative to the signal gain of 21 dB at 200 MHz, this is already close to the 54 dB dynamic range of an ideal 10 bit digitiser. The rejection is well above this level at the first aliasing frequency of the IF, which is 1.8 GHz. Aliasing should therefore not be the cause of additional phase noise. The impulse response of the intermediate frequency section is also shown in Figure 3.15. It shows extra features that were not seen in the impulse response of the filters when measured individually and rings out for 10s of nanoseconds.

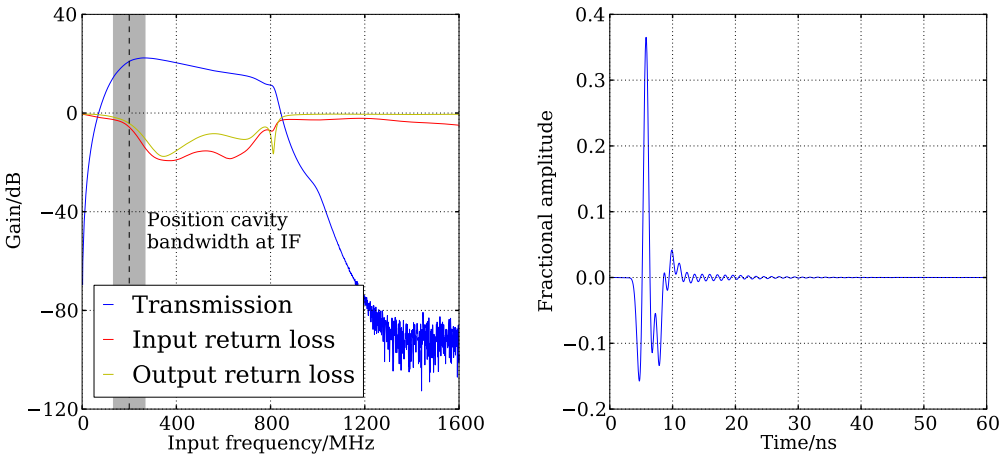

Figure 3.15: Two port measurement of the intermediate frequency section (left) and the corresponding impulse response (right).

This is the influence of the amplifier which also completely defines the power response shown in Figure 3.16.

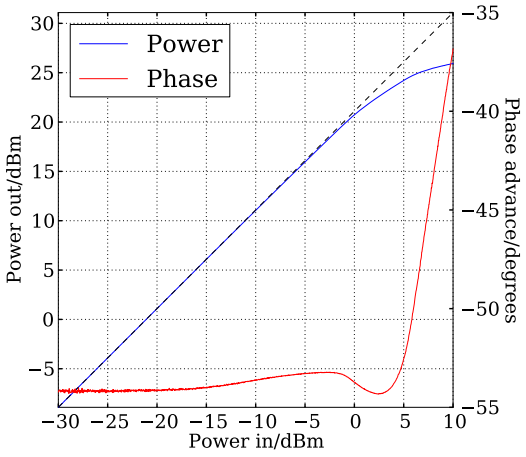

Figure 3.16: Output power as a function of input power for the intermediate frequency section.

3.2.3 Full Channel Measurements

Full channel measurements of the down-converter box had to be made in order to determine the important overall features of the electronics so that they could be specified in future. In order to measure a full channel, two signal generators were used, one for the input signal and one for the local oscillator. The output signal was measured using a spectrum analyser. In all cases, the input frequency used to measure the power response curves was always 15 GHz. As discussed in Section 3.2.2, it was not possible to use a network analyser for these measurements.

3.2.3.1 Frequency and Power Response

The frequency and power responses of all three channels (labelled X, Y and REF for the horizontal position, vertical position and reference signals respectively) were measured in order to ensure that they were all performing similarly and to get an idea of the variability between them. The frequency response of each

Channel	3 dB points/GHz		Bandwidth/MHz	Flatness/±dB
REF	14.968	15.173	205	1.2
X	14.964	15.181	216	0.6
Y	14.972	15.176	204	0.8
REF Coupler	14.820	15.172	352	1.0

Table 3.2: Summary of the frequency responses of the three channels and the output of the 20 dB coupler in the reference channel where the flatness is relative to the power at 15 GHz input in the 150 MHz frequency range from 14.99-15.14 GHz input.

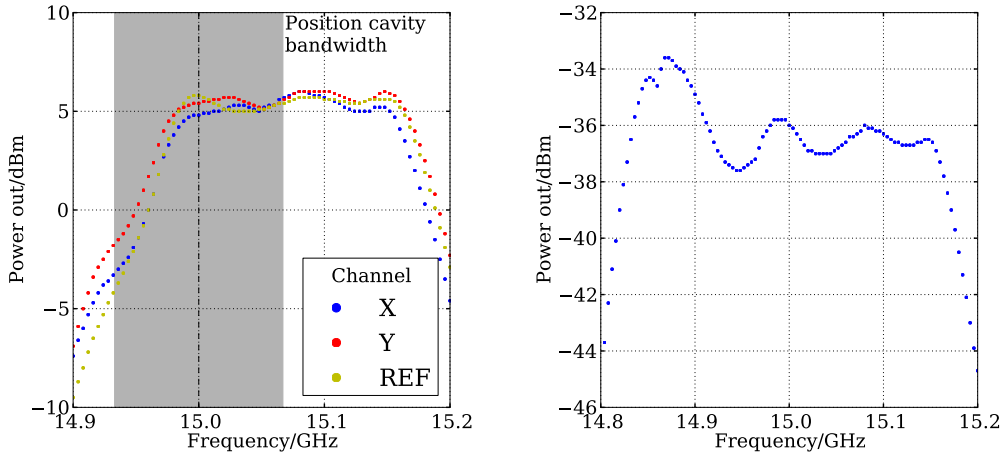

Figure 3.17: Frequency response of the three channels (left) and the coupler of the reference channel (right).

channel was measured with a 0 dBm input signal. The input frequency was swept by hand from 14.9 GHz to 15.1 GHz and the peak in the output power spectrum was measured using the spectrum analyser. The results are shown in Figure 3.17 and the important values are listed in Table 3.2. The high frequency cut-off is defined by the front end bandpass filter while the low frequency cut-off is defined by the IF amplifier which provides less gain at lower frequencies. The response curves of the three channels do not follow each other exactly but their lower and upper 3 dB points are similar to within a few MHz. The frequency response was also measured at the coupler output of the reference channel. Since this comes before the amplifier, the bandwidth is larger. The flatness is comparable to that of the reference channel in the passband.

Power sweeps were performed with a 15 GHz input frequency for the three channels and are shown in Figure 3.18. More points were taken close to saturation in order to determine the output power at 1 dB compression. This came to +13 dBm for all three channels. The coupling between the channels was also determined by measuring the output of one channel while varying the input power of another channel between -10 dBm and +10 dBm. A local oscillator signal was provided to both channels during the measurement and all unused ports were terminated with matched loads. The results are summarised in Table 3.3 along with the gain of each channel, taken as the mean difference between the output and input powers in the linear region.

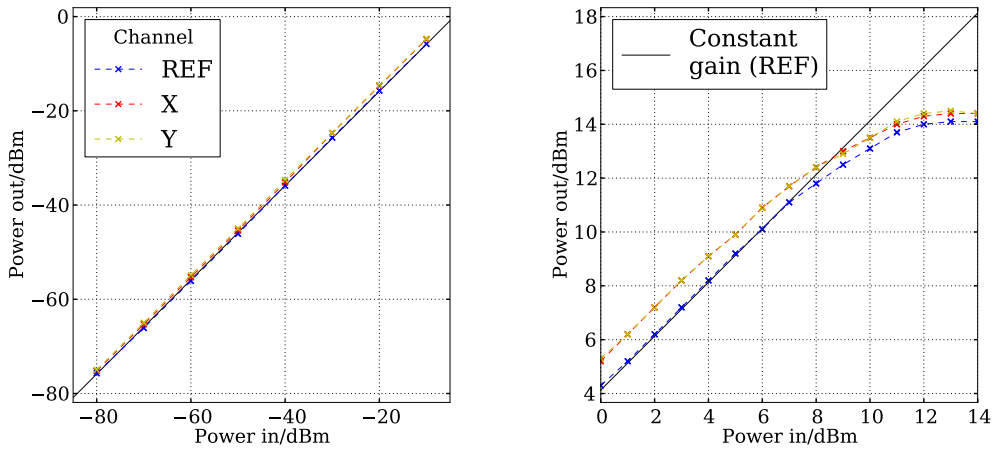

Figure 3.18: Output power against input power for the three channels over a broad range (left) and at high power (right) at 15 GHz input frequency.

In \ Out	REF/dB	X/dB	Y/dB
REF	4.13±0.06	-90.4±0.5	-87.0±0.5
X	-96.8±0.5	5.0±0.1	-82.2±0.5
Y	-81.0±0.5	-87.4±0.5	5.19±0.05

Table 3.3: Summary of gain and coupling of the down-converter box where each cell shows the output power relative to the input power of another channel.

The measured gain varies by just over 1 dB between channels and is around 5 dB. This compares well to the 8.4 dB prediction in Table 3.1 when limited com-

ponent reproducibility and the additional losses of a real system in connections and cables are assumed. The isolation has some dependence on the physical separation between the two channels and the routing of the LO cables but is always better than -80 dB which is negligible compared to the cross-coupling between the pick-up output couplers and the digitiser dynamic range.

3.2.3.2 Local Oscillator Power

The effect of varying the local oscillator power to the reference channel was investigated, first over a broad input power range and then, only at higher input powers. The effect on the linearity is more clear in Figure 3.19 than when the mixer was measured individually in Section 3.2.1.2. The larger scale plot of the nonlinear region shows again that for all LO power levels, the output reaches the same compression level defined by the amplifier.

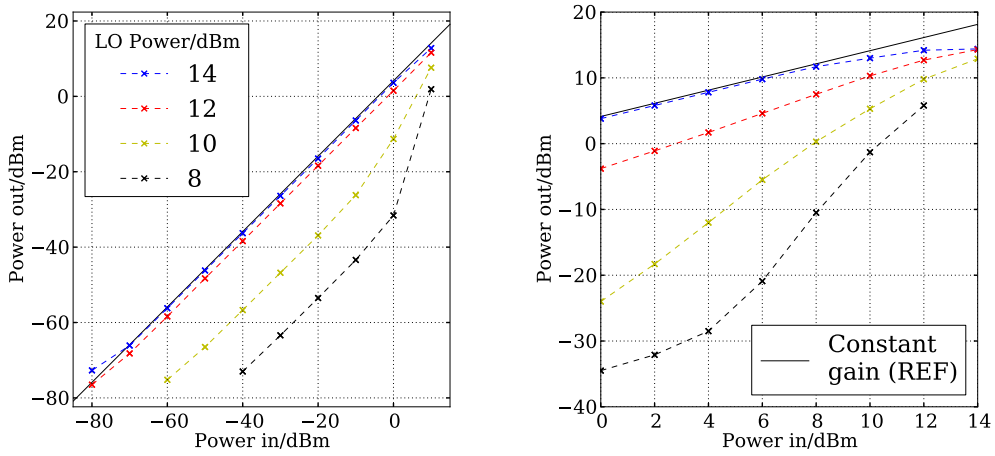

Figure 3.19: Output power as a function of input power for different settings of the LO power over a broad range (left) and at high power (right).

The LO power was also varied at finer intervals with a constant input power of 0 dBm as shown in Figure 3.20. This again shows the mixer saturation at the design LO power of 13 dBm. No further effects are expected since the LO power only has an influence on the performance of the mixer and none of the other components.

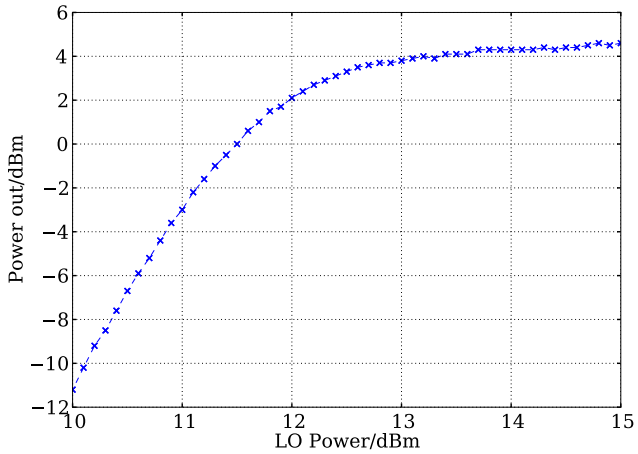

Figure 3.20: Output power as a function of local oscillator power for 0 dBm input power.

3.2.3.3 Attenuation Setting

A power sweep was performed with five different settings on the variable attenuator, again using the reference channel. The seven points from -60 dBm to 0 dBm were used to determine the overall gain at each attenuation setting. The results are shown in Figure 3.21 and summarised in Table 3.4. The compression at high power is only visible with less than 5 dB of attenuation up to 14 dBm of input power. The change in gain is slightly less than the change in attenuation setting with a gradient of -0.911 ± 0.007 obtained from a linear fit of the data in Table 3.4. This is once again, smaller than the ideal value of -1.0. Table 3.4 can be a useful lookup table for the electronics gain at several key attenuation settings if one is required.

Attenuation/dB	Gain/dB
0	4.11±0.06
2.5	2.37±0.06
5	0.11±0.03
10	-4.83±0.07
15.5	-9.93±0.12

Table 3.4: Summary of overall gain for several attenuation settings.

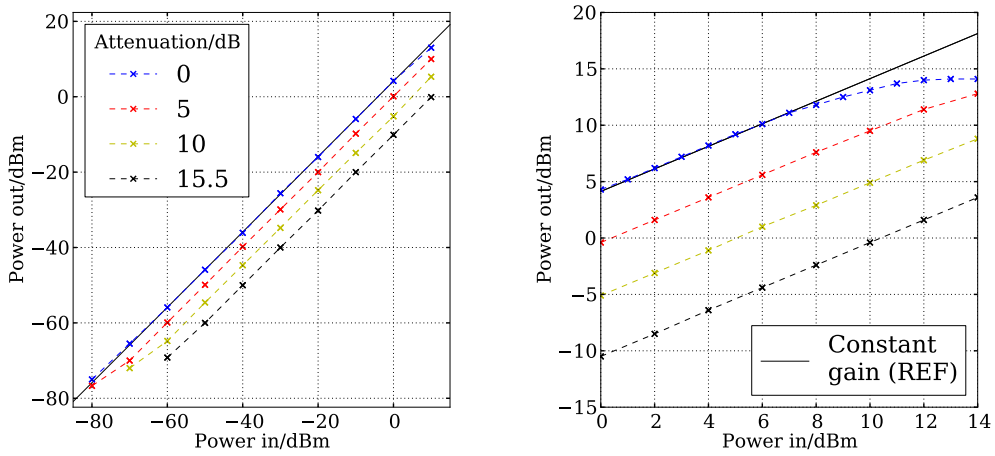

Figure 3.21: Power sweep of the reference channel with different attenuation settings (left) and at higher power (right).

3.2.3.4 Linearity

As discussed in Section 1.2.2.2, the powers for 1 dB compression are useful for specifying the input power limits but the linearity over the whole input range is better quantified by the powers at the third order intercept points. These were determined by performing two-tone measurements, also outlined in Section 1.2.2.2.

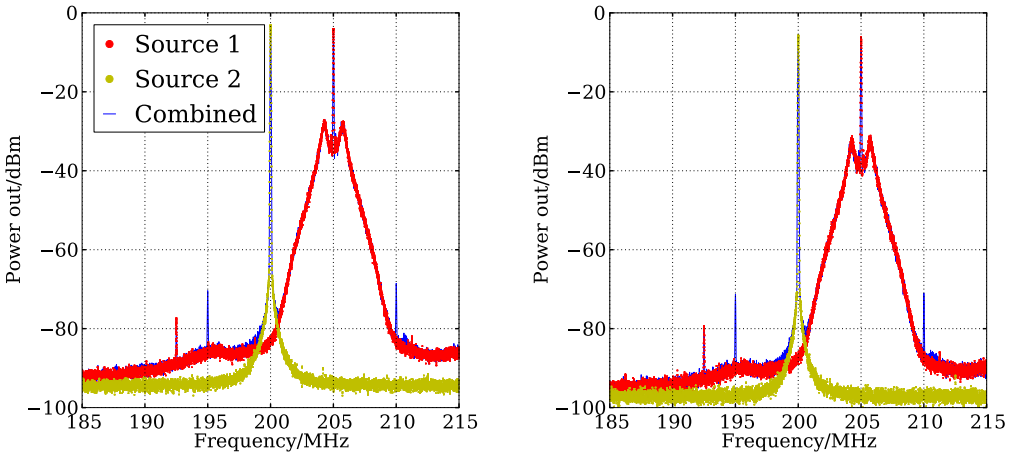

Figure 3.22: Results of measurements of third order intercept points with the first amplifier (left) and the second amplifier (right).

Two signal generators were used to provide -10 dBm input signals at frequencies offset by 5 MHz. The signals from the two sources were added in a 3 dB resistive combiner before being input to the reference channel of the down-converter. They were then measured individually, without the presence of the other, before the two tone measurement was made. The results are shown in Figure 3.22 for both the first choice amplifier and the second amplifier. The two large peaks in the output spectra separated by 5 MHz are the two input signals while the small spurs, 5 MHz above the higher frequency input signal and 5 MHz below the lower one, are caused by the nonlinear response of the electronics. The output power at these frequencies was measured and Equation 1.69 was used to determine the output power at the third order intercept point. This came to 30 dBm with the first amplifier and 27 dBm with the second. The difference in the two is the same as the difference in gain. This suggests that the dominant source of nonlinear effects in the electronics is not the amplifier and so is almost certainly the mixer.

The two-tone measurement, while reasonably simple to set up, requires careful consideration to perform successfully. As mentioned in Section 1.2.2.2, it is important to keep the input powers low for the assumptions used in the calculation of the $OIP3$ to be justified. A compromise must be found, however, so that the signal spurs are resolvable above the noise of the spectrum analyser. The two input signals must also be kept close in frequency so that the result is not effected by the bandwidth of the electronics. As is clear in Figure 3.22, this is not always easy when the signal generator outputs have their own widths in the frequency domain. Figure 3.22 also shows a spur at 192.5 MHz that appears in the output of the second source and the combined spectrum. Care must be taken so that the spurs produced by the nonlinear components in the electronics do not fall on one of these signal generator produced spurs which would give a false result.

3.2.3.5 Performance with Second Amplifier

The reference channel was measured with the second amplifier in the same way as with the first amplifier. The same power sweep was performed (Figure 3.23) and the gain measured to be 2.87±0.12 dB. The output power at 1 dB compression

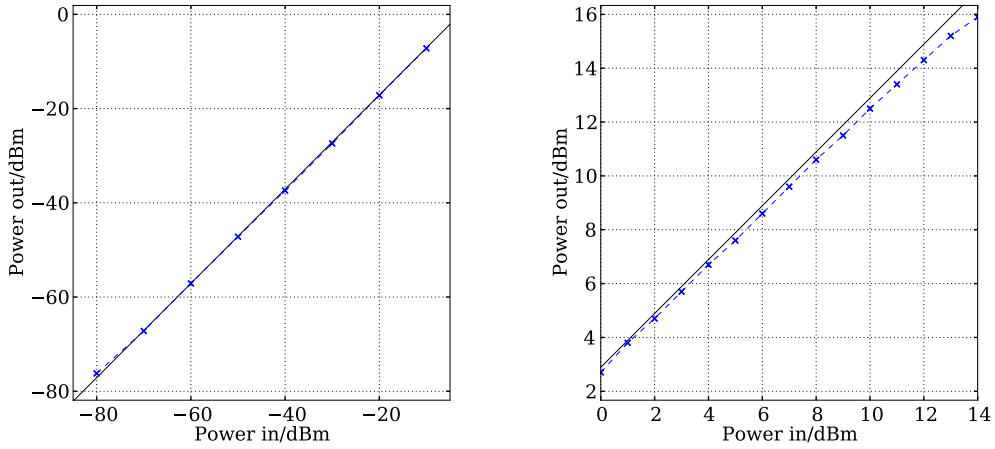

Figure 3.23: Output power against input power for the reference channel with the second amplifier over a broad range (left) and at high power (right).

was not reached and is therefore, higher than 15.9 dBm.

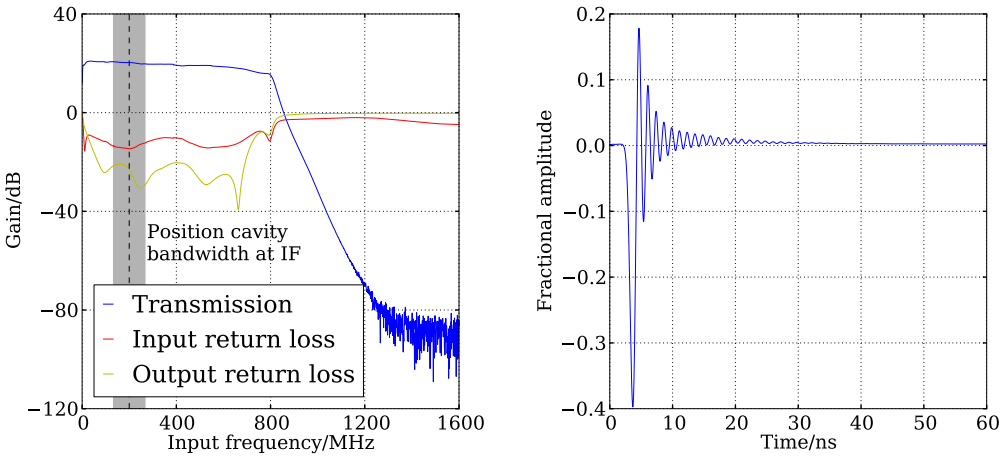

Figure 3.24: Two port measurement of the intermediate frequency section with the second amplifier (left) and the corresponding impulse response (right).

The transfer function and impulse response of the IF section with the second amplifier were also measured. These are shown in Figure 3.24. The output return loss is improved by about 10 dB while the bandwidth is hugely increased by about 370 MHz thanks to the higher gain of the second amplifier at lower frequencies.

The impulse response is also more similar to the impulse responses of the filters in Figure 3.6 since the amplifier has a smaller influence.

3.2.4 Time Domain

Since the BPM signals are principally processed digitally in the time domain, the time domain response of the electronics should be known. Time domain measurements were made with a signal generator producing 200 ns long CW pulses at 15 GHz, each separated by a further 200 ns. The channel output was measured with an oscilloscope. The positive and negative waveform peaks were measured in order to reconstruct the signal envelope. This was then averaged over 5 pulses. This envelope reconstruction results in a periodic ripple across the output because of the procession of the source in phase relative to the oscilloscope clock but has minimal influence on the overall signal shape because it involves no filtering. The results are shown in Figure 3.25 for a full channel with the first and second amplifiers at different intermediate frequencies. Measurements were also made at the output of the mixer and with different attenuation settings.

It is evident that the choice of amplifier has an effect on the rise of the signal. For the first amplifier, when the IF is close to the lower 3 dB point, there is some smoothing of the pulse rise. As the intermediate frequency is increased to the centre of the IF section passband, a small amount of signal droop is evident at the beginning of the pulse. With the second amplifier, the droop is present at all intermediate frequencies because the low frequency cut-off is now being defined in the microwave front end and not by the amplifier. With both amplifiers, there is some dependence of the signal amplitude on the choice of intermediate frequency but not as would be expected from the measurement of the IF sections in Sections 3.2.2.2 and 3.2.3.5. The measurements at the output of the mixer show a similar change in amplitude. It can therefore be concluded that this change is caused by the mixer, since the signal frequency at the microwave front end is fixed. Changes in the signal shape that arise from reflections in the microwave front end are also likely to be affected by the attenuation setting because, as mentioned in Section 3.2.2.1, most reflections occur upstream of the attenuator. No change in the signal shape is seen when the attenuation setting is varied so

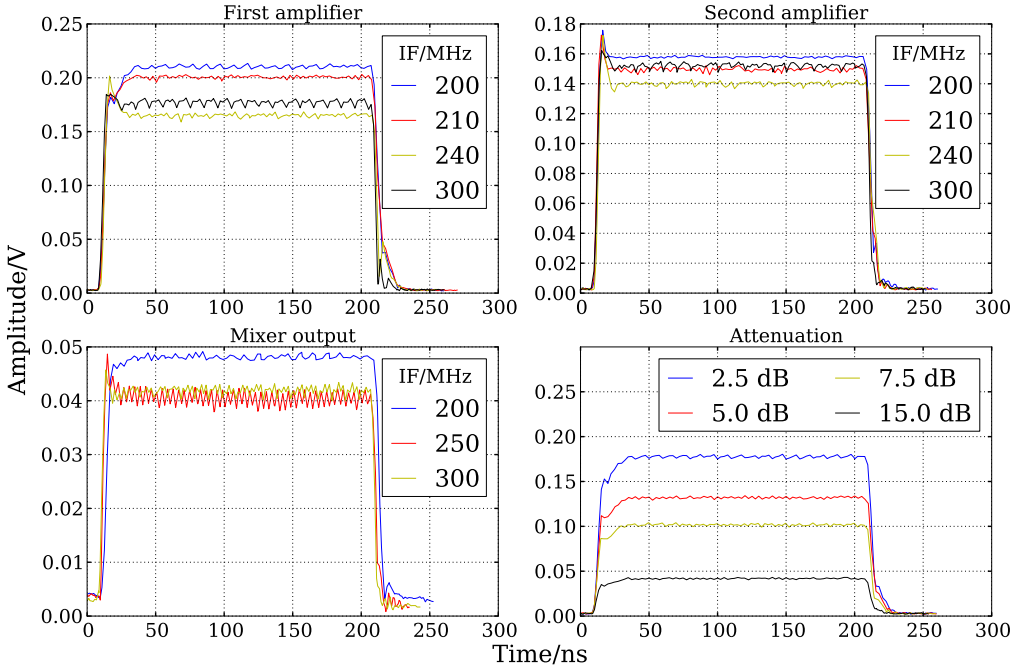

Figure 3.25: Time domain measurements of a full channel with the first amplifier (top left) and second amplifier (top right) and at the output of the mixer (bottom left) at different intermediate frequencies. The bottom right plot shows the signal shape for different attenuation settings at an IF of 200 MHz.

reflections in the microwave front end are not a major cause of changes to the signal shape.

3.2.5 Digitiser

The digitiser is a DC282 from Acqiris [34] with a compact PCI interface. It has four channels, each with a sampling rate of 2 GS s^{-1} and 10 bit resolution. A pre-amp allows the measurement range to be varied between ± 0.025 V and ± 2.5 V. The digitiser noise was determined by terminating the channel inputs with 50 Ω loads and recording 70 waveforms, each 400 samples long. The root-mean-square voltage of each waveform was then used as a measure of the noise. During the measurement, the setting of the pre-amp was varied to change the range of the digitiser. Figure 3.26 shows the measured noise in terms of digitiser

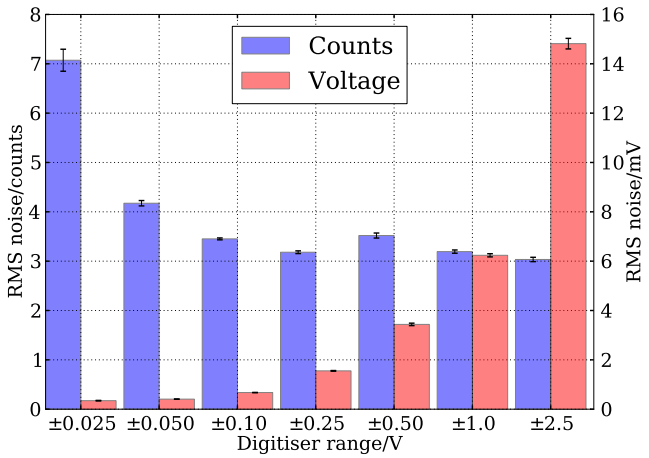

Figure 3.26: Measured RMS noise for the different digitiser voltage range settings.

counts and voltage for the different voltage range settings. For the two widest ranges, the noise is constant at 3.13 ± 0.03 counts and it is close to this number for the next three smaller ranges. The effective resolution of the digitiser is therefore $10 - \log_2(3.13) = 8.4$ bits. At smaller digitiser ranges, the thermal noise overtakes the digitiser noise and it becomes almost constant in voltage. At the smallest digitiser range, the noise is largest at 7.1 ± 0.2 counts which corresponds to 0.345 ± 0.011 V or -56.2 $\pm$ 0.3 dBm. For a signal spanning the full range of the digitiser, this gives a signal to noise ratio of 34 dB. This is to specification but the noise is much larger than the expected thermal noise calculated using Equation 1.56 which is -78 dBm. This could be due to aliasing or the noise figure of the digitiser's analogue front end, which is not quoted in the documentation.

When the digitiser is used to sample alternating current (AC) signals, because the signal is no longer constant in time, there is additional noise that arises from jitter of the digitiser clock. This would add to the noise power at DC N_{DC} so that the total noise power N_{AC} is given by

$$N_{\mathrm{AC}} = N_{\mathrm{DC}} + \sigma_t^2 \frac{(2\pi f_{\mathrm{sig}} A_{\mathrm{sig}})^2}{2Z} \tag{3.1}$$

where A_{sig} is the voltage amplitude of the signal, f_{sig} is the signal frequency, σ_t is the RMS jitter of the digitiser clock and Z is the load impedance [40]. For a signal

amplitude of 1 V and the specified digitiser clock jitter of 1.2 ps, the additional voltage noise for a 200 MHz signal would be 0.0015 V. This corresponds to 0.77 counts at the ±1.0 V digitiser range setting and only reduces the effective digitiser resolution by less than 0.1 bits.

3.2.6 Local Oscillator Source

Because the local oscillator has a high frequency of 14.8 GHz, a lot of power would be lost if it were transmitted over a long distance. The LO source would therefore have to be installed inside the accelerator tunnel and so could not be a commercial signal generator, which may not be radiation resistant. A purpose-built source was constructed at Royal Holloway in a separate housing from the down-converter electronics. It is based on a voltage controlled oscillator (VCO), HMC398QS16G, and a power amplifier, HMC965LP5E, both from Hittite Microwave [36]. The power from its eight output ports is sufficient to supply one down-converter channel each, allowing for additional cable losses.

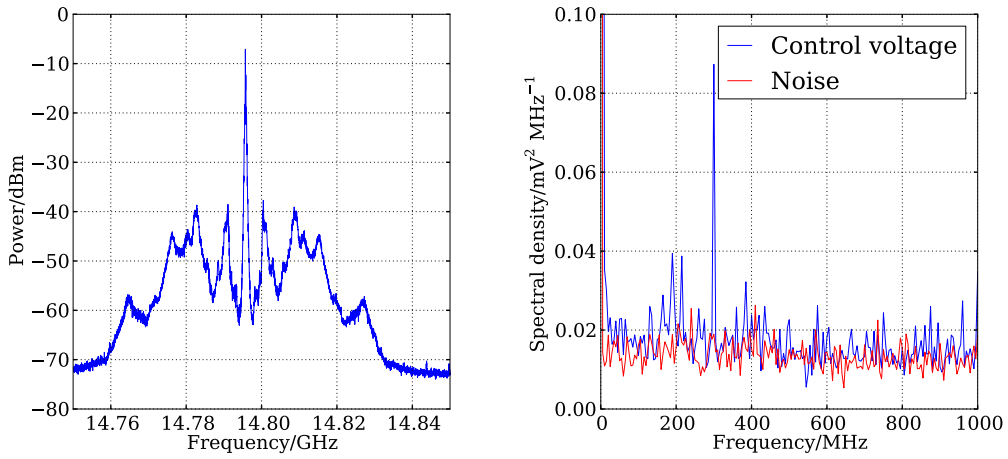

Figure 3.27: Frequency spectrum of the LO source output as measured using a spectrum analyser (left) and the frequency spectrum of the LO control voltage (right) as measured using the digitiser compared to a spectrum of the digitiser noise.

The LO source was measured both directly, using a spectrum analyser, and in conjunction with the down-converter using an oscilloscope. The plot on the

left of Figure 3.27 shows the output spectrum of the LO source. The phase noise appears to be much larger than the value of 105 dBc Hz^{-1} at 100 kHz offset specified for the VCO. This may be due to variations in the control voltage. This was measured in the CTF klystron gallery using the digitiser as a scope. Figure 3.27 also shows the spectrum of the digitiser output with and without the control voltage at the input. It can be seen that there is some additional noise coming from the control voltage power supply with a peak at a frequency of about 300 MHz. This is faster than the spectrum analyser sweep rate so it would smear out the measured spectrum. The measured RMS variation in the control voltage is 7 mV which corresponds to a 1 MHz change in LO frequency based on the beam-based measurements in Section 5.3.2. This is about the same as the 3 dB width of the peak in the LO spectrum.

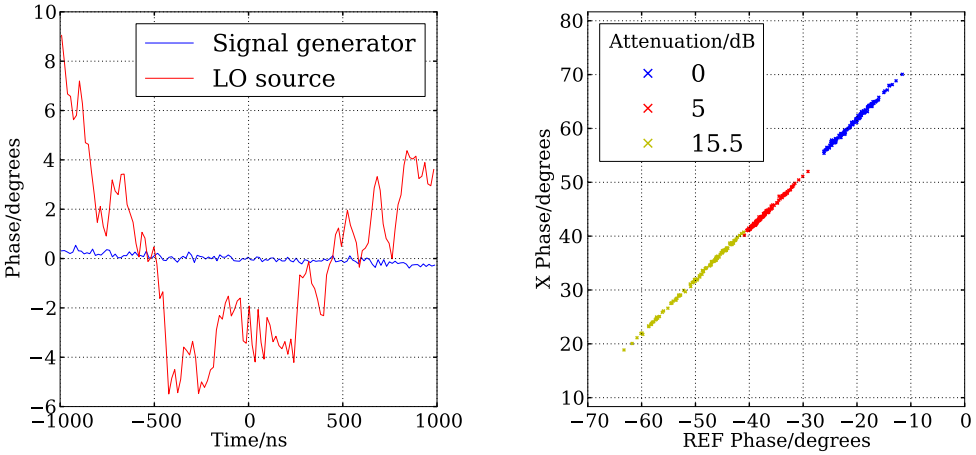

Figure 3.28: Phase variation over time of a down-converter channel (left) and correlation of the output phases of two channels (right).

The effect of the LO source on the phase stability of the down-converter output was investigated. The source was connected to the LO input of the down-converter while a signal generator was used for the 15 GHz signal input. The 200 MHz output was then measured using an oscilloscope. The scope trace was divided into sections, each about 15 ns in length (about 3 oscillation periods at

200 MHz), and a fit was made to each section. The fit function is of the form

$$V(t) = A\sin(2\pi f_{\mathrm{IF}} t + \phi) \tag{3.2}$$

where f_{IF} is the signal frequency measured using numerical analysis of the fundamental frequencies (NAFF, see Section 5.3.1) and A and ϕ are fit parameters. The time-base of the oscilloscope is used for the time t so that the phase advance between each section is accounted for. The results for the measured phase ϕ are shown in the plot on the left of Figure 3.28. It can be seen that, with the LO source connected, the phase changes by more than 10° over 2 μs. The phase stability is significantly better when a signal generator is used for the LO signal. The input signal was also split and sent to two of the down-converter channels so that two phases could be measured. Figure 3.28 shows that the phases of the two channels are well correlated, including the step changes in phase expected for different attenuation settings. During the beam tests, the variation in phase coming from the local oscillator is expected to be visible in the long signals excited by many bunches.

3.3 Measurements In-Situ

The electronics were installed in the accelerator hall close to the installation of the cavity BPM pick-up. More details of the installation are given in Section 5.2. The electronics were then measured in-situ so that the response of the system as a whole from pick-up output to digitiser readout was known. The total gain could then be used to determine the output power of the pick-up and the time domain result could be compared directly with the beam excited signals. A signal generator and cable of known loss were used to provide a signal at the input of the rigid cables that run between the pick-up and the down-converter. The local oscillator source was used and the amplitude of the output signal was measured using the digitiser.

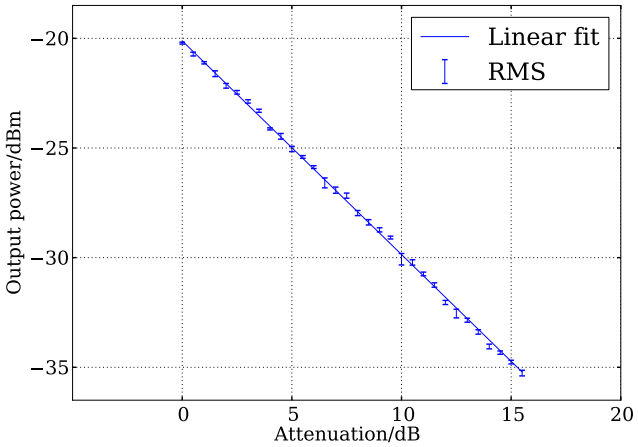

Figure 3.29: Output power of the vertical position channel for different attenuation settings.

3.3.1 System Gain

The total system gain from the cavity output to the digital sampling was measured using a CW signal. A signal power of 0 dBm was used for both the position cavity channels which have 6 dB of fixed attenuation before the input of the down-converter. For the reference cavity channel, there is 20 dB fixed attenuation before the down-converter input so the signal power used was 10 dBm. The power of the output signal was then determined from the root-mean-square of the signal recorded by the digitiser. The results are summarised in Table 3.5.

The effect of varying the attenuation in the horizontal channel was also checked at this point with an input signal of -10 dBm from the signal generator. The result is shown in Figure 3.29. The fit gradient is -0.974 ± 0.003 with a χ^2 per degree of freedom of 2.77. This is again, slightly higher than -1, the value which would

Channel	Input/dBm	Output/dBm	Total gain/dB
X	0	-10.71 ± 0.09	-6.1
Y	0	-10.55 ± 0.05	-5.9
Reference	+10	-12.348 ± 0.017	-17.7

Table 3.5: Results of gain measurements of the full processing electronics and signal transmission performed using a continuous wave signal generator.

be expected if the variable attenuator was ideal. This differs from the results of Section 3.2.1.4 and Section 3.2.3.3 because of variation between the individual attenuators.

3.3.2 Pulsed Input

Figure 3.30 shows the waveform measured when pulses of 50 ns are input to the down-converter in-situ. Digital down conversion (see Section 4.1.1) has been applied in order to measure the signal envelope and the phase and the frequency was measured using NAFF (see Section 5.3.1). The phase during the rise and fall of the pulse varies nonlinearly. This is significant for the measurement of short pulses produced by a single bunch excitation. Similar behaviour can be seen in the phase at the start of the second pulse which is also visible in the digitiser window. During the constant amplitude portion of the waveform, the phase is almost constant but demonstrates the variation expected from the finite phase stability of the LO signal.

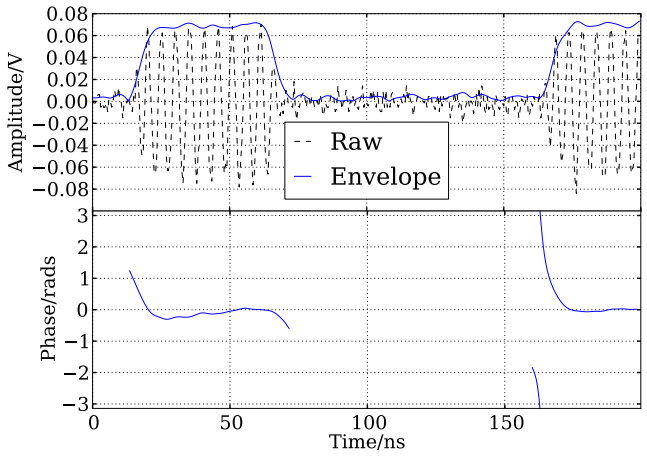

Figure 3.30: Measurement of the electronics in-situ with a 50 ns long input pulse.

Chapter 4

Digital Signal Processing

4.1 Theory

As discussed in Section 1.2.1.4, the relative phase between the position and reference cavity signals must be measured in order to remove the beam angle and bunch tilt components and determine the sign of the position offset. The phase of an RF signal is usually measured by splitting and then multiplying by two more signals at the same frequency that are 90 degrees apart in phase. The arctangent of the ratio of output amplitudes after multiplication then gives the phase while combining the amplitudes in quadrature gives the amplitude of the signal. The multiplication can either be done in analogue electronics using an IQ mixer or digitally using a number of different algorithms. The methods described here are based on an algorithm known as digital down conversion but others based on the same principle, such as fast Fourier transform or principle component analysis have been used.

4.1.1 Digital Down Conversion

The signal from the prototype cavity BPM pick-up is processed by the electronics described in the previous chapter and then digitised. An algorithm known as digital down conversion (DDC) can then be applied to convert the signal from the intermediate frequency of the electronics to baseband in order to separate the signal amplitude and phase. Assuming an exponentially decaying sinusoidal sig-

nal from a single bunch excitation of the prototype pick-up, the discrete digitised waveform in a single channel V_i is of the form

$$V_\mathrm{i}[n] = u(n-n_0)Ae^{-\frac{n-n_0}{\tau f_\mathrm{s}}}\cos\left(2\pi\frac{f_\mathrm{IF}}{f_\mathrm{s}}(n-n_0)+\phi\right) \tag{4.1}$$

where A is the signal amplitude, f_IF is the signal intermediate frequency, τ is the signal decay time, f_s is the sampling frequency, ϕ is an unknown phase offset, n is the sample number and n_0 is the time in samples, not necessarily an integer, at which the beam-induced signal starts. $u(k)$ is the unit step function and is equal to 0 for negative values of k and 1 for values greater than or equal to 0. The square bracket convention has been adopted here to indicate that the function is discrete and only exist for integer input values. Expressing Equation 4.1 for $n \geq n_0$ in exponential form gives

$$V_\mathrm{i}[n] = u(n-n_0)Ae^{-\frac{n-n_0}{\tau f_\mathrm{s}}}\left[\frac{e^{i\left(2\pi\frac{f_\mathrm{IF}}{f_\mathrm{s}}(n-n_0)+\phi\right)}+e^{-i\left(2\pi\frac{f_\mathrm{IF}}{f_\mathrm{s}}(n-n_0)+\phi\right)}}{2}\right]. \tag{4.2}$$

For conversion to baseband, the signal is multiplied by a digital complex local oscillator given by

$$L[n] = 2e^{i2\pi\frac{f_\mathrm{IF}}{f_\mathrm{s}}n} \tag{4.3}$$

where the factor of 2 is included so that the amplitude of the demodulated signal is the same as the amplitude of the raw signal. Multiplying Equation 4.1 by Equation 4.3 gives

$$L[n]V_\mathrm{i}[n] = u(n-n_0)Ae^{-\frac{n-n_0}{\tau f_\mathrm{s}}}\left[e^{i(2\pi\frac{2f_\mathrm{IF}}{f_\mathrm{s}}n-2\pi\frac{f_\mathrm{IF}}{f_\mathrm{s}}n_0+\phi)}+e^{-i\left(\phi-2\pi\frac{f_\mathrm{IF}}{f_\mathrm{s}}n_0\right)}\right] \tag{4.4}$$

There is now a sum frequency component at $2f_\mathrm{IF}$ and a zero frequency component. In order to remove the sum frequency component, a digital Gaussian filter is used whose impulse response $h[n]$ is given by

$$h[n] = \frac{1}{w\sqrt{2\pi}}e^{-\frac{n^2}{2w^2}} \tag{4.5}$$

where w is a width in number of samples. Equation 4.5 is convolved (see Section 4.1.3) with the down-mixed waveform. The frequency response of the filter is also a Gaussian and is centred at zero. The width of the Gaussian in the frequency domain Δf is given by

$$\Delta f = \frac{f_\mathrm{s}}{2\pi w}\,. \tag{4.6}$$

Each width Δf corresponds to a suppression of $e^{-\frac{1}{2}}$ in voltage or -4.3 dB. After filtering, assuming total removal of the sum frequency component and perfect preservation of the zero frequency component, the final processed signal V_f is complex and given by

$$V_\mathrm{f}[n] = Ae^{-\frac{n-n_0}{\tau f_\mathrm{s}}} e^{-i\left(\phi - 2\pi\frac{f_\mathrm{IF}}{f_\mathrm{s}}n_0\right)}\,. \tag{4.7}$$

The signal amplitude and phase at each sample are then easy to separate as the modulus and argument of the complex number [41].

4.1.2 Calibration

Once the position cavity signal has been down-converted, the reference cavity signal, processed in the same way, must be used to remove charge dependence and provide a reference phase that is fixed with respect to the time of the beam arrival. This is done by sampling the result of Equation 4.7 for the position cavity signal at a single point n_p and dividing it by the processed reference cavity signal sampled at point n_r. The result can be expressed as a complex phasor whose real part is known as the in-phase component I and whose imaginary part is the quadrature-phase component Q.

$$I + iQ = \frac{A_\mathrm{p}}{A_\mathrm{r}} e^{\left(\frac{n_\mathrm{r}-n_0}{\tau_\mathrm{r} f_\mathrm{s}} - \frac{n_\mathrm{p}-n_0}{\tau_\mathrm{p} f_\mathrm{s}}\right)} e^{i\left(2\pi\left(\frac{f_\mathrm{IF,p}-f_\mathrm{IF,r}}{f_\mathrm{s}}\right)n_0 + \phi_\mathrm{r} - \phi_\mathrm{p}\right)} \tag{4.8}$$

where the subscript p indicates that the parameter refers to the position cavity channel while the subscript r refers to the reference cavity channel.

Provided that the sampling points relative to the beam arrival $n_\mathrm{r} - n_0$ and $n_\mathrm{p} - n_0$ are fixed as well as the decay times τ_r and τ_p and frequencies f_r and f_p,

Equation 4.8 gives a complex quantity that is only dependent on beam position, trajectory angle and bunch tilt. Dividing by the reference signal amplitude A_r removes the charge dependence and subtracting the reference signal phase ϕ_r removes the influence the analogue LO has on the phase. The phase can then be used to extract the contribution of the beam position offset, which is in quadrature phase with the contributions from the beam trajectory angle and the bunch tilt.

Once the signal is represented in the form given by Equation 4.8, a rotation must be applied so that changes in position show up only in one component. This is equivalent to multiplication by a complex number of modulus 1 whose argument is equal to a calibration constant known as the IQ rotation angle θ_{IQ}. The real component then corresponds to the changes in position and can be scaled to physical units of position with a second calibration constant known as the position scale factor S_P. The position P is then given by

$$P = S_\mathrm{P}\mathrm{Re}\left[(I + iQ)e^{-i\theta_{IQ}}\right] . \tag{4.9}$$

The imaginary part of the complex number in square brackets is the contribution from the bunch tilt and the beam trajectory angle [41].

4.1.3 Deconvolution

The signal from any linear processing system is the convolution of the impulse response of the system with the input. In the case of a cavity BPM, the input is the bunch structure (scaled by the position offset in the case of the position cavity) and the impulse response of the pick-up has the form of an exponentially decaying sinusoid. Therefore, if it were possible to deconvolve the impulse response of the cavity BPM from the measured signal, the bunch structure of the beam scaled by the beam position would remain. The convolution $z(t)$ of two functions of t, $x(t)$ and $y(t)$, is given by

$$z(t) = \int_{-\infty}^{\infty} x(t')y(t - t')dt' . \tag{4.10}$$

It can be shown that in the frequency domain, this is equivalent to the product of the Fourier transforms of $x(t)$ and $y(t)$ ($\tilde{x}(k)$ and $\tilde{y}(k)$ respectively).

$$\tilde{z}(k) = \tilde{x}(k)\tilde{y}(k) \tag{4.11}$$

where the variable k is proportional to $1/t$. Deconvolution is a method of obtaining one of the original functions from the convolution given knowledge of the other function using

$$x(t) = \mathcal{F}^{-1}\left(\frac{\tilde{z}(k)}{\tilde{y}(k)}\right) \tag{4.12}$$

where $\mathcal{F}^{-1}$ represents the inverse Fourier transform.

In the case where the signals are discrete and of finite length, the discrete Fourier transform (DFT) must be used instead of the continuous Fourier transform in Equation 4.12. It can be shown that, analogous to Equation 4.11, the product of two DFTs is equivalent to

$$\tilde{x}[k]\tilde{y}[k] = \mathcal{D}(z[n]) = \mathcal{D}\left[\sum_{m=0}^{N-1} x[m]y[n-m]\right] \tag{4.13}$$

where m, n and k are integers between 0 and $N-1$ and $\mathcal{D}$ represents the DFT. $x[n]$ and $y[n]$ are periodic signals of period N and $\tilde{x}[k]$ and $\tilde{y}[k]$ are their respective DFTs calculated over the length of one period. The term in brackets on the right hand side of Equation 4.13 is known as the circular convolution of the two periodic signals. It is similar to the discrete convolution defined as

$$z[n] = \sum_{m=-\infty}^{\infty} x[m]y[n-m]\,. \tag{4.14}$$

If two finite-length signals $x[n]$ and $y[n]$ are only nonzero over the finite ranges $0 < n < N_x$ and $0 < n < N_y$ respectively and their periodic extensions are found by repeating after every $N_x + N_y - 1$ samples, then their discrete convolution is equal to the circular convolution of their periodic extensions between $n = 0$ and $n = N_x + N_y - 1$ [42].

Deconvolution of discrete signals involves dividing by the DFT of one of the convolved waveforms or multiplication by it's inverse frequency response $1/\tilde{y}[k]$.

When y contains a real signal spread over a limited set of frequencies, the inverse frequency response will be small at these frequencies and large elsewhere where there may be noise. In order to mitigate this noise amplification, the inverse frequency response may be filtered by multiplying it by the expression

$$g[k] = \frac{k_2 - k}{k_2 - k_1} \tag{4.15}$$

where k_1 and k_2 are constant integers. This expression ensures that the inverse frequency response is smoothly reduced to zero between the frequencies of k_1/N and k_2/N which would be above the signal frequencies [43]. It should be noted that from Equation 4.14 that if there is a delay in one of the signals, ie. $x[n] \rightarrow x[n - \delta n]$, then the convolved signal will be delayed by the same amount. Therefore, the delay in the deconvolved waveform δn_x is related to a delay in the other waveform δn_y and the delay of their convolution δn_z by the following expression

$$\delta n_x = \delta n_z - \delta n_y \,. \tag{4.16}$$

The signal excited in a cavity BPM by multiple bunches is the single bunch excitation convolved with a series of delta functions at the bunching frequency, scaled by the bunch charge in the case of the reference cavity and charge and position in the case of the position cavity. Therefore, information about the bunch positions can be extracted by deconvolving the single bunch response from the multiple bunch signal. Equation 4.13 is used to do this digitally. The DFT of the multiple bunch signal is multiplied by the inverse frequency response, filtered using Equation 4.15, of the single bunch signal . An inverse DFT is then applied to obtain the deconvolved waveform. Because the convolution in Equation 4.13 is circular, care must be taken with the signal timing. If the start of the single bunch waveform comes after the multiple bunch waveform within the finite digitiser window, following Equation 4.16, the start of the deconvolved waveform will be negative. The periodicity in N means that the deconvolved waveform may be split and appear at the start and end of the window and problems with edge effects would arise if further digital filtering were applied.

There is an additional effect since the signal phase relative to the analogue LO depends on the bunch arrival time and, therefore, the number of the bunch within

the train. This phase processes at the frequency difference between the analogue LO and its nearest bunching frequency harmonic. The instant change in phase at the arrival of each bunch can only be accounted for in the convolution by using complex signals and complex scalings at the arrival of each bunch. However, the digitised signals are real and discrete. In this case, the change in phase can be approximated by adding the single bunch waveform at different amplitudes and offset in time. This is what appears in the deconvolved discrete signal. In the case of CLIC and CTF3 however, because the bunch arrival frequency is larger than the digitiser sampling frequency, the changes in phase are close together and form a sinusoidal signal with constant amplitude that oscillates at the frequency difference between the bunching frequency harmonic and the LO and lasts as long as the bunch train.

4.2 Simulation

In order to test the deconvolution algorithm, a simulation was run using the DiagSim diagnostics simulation package written in python [44]. The simulation parameters are given in Table 4.1. The electronics parameters are the same as those calculated for the baseline electronics design with extra attenuation on the reference cavity channel to prevent digitiser saturation. The pick-up sensitivities

Parameter	Reference cavity	Position cavity
Mode resonant frequency/GHz	15	15
Pick-up sensitivity/V nC^{-1}	50	14 mm^{-1}
Signal decay time/ns	5	5
Electronics gain/dB	-18.6	8.4
Electronics noise figure/dB	32.3	16.0
Electronics bandwidth/MHz	300	300
Intermediate frequency/MHz	200	200
Digitiser sampling rate/GS s^{-1}	2	2
Digitiser resolution/bits	32	32
Digitiser range/V	± 1	± 1

Table 4.1: Parameters used for simulation of different digital processing algorithms.

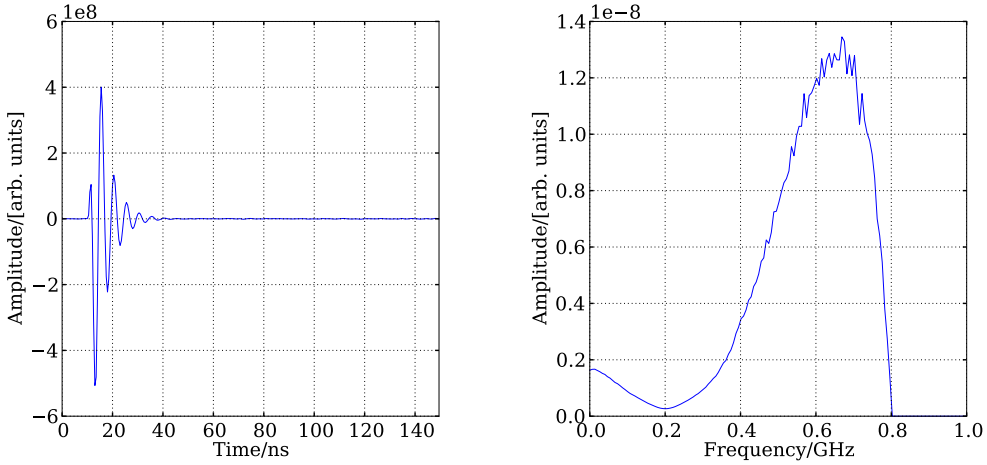

Figure 4.1: Simulated single bunch response as measured by a 32 bit digitiser (left) and its filtered inverse frequency response (right).

were selected within an order of magnitude of the predicted values, which should be sufficient since the simulation is not expected to be accurate in terms of absolute resolution. The simulation is kept as linear as possible with linear responses in all the electronic components and a high digitiser resolution. Random thermal noise is also included. The bunch arrival frequency used is the 1.5 GHz of the CTF3 probe beam and the bunch charge is 0.6 nC in the case of a single bunch and 0.05 pC for the simulated trains of 150 bunches.

In order to perform the deconvolution, some single bunch waveforms were first generated such as the one in Figure 4.1. For this, DiagSim assumes a perfect decaying sine wave at the output of the pick-up and then simulates the analogue signal processing and digitisation. Figure 4.1 also shows the inverse frequency response which was determined with a filter as described by Equation 4.15 with $k_1 = 400$ MHz and $k_2 = 800$ MHz. This was then applied to the raw multiple bunch signal shown on the left of Figure 4.2. The start time of the single bunch waveform, 10 ns, is sooner than the start time of the multiple bunch waveform, 20 ns. The start time of the deconvolved waveform is then given by Equation 4.16 as 10 ns so the deconvolved waveform does not cross the ends of the digitiser sampling window. Deconvolution was applied to both the position cavity signal and the reference cavity signal.

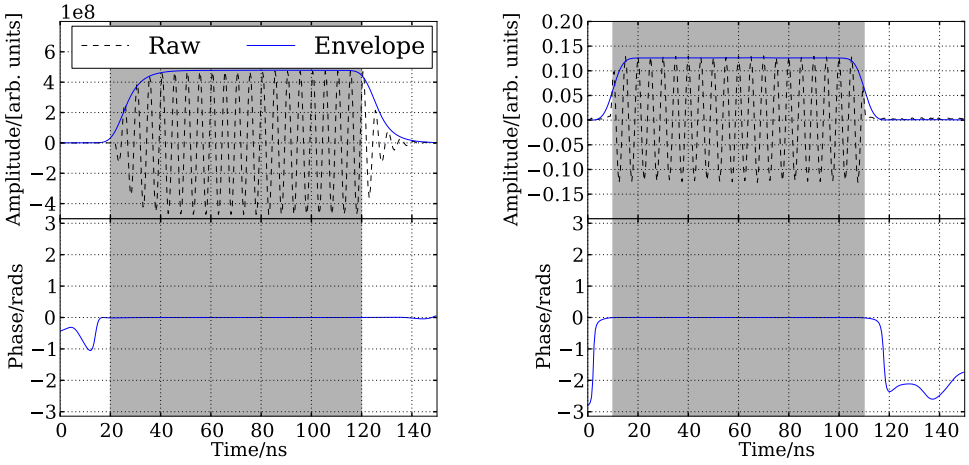

Figure 4.2: Simulated raw digitised waveform from a train of 150 bunches and its amplitude and phase as determined using DDC (left) and the same waveform where deconvolution and has been applied along with its amplitude and phase from DDC (right).

Digital down-conversion was applied to the raw multiple bunch signal and the deconvolved signal. The down-converted waveforms are shown in Figure 4.2. It is immediately clear that the deconvolution has steepened the rise and decay of the multiple bunch waveform and leaves a signal at the intermediate frequency (which here is the same as the bunch arrival harmonic minus the LO frequency) of almost constant amplitude as explained in Section 4.1.3.

Position values were calculated for every sample within the 120 ns over which the integrated amplitude of the deconvolved reference cavity signal was greatest. The same time period shifted by 10 ns was used for the raw multibunch waveforms. A calibration using 5 step changes in position, each of 10 μm, was performed in the horizontal direction to determine the IQ rotation angle and the position scale factor and these same constants were used for the vertical direction. The method used for the calibration is the same as described in Section 5.6.1. No bunch to bunch position jitter was added in the horizontal direction so that the calibration would be accurate and the spatial resolution of the simulated processing could be estimated. The nature of the vertical beam position variation along the bunch train was adjusted to probe the time resolution in different ways.

4.2.1 Resolution

The spatial resolution from the simulation was determined as the root mean square (RMS) of the measured position minus the simulated vertical beam position. Since there is no bunch to bunch position jitter, the simulated position is the same for the whole bunch train. Figure 4.3 shows the calculated position along one train in this case and the resolution determined from 250 pulses. The quantisation of the digitiser is visible in the train position measurements as the readings jump between different values. The deconvolution degrades the spatial resolution in this case because it attenuates the signal frequency and amplifies the out of band noise.

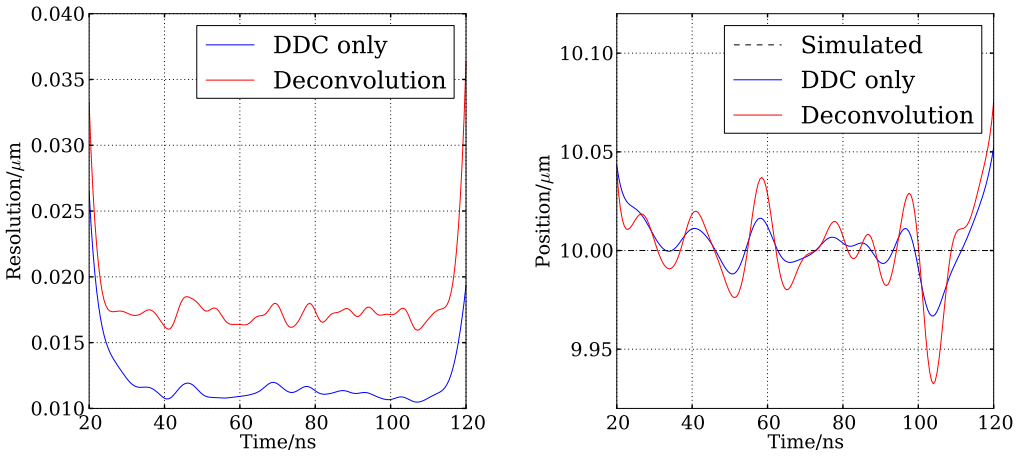

Figure 4.3: Spatial resolution of position measurements (left) and the position measurements themselves (right) along a bunch train with no change in bunch position.

Figure 4.4 shows the same plots but this time, variation in the vertical beam position has been added. The bunch to bunch position jitter is 1 μm and the position of each bunch is drawn from a normal distribution centred around the position of the previous one. This results in a random walk in position along the bunch train. It can be seen qualitatively that the positions calculated after deconvolution follow the evolution of the random walk more closely than when no deconvolution is applied suggesting that the time resolution has been improved. This also results in a better position resolution because the deconvolution algo-

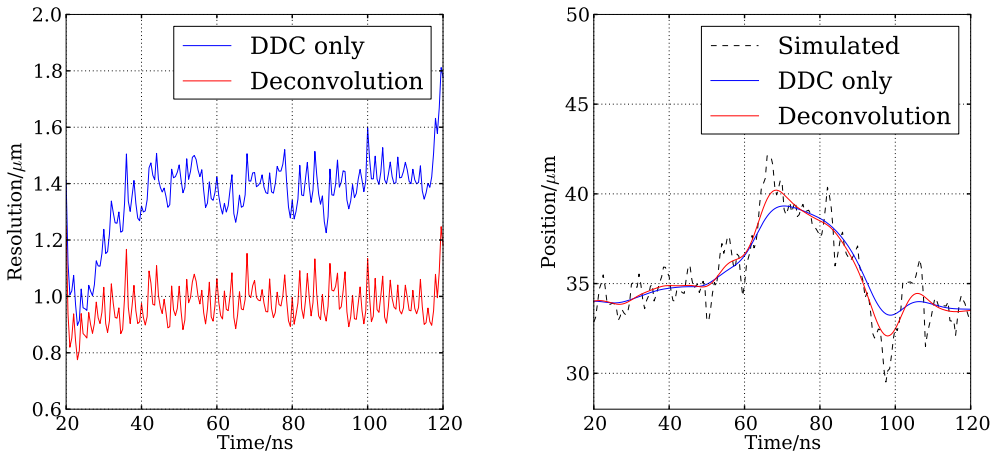

Figure 4.4: Spatial resolution of position measurements (left) and the position measurements themselves (right) along a bunch train with a random walk in bunch position.

rithm is able to follow the changes in position more closely.

Next, the vertical bunch to bunch position jitter was increased to 10 μm and was made completely random. The correlations between the positions calculated for each sample with the positions calculated for all other samples over 250 pulses were then determined. The correlations are shown for a single sample at the top of Figure 4.5 and for all samples, in correlation matrices represented by colour plots at the bottom of Figure 4.5. Despite the random motion, there is some correlation between nearby samples because of the finite measurement bandwidth. This correlation is a near symmetric peak of almost constant width along the 120 ns train. Deconvolution reduces the average width of this peak from a FWHM of 16 ns to 10 ns, thus improving the time resolution.

4.2.2 Average Train Position

As well as making several position measurements within a single bunch train, the CLIC main beam BPM must be able to make an average position measurement of the whole train. Further simulations were performed in order to find the best way to do this. Figure 4.6 shows the results for three different methods, for which, the signals were averaged over the same 120 ns used in the previous

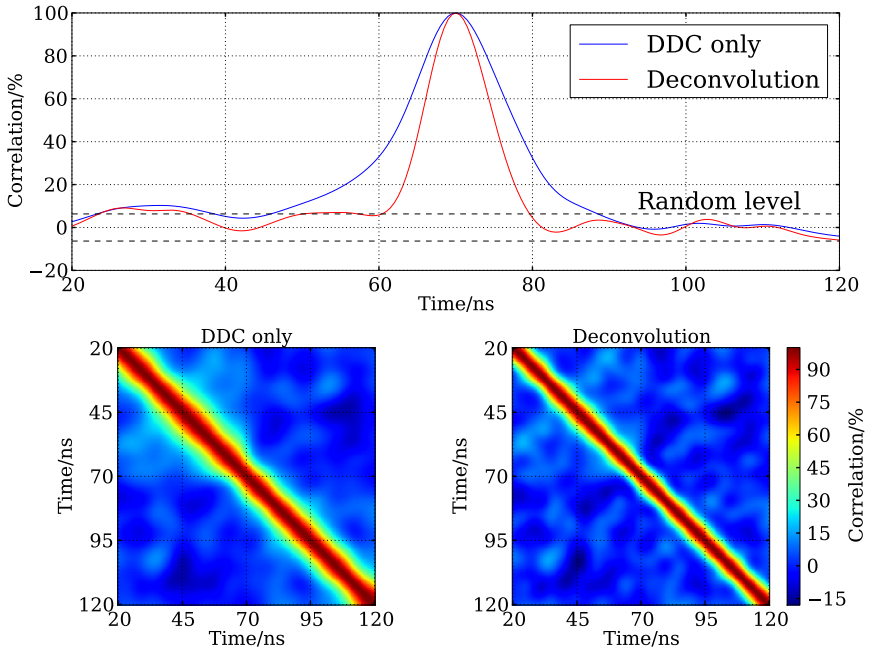

Figure 4.5: Correlation between a central position measurement and the rest within a bunch train (top) and the correlation matrices with digital down-conversion only (bottom left) and after deconvolution (bottom right).

two sections. The difference between the position calculated from the average signal and the average simulated position for 250 trains was used to calculate the resolution. For the first two methods (one without deconvolution and one with), each sample is individually normalised for charge and referenced in phase using its corresponding sample from the reference cavity signal. The average is then taken last. For the third method, the samples from each signal are averaged first and then the charge normalisation and phase referencing is done using the average signals. Deconvolution is not applied in this case.

For all three methods, the resolution is degraded with increased bunch to bunch position jitter. The third method performs worst, however, because the shape of the multiple bunch signal means that not all samples in the average are weighted equally. Performing the normalisation first makes the weighting in the average more uniform along the train. Deconvolution improves the resolution

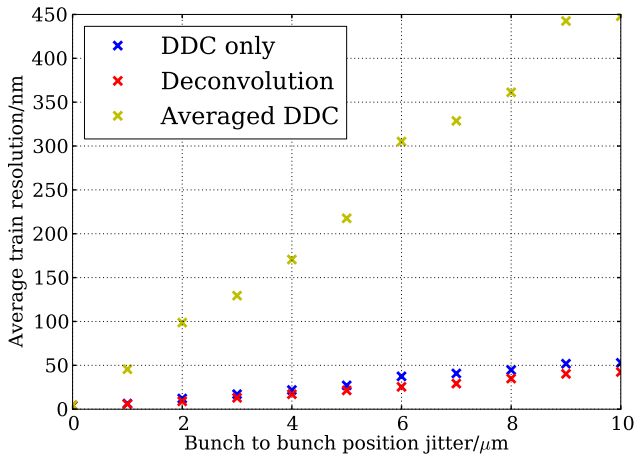

Figure 4.6: Resolution of average train position measurements for different levels of bunch to bunch position jitter.

further since it removes the signal rise and decay as can be seen in Figure 4.2, further improving the uniformity of the sample weights. As a further check, the simulation was performed again but with a shorter decay time for the reference cavity and the best resolution was still obtained when the averaging was done after the normalisation. Challenges such as timing and correlated position variation that are not present in the simulation may have an effect that improves the performance of the third method over the other two. Results of actual beam measurements will give the best indication.

Chapter 5

Beam Tests

5.1 CALIFES

CALIFES (Concept d'Accélérateur Linéar pour Faisceau d'Électron Sonde) is the name given to the probe beamline of the 3rd CLIC Test Facility (CTF3). Its principle role is described in Section 1.1.2. The prototype cavity BPM is currently installed at the end of the beamline and beam data was taken during the running period from February through to May 2013.

The nominal beam parameters are listed in Table 5.1 but large variation from these was observed during beam operation. The energy has been seen to go as low as 180 MeV. At the time of the beam tests, it was not possible to achieve a single bunch charge of 0.6 nC due to degradation of the photocathode. The maximum charge of long beam pulses is also limited by beam loading, extraction of more RF energy from the accelerating structures than can be replenished for the later bunches. Typically, 0.05 nC was used as a single bunch charge and 60 ns long trains of 3.5 nC (0.04 nC bunch^{-1}) were used. The repetition rate commonly used, 1.6 Hz, is twice the repetition rate of the high current drive beam. On the occasions when there was drive beam in the CLIC Experimental Area (CLEX), electromagnetic interference could be seen on many monitors, including the cavity BPM prototype. The data presented here was taken with no drive beam present.

Parameter	Value
Energy/MeV	200
Repetition rate/Hz	1.6
Bunch arrival frequency/GHz	1.5
Bunch charge/nC	0.6
Number of bunches	1-226

Table 5.1: Parameters of the CTF3 probe beam CALIFES [6].

5.1.1 Beam Generation and Transport

Figure 5.1 shows a diagram of the full CALIFES beamline. The source of the probe beam is a photocathode injector where a high power laser incident on a cathode causes the photo-emission of electrons. The 1.5 GHz laser oscillator defines the bunch arrival frequency and a pulse picker, using two Pockels cells and three polarisers, defines the number of bunches. Each Pockels cell can be switched in one direction in under 400 ps using a high voltage pulser. The fast switching of each Pockels cell defines one edge of the laser pulse going to the photo-cathode and so by adjusting the relative timing of the two, the beam pulse length can be varied with high precision. An RF gun and three 3 GHz accelerating structures, all powered by the same klystron, are then used to accelerate the beam to 200 MeV. Mechanical phase shifters allow the phase of the RF power to the gun and to the final accelerating structure to be controlled independently. This is crucial for stable operation. Additionally, by operating the final accelerating structure off crest, some bunch length compression can be achieved. Solenoid magnets are used for focusing of the beam at low energy and during acceleration. After the accelerating structures, there is a quadrupole triplet and a beam energy spectrometer and then another quadrupole triplet before the beam enters TBTS. At this point, CLIC accelerating structures, powered by the drive beam, are tested at gradients exceeding 150 MV m^{-1} [17]. After TBTS, there is a final quadrupole triplet and another spectrometer to measure the full energy of the beam. The probe beamline is instrumented with screens for transverse beam profile measurements and several BPMs.

During a typical start-up procedure, the klystron must first be conditioned so that it can produce the required output power without electric breakdown. Next,

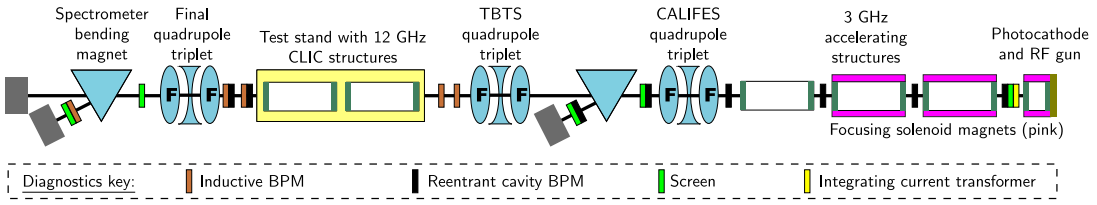

Figure 5.1: Diagram of the full CALIFES beamline [1].

a shutter is opened, allowing the laser onto the photocathode. The condition of the beam coming out of the gun can be inspected using a screen (described in Section 5.1.2.2) just downstream and tuned by adjusting the phase of the RF gun. The beam must then be steered along the line to ensure transmission as far as the beam dump of the final spectrometer. Beam position monitors and several screens along the line can be used to do this. There is a screen immediately after the first quadrupole triplet in CALIFES. By measuring the beam size while varying the strength of two quadrupole magnets and using the energy measured at the spectrometer, the Twiss parameters and the emittance, described in Section A.1 of the Appendix, can be measured in both transverse directions.. A model of the beamline is then employed to focus the beam both horizontally and vertically in a chosen location using any of the three quadrupole triplets [4].

5.1.2 Diagnostics

As mentioned in Section 5.1.1, beam diagnostics are essential to be able to steer and tune the beam. They can also be used to make measurements that can be compared with the prototype cavity BPM output. Only a few of the beam monitors in CALIFES give processed and calibrated beam readings and each may be used to measure more than one beam property. It is therefore necessary and advantageous to develop user-specific processing of the raw digital readout from each monitor.

5.1.2.1 Inductive BPMs

The pick-up of an inductive BPM is different from that of a button or stripline BPM described in Section 1.1.1. The metallic beam pipe is interrupted along

its length with a ceramic section. Equally spaced on the outside of the ceramic are eight long, flat electrodes that have a continuous electrical connection to the beam pipe at both ends. The electrodes are surrounded by ferrite and are designed to have a high inductance and a low cut-off frequency to minimise signal droop. At one end of each electrode is a screw that passes through a transformer on a printed circuit board (PCB). As the beam passes through the pick-up, its image current must pass through the transformers and along the electrodes. The signals induced in the electrodes are in pairs and the difference between opposing pairs is the difference signal Δ, which is sensitive to position, while all of the electrodes together provide the sum signal Σ. Because the whole image current must pass through the transformers, the sum signal of an inductive BPM gives a good absolute measurement of the beam current. Each transformer also has a calibration turn where a current known to 0.1 % can be injected to calibrate the beam current measurement [45].

The readout of the inductive BPMs along the probe beamline is such that treated signals are directly available which give the position and current profile of the bunch train at a sampling rate of 200 MHz. An example of a treated set of signals from one of the inductive BPMs (CA.BPM0750) is shown in Figure 5.2. In order to extract a single position reading, the start and end of the beam pulse are determined from the maximum and minimum difference between consecutive samples of the sum signal. The position signal is then sampled half way between these two points.

5.1.2.2 Beam Profile Monitors

The probe beamline includes 6 screen monitors that can be inserted into the beam path at a 45° angle in order to inspect the transverse beam profile. The screens are either made from yttrium aluminium garnet (YAG) or silicon and emit light through scintillation or optical transition radiation (OTR) respectively when the beam passes through. After each quadrupole triplet, there is a screen that can be used for beam size measurements in quadupole scans. The two spectrometers also have screens so that the beam average energy and longitudinal energy profile can be determined from the position and width of the image on the

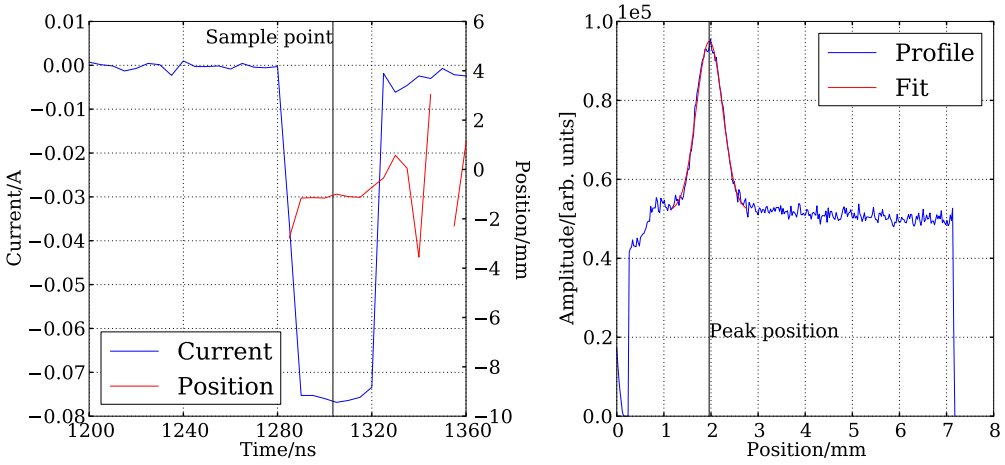

Figure 5.2: Example signals from the sum channel and one position channel of an inductive BPM (left) and a YAG screen image projection (right).

screen respectively [46]. For these tests, two of the screens are used as a position diagnostic for comparison with the cavity BPM.

The light from each screen is emitted at a 90° angle to the beam and is detected by a charge coupled device (CCD) after some optics. Figure 5.2 shows an example of a projected image profile from the screen (CA.MTV0790) that is after the final quadrupole triplet. The amplitude of each bin in the image profile is obtained from the summation of one line of pixel values. Summation along the vertical axis of the image gives the horizontal image profile and vice versa. The location of the beam in the profile is determined from the bins whose values are above 75 % of the profile maximum. These bins and 20 bins either side are then used for a Gaussian fit to the peak in the beam profile. The location of the maximum from the Gaussian fit is used as the beam position.

5.1.2.3 Reentrant Cavity BPMs

The second type of beam position monitor used in the probe beamline is based on a pick-up with a reentrant cavity that has four magnetically coupled antennae around the beam-pipe. As with the cavity BPM prototype, resonant electromagnetic modes excited in the cavity are used to determine the beam position. The signals from opposite ports are combined in a 180° hybrid where the signals from

the cavity monopole mode adds constructively in the hybrid sum port and destructively in the difference port and vice versa for the cavity dipole mode signal. The monopole mode signals from each pair of output ports are combined and rectified using a diode while the dipole mode signals are demodulated using an IQ mixer and a local oscillator locked to the machine RF. The signals are then digitised [47]. Some examples of digitised waveforms are shown in Figure 5.3 for a 20 ns pulse train. The slope of the monopole mode signal is used to determine the signal start time and the delay between the different channels is measured once and kept fixed. The signal from each channel is then averaged over the 15 ns that corresponds to the maximum monopole signal level and the unscaled position signal P is determined as

$$P = \frac{\sqrt{I^1 + Q^2}}{\Sigma} \text{sign}\left(\tan^{-1}\left(\frac{\text{Q}}{\text{I}}\right)\right) \tag{5.1}$$

where I and Q are the in-phase and quadrature-phase signals of a position channel respectively and Σ is the sum signal level. These BPMs are more sensitive to beam charge and position than the inductive type BPMs but, like the prototype cavity BPM, require a beam based calibration to be able to give position readings in physical units.

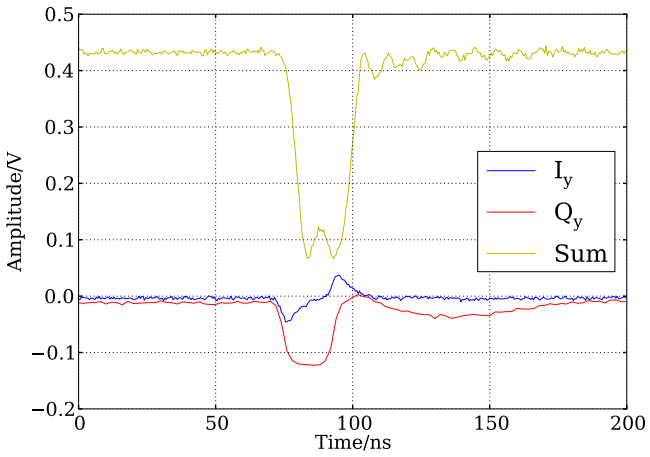

Figure 5.3: Example of signals from the sum and vertical channels of a reentrant cavity BPM.

5.2 Cavity BPM Installation

5.2.1 Location

The cavity BPM is installed downstream of the final spectrometer magnet such that the beam passes through it when the spectrometer magnet is not being used. This means that the beam is not passing through the BPM pick-up during normal operation which protects the electronics from saturation. This location was chosen because it is downstream of all other experiments so the small beam-pipe aperture of the BPM (8 mm diameter) has no effect on the beam transmission. A quadrupole triplet is just upstream and this can be used to focus the beam in the cavity pick-up to minimise the beam position jitter.

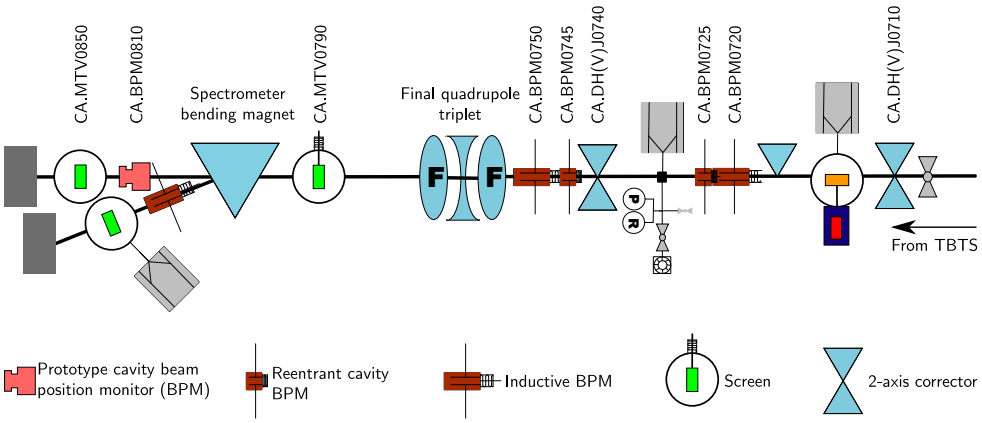

Figure 5.4: Beamline diagram of the end of the probe beamline where the cavity BPM prototype is installed.

Figure 5.4 is a diagram of the end of the probe beamline where the prototype cavity BPM is installed. For these tests, either the CALIFES quadrupole triplet or the final one (visible in Figure 5.4) was used to focus the beam in the centre of the cavity BPM. The two dual-axis dipole corrector magnets that are labelled in Figure 5.4 (CA.DH(V)J0740 and CA.DH(V)J0710) were used to vary the beam position in order to calibrate the BPM and measure the sensitivity of the cavity pick-up. As well as being used individually, the two correctors could be used antagonistically as a pair to shift the beam to a different position without causing

a change in beam trajectory angle, assuming the response of the beam to both correctors is the same.

By using the final quadrupole triplet, a smaller beam size and therefore less beam position jitter can be achieved. However, an optical model, described in Section A.2 of the Appendix, is then required to determine the change in the beam position at the BPM due to a change in current in one of the corrector magnets. When the final quadupole triplet is switched off, a simple ballistic model can be used to do this. This is where the change in transverse beam position at a specific location downstream of a corrector is simply the angle by which the beam is deflected in the corrector multiplied by the distance to that location. Measurements made by the cavity BPM were compared with those made using the other diagnostics that are labelled in Figure 5.4.

5.2.2 Layout

A diagram and photo of the prototype cavity BPM installation are shown in Figure 5.5. The pick-up assembly is mounted on mover stages from Newport [48]. The motion of the BPM is accommodated for by bellows and 40 mm long beam pipe spool pieces which are of the same aperture as the BPM. There are then tapered transitions to the 40 mm diameter aperture of the upstream and downstream vacuum chambers. A fixed OTR screen (CA.MTV0850) is installed downstream so that good beam transmission can be ensured. The down-converter

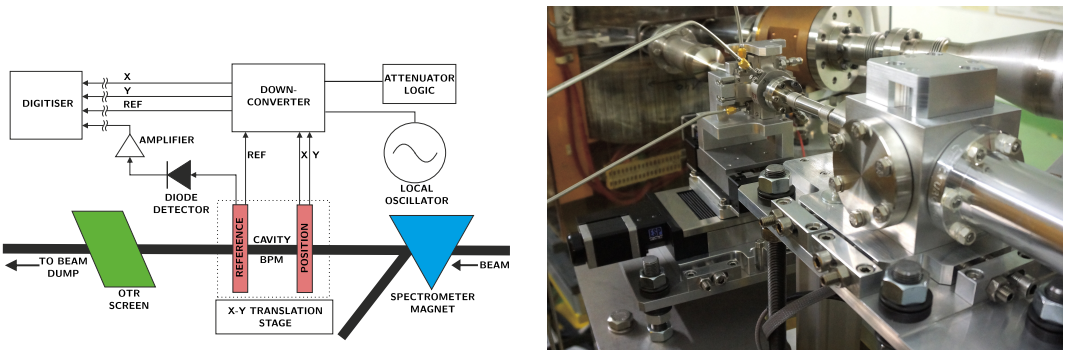

Figure 5.5: Diagram (left) [3] and image (right) of the prototype installation where the pick-up, the translation stages and the vacuum chamber for the downstream OTR screen are visible.

electronics are mounted on the legs of the supporting table and 700 mm long rigid coaxial cables connect the three output ports of the cavity BPM (one for each channel) to the input ports of the down-converter. For the two position cavity channels, 3 dB fixed attenuation is added at each end of the cable giving a total of 6 dB fixed attenuation while for the reference cavity, 10 dB is installed at each cable end. There is a diode rectifier and amplifier installed on the second output port of the reference cavity [3] which can be used for signal arrival time measurements.

The local oscillator box and a control board for the digitally controlled attenuators are installed on the other side of the spectrometer beam dump. Corrugated 1/4" coaxial cables, 3 m long, are used for the transmission of the LO signal to the down-converter electronics. The same type of cable is used to take the signals from the output of the down-converter to the klystron gallery where the digitiser is located. The control voltage for the VCO that produces the LO signal comes from a power supply in the gallery. This can be controlled remotely via a laptop running LabVIEW [49] that is also in the gallery and has a network connection. The control board for the digitally controlled attenuator has an RS232 interface and a LabVIEW driver and is controlled remotely via a serial device server in the tunnel that is connected to the CERN technical network. Both the VCO control voltage and the attenuation setting can be adjusted from the control room.

5.2.3 Data Acquisition

The CERN Open Analogue Signal Information System (OASIS) [50] is used to acquire the data from the digitiser. This allows the signals to be connected to, viewed and controlled through a virtual scope interface. Once connected, the signals can also be acquired independently through the controls middleware [51]. The control available through the virtual scope includes the digitiser range (through the digitiser pre-amp), the length of the digitiser window, the pedestal value and the sampling frequency. For all four digitiser channels, the sampling frequency was set to its maximum value of 2 GS s^{-1} and the length of the digitiser window was commonly set to 400 samples (200 ns). For the down-converter channels, the digitiser range was mostly kept at ± 1 V and the pedestal

value was left at 0 V. For the diode channel, the digitiser range was adjusted so that all of the signal was clearly visible and the pedestal level was set at about 80 % of the range maximum since the diode response is unipolar and negative in voltage. The output of the digitiser is cast to a 16-bit signed integer which can be converted to a voltage using a sensitivity value that is dependent on the digitiser range. The Front-End Software Architecture (FESA) [52], on which OASIS is based, provides a hardware independent interface and is universal across CTF. This allows signals from multiple sources to be acquired synchronously in the same way. The signals from the reentrant cavity and inductive BPMs as well as the beam pulse charge information and screen images can therefore all be recorded at the same time as the signals from the prototype cavity BPM.

5.2.4 Synchronisation

The ability to compare the cavity BPM signals with the other diagnostics was tested during two beam position scans performed using the pair of corrector magnets. For the first scan in vertical position, a 20 ns long beam pulse (about 30 bunches) was used and the measurements were compared with the signals from the reentrant BPM that is closest to the cavity BPM installation and after the

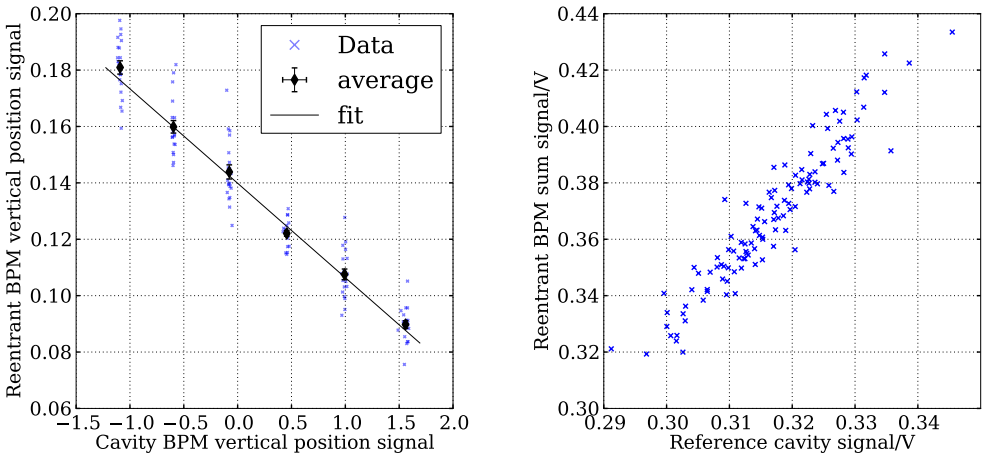

Figure 5.6: Correlation of the prototype cavity BPM position (left) and charge (right) signals with the signals from one of the reentrant cavity BPMs.

two correctors. The processing of the cavity BPM signals is as described later in Section 5.6.1 and the results are shown in Figure 5.6. In this configuration, both BPMs are able to measure the position shifts, which are about 1 mm, but between the position steps, no correlation can be seen. Although it was not measured, it is clear that the noise from the reentrant BPM is significantly worse than the noise from the cavity BPM. It is also well above the specified resolution of 4 μm so the performance of the reentrant BPM is clearly not optimal. The monitors are synchronised and measuring the same beam pulses. This is evident in the 95 % correlation between the reentrant BPM sum signal and the cavity BPM reference signal shown on the right side of Figure 5.6.

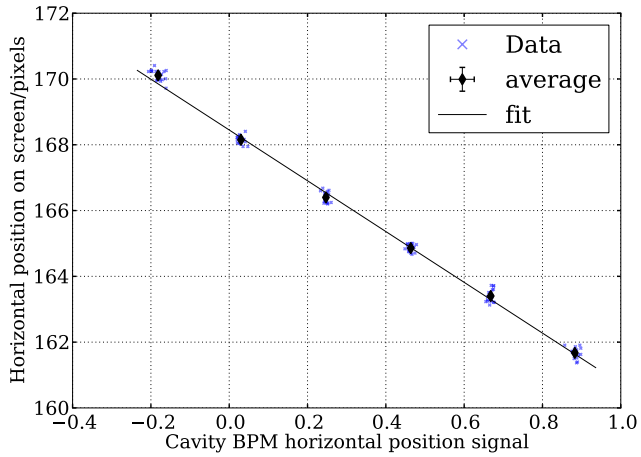

Figure 5.7: Correlation of the prototype cavity BPM position signal with the position of the beam measured on the downstream OTR screen.

Figure 5.7 shows the correlation between the cavity BPM output and the measured positions of the beam on the downstream OTR screen (CA.MTV0850) which were made during a horizontal position scan with 60 ns long beam pulses. Because this screen has a lower magnification than the one discussed in Section 5.1.2.2, there were not enough points over the beam profile to perform a Gaussian fit. The average position of the bins whose values were over 95 % of the image profile maximum was therefore used for the beam position, given in pixels since the screen does not have a calibrated position scale. There appears to be less noise than from the reentrant BPM but still no apparent correlation between

the large step changes in position. Since there is no correlation between position steps for either monitor, neither can be used to subtract the pulse to pulse beam position jitter.

5.2.5 Corrector Angular Response

In order to make sensitivity measurements, it must be possible to change the position of the beam at the location of the cavity BPM by a known amount. This was done using the two two-axis correctors upstream, CA.DH(V)J0710 and CA.DH(V)J0740 in Figure 5.4, where H or V refers to the horizontal or vertical axis of the corrector respectively. There are two types of diagnostic downstream of the correctors that are able to give calibrated position readings. These are the inductive BPMs, CA.BPM0720 and CA.BPM0750, and the YAG screen CA.MTV0790. The inductive BPMs are not sensitive enough to effectively measure changes in position within the aperture of the cavity BPM, particularly at low charge, and the screen intercepts the beam. Therefore, neither of these are able to give position readings during the sensitivity measurement itself so the response of the beam to the dipole correctors was measured separately.

Since the first corrector is upstream of the screen and both BPMs, its effect on the beam trajectory can be measured using all three monitors. For the second corrector, only the second BPM and the screen are downstream and can be used. As well as being measured individually, the response of the beam to two correctors acting antagonistically as a pair was measured, also using the second BPM and the screen.

The corrector current was varied in steps and at each step, 20 pulses were recorded and the processing described in Sections 5.1.2.1 and 5.1.2.2 was applied. For each step, the median and upper and lower quartiles were found. The position readings that were more than three times further from the median than the most distant quartile were cut. A weighted linear least squares fit was then performed to measure the response of the beam in position at the monitor location. Since there are no strong focusing or nonlinear elements between the correctors and the BPMs and the quadrupole triplet was switched off for the measurements using the screen, the response in beam angle could be determined by dividing the

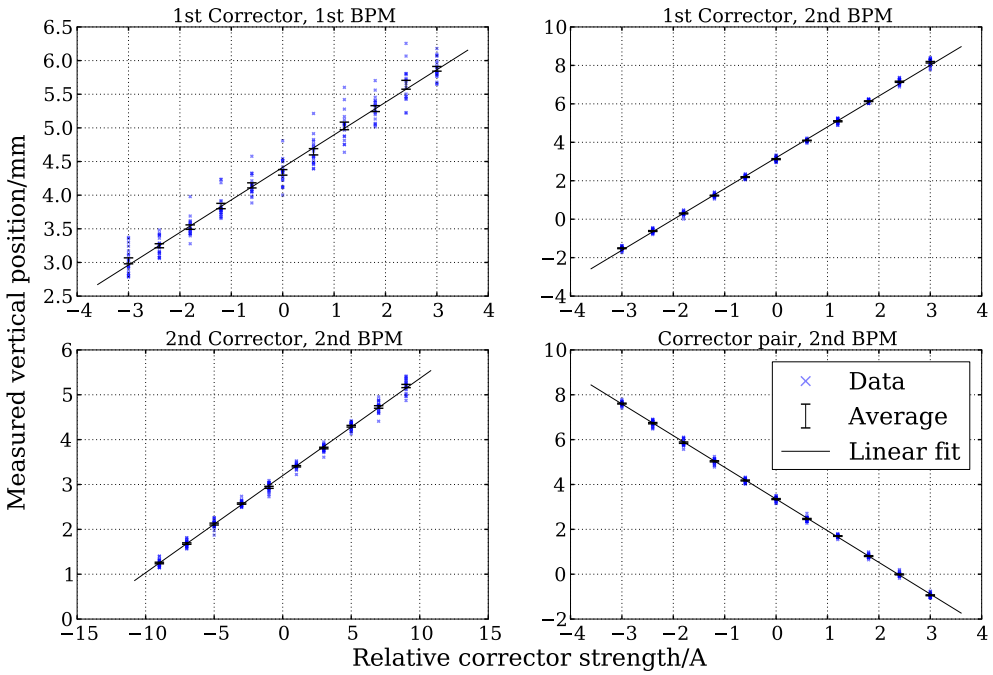

Figure 5.8: Measured vertical position against current in the corrector for the first corrector (CA.DVJ0710) and the first inductive BPM (CA.BPM0720) (top left), the first corrector and the second BPM (CA.BPM0720) (top right), the second corrector (CA.DVJ0740) and the second BPM (bottom left) and both correctors acting as a pair and the second BPM (bottom right).

response in position by the distance between the corrector and the monitor used. In each case, the measurement was performed twice to get an impression of the repeatability.

The fits to the vertical BPM data are shown in Figure 5.8 and the results for both transverse directions are summarised in Table 5.2. In all cases, the two results are within 2 % of each other except for the measurement of the vertical beam response to the first corrector using the first BPM. The vertical channel of this BPM was faulty and giving many anomalous readings. Only the two results for the pair of correctors in the vertical direction are further apart than three combined errors.

For the screen, the fits for the vertical direction can be seen in Figure 5.9 and

Axis	Corrector	BPM	Cut pulses	$\frac{\chi^2}{\text{DOF}}$	Position/ mm A^{-1}	Angle/ mm m^{-1} A^{-1}
Horizontal	1st	1st	0	1.40	-0.477 ± 0.003	-0.681 ± 0.004
			1	2.46	-0.486 ± 0.003	-0.694 ± 0.004
	1st	2nd	2	1.09	-1.624 ± 0.003	-0.7061 ± 0.0014
			5	2.51	-1.629 ± 0.003	-0.7081 ± 0.0013
	2nd	2nd	1	1.43	-0.2177 ± 0.0010	-0.726 ± 0.003
			3	2.90	-0.2192 ± 0.0010	-0.731 ± 0.003
	Pair	2nd	4	4.10	1.402 ± 0.003	N/A
			4	5.53	1.407 ± 0.003	N/A
Vertical	1st	1st	17	2.08	0.460 ± 0.005	0.657 ± 0.007
			13	1.00	0.484 ± 0.006	0.691 ± 0.009
	1st	2nd	3	9.38	1.593 ± 0.003	0.6926 ± 0.0015
			1	11.03	1.604 ± 0.004	0.6975 ± 0.0017
	2nd	2nd	1	1.45	0.2164 ± 0.0011	0.721 ± 0.004
			0	1.81	0.2165 ± 0.0012	0.722 ± 0.004
	Pair	2nd	3	1.06	-1.395 ± 0.004	N/A
			0	2.80	-1.413 ± 0.003	N/A

Table 5.2: Response of the beam in position and tilt to the two correctors measured using the inductive BPMs.

the results are summarised in Table 5.3. Each measurement made using the screen is consistent with the repeated measurement. However, the values for the χ^2/DOF are very high. This is likely to be because of systematics in the screen monitor itself. Dead pixels would reduce the variation seen between repeat measurements and this variation is used to estimate the error. Local deformations in the screen material may also introduce a systematic to the position measurement at each separate corrector current setting. In any case, the overall trend appears linear and so the results of the fits are still used.

A summary of the measured responses is shown in Table 5.4 where the weighted average of each pair of measurements has been used with the associated error. The measurements in the horizontal direction are self-consistent for both types of monitor since combining the calibrations of the individual correctors using the ballistic model gives a prediction that is consistent with the calibration of the corrector pair. In the vertical direction, the pair calibration is close to the two individual results combined but not quite consistent. Furthermore, for the screen, the prediction for the response in beam angle in the vertical direction with two correctors is also zero to within one error. The difference between the results for the two types of monitor is less than 10 %.

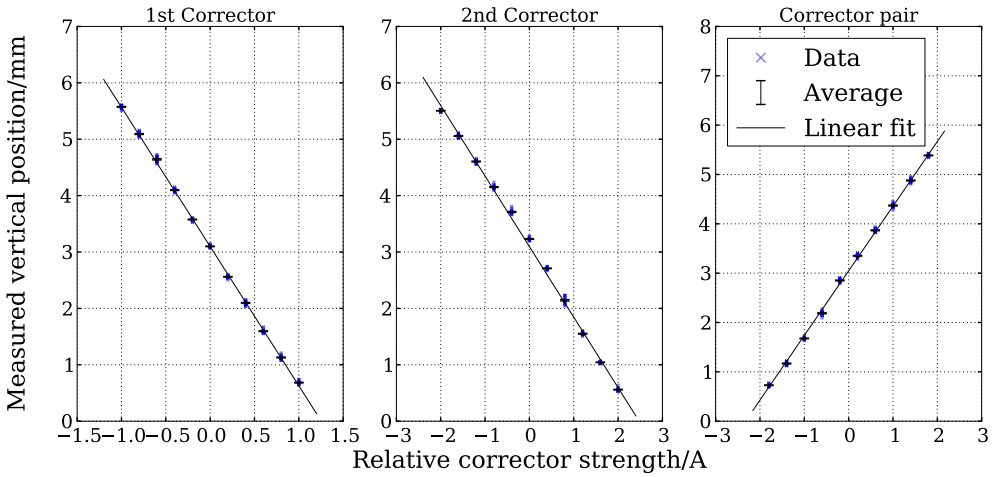

Figure 5.9: Vertical position as measured on the screen (CA.MTV0790) against current in the first corrector (CA.DVJ0710) (left), the second corrector (CA.DVJ0740) (middle) and the two acting as a pair (right).

The values used for the sensitivity measurements are based on those obtained using the screen as its calibration is likely have a smaller theoretical uncertainty. When the second corrector is used alone, the value used for the angular response in the horizontal direction is -0.64 mm m^{-1} A^{-1}. This must be multiplied by the distance from the corrector to the cavity BPM (3.034 m) and the change in corrector magnet current in order to obtain the resulting change in horizontal beam position at the cavity BPM. When the two correctors are used as a pair, 1.30 mm A^{-1} is used. It is assumed that in this case, there is no change to the beam trajectory angle and so this only needs to be multiplied by the corrector magnet current to obtain the position change at the cavity BPM location. The same two values are used for the vertical direction but with opposite sign to match the convention of the other BPMs in the CALIFES line.

Axis	Corrector	Cut pulses	$\frac{\chi^2}{\mathrm{DOF}}$	Position/ mm A^{-1}	Angle/ mm m^{-1} A^{-1}
Horizontal	1st	7	15.53	-2.5631 ± 0.0016	-0.6487 ± 0.0004
		6	26.06	-2.5709 ± 0.0017	-0.6507 ± 0.0004
	2nd	7	21.89	-1.2506 ± 0.0008	-0.6410 ± 0.0004
		7	21.58	-1.2428 ± 0.0010	-0.6370 ± 0.0005
	Pair	2	66.71	1.3057 ± 0.0009	N/A
		4	36.79	1.3212 ± 0.0007	N/A
Vertical	1st	1	31.53	-2.553 ± 0.003	-0.6461 ± 0.0007
		1	25.71	-2.469 ± 0.003	-0.6250 ± 0.0008
	2nd	2	51.04	-1.2348 ± 0.0016	-0.6329 ± 0.0008
		4	296.90	-1.2507 ± 0.0011	-0.6411 ± 0.0006
	Pair	4	21.13	1.2958 ± 0.0017	N/A
		2	55.71	1.3101 ± 0.0017	N/A

Table 5.3: Response of the beam in position and tilt to the two correctors as measured using the screen.

	Corrector	Position/mm A^{-1}		Angle/mm m^{-1} A^{-1}	
		Horizontal	Vertical	Horizontal	Vertical
2nd BPM	1st	-1.626 ± 0.002	1.599 ± 0.002	-0.7071 ± 0.0009	0.6951 ± 0.0009
	2nd	-0.2185 ± 0.0007	0.2165 ± 0.0007	-0.728 ± 0.002	0.722 ± 0.002
	Combined	1.408 ± 0.002	-1.382 ± 0.002	-0.021 ± 0.002	0.026 ± 0.002
	Pair	1.404 ± 0.002	-1.403 ± 0.002	0	0
Screen	1st	-2.5668 ± 0.0012	-2.5131 ± 0.0012	-0.6497 ± 0.0003	-0.6361 ± 0.0003
	2nd	-1.2475 ± 0.0006	-1.2411 ± 0.0006	-0.6394 ± 0.0003	-0.6361 ± 0.0003
	Combined	1.3193 ± 0.0013	1.2720 ± 0.0013	0.0102 ± 0.0004	-0.0001 ± 0.0004
	Pair	1.3156 ± 0.0006	1.3049 ± 0.0006	0	0

Table 5.4: Summary of measured responses to the first and second correctors with the individual calibrations combined and with both correctors measured together as a pair.

5.3 Raw Signal Characteristics

Figure 5.10 shows examples of beam excited signals measured at the digitiser after the analogue signal processing. It can be seen that the signals from short 2.1 ns beam pulses, have a finite rise time that is longer than the pulse length. For the diode signal, this is consistent with the time constant defined by the capacitance and load impedance of the installation [3]. The cause for the down-converted signals is less clear and could be from the electronics or the excitation of the cavity itself. The position profile of the long bunch train is clearly visible in the amplitude of the position channel signals while the reference cavity signals, both

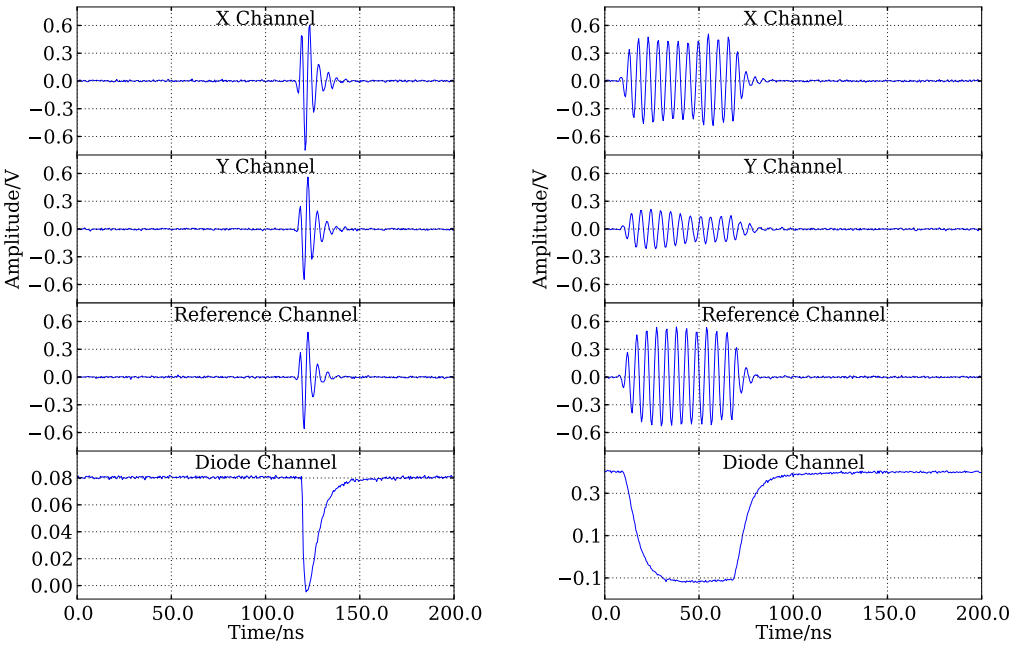

Figure 5.10: Examples of raw digitised signals from the cavity BPM prototype after the analogue signal processing for short beam pulses (left) and long beam pulses (right).

down-converted and diode rectified, show a clear exponential rise and decay.

5.3.1 Frequency and Decay Time

Because the beam induced signals are unique and not reproducible in the lab, their frequency and decay time were measured. Because of factors such as variation in bunch charge and position within a train or the response of the electronics, the recorded beam signals may not have the exact form discussed in Section 2.3 and expected from a cavity BPM pick-up. Therefore, several techniques were used to measure the frequency and decay time so that the best technique could be identified. The decay time measurements are based on fits, some of which, also give the frequency. The signal frequency can also be measured using a DFT but better precision can be achieved through methods such as numerical analysis of the fundamental frequencies (NAFF) [53] or phase-flattening [41].

The signal decay times were measured by fitting a decaying sine wave to the relevant part of the recorded signals. The start of the signal decay was taken from the signal maximum in the case of a short beam pulse and from the end of the signal flat-top in the case of a long beam pulse. The latter was located by taking the absolute signal level at the positive and negative peaks and then applying the same algorithm as for the inductive BPMs (Section 5.1.2.1) to find the timing of the end of the beam pulse. In both cases, the following function was fitted to the 30 samples after the start of the signal decay, which corresponds to 15 ns or about 3 signal decay times:

$$S(t) = Ae^{-\frac{t}{\tau}} \sin(2\pi f t + \phi) \tag{5.2}$$

where t is the time of the sample and the other symbols represent free parameters in the fit: A is the signal amplitude, τ is the decay time, f is the frequency of the recorded signal and ϕ is the phase. In the case of long beam pulses, fits were also made to the signal rise using a similar equation but with a charging exponential factor:

$$S(t) = A\left(1 - e^{-\frac{t}{\tau}}\right) \sin(2\pi f t + \phi)\,. \tag{5.3}$$

Examples of fits to raw signals are shown in Figure 5.11. In addition, the signals were down converted digitally using the algorithm described in Section 4.1.1 and fits were made to the amplitude envelope, also shown in Figure 5.11. The fitting was done using the python module, PyMinuit [54], an interface to the Minuit numerical function minimisation package [55]. 200 beam pulses were recorded, 100 short and 100 long and the variation in the fit results was used to estimate the error. The results are summarised in Table 5.5 and in histograms in Figure 5.12. Unsuccessful fits and fits where the sum of the squared fit residuals exceeded 5×10^7 (digitiser units)^2 and which therefore gave meaningless results, were ignored. Results from pulses where the digitiser was saturated were also ignored.

The resulting values are not consistent with those expected from the quality factor measurements in Table 2.7 because they are a characteristic of the full system, not just the pick-up. It is clear from Figure 5.11 that the signal decay is not a perfect exponential but includes additional features coming from the electronics. The decay time of the signal in the vertical channel from a long pulse

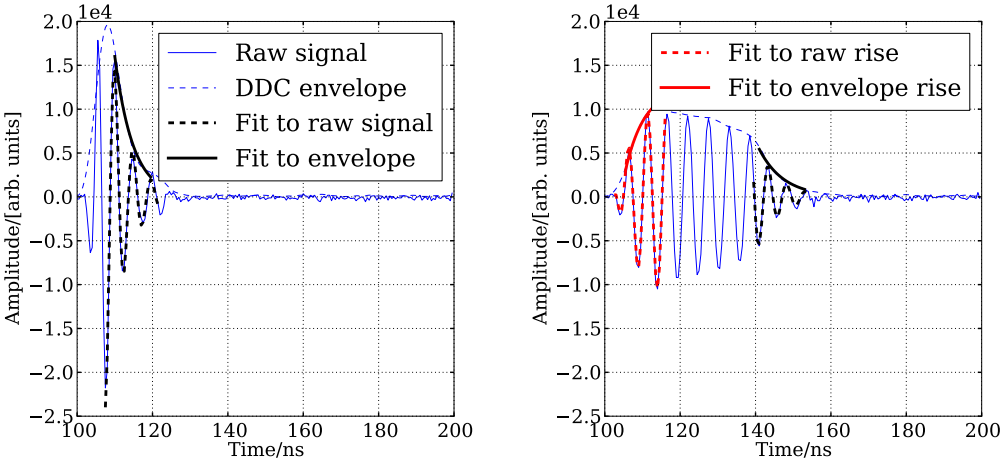

Figure 5.11: Raw signals for short (left) and long (right) beam pulses in the horizontal channel with the fits to the data also shown. The vertical scale corresponds to the 16 bit representation of the digitiser output.

Channel	Decay time/ns					
	Short Train		Long Train		Rise	
	Raw	DDC	Raw	DDC	Raw	DDC
X	4.77 ± 0.04	4.79 ± 0.03	5.68 ± 0.05	5.86 ± 0.04	3.49 ± 0.13	3.70 ± 0.12
Y	4.85 ± 0.03	5.266 ± 0.014	5.05 ± 0.04	5.00 ± 0.02	4.92 ± 0.15	5.36 ± 0.14
REF	4.41 ± 0.04	4.591 ± 0.013	4.82 ± 0.03	4.66 ± 0.02	5.4 ± 0.2	5.78 ± 0.19

Table 5.5: Results of signal decay time measurements for different pulse lengths from fits to raw digitised signals and their envelopes.

excitation is also inconsistent with the decay time measured from the short pulse excitation. This inconsistency would not be expected from a measurement of the cavity alone. The fits to the DDC amplitude are consistent with the fits to the raw amplitude in each case and have smaller statistical errors.

The same signals as used for the decay time measurement were used to determine the signal frequency. This was done in three different ways. The first is the numerical analysis of the fundamental frequencies [53], the second is phase flattening [41] and the third is to simply take the frequency f from the fits given by Equations 5.2 and 5.3. Each of these methods provides a higher precision than just a discrete Fourier transform alone.

A discrete Fourier transform (DFT) is a representation of the data in the

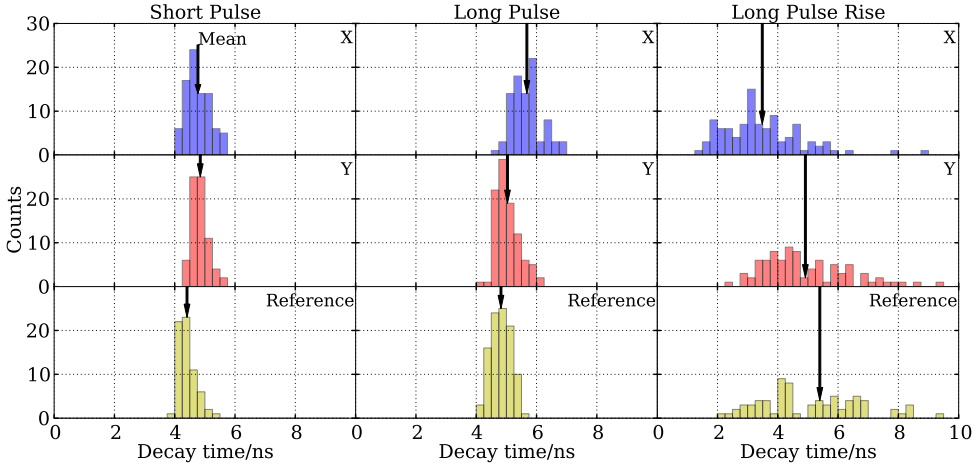

Figure 5.12: Distribution of the different decay time measurements as measured from fits to the raw digital signals. The downward arrow on each histogram points to the mean.

frequency domain and is calculated as

$$\tilde{x}[k] = \sum_{n=0}^{N-1} x[n] e^{i2\pi \frac{k}{N} n} \tag{5.4}$$

where $x[n]$ is the signal level at sample number n, N is the signal length in samples, and k is an integer between 0 and $N-1$. Each value of k/N is a frequency as a fraction of the digitiser sampling frequency f_s. The value of k for which $\tilde{x}[k]$ is at maximum amplitude corresponds to the dominant signal frequency. NAFF improves the precision of the DFT and is computed for any frequency around the peaks in the DFT as

$$\tilde{x}(f) = \sum_{n=0}^{N-1} x[n] e^{i2\pi \frac{f}{f_s} n} \tag{5.5}$$

where f is no longer constrained to being an integer [53]. For each beam pulse, Equation 5.5 is computed for 500 frequencies within half a DFT bin-width of the maximum in the DFT spectrum. The frequency for which $\tilde{x}(f)$ is at maximum amplitude is taken as the measured signal frequency.

The second method is known as phase flattening. Using a crude initial esti-

mate for the frequency based on knowledge of the system, the signal is digitally down-mixed and filtered using an 80 MHz bandwidth Gaussian filter as described in Section 4.1.1. The phases ϕ of several samples (10 or more depending on the pulse length) during the period of maximum signal level are then fitted linearly against sample number n and the change in frequency Δf that would set the gradient of the fit to zero is determined as

$$\Delta f = \frac{f_s}{2\pi} \frac{\Delta\phi}{\Delta n} \tag{5.6}$$

where $\Delta\phi/\Delta n$ is the gradient of the linear fit. In order to do this, the phase must first be unrolled so that there are no jumps between $-\pi$ and π that would invalidate the measurement [41]. In the case of a short pulse excitation, 10 samples of the phase starting from 5 samples after the signal maximum amplitude are chosen for the linear fit. In the case of the long pulse excitation, the start and end of the constant amplitude portion of the signal is determined from the signal rise and fall and the phase during this portion is fitted.

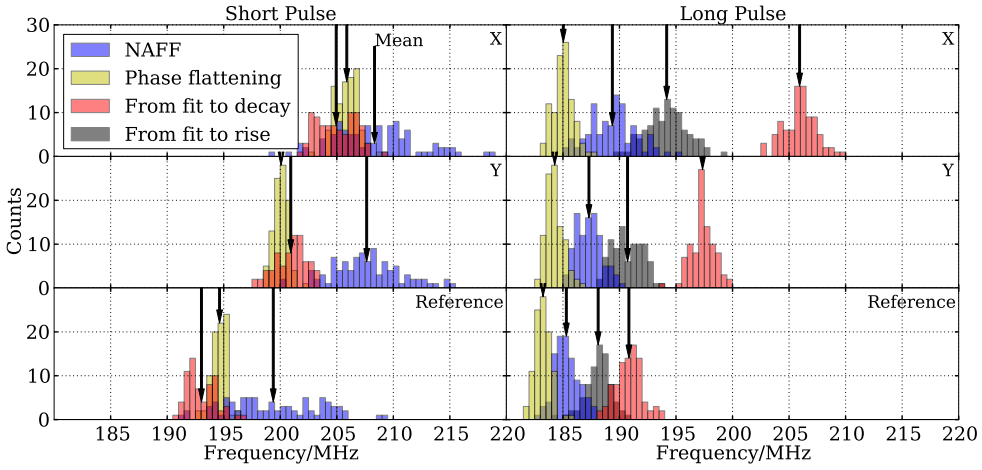

Figure 5.13: Distribution of the frequency as measured using different methods for short and long pulse trains. The arrows in each plot point to the mean in each distribution.

The results of the frequency measurements for short and long beam pulses are shown in Table 5.7 and Table 5.6 respectively and the distributions are shown in Figure 5.13. For the fit method, the same cuts based on the fit quality were

Channel	Frequency/MHz			
	NAFF	Phase flattening	Fit to decay	Fit to rise
X	189.4 ± 0.2	185.03 ± 0.09	205.93 ± 0.15	194.15 ± 0.18
Y	187.30 ± 0.11	184.27 ± 0.08	197.34 ± 0.10	190.69 ± 0.14
Reference	185.30 ± 0.11	183.24 ± 0.07	190.82 ± 0.11	188.11 ± 0.13

Table 5.6: Results summary of signal frequency measurements for four different methods applied to long beam pulses.

Channel	Frequency/MHz		
	NAFF	Phase flattening	Fit to decay
X	208.3 ± 0.4	205.89 ± 0.10	204.95 ± 0.18
Y	207.6 ± 0.3	200.07 ± 0.07	200.93 ± 0.16
Reference	199.4 ± 0.5	194.63 ± 0.07	193.03 ± 0.16

Table 5.7: Results summary of signal frequency measurements for three different methods applied to short beam pulses.

applied as were applied to the decay time measurements. For the other two frequency measurement methods, only saturated pulses were ignored. It is clear that for short pulses, the result using NAFF is less precise than the other two methods since it varied more pulse to pulse while phase flattening is the most precise. NAFF and phase flattening show better agreement in the case of long pulse excitations. In this case, there is also a clear difference in the frequency measured by the fit method from the others. This is because the fit method is applied only to the decaying and rising portions of the waveform. The signal frequency is expected to differ for these two portions of the waveform as discussed in Section 5.5.1.

The ultimate purpose of measuring the signal frequency is to obtain the frequency of a digital complex LO for a digital IQ mixer that can be used to demodulate the signal so that the phase around the sample point, just after the signal peak, is constant in time. The DDC amplitude then just represents the signal envelope. The phase flattening method therefore has the obvious advantage in that the method minimises the phase gradient of the signal around the peak which is exactly the desired outcome. The advantage of the NAFF method is that the measurement can be made without knowing anything about the form of the signal in terms of pulse length or signal arrival time. One of its disadvantages

is that it is sensitive to variations of the signal amplitude in time and so may not give the most accurate value for the signal frequency. The fit method serves as a comparison between the other two and is not intended for routine use because the fits sometimes fail or give anomalous results that in general, cannot be easily identified automatically and removed.

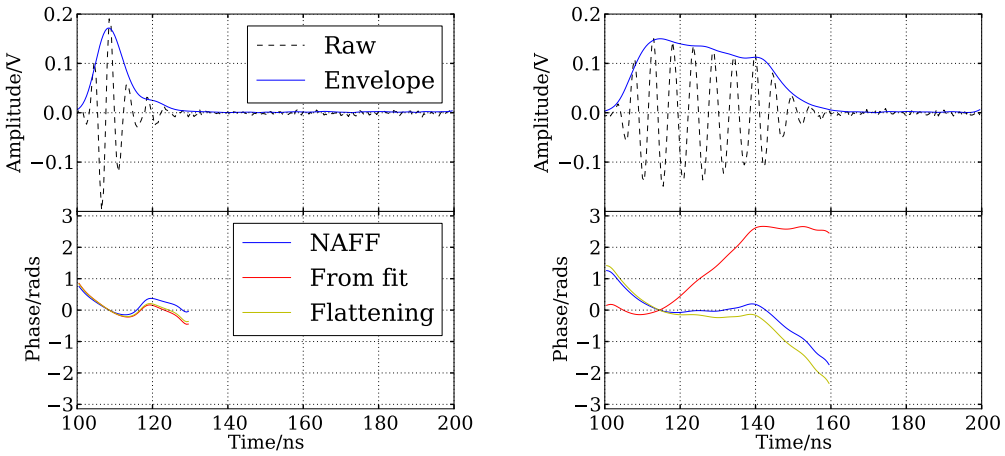

Figure 5.14: Digitally down-converted short (left) and long (right) waveforms from the horizontal channel where the frequency used in the algorithm has been determined using three different methods.

Figure 5.14 and Figure 5.15 show the effectiveness of the three methods on flattening the phase in the horizontal and reference channels respectively for both a short and long beam pulse excitation. The raw signals have been mixed down by multiplying by three digital LOs, each at the frequency measured using a different one of the three methods. The same Gaussian filter was then applied to each. It can be seen that in the horizontal channel with a short pulse excitation, the phase seems to vary nonlinearly and is therefore, not possible to flatten over a time period longer than several samples. This variation in phase over time is the same from pulse to pulse and is most likely, a feature of the electronics. It was also present in the results of the in-situ electronics measurements in Section 3.3.2. The reference channel, on the other hand, does show a near-constant phase for a period after the peak for the phase flattening and fit methods. The difference between the two channels in this respect is most likely down to the difference in

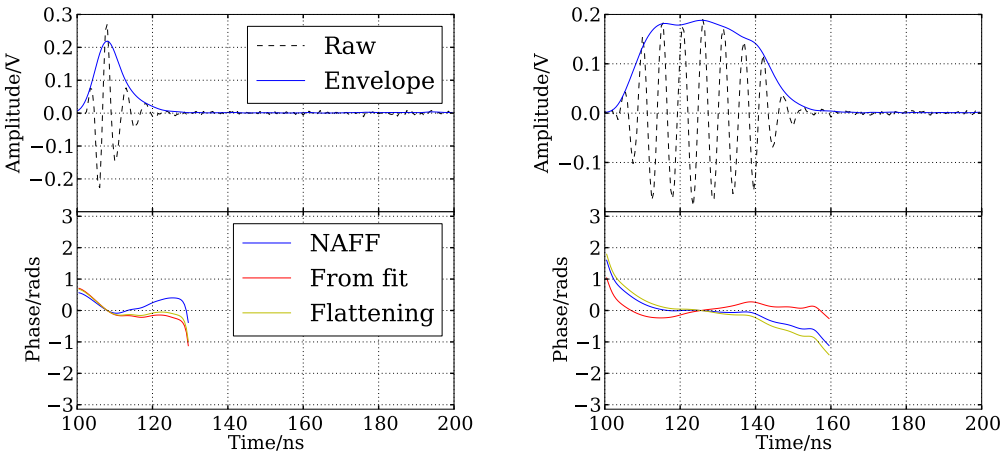

Figure 5.15: Digitally down-converted short (left) and long (right) waveforms from the reference channel where the frequency used in the algorithm has been determined using three different methods.

the quality factor of the two cavity modes.

For the long pulse excitation, both the NAFF and phase flattening methods are effective at producing a constant phase during the steady state region of the signal where as the fit method, which is applied to the decaying portion of the waveform, only produces a constant phase in this period. It is clear from the variation in phase that the frequency is different during the decaying portion of the waveform.

5.3.2 Local Oscillator Frequency

The LO signal supplied to the down-converter electronics is based on a independent voltage controlled oscillator and is not phase-locked to the main accelerator RF. This leads to greater variation in the phase of the BPM signals with the minor advantage that any systematic effects that are phase dependent are averaged out. It also allows the intermediate frequency to be varied by changing the control voltage. This could be used to further probe the signal characteristics. Figure 5.16 shows how the pulse shape and phase advance along the pulse is affected by the choice of intermediate frequency after digital down conversion, for which the frequency has been measured using the phase flattening method. Both

the signal amplitude and phase advance, whose behaviour does not change much from pulse to pulse, is affected, especially after the signal rise. This proves that the electronics have an effect on the signal shape and decay time measurements since the change is introduced after the cavity output. The changes can be more specifically attributed to the intermediate frequency section of the electronics since the frequency in the microwave front end remains the same. Similar features in amplitude to those seen in the signal decay were seen in the time domain measurements of the electronics in Section 3.2.4.

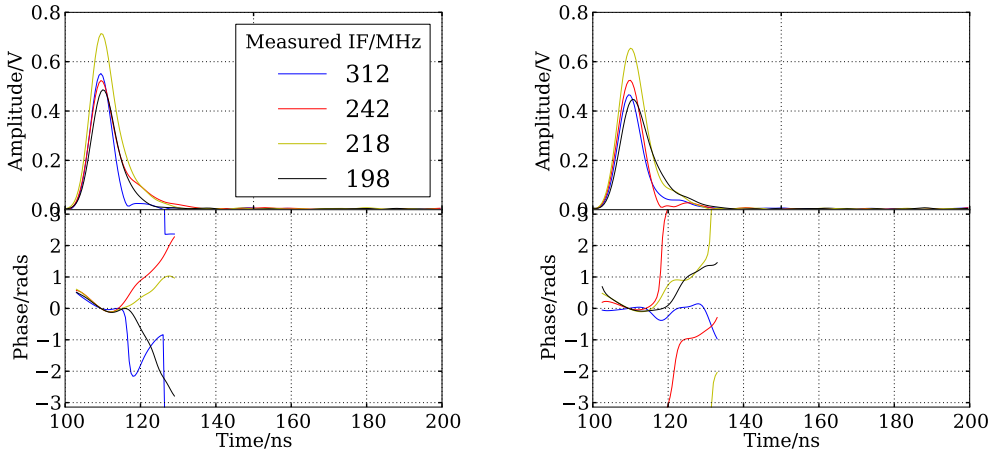

Figure 5.16: DDC amplitude and phase of signals from the reference channel (left) and the horizontal channel (right) for different intermediate frequencies.

Figure 5.17 shows how the intermediate frequency, as measured using the phase flattening method, changes with the oscillator control voltage for both long and short pulses. The long pulses show less erratic behaviour. This is expected since the pulse length is long compared to the settling time of the electronics and so the frequency can be measured where the amplitude of the signal is constant. The relationship between the output frequency of the VCO and the control voltage is not expected to be linear over large ranges but appears to be for this small range. A fit was therefore made to the reference channel frequency since this is the most reliable as discussed in Section 5.5.1. The change in frequency was found to be -147.41 ± 0.15 MHz V^{-1} and for long pulses, an intermediate frequency of 200 MHz would be expected for a control voltage of

4.86 V.

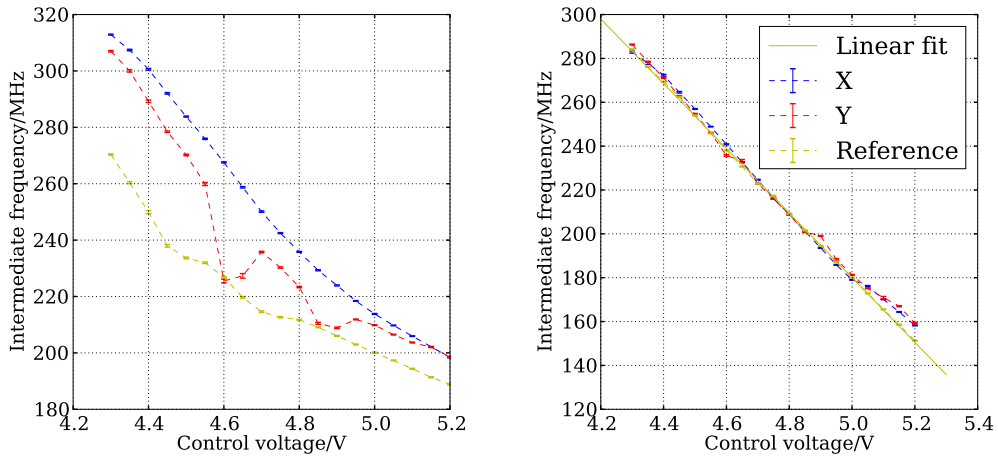

Figure 5.17: Frequency as measured using the phase flattening method for short beam pulses (left) and long beam pulses (right).

5.3.3 Saturation Detection

Figure 5.18 shows a pulse that has saturated the digitiser. One of the samples has driven the analogue to digital converter (ADC) to it's maximum value which always appears negative. Unfortunately, because of how the digitiser interacts with the front end controls software, which casts the digitiser output to a signed 16 bit short, it is not always clear what this value will be so it is not possible to simply search for occurrences of this value to detect saturated pulses. However, the effect of these repeated negative values on the signal frequency spectrum is clear. It can be seen in the spectra in Figure 5.18 that for the saturated pulse, as well as the clearly defined signal peak, there is a large background throughout the rest of the spectrum that is not there for an unsaturated pulse. Therefore, the signal is said to be saturated if the median amplitude of the frequency spectrum, obtained by a DFT of the raw digitised signal, is greater than 2^{15}. From experience, this method is very efficient at identifying saturated pulses but it does not provide any information about where the signal saturation ends. This would be required if any useful information were to be obtained from saturated signals.

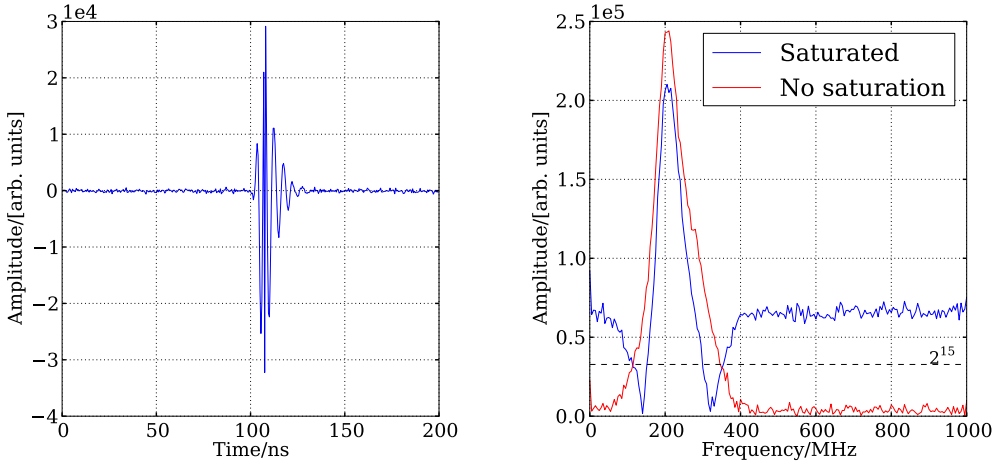

Figure 5.18: Example of a saturated beam pulse signal (right) and its discrete Fourier transform (right) compared with the transform of an unsaturated pulse.

5.3.4 Arrival Time Measurement

The start time of the cavity BPM signal within the digitiser sampling window will change when users vary the timing of the photo-injector laser and if the trigger drifts in time relative to the machine RF. It is also dependent on the phase of the digitiser clock relative to the trigger. Therefore, rather than using the same samples in the window for the measurement of each pulse, a timing measurement should be made of each pulse individually and the samples used for the beam measurement chosen relative to that.

The timing of the beam signal within the digitiser window is measured using the diode rectified signal from the reference cavity. The signals from the down-converter channels are not as suitable because their rise has a noticeable dependence on their phase relative to the analogue LO, which is not locked to the main accelerator RF. The method used for the measurement must be more precise than the beam arrival time jitter (or the clock jitter, whichever is larger) so that the variation in the beam measurements is reduced and not increased. The result of the timing measurement does not have to correspond to any specific part of the waveform but must provide a consistent reference time from which the signal can be sampled at a fixed offset.

Two methods have been tested. The first involves finding the largest difference

between two adjacent samples of the diode rectified signal. This is always on the signal rise. The arrival time is then calculated as the time at which a line going through these two points crosses the axis where the signal level is zero [41]. In order to do this, the pedestal must first be subtracted so that the zero corresponds to where there is no signal. The pedestal value is determined from the mean sample amplitude during a portion of the waveform, about 100 samples long, where there is no signal.

The second method involves taking an average of the time of each sample that is weighted by the sample amplitude. This can be written as

$$t_0 = \frac{\sum_i^N (t_i V_i^k)}{\sum_i^N V_i^k} \tag{5.7}$$

where t_i is the time of sample number i, V_i is the sample amplitude, N is the number of samples in the digitiser window and k is an odd integer. Again, the pedestal is subtracted beforehand so that the samples where there is signal have the largest weights.

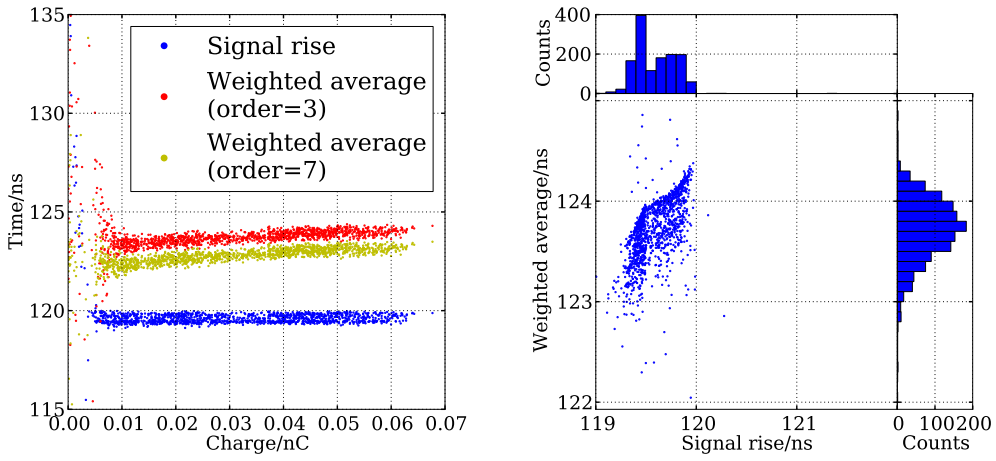

Figure 5.19: Results of beam arrival time measurements against the beam charge (left) and a correlation of two different methods (right) for short beam pulses.

Figure 5.19 shows the measurement of the arrival time of short beam pulses of different charge using the two methods and with two orders (values for k in Equation 5.7) for the second method. It can be seen in both plots that the first

method shows less statistical variation. However, it is affected by some systematic because of which, the measurements are not distributed normally but in some two peak distribution. This is certainly due to the measurement method because the digitiser clock is not locked to any machine timing and this feature is not seen using the other method, even at high charge. The second method shows some charge dependence, particularly at low charge. This is likely to be because the response of the diode rectifier becomes nonlinear for charges below 50 pC and the short beam pulses used here are below this charge. Therefore, the first method was chosen for short pulses since low charges would be used.

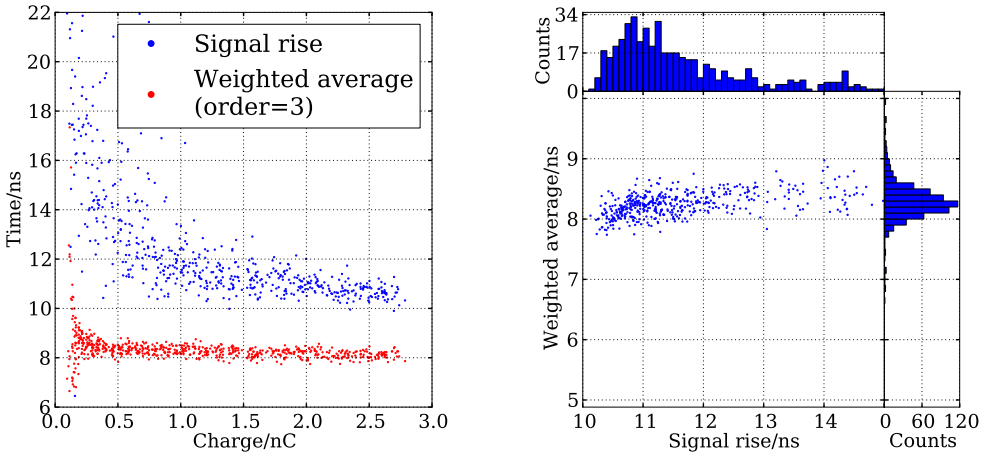

Figure 5.20: Results of beam arrival time measurements against the beam charge (left) and a correlation of two different methods (right) for long beam pulses.

The two methods were also applied to long beam pulses. Figure 5.20 shows the results for a 30 ns pulse length. The performance now appears to be worse for the first method, which shows large statistical variation, especially at low charges, and no correlation with the weighted average method that shows less variation for all charges. The weighted average method was therefore chosen for the timing measurements of long beam pulses.

5.4 Beam Based Electronics Measurements

The output signals from the cavity modes excited by the beam are not easily reproducible in the laboratory. Therefore, measurements were made in order to determine the behaviour of the electronics with beam-induced signals at the input. Unlike for the electronics measurements in the laboratory, the input power was not known because it would have to be measured and the hardware required to do this was not available. One way in which the input power could be measured is to install a coupler upstream of the electronics and measure the coupled power using a diode detector. Uncertainties in the response of the diode, the strength of the coupling and the cable losses from the output of the diode would lead to systematic errors in the measured input power of the electronics.

5.4.1 Remotely Controlled Attenuation

The effect of the attenuator during beam operation was tested by simply varying the attenuation in steps of 0.5 dB while keeping the beam conditions constant. Short beam pulses were used and digital down conversion was applied with a 80 MHz Gaussian filter bandwidth. The peak signal was then measured for each channel. 20 pulses were recorded for each attenuation setting so that the mean

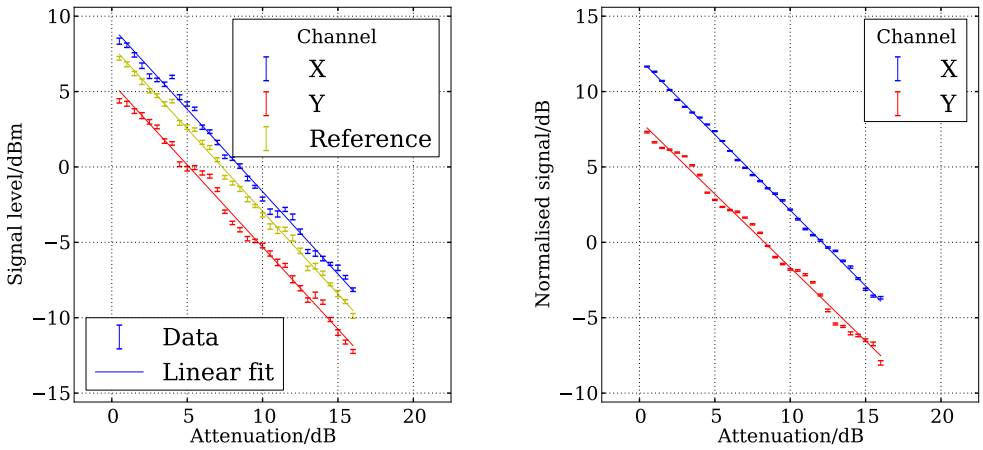

Figure 5.21: Signal level against the setting of the variable attenuator for all three channels (left) and charge normalised signals (right).

and standard error could be calculated. The results are shown in Figure 5.21. The signals are not CW but, so that linear fits could be performed, the power in dBm was estimated by dividing the peak voltage by $\sqrt{2}$. The gradients of the linear fits are in Table 5.8.

Pulse length/ns	Channel	Normalised	χ^2/DOF	Absolute/ dBm dB^{-1}	χ^2/DOF
2.0	X	-1.0034 ± 0.0016	17.4	-1.092 ± 0.012	13.6
	Y	-0.974 ± 0.002	74.2	-1.090 ± 0.011	17.3
2.7	X	-	-	-1.049 ± 0.010	6.0
	Y	-	-	-1.069 ± 0.011	8.4
	Reference	-	-	-1.004 ± 0.016	3.9

Table 5.8: Gradients of linear fits to the absolute signal level against attenuation setting and to the normalised signal level when the reference cavity attenuation was kept constant.

The test was first performed with the attenuation on the reference channel kept constant at a nominal value of 0 dB and with a pulse length of 2 ns. This allowed the charge variation during the scan to be removed by normalising the horizontal (X) and vertical (Y) signals by the reference signal. Splitting the signal from a single pick-up output port into all three electronics channels and then performing the test in the same way would have better removed the effects of changes in the beam position and charge but this was not done due to limited access to the accelerator tunnel during the beam tests.

The slow orbit position drift is still visible and so the values of χ^2 per degree of freedom (DOF) are high. It can clearly be seen, however, that the overall trend is linear. The gradients from all the fits are close to the expected value of -1 but only the normalised result from the horizontal channel is within three standard errors. The result for the vertical channel is consistent with the measurements made of the horizontal channel using the signal generator in Section 3.3.1. The phase relative to the reference cavity signal was also recorded for different attenuation settings (Figure 5.22). For both position channels, the change in phase is of the same magnitude and follows the same pattern as the measurement of the variable attenuator itself in Section 3.2.1.4. The phase change is in the opposite sense because the digital IQ mixer converts the negative frequency component to baseband.

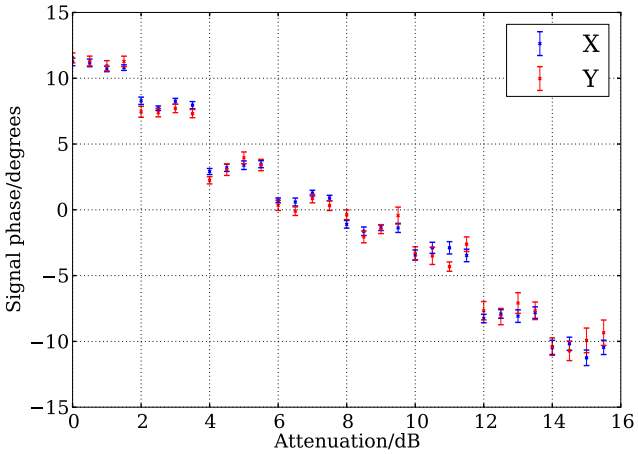

Figure 5.22: Signal phase against setting of the variable attenuator for the two position channels with the reference signal used as a reference phase.

The test was then repeated with a pulse length of 2.7 ns which corresponds to one additional bunch and therefore, a higher beam intensity. The attenuation was then scanned on all three channels. The higher intensity reduces the overall variation in charge and so the values of χ^2/DOF are lower and in this case, only the gradient from the reference channel, which is not affected by beam position jitter, is within three standard errors of the expected result of -1. Because the behaviour of the attenuator is close to ideal for all three channels, it is acceptable to convert between measurements made at different attenuation settings where necessary. The calibration scale factors should also be independent of the attenuation setting as long as the setting is the same for the reference and position channels.

5.4.2 Electronics Saturation

The electronics saturation was also tested by varying the attenuation and measuring the peak signal. This time, a 12 ns pulse length was used followed by a 20 ns pulse length. The results are shown in Figure 5.23. A measurement of the output power at 1 dB compression was also attempted using the signals from the 12 ns pulse. This was done by first determining the difference between the measured signal level and the nominal attenuation for the five highest attenua-

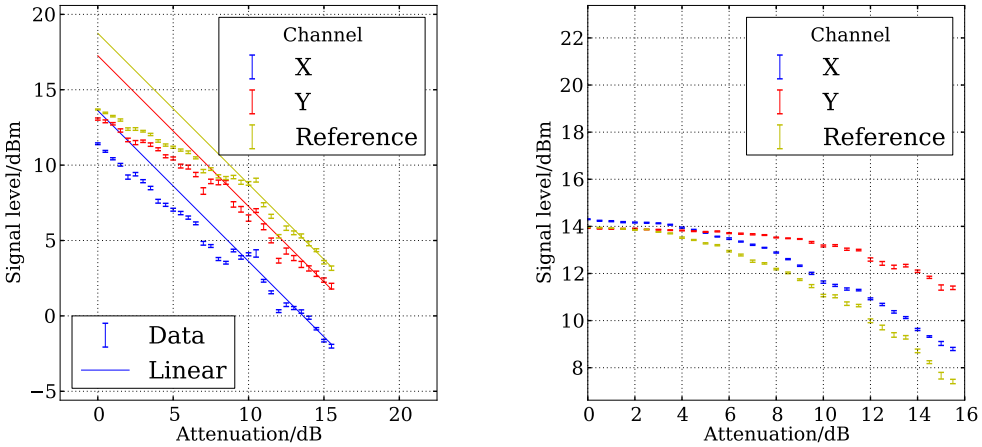

Figure 5.23: Measurement of the electronics saturation by varying the signal attenuation with 12 ns (left) and 20 ns (right) pulse lengths: the linear fits in the left plot have been used to measure the 1 dB compression point.

tion settings. This was used to produce the line on the graph with gradient -1. Measured signals at lower attenuation were than subtracted from points along this line until 3 consecutive points were found to be more than 1 dB below. The first of these three points with its error was then used as the estimated value for the output at 1 dB compression. The results are shown in Table 5.9. The results

Channel	Output Power for 1 dB Compression/dBm	χ^2/DOF
X	3.53 ± 0.12	4.96
Y	8.28 ± 0.12	1.33
Reference	9.14 ± 0.11	1.16

Table 5.9: Results of the compression point measurements using beam data for the three channels.

for the reference and vertical channels are close to each other and are close to the result of +13 dBm measured in Section 3.2.3.1 after the cable loss at IF of -2.6 dB has been accounted for. However, the systematic error due to drifts in beam position and charge is clearly large. The result for the horizontal channel is low but this is probably due to a large horizontal position shift since the compression is not as clear as for the other two channels. The value corresponding to

a continuous-wave signal of ±1 V, 10 dBm, is above the 1 dB compression point. Therefore, care must be taken when using the digitiser at its ±1 V setting.

The results for the 20 ns pulse length could not be used for a measurement of the electronics compression because there are no points inside the linear range of the electronics. It can be seen from the right plot of Figure 5.23, however, that at complete saturation, the output power of all three channels is the same. This confirms that the three down-converter channels are behaving similarly in terms of their operating input power range.

5.5 Cavity BPM Response to Beam

The main goal of the beam tests was to measure the response of the cavity BPM pick-up to the beam to ensure that it is working as predicted and that it will be possible to achieve the desired performance without significant alterations to the design. For this purpose, the sensitivity of the reference cavity signal to beam charge and the sensitivity of the position cavity signal to beam charge and position were measured. The response of the pick-up signals to pulse length was also measured for a comparison with the predictions in Section 2.3.

5.5.1 Pulse Length

The beam pulse length was varied in order to further investigate the predictions made in Section 2.3 about the signals from multiple bunches. This could be achieved by changing the timings of the pulse picker of the photo-injector laser. The relative timings were changed in steps of 2 ns and 20 pulses were recorded at each laser pulse length. The digitally down converted signal for each channel was sampled at its maximum level in the first 30 ns after the signal rise. In this time window, the signal is able to reach an essentially steady state and the measurement is not affected by position or charge variation towards the end of the train.

The results of the pulse length scan are shown in Figure 5.24. The form of the curve is as expected from Equation 2.21. Fits were therefore made to determine the decay time τ that describes the rise in the signals from the three channels.

Channel	X	Y	Reference
Decay time from pulse length scan/ns	3.49 ± 0.11	3.5 ± 0.3	6.13 ± 0.18

Table 5.10: Decay time measurements from the rise of the signal level against pulse length.

The values obtained are shown in Table 5.10. The result for the reference cavity is longer than the measurements summarised in Table 5.5 because of the shaping of the signal by the analogue processing. However, the information is still useful for converting between measurements at different pulse lengths when the reference cavity is used to measure the beam charge. The decay time results for the position cavity signals are much shorter than expected. This may be due to differences in the cavity resonant frequency from 15 GHz, the closest harmonic of the bunch arrival frequency. This would lead to a pulse shape with a sharper rise and then beating at the frequency offset. The sharper rise would result in a shorter result for the measured decay time but the signal amplitude at its steady state would not be reduced by the corresponding amount. Because the decay time and frequency both have a similar effect on the fitted curve, it is impossible to perform a fit in which both are free parameters. The small value for the decay time of the position channels could also be due to position variation at the beginning of the bunch train.

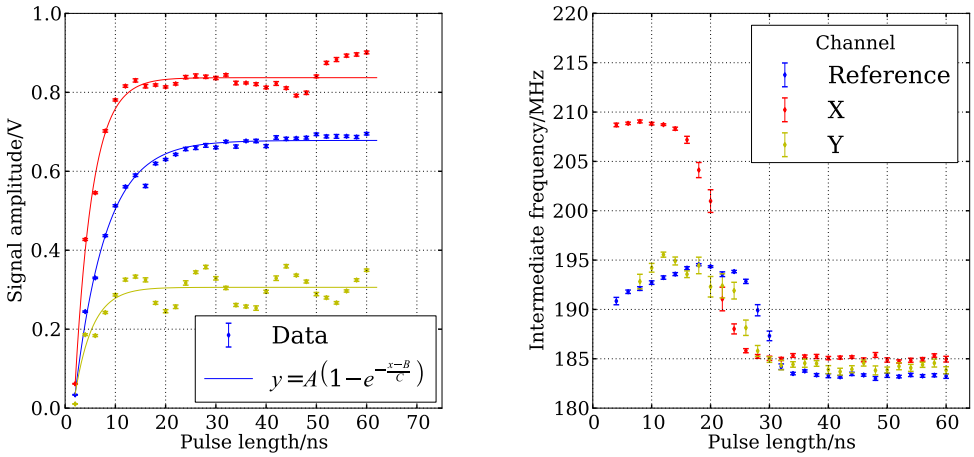

Figure 5.24: Signal amplitude (left) and measured frequency (right) for different beam pulse lengths.

Equation 2.17 suggests that for long pulse trains, the signal becomes periodic with the bunch arrival frequency. This was tested by measuring the frequency of the signals of different pulse lengths using the phase flattening method. A different data set from the amplitude measurement was used where there was less position variation along the trains. It can be seen from Figure 5.24 that for short pulses, since the resonant frequency is defined by the geometry of the cavity modes, the measured frequency is different for all three channels. For long pulses, however, the signal becomes periodic with the bunch arrival frequency and the dominant signal frequency is the bunch arrival harmonic that is closest to the cavity mode resonant frequency. The measured frequency is therefore the same for all three channels.

5.5.2 Charge Sensitivity

The sensitivity of the reference cavity signal to the bunch charge was measured by varying the attenuation of the photo-injector laser and measuring the charge with an integrating current transformer (ICT). The BPM signal was digitally down converted, this time with a narrower 45 MHz bandwidth for the Gaussian filter so that variation in charge along the train was averaged out without affecting the

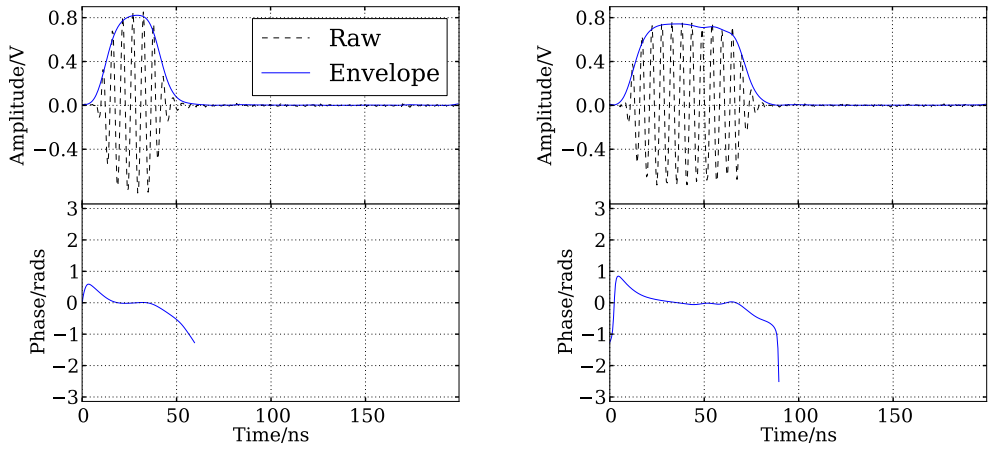

Figure 5.25: Example of a digitally down-converted waveform where the Gaussian filter bandwidth is 45 MHz instead of the 80 MHz used previously.

maximum signal level. Results were taken with two different pulse lengths: 30 ns (about 45 bunches) and 60 ns (about 90 bunches) and the variable attenuator was at its maximum setting, 15.5 dB. Long pulses were chosen in order to get a high charge and a large signal to noise ratio for the ICT. The maximum attenuation was necessary to match the operating range of the ICT to that of the cavity BPM.

Examples of signals processed with the narrower digital filter bandwidth are shown in Figure 5.25. The reference cavity signal for each pulse was sampled at its maximum amplitude and 20 pulses were taken for each charge setting. The ICT measurement is of the total charge of the whole bunch train. This was divided by the number of bunches so that the results for the two different pulse lengths are comparable. A weighted linear least-squares fit was then applied to obtain the sensitivity as shown in Figure 5.26. A nonzero intercept can be seen that is different for the two train lengths. This is due to an offset that is dependent on the ICT gain setting and does not affect the final value of the measured sensitivity.

As well as sampling the peak signal, the total energy E in the digitised waveforms was calculated using the equation

$$E = \frac{\Delta t}{Z} \sum_{i=0}^{N} V_i^2 \tag{5.8}$$

where Z is the line impedance, Δt is the time between samples, N is the total number of samples and V_i is the signal level measured for sample i. A quadratic fit was performed to obtain the energy per square unit of bunch charge.

The results of the fits are displayed in Table 5.11. The sensitivity due to a single bunch excitation can be obtained by dividing by the multiple bunch gain factor given by Equation 2.23, assuming a mode resonant frequency of exactly 15 GHz. This was done for two different decay times: 2.8 ns, which corresponds to the measured loaded quality factor of the reference cavity in Table 2.7 and 6.1 ns which was obtained from the fit of signal amplitude against pulse length in Section 5.5.1. The former may be closer to the physical value because, as discussed in Section 5.3, it is likely that the analogue processing electronics have an effect on the signal shape. The single bunch sensitivity estimated for $\tau =$ 6.1 ns, 63.8 V nC^{-1}, is what would be measured if the test were performed with

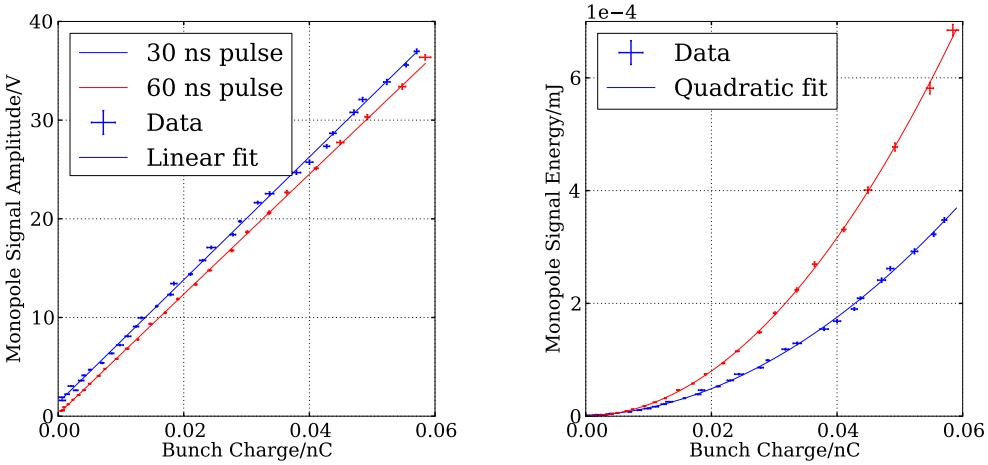

Figure 5.26: Maximum amplitude of the reference cavity monopole mode signal for different values of the bunch charge as measured by the ICT (left) and the total energy in the waveform (right). The vertical scale shows the value expected at the output of the cavity after the cable loss and electronics gain have been accounted for.

Pulse length/ns		30	60
Long pulse sensitivity/V nC^{-1}		622 ± 3	608 ± 2
χ^2/DOF of linear fit		1.39	0.65
Single bunch sensitivity/V nC^{-1}	$\tau = 2.8$ ns	131.9 ± 0.6	128.8 ± 0.5
	$\tau = 6.1$ ns	64.5 ± 0.3	63.0 ± 0.2
Energy gradient/mJ nC^{-2}		0.1115 ± 0.0006	0.1989 ± 0.0012

Table 5.11: Results of the reference cavity sensitivity measurements in terms of peak signal amplitude and total energy along with calculations of the estimated single bunch response.

a short bunch train. The sensitivity predicted from measurements of the external quality factor and the normalised shunt impedance as described in Section 2.2.7 is 165 V nC^{-1}. As expected, this is closer to the sensitivity estimated for the 2.8 ns decay time, 130.4 V nC^{-1}. The predicted energy outputs, using Equation 2.31, are 0.129 mJ nC^{-2} for the 30 ns pulse length and 0.271 mJ nC^{-2} for the 60 ns pulse length. These values are both larger than the measured values of 0.1115 ± 0.0006 mJ nC^{-2} and 0.1989 ± 0.0012 mJ nC^{-2} by 16 % and 36 % respectively.

In order to further check the validity of the ICT measurement, an independent measurement of the charge was made using the two inductive BPMs closest to the

Pulse length/ns	CA.BPM0720	χ^2/DOF	CA.BPM0750	χ^2/DOF
30	0.929 ± 0.018	0.8	0.787 ± 0.014	1.6
60	0.909 ± 0.011	1.2	0.763 ± 0.009	0.8

Table 5.12: Results of the fits to the charge as measured by the inductive BPMs against the charge as measured by the ICT.

cavity BPM. For each, the charge was determined by integrating the pedestal-subtracted current signal. This was fitted against the ICT measurement as shown in Figure 5.27. The results of the fits are summarised in Table 5.12. The gradient is not exactly 1 because there is beam loss between the ICT, which is at the beginning of the CALIFES line, and the two BPMs which come after the two-beam test stand. The gradients therefore, give an estimate for the fraction of the beam charge that is transmitted down the beamline and may be used to give a more accurate estimate for the charge at the cavity BPM.

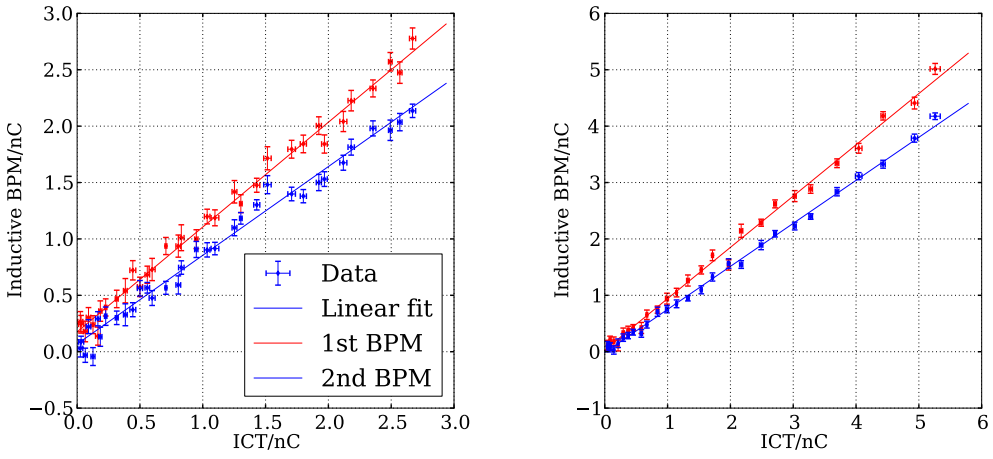

Figure 5.27: The beam pulse charge as measured by the inductive BPMs against the charge as measured by the ICT for the 30 ns pulse length (left) and the 60 ns pulse length (right).

Charge scan data was taken on other days but was not as successful because of poor beam charge stability. Another measurement was taken where the charge did appear to be stable but the ICT measurement did not agree with the monopole signals of the reentrant cavity BPMs, which may be due to a malfunctioning gain setting on the ICT. This result was therefore also discarded. All the results agree

with the same measurement using the rectifying diode when it was available. Therefore, any discrepancies are down to factors external to the cavity BPM pick-up and down-converter electronics.

5.5.3 Position Sensitivity

The sensitivity of the position cavity signal to the beam position was measured by varying the position of the beam within the cavity pick-up using the dipole corrector magnets. The beam was focused in the cavity BPM using the first quadrupole triplet in the CALIFES beamline so that the results were not affected by the model uncertainties discussed in Section 5.2.1. The test was performed in both vertical and horizontal directions, first using the second corrector only, and then using both correctors as a pair to shift the beam position without affecting its trajectory angle. The beam pulse length used was 2.1 ns since this produces a response that is similar to a single bunch excitation while maintaining good charge stability. In the worst possible case, 2.1 ns corresponds to three bunches of equal charge. It can be shown using Equation 2.21 that in this case, for decay times as short as 4 ns, the maximum signal amplitude is no less than 85 % of the amplitude expected from a single bunch excitation.

In order to ensure that the signal amplitude was always sampled at the peak, even for low amplitude pulses where the peak may be difficult to distinguish from the background noise, the signal arrival time was measured from the rise of the diode rectified monopole mode signal as described in Section 5.3.4. The timing offset from the diode signal rise to the peak of the dipole mode signals was determined beforehand. The sampled peak amplitude was normalised by the pulse charge to account for fluctuations in beam intensity. The larger the beam offset, the more these fluctuations would show up in the position cavity signals.

Since the beam pulse charge was too small for the ICT to measure and since the ICT does not provide a pulse to pulse measurement, the reference cavity had to be used to measure the charge. The charge calibration of the reference cavity signal was determined from the charge sensitivity data in Section 5.5.2. The measured amplitude as a function of pulse length, Figure 5.24, was used to convert the calibration to the right value for single bunch excitation. The

sensitivity that corresponds to the calibration scale used is 64.5 V nC^{-1}, the result for reference cavity decay time $\tau = 6.1$ ns in Table 5.11. 20 pulses were averaged for each position step.

When the beam is at the cavity centre, there is a nonzero minimum signal coming from the beam angle and pollution from monopole cavity modes. It is therefore not possible to simply reverse the sign when the beam crosses the centre of the cavity and perform a linear fit to determine the sensitivity. Instead, a three parameter fit of the form

$$\frac{\hat{V}(x)}{q} = S_x|x - x_m| + C \tag{5.9}$$

was used. Here, $\hat{V}$ is the peak voltage as a function of beam position x, and q is the beam pulse charge. The three fit parameters are the position sensitivity S_x, the position that corresponds to the electrical centre of the position cavity x_m and the signal level measured at this point C. The sensitivity can also be determined from the quadratic fit to the peak output power $\hat{P}$ since

$$\hat{P} = \frac{\hat{V}^2}{Z_\mathrm{L}} = \frac{[S_x q(x - x_m)]^2}{Z_\mathrm{L}} \tag{5.10}$$

where Z_L is the load impedance, 50 Ω in this case. Examples of the two fits are shown in Figure 5.28 and Figure 5.29 for the horizontal and vertical channels respectively.

As with the charge sensitivity measurement, the waveforms were integrated using Equation 5.8 to obtain the total energy detected by the digitiser. Example fits to the output energy in both horizontal and vertical directions are shown in Figure 5.30.

Several beam position scans were performed on the same day. The results for each individual scan can be found in Section A.3 of the Appendix. The average sensitivity results are given in Table 5.13 for the signal amplitude and Table 5.14 for the signal energy. The gain of the system, as measured from the cavity output to the input of the digitiser, as summarised in Section 3.3.1 was used to scale the fit results to the output of the cavity pick-up.

The sensitivity measurements made using the linear fits are consistent with

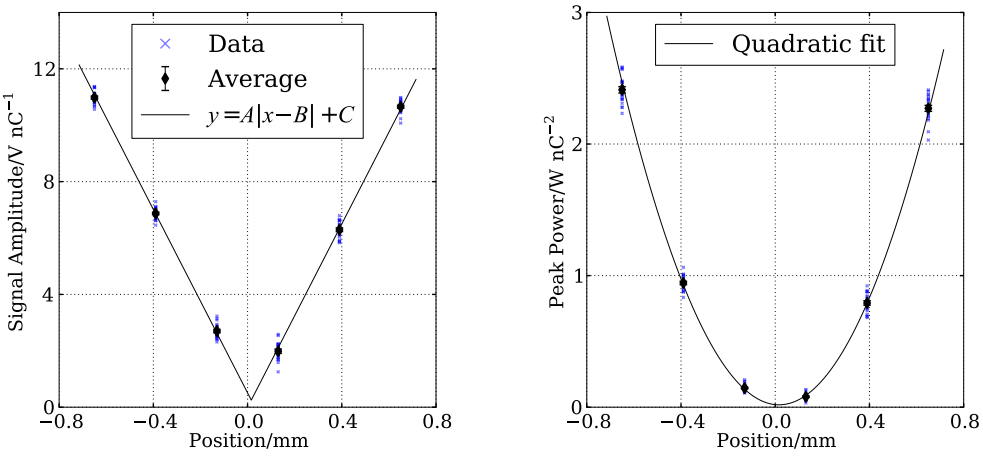

Figure 5.28: Signal amplitude (left) and peak power (right) at the output of the cavity pick-up for different relative beam positions in the horizontal direction.

	Quantity	Sensitivity/V nC^{-1} mm^{-1}		
		X	Y	Long pulses (X)
Single Corrector	Amplitude	16.74 ± 0.18	15.5 ± 0.5	133.7 ± 0.7
	Peak Power	16.82 ± 0.09	16.5 ± 0.3	130.6 ± 0.5
Corrector pair	Amplitude	16.49 ± 0.10	15.9 ± 0.4	128.7 ± 0.5
	Peak Power	16.74 ± 0.06	16.4 ± 0.6	128.9 ± 0.6
Average	Amplitude	16.61 ± 0.11	15.6 ± 0.3	131.2 ± 0.8
	Peak Power	16.78 ± 0.05	16.5 ± 0.3	129.8 ± 0.4

Table 5.13: Average results for the fits of maximum signal amplitude against beam position.

the quadratic fits, which is expected since they are based on the same data. The measurements made using one corrector are also consistent with the measurements made using the corrector pair. This shows that the ballistic model used to track the beam between the correctors and the cavity BPM is valid and that the longitudinal positions of the different correctors and beam monitors are correct. Added to the measurements of the corrector responses in Section 5.2.5, the ballistic model has been used to produce consistent results for three separate measurements using three different monitors. The average measured sensitivity of 16.4 V nC^{-1} mm^{-1} is 70 % of the theoretical prediction of 24.2 V nC^{-1} mm^{-1} (based on the measured RF characteristics of the cavity and the shunt impedance determined from simulation as described in Section 2.2.7). This level of agreement

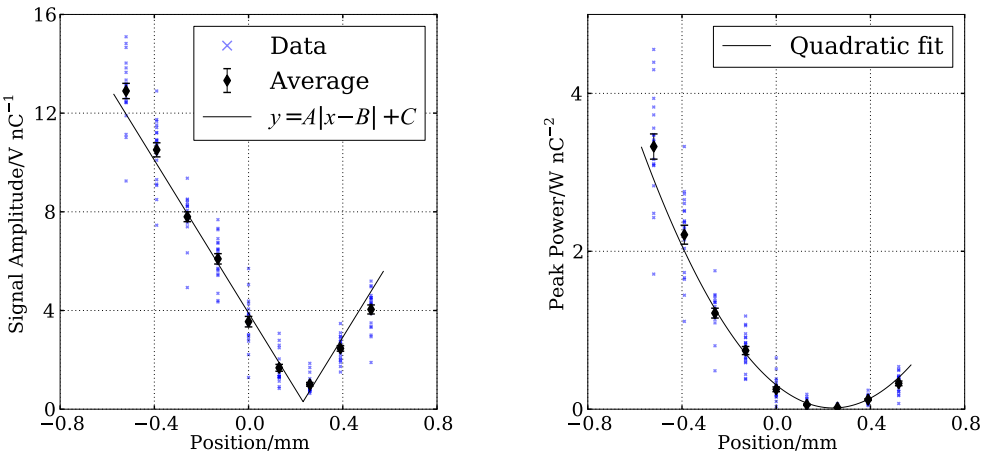

Figure 5.29: Signal amplitude (left) and peak power (right) at the output of the cavity pick-up for different relative beam positions in the vertical direction.

Axis	Energy gradient/mJ nC^{-2} mm^{-2}		
	Single	Pair	Average
X	$(2.36 \pm 0.02) \times 10^{-5}$	$(2.332 \pm 0.016) \times 10^{-5}$	$(2.345 \pm 0.014) \times 10^{-5}$
Y	$(2.40 \pm 0.09) \times 10^{-5}$	$(2.39 \pm 0.16) \times 10^{-5}$	$(2.39 \pm 0.07) \times 10^{-5}$
Multibunch (X)	$(9.85 \pm 0.09) \times 10^{-3}$	$(9.72 \pm 0.09) \times 10^{-3}$	$(9.79 \pm 0.06) \times 10^{-3}$

Table 5.14: Average results for the fits of integrated signal energy against beam position.

is reasonable considering the large number of systematic effects, outlined explicitly in Section 5.5.5. Equation 2.29 was used to predict the sensitivity in terms of signal energy for a single bunch. The predicted result is 1.23×10^{-5} mJ nC^{-2} mm^{-2} which is about half of the measured value of 2.4×10^{-5} mJ nC^{-2} mm^{-2}.

5.5.4 Sensitivity for Long Beam Pulses

The position sensitivity for long beam pulses was measured in a similar way to the sensitivity measured using short pulses except instead of sampling at the peak, the signal was averaged over its steady state region. This is illustrated in Figure 5.31. The window used for the average starts at a fixed offset of 15 ns after the signal rise and is 40 ns long. This is also done for the reference cavity which is used to normalise for charge. Since the beam pulse length used, 60 ns, was also

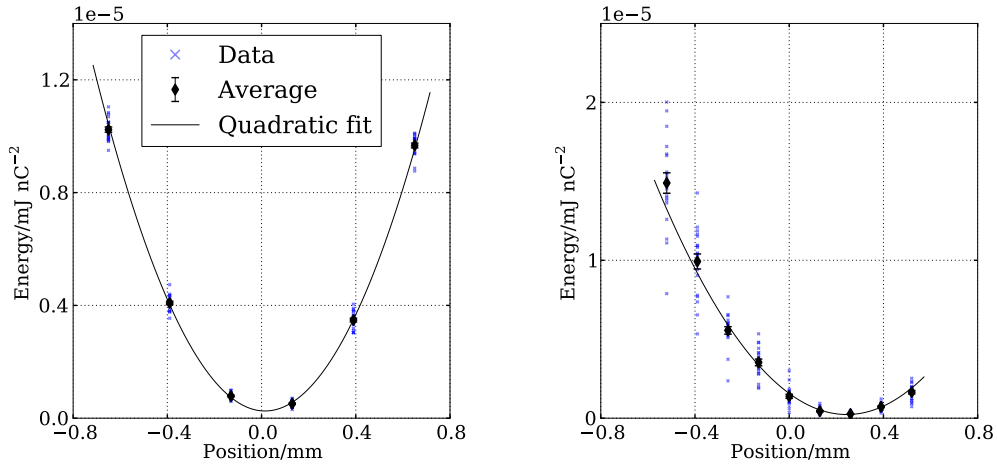

Figure 5.30: Total waveform energy for different relative beam positions in the horizontal (left) and vertical (right) directions.

used for a charge sensitivity measurement, no conversion of the calibration was necessary. The total pulse charge was also sufficient to be measured using the ICT. The average charge measurements for each position scan using the reference cavity and the ICT are compared on the right in Figure 5.31 which shows that the two measurements are correlated as expected and do not differ by more than 0.1 nC, about 3 % of the total charge. There are other differences in the processing compared with the short pulse measurements. Firstly, the beam arrival time, measured using the diode signal, was estimated using the weighted average method described in Section 5.3.4 instead of detecting the signal rise. Secondly, the digital local oscillator frequency was the same for all three channels and was measured from the reference cavity signal. This is suitable for long pulses because of the signal periodicity with the bunch arrival frequency harmonic seen in Section 5.5.1.

Examples of the each fit are shown in Figure 5.32 and the results are summarised in Table 5.13 with the short pulse measurements. Again, the results of each individual position scan performed with long beam pulses can be found in Section A.3 of the Appendix. Data with the quadrupole triplet switched off was only taken for the horizontal direction but the short train measurements show that the position cavity behaves very similarly in the vertical direction. As with

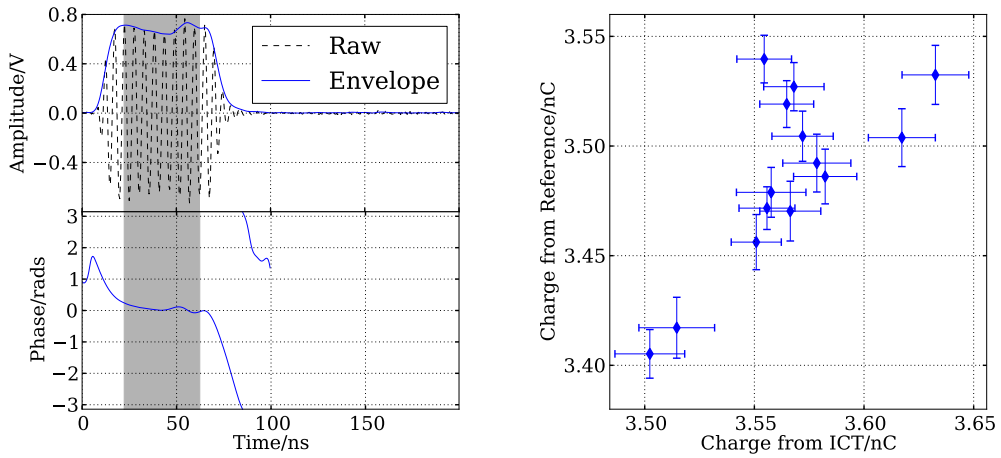

Figure 5.31: Example of a processed waveform where the averaging window is shaded (left) and correlation of the beam charge measurements using the reference cavity with the measurements made using the ICT.

the short pulses, the signals were integrated to obtain the total energy. The results are plotted in Figure 5.33.

The measured sensitivity in terms of peak signal amplitude, 130 V nC^{-1} mm^{-1}, is about 8 times greater than the result for the short pulse excitation, 16.4 V nC^{-1} mm^{-1}. Using Equation 2.25, this corresponds to a decay time of 5 ns. This is close to the value of 4.2 ns expected from the quality factor measurements. The agreement with the RF measurements is much better than for the reference cavity in Table 5.11 because the bandwidth of the position cavity is smaller and so less affected by the frequency response of the electronics. The measured sensitivity in terms of waveform energy, 0.0098 mJ nC^{-2} mm^{-2}, is in comparison with the theoretical value of 0.0130 mJ nC^{-2} mm^{-2} calculated using Equation 2.31 for a pulse length of 60 ns. The theoretical value is 30 % larger than the measured value as with the charge sensitivity measurements for this pulse length. Equation 2.31 is expected to give an overestimate since it assumes that the resonant frequency of the cavity is an exact multiple of the bunching frequency. The measured value is just over 400 times greater than the value measured with short pulses. The difference between this and the expected factor of just over 1000 with a 4.2 ns signal decay time can be almost entirely accounted

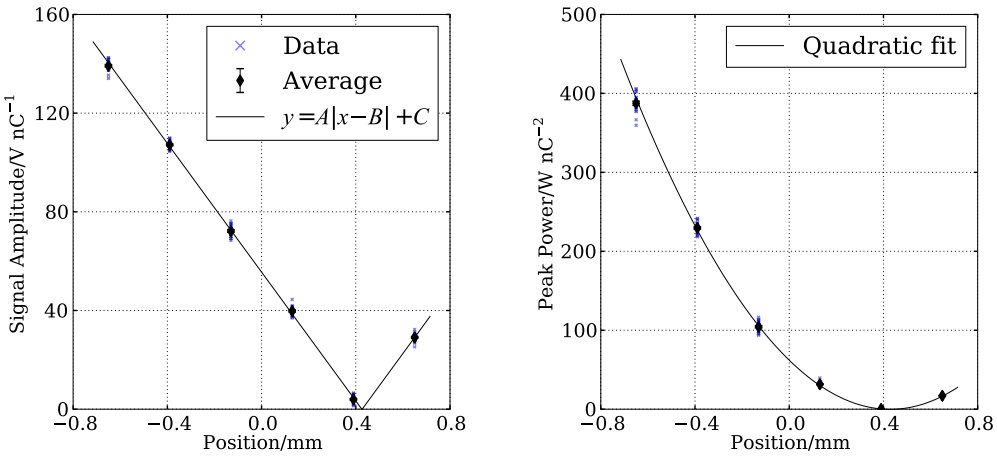

Figure 5.32: Example of fits made to average signal amplitude (left) and peak power (right) during the steady state portion of the waveform from a long beam pulse.

for by the fact that value for the short pulse measurement is half the predicted value.

5.5.5 Systematics Affecting Sensitivity Measurements

There are many systematic effects which make it hard to achieve accurate results from sensitivity measurements performed in this way. For the single bunch sensitivity, the two largest effects are the measurement of the charge and the response of the electronics. The charge during the single bunch sensitivity measurements was determined using the reference cavity signal whose calibration had been extrapolated from measurements using long beam pulses. The decay time used in the extrapolation was determined from a scan of the beam pulse length. If the decay time measured from the short pulse signal, 4.5 ns in Table 5.5, was used, the charge would have been measured to be 33 % lower and the measured position sensitivity would be higher by the same percentage. The gain of the electronics was measured, as described in Section 3.3.1, using continuous wave signals. In this case, standing waves caused by reflections in the electronics lead to a constant offset in amplitude. The gain for pulsed signals, where the reflections simply result in secondary output peaks, will be slightly different. Indeed, the

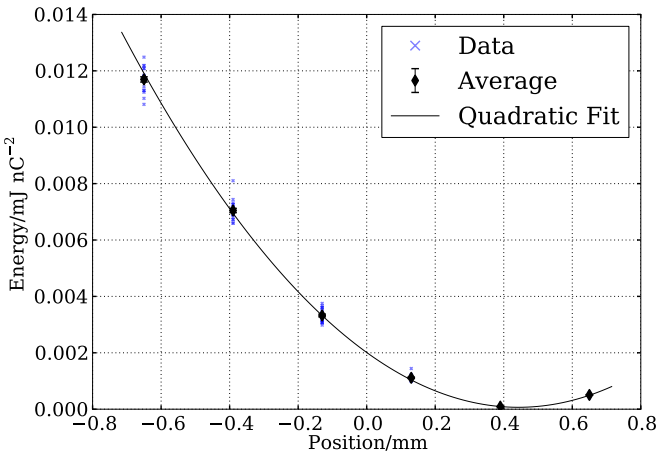

Figure 5.33: Total signal energy for different relative beam positions.

time domain measurements in Section 3.2.4 show that the response of the electronics varies by at least 15 % during the settling time. Furthermore, for short pulses, the finite electronics bandwidth means that the peak signal at the output of the electronics will occur later in the waveform compared with the signal at the input due to a finite electronics rise time. The form of the excitation coming from the cavity pick-up with short beam pulses is likely to produce a response from the electronics that is different to the CW case and will probably result in a lower peak output voltage. Furthermore, as discussed in Section 5.5.3, the use of a 2.1 ns pulse length to approximate a single bunch will lead to a slight underestimate of the measured sensitivity, at most by about 15 %.

The results of the short pulse measurements have been compared with the theoretical sensitivity calculated using Equation 1.41. This equation is based on instantaneous excitation and does not account for the signal shape seen which clearly peaks after the beam has left the cavity. This may be due to the filling of the cavity or the rise time of the electronics which, as can be seen in Figure 3.25, is not instantaneous. It was for this reason that the energy of the waveforms was calculated to be compared with theoretical estimates. These results are, however, also affected slightly by the response of the electronics because power at frequencies outside of the electronics bandwidth that is reflected is not accounted for.

For the long pulse sensitivity measurements, the response of the electronics is less important since the signal is close to CW. The results were compared with theoretical estimates of the waveform energies using Equation 2.31. These are inaccurate if the actual decay time of the cavity is different from the prediction based on the quality factor measurements. Equation 2.31 is also based on the assumption that the resonant frequency of both cavities is a harmonic of the bunch arrival frequency. The result will be lower in the case where the frequencies are offset, up to 20 % for a 20 MHz frequency offset and 60 ns pulse length.

There are further systematic effects that are accounted for or not expected to be as large. Slow drifts in the beam position during the measurement in addition to the changes introduced intentionally will affect the results. By repeating the measurement several times over a period that is longer than the time scale of the drifts, this effect should be averaged out. Hysteresis in the corrector magnets means that the deflection of the beam is not always as expected. This may contribute to the variation in the results of the corrector response measurement. The corrector response measurement itself may also have a systematic offset that will propagate to the sensitivity measurements. Finally, the reference cavity is calibrated in terms of beam charge to an ICT that is about 35 m upstream. There are likely to be some losses along the beamline which would make the charge at the cavity BPM lower than expected. This would lead to an underestimate of the pick-up sensitivity.

5.6 Position Measurement

The ultimate purpose of a cavity BPM is to measure the transverse beam position at its location to help optimise the performance of the accelerator. To be able to use a cavity BPM for this, first a beam-based calibration must be carried out. After this, the BPM output can be converted into a beam position in physical units and used to measure characteristics of the accelerator such as the beam position jitter.

5.6.1 Calibration

Calibrations were performed in order to obtain position readings from the cavity BPM output signals. The calibration determines the IQ rotation angle and the position scale factor described in Section 4.1.2. The position of the beam was scanned across the cavity BPM in known steps as in the sensitivity measurements described in Section 5.5.3. The timing was once again determined using the diode signal but because the phase is now important, the signals from the down-converter channels were sampled just after the peak in amplitude where the phase is most constant. The sampled signal is then expressed as a complex number. The changes in the signals during horizontal and vertical position scans are shown in Figure 5.34. The steps in the position scan, 20 pulses apart, are clearly visible as is the phase change when the beam crosses the electrical centre of the position cavity. The phase change is not instant or exactly equal to π because of the remnant signal from beam tilt and other modes as discussed in Section 5.5.3. The phase difference between signals from bunch trains either side of the cavity centre will tend to π for large offsets, however, as the influence of this remnant signal becomes negligible.

The complex position cavity signals are divided by the complex reference

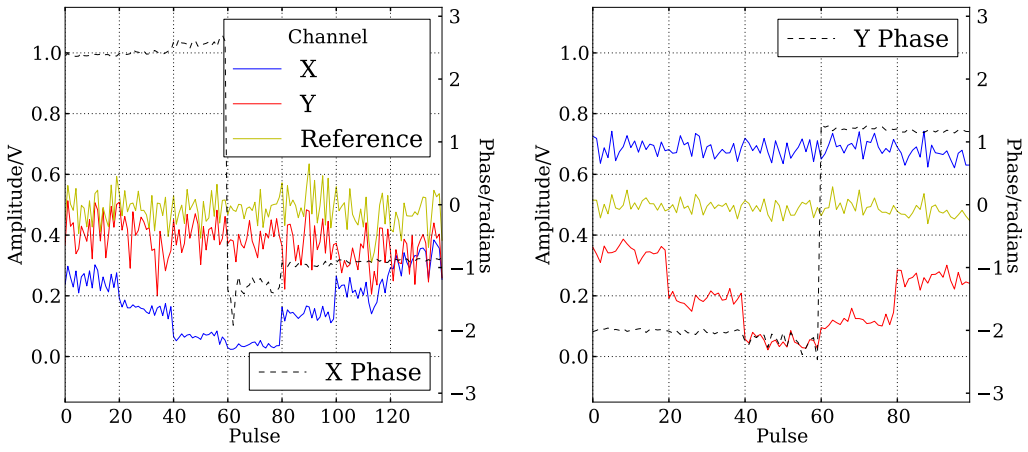

Figure 5.34: Change in amplitude and phase of the three channels during position scans in the horizontal (left) and vertical (right) directions.

cavity signals, giving a complex number whose real and imaginary components are the in-phase I and quadrature-phase Q components given by Equation 4.8. As discussed in Section 4.1.2, this complex quantity is dependent on the beam position, trajectory angle and bunch tilt only. The IQ rotation angle θ_{IQ} is determined by fitting Q against I under the assumption that any changes in signal are only due to changes in beam position. Next, the IQ rotation is applied so that all changes in position show up in one component, the position signal, while the other component becomes the tilt signal. Fitting the position signal against the known changes in position during the position scan gives a scale factor S_{P} to convert to a position in physical units, in this case, millimetres. Example scans in the horizontal and vertical directions are shown in Figure 5.35 and Figure 5.36 respectively where the fits used to determine the IQ rotation angle and the position scale factor are included.

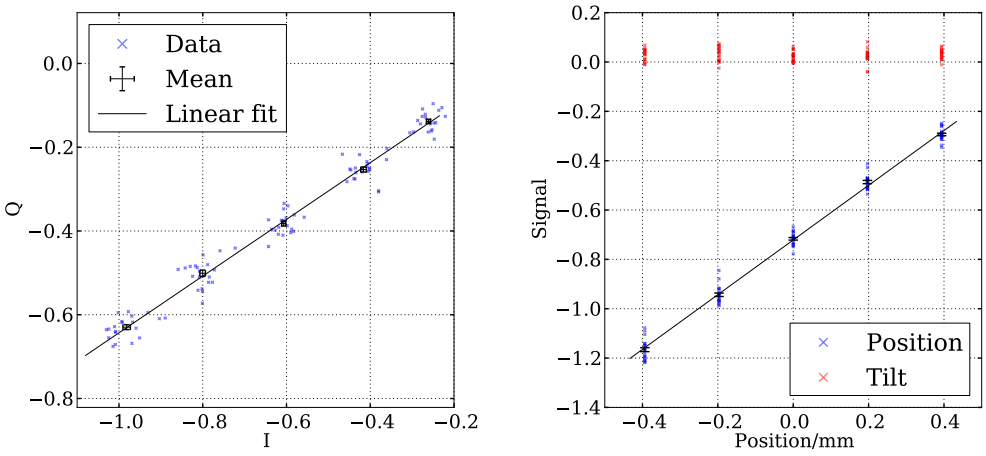

Figure 5.35: Measurement of the IQ rotation angle (left) and position scale factor (right) using a position scan in the horizontal direction.

In order to be able to compare calibrations from different days, the processing variables have to be kept constant. The sampling time of a signal is dependent on the timing offset used for that channel. The time at which the signal is sampled affects the measured phase because the phase is not perfectly constant around the signal peak. It also affects the measured amplitude. The phase is also affected by the frequencies used in the digital down conversion of the different channels.

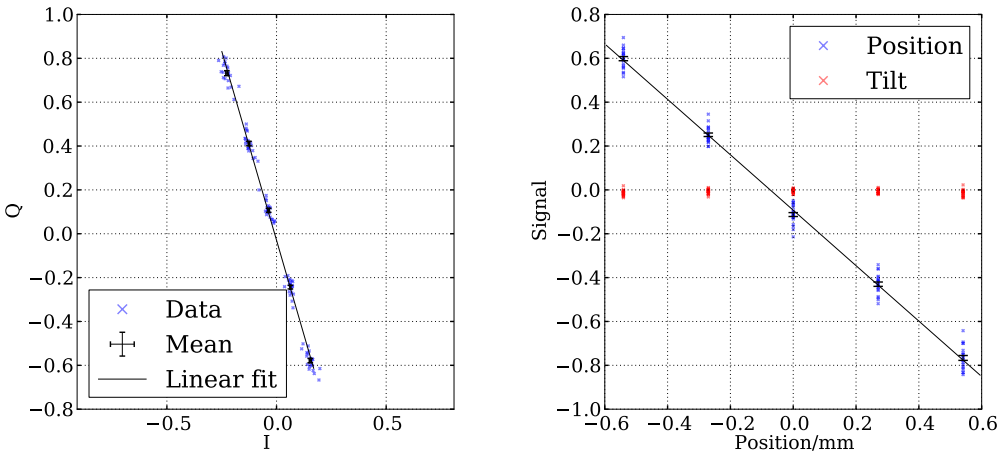

Figure 5.36: Measurement of the IQ rotation angle (left) and position scale factor (right) using a position scan in the vertical direction.

Therefore, the frequencies and timing offsets were measured from a single data file recorded on the 29th April.

Calibrations were performed on six different days and on each day, several short pulse position scans were taken that could be used as calibrations. The corrector responses as measured in Section 5.2.5 were used but were scaled to the beam energy, which was measured on each day using the energy spectrometer at the end of the probe beamline.

	Date (2013)	Position Scale Factor S_P/mm		
		Single Corrector	Corrector Pair	Average
X	29th April	0.856±0.008	0.882±0.011	0.872±0.008
	6th May	-	0.901±0.007	0.901±0.007
	13th May	−0.888±0.006	−0.884±0.003	−0.886±0.003
	6th May	-	1.470±0.016	1.470±0.016
Y	16th April	−0.802±0.003	−0.794±0.006	−0.797±0.004
	6th May	0.988±0.008	-	0.988±0.008
	13th May	−0.918±0.018	−0.931±0.020	−0.923±0.013
	15th April	−0.738±0.004	−0.677±0.007	−0.705±0.008
	6th May	0.954±0.013	0.93±0.05	0.949±0.015

Table 5.15: Position scale factors measured on different days. The highlighted results were measured with the final quadupole triplet switched on.

The results of the calibrations in terms of position scale factor and IQ rotation

	Date (2013)	Rotation Angle θ_{IQ}/degrees			
		Single Corrector	Corrector Pair	Average	$n_0 = 0$
X	29th April	33.6±0.7	33.77±0.13	33.7±0.3	−2.3±0.3
	6th May	-	26.88±0.13	26.88±0.13	−3.01±0.13
	13th May	71.4±0.2	71.58±0.07	71.51±0.10	−4.36±0.16
	6th May	-	26.76±0.07	26.76±0.07	−3.44±0.07
Y	16th April	−73.10±0.08	−73.70±0.10	−73.42±0.09	73.17±0.09
	6th May	1.0±0.4	-	1.0±0.4	71.8±0.4
	13th May	−84.3±0.2	−84.29±0.07	−84.31±0.12	72.75±0.12
	15th April	−73.4±0.2	−74.28±0.04	−73.6±0.3	72.70±0.14
	6th May	1.38±0.15	1.0±0.3	1.30±0.13	72.11±0.13

Table 5.16: IQ rotation angles measured on different days. The highlighted results were measured with the final quadupole triplet switched on.

angle are shown in Table 5.15 and Table 5.16 respectively. The horizontal results start at a later date because, due to problems with the field in the upstream spectrometer dipole magnet, this was when the beam was first centred in the position cavity in the horizontal direction. In both tables, the last row in the horizontal direction and the last two in the case of the vertical direction show results taken where the focusing was done using the final quadrupole triplet. This means that the beam size, and therefore, the beam position jitter is smaller but also means that the ballistic model is no longer valid and the model described in Section A.2 of the Appendix must be used to propagate the beam position from the correctors to the cavity BPM. This tracking is not successful which is why, for these runs only, the scale factors from the calibrations using one corrector are inconsistent with those where two correctors were used. No time was dedicated to optimisation of the model and in terms of the position scale, these results can be ignored. Significant variation of up to 20 % between the different dates can be seen. It is suspected that this is partly due to drifts in the relative timings of the signals. When a data file from the 15th April is used to determine the frequencies and timing offsets for the vertical calibrations on the 16th, an average scale factor of 0.881 ± 0.004 mm is obtained which is closer to the values obtained on other days.

From inspection of Equation 4.8, the relationship between the IQ rotation

angle θ_{IQ} and the arrival time n_0 (in samples) is

$$\theta_{IQ} = 2\pi \left(\frac{f_{\mathrm{IF,p}} - f_{\mathrm{IF,r}}}{f_{\mathrm{s}}} \right) n_0 + \phi \tag{5.11}$$

where $f_{\mathrm{IF,p}}$ and $f_{\mathrm{IF,r}}$ are the intermediate frequencies of the position and reference channels respectively, f_{s} is the sampling frequency and ϕ is some unknown offset. Because both channels use the same analogue LO, the difference in the intermediate frequencies is the same as the difference in the resonant frequencies of the modes of interest in each cavity. In order to remove this dependence, the measured IQ rotation angles were extrapolated back to a beam arrival time of $n_0 = 0$. This allowed the calibrations taken on different days, when the timings were different, to be directly compared. In this case, the results from all the calibrations in each transverse direction are within 1.1° of the mean across the different days. This includes the results taken with the final quadrupole triplet switched on because no tracking is required for the determination of the IQ rotation angle. The only assumption is that any observed changes in signal are purely due to changes in the beam position and this holds in both cases.

5.6.2 Beam Jitter

Once a calibration was obtained, the BPM could be used to give position readings in mm. This could then be used to measure the beam position jitter from pulse to pulse. Because no dedicated data was taken for this purpose, the beam jitter in each direction was taken from 100 pulses of a position scan in the perpendicular direction. The results are shown in Table 5.17 and include the effects of any slow orbit drifts that occur on a time scale less than 100 pulses. The beam jitter was measured both with strong focusing using the final quadrupole triplet and without. The IQ rotation angles used were from the date when the data was taken while the position scales were taken from calibrations on the 13th May. The beam jitter is smaller when there is stronger focusing as it is expected to be a constant fraction of the beam size. The beam size was predicted using the beamline model based on Twiss parameters measured during a quadrupole scan taken earlier in the day [4]. The beam size on the 13th May was not recorded.

Date	Position Jitter/μm		Beam Size/μm	
	X	Y	X	Y
06/05 (Distant focusing)	20	34	190	390
06/05 (Strong focusing)	13	27	60	260
13/05 (Distant focusing)	18	78	-	-

Table 5.17: Pulse to pulse beam jitter measured from 100 pulses and compared to the predicted beam size.

5.6.3 Intra-Train Motion

Position measurements within a single bunch train were attempted over the constant amplitude portion of the signal. Although it is not possible to verify the results using a single BPM, this allows any obvious systematic effects to be identified and should give an estimate for the extent of the intra-train beam motion. The deconvolution algorithm described in Section 4.1.3 was applied and compared with digital down conversion.

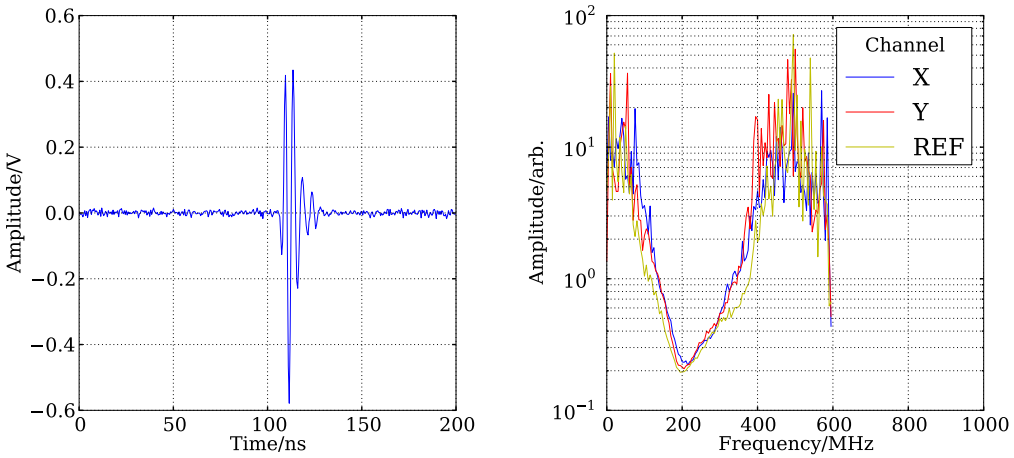

Figure 5.37: Single bunch waveform from the horizontal channel (left) used to determine the inverse frequency response, which is shown with the inverse responses of the other two channels (right). These are used in the deconvolution algorithm.

A single bunch signal from the horizontal channel and the corresponding inverse response are shown in Figure 5.37. The values used for k_1 and k_2 in Equation 4.15 were 500 MHz and 600 MHz respectively. The inverse frequency response

is 0 for frequencies above k_2. It is clear that there is still a lot of noise in the inverse response either side of the signal frequency.

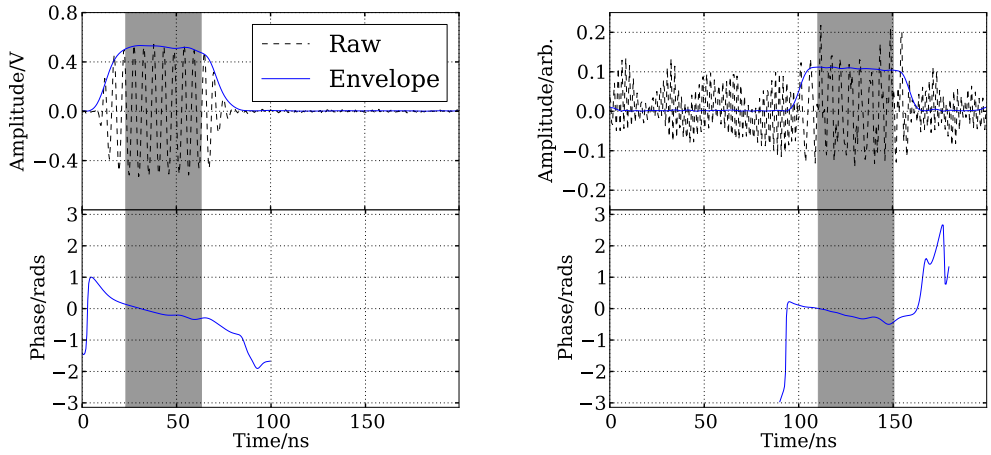

Figure 5.38: Example waveforms with no deconvolution (left) and after deconvolution (right). The same down-conversion processing has been applied to both to determine the amplitude envelope and phase.

Figure 5.38 shows a raw digitised waveform from the reference channel along with a waveform where deconvolution has been applied. As expected, the deconvolution has made the signal's rise and fall steeper but has also significantly increased the noise outside of the signal bandwidth. The same digital down conversion was applied before and after the deconvolution. The Gaussian filter bandwidth was chosen to be 53 MHz to further attenuate the additional noise in the deconvolved waveform. It appears from the signal amplitude envelopes that the deconvolution has also flattened the top of the signal. This suggests that the variation in the amplitude of the monopole signal is not simply due to charge variation but also the addition of signals from the different bunches.

After the processing, the waveforms from all three channels were sampled over a 40 ns window within the constant amplitude portion. The timing of the sampling windows was kept fixed relative to the timing calculated using the weighted mean of the diode signal. In order to do this, the timing offsets had to be measured beforehand. The start time of the single bunch signal in Figure 5.38 was then subtracted from these offsets to obtain the offsets of the deconvolved wave-

forms (negative values would have to be added to the digitiser window length, see Section 4.1.3).

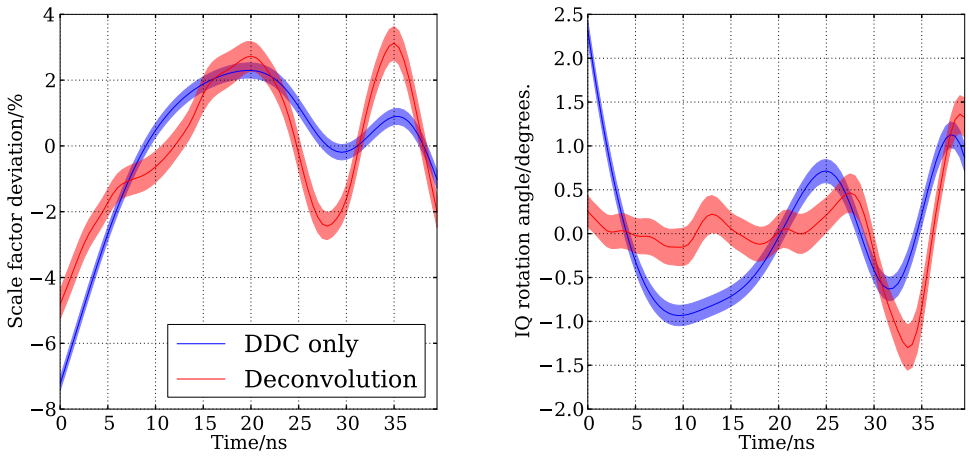

Figure 5.39: Variation of the position scale factor (left) and IQ rotation angle (right) along the sampling window during the long pulse signal waveform. The coloured band represents the fit error.

Each sample in the 40 ns window of the position signal was individually normalised for charge and referenced in phase using the corresponding sample in the window of the reference signal. A horizontal position scan using two correctors was then used to calibrate each sample individually. The variation in the calibration constants along the train are shown in Figure 5.39. There are small oscillations in the IQ rotation angle with both processing methods. These are due to the relative arrival times of the two signals. Small bumps in the measured phase appear consistently in every pulse and in all channels (some are visible at the end of the waveforms in Figure 5.38). If the timing of the position channel and the reference channel waveforms are offset, these bumps occur at different times so that when the reference signal phase is subtracted from the position signal phase, the bumps don't cancel out. Such oscillations in phase could be caused by variations in bunch intensity along the train. Because of the high resonant cavity frequency, small changes in the separation between consecutive bunches (< 0.1 %) could also be the cause such phase changes.

A similar effect appears in the position scale factors. If the position and

reference cavity sample windows are not lined up in time, small changes in bunch intensity in the middle of the train will not be normalised out correctly and will show up as oscillations in the position scale factor. Figure 5.40 shows calibration constants based on the same data but where the reference cavity sampling window was adjusted by eye so that the bumps in amplitude and phase line up with the correlated bumps in the position cavity signals. This better synchronises the two channels and so the oscillations in both the calibration constants are smaller.

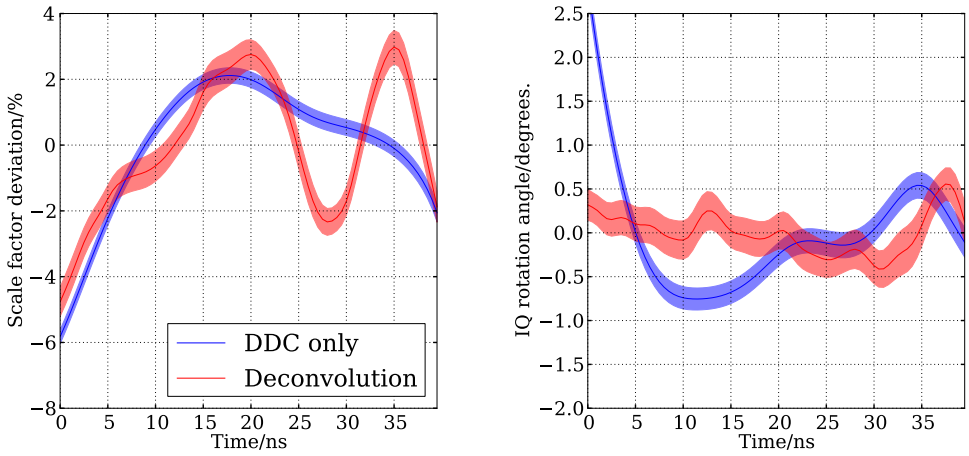

Figure 5.40: Variation of the position scale factor (left) and IQ rotation angle (right) along the sampling window after the timing offsets have been adjusted so that the sampling windows of the position and reference signals are closer in real time. The coloured band represents the fit error.

Slower variations in the position scale factor of up to 8 % are also seen along the bunch train after the adjustment to the timing has been made. The first part of the sampling window is at the end of the signal rise. Since, as seen in Section 5.5.1, the position channel signals have a steeper rise than the reference signal, the measured scale factor is smaller at the beginning of the waveform. The deconvolution is expected to steepen the rise but only has a small effect. It has a much clearer effect on the IQ rotation angle which is clearly more constant along the train after the deconvolution has been applied. Differences in the shape of the signals from the position and reference channels can also explain the slow decrease in the scale factor throughout the rest of the waveform. The outstanding

oscillation in the position scale factor with deconvolution is due to a feature in the horizontal position channel signal that does not appear in the reference cavity signal. Since there is no trace of such a variation before deconvolution, it must be a feature the deconvolution algorithm. Without resolution measurements, it is not possible to determine whether this affects the accuracy of the position measurements.

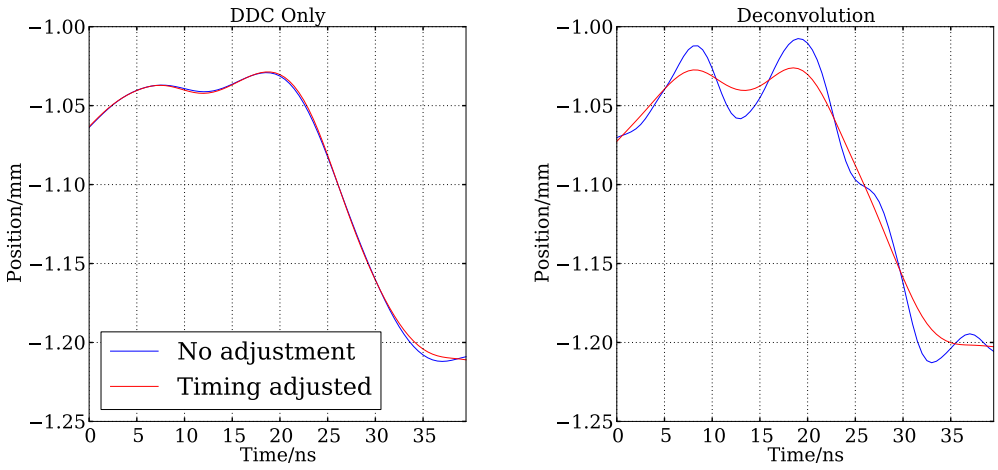

Figure 5.41: Measurement of the beam position along a single bunch train without deconvolution (left) and with (right) before and after the adjustment to the timing of the sampling windows.

Figure 5.41 shows a horizontal position measurement along the sampling window from one beam pulse in the same data set used for the calibration. The variation in position does not follow the variation in the calibration constants which means that the calibration has successfully removed the systematic differences between the different samples. A drift in the transverse position of 200 μm is seen towards the end of the bunch train. The timing adjustment that was used to remove some of the variation in the calibration constants has a large effect on the measurements made using the deconvolution algorithm. It is therefore clear that the synchronisation of the position and reference cavity channels is important if measurements are to be performed in this way. The synchronisation could be improved in hardware through adjustments to the clock phases of the digitiser channels or through time domain measurements to match the delay of the three

analogue processing channels.

The position measurement at each sample is a weighted average of the positions of the previous bunches where the weights are a function of the time between the bunch and sample and the bunch charge. In the case where the bandwidth of the position channel is not the same as that of the reference channel, this is no longer the case so that the weighted average is no longer correctly applied and can lead to apparent beam position variation along the train where there is none. The bandwidth of the two position channels should therefore be closely matched to the reference channel. In some parts of the train shown in Figure 5.41, the difference in the position measurements with deconvolution and without are at the level of 20 μm. Differences in the bandwidths of the position and reference channels could be the cause of such a systematic uncertainty.

Chapter 6

Conclusions and Discussion

6.1 Results

A prototype cavity beam position monitor pick-up has been characterised in the laboratory using a network analyser and then installed on the probe beamline at CTF3. Analogue signal processing electronics were designed, constructed and measured in order to determine the prototype's performance. A working beam test setup including software and controls has been put together in order to perform the tests. Measurements performed using the beam include signal decay time and frequency, the response of the monopole reference cavity to beam charge and the response of the position cavity to beam charge and position offset. A beam based calibration was performed so that the BPM signals could be converted into position readings in physical units. The single bunch position jitter could then be measured and the minimum value observed was 13 μm. Position measurements were also made across the constant amplitude portion of a multiple bunch signal and the challenges that this presents have been identified.

Analytical expressions for the form of the signals excited by multiple bunches have been derived. These predicted a convergence of the signals in terms of amplitude and phase and a signal component at the harmonic of the bunching frequency that is closest to the cavity mode resonant frequency. The predictions were confirmed in simulation and in beam test measurements, where it was found that the bunching frequency harmonic becomes the dominant signal frequency for long bunch trains, as might be expected for such an externally excited oscillator.

A deconvolution algorithm has also been developed in order to reduce the time between uncorrelated position measurements. This has been applied to both real and simulated data but resolution measurements, discussed in Section 6.3, must be made to determine whether it provides any real advantage over digital down conversion alone.

6.1.1 Sensitivity

The sensitivity of the position cavity output to beam position and charge and the sensitivity of the reference cavity output to charge were measured both in terms of peak signal and the total signal energy. For the position cavity, the measured sensitivities of 16.70 ± 0.09 V nC^{-1} mm^{-1} in x and 16.1 ± 0.5 V nC^{-1} mm^{-1} in y are about 70 % of the predicted value of 24.2 V nC^{-1} but there are a lot of systematic effects as described in Section 5.5.5. It is hard to determine the quantitative impact of all these effects, particularly the response of the electronics to the short beam pulse signal. Furthermore, the predicted position sensitivities are based on microwave measurements of the external quality factor that rely on the assumption that a short-circuited port has no influence on the power loss in the pick-up. The result of these measurements is likely to be an overestimate, which would lead to an underestimate for the predicted sensitivity. The signal energy was used as a way of removing the dependence on the signal shape but it still requires the gain of the electronics to be known. The results of the measurements, $(2.345 \pm 0.014) \times 10^{-5}$ mJ nC^{-2} mm^{-2} in x and $(2.39 \pm 0.07) \times 10^{-5}$ mJ nC^{-2} mm^{-2} in y, are close to twice the predicted value of 1.23×10^{-5} mJ nC^{-2} mm^{-2}. Conversely, the multiple bunch sensitivity measurement of $(9.79 \pm 0.06) \times 10^{-3}$ mJ nC^{-2} mm^{-2} is 25 % lower than the predicted value of 0.0130 mJ nC^{-2} mm^{-2}. Part of this may be down to the offset in frequency from the nearest harmonic of the bunch arrival frequency, 15 GHz. The resonant frequency from the laboratory measurements in Section 2.2.5 is offset from 15 GHz by 12 MHz which would lead to a predicted decrease in signal energy of 8 %.

The measured charge sensitivity of the reference cavity is also affected by systematics, particularly because the measurement could only be made with long

pulse lengths. The beam pulse length scan was used to scale the measured sensitivity for the purposes of single bunch charge measurements and gave a calibration constant of around 64.5 V nC^{-1}. Using the decay time from the RF measurements resulted in a measured single bunch sensitivity of 132 V nC^{-1} which is twice as large and closer to the prediction of 165 V nC^{-1}. The sensitivity in terms of signal energy does not depend on the signal shape and the response of the electronics and so should only depend on the cavity parameters determined from the RF measurements and the electronics gain. The measured results of 0.1115 ± 0.0006 mJ nC^{-2} for the 30 ns pulse length and 0.1989 ± 0.0012 mJ nC^{-2} for the 60 ns pulse length are lower than their respective predicted values of 0.129 mJ nC^{-2} and 0.271 mJ nC^{-2}. For the 13 MHz offset of the cavity resonant frequency from 15 GHz measured in Section 2.2.5, only a 4 % decrease in the signal energy is expected for the reference cavity. The effect is smaller than for the position cavity because of the shorter resonant mode decay time.

6.2 System Performance

6.2.1 Microwave Pick-Up

The beam-based sensitivity measurements of the position cavity to beam charge and position are close enough to the predicted values to say that the design is successful. The cross-coupling between adjacent output ports of the position cavity increased significantly after brazing but with modified antenna designs resulting in looser tolerances, the necessary improvements seem feasible without modifying the pick-up geometry. The reference cavity had to be modified so that the resonant frequency was correct which meant that its loaded quality factor ended up being much lower than that of the position cavity. The reference cavity must therefore be redesigned so that it not only has the correct resonant frequency but is better matched to the position cavity in terms of signal decay time. This is important in terms for the multiple bunch position measurements where, as mentioned in Section 5.6.3, if this is not the case, variation in charge along a bunch train can lead to apparent variation in the beam position when there is

none.

The pick-up geometry for the second iteration has recently been finalised at CERN. The decision was taken to build the assembly from copper so that the internal quality factor is higher. Using the similar geometry still gives a low external quality factor. This means that less energy is lost in the cavity walls and more is extracted over the longer signal decay. It also means that the bandwidth of the signal is narrower so that rejection of the image frequency is easier in the electronics [35] and a lower IF can be used as discussed in Section 6.2.3. The final predicted value for the loaded quality factor is 520. The requirement for the signal decay has been relaxed by a factor of 10 so that over 50 ns, the signal will have decayed to 1 % of its initial value. Some signal subtraction may therefore have to be applied if independent measurements are to be made to a high precision with this separation in time. A beaded feedthrough antenna has been designed with a bead of length 0.5 mm, radius 0.155 mm and curvature 0.25 mm. The optimum spacing between the antenna and the wall of the waveguide is indeed larger at 0.516 mm. Due to a change in manufacturer, the antenna radius has also changed along with the radius of the antenna holes to keep the 50 Ω impedance of the coaxial sections. The geometry of the reference cavity, which will not have beaded antennas, was also changed, first to ensure a 15 GHz resonant frequency with the new feedthrough antennas and then to match the loaded quality factor of the position cavity. This will ensure a similar signal build up in both cavities and make both the analogue and digital processing of the multiple bunch signals simpler [3].

All dipole like modes excited in a cavity, including those used to determine the beam position, contribute to the transverse wakefields. Axially symmetric modes can also degrade the beam by extracting energy, a longitudinal wakefield effect. The installation of so many cavity BPMs in CLIC must be assessed in terms of the wakefields. If the degradation of the beam is too much, other solutions must be considered. One proposed solution is a BPM with an RF choke that traps the first dipole mode and uses RF absorbing materials to damp the unwanted modes [56].

6.2.2 Electronics

A set of receiver electronics with three channels and a single stage of down-conversion has been designed and constructed for the beam tests of the prototype pick-up. A scheme with one gain stage was chosen after estimating the noise figures, total gain and dynamic range of several schemes. Each component in the final layout was measured individually before partial and full channel measurements were made. The total gain of each channel is 4.8 dB and their average bandwidth is 208 MHz. The output power at 1 dB compression is 13 dBm for all three channels, which, combined with the predicted noise figure of 16 dB and the estimated thermal noise floor of -108 dBm, gives a conservative estimate of 74 dB for the dynamic range.

Experience gained during the first beam tests has highlighted the features of the electronics that are most important for these measurements. The response in the time domain was found to be important in both the decay time measurements and the sensitivity measurements, particularly in phase. Conversely, the pick-up has been found to be sufficiently sensitive that a small amount of gain can be sacrificed. From the results of Section 3.2.3.5 and Section 3.2.4, it may therefore be better to use the more broadband, spare amplifier if the same down-converter electronics are to be used in further beam tests. Future iterations of the electronics may also include a small amount of gain (less than 10 dB) at the front end to improve the noise figure. The maximum input power could be raised slightly before being limited by the input 1 dB compression point of the mixer.

Other electronics schemes have been used for cavity BPM signal processing that are worth considering. A homodyne system, which has been used elsewhere [57], in which an IQ mixer is used to convert the BPM signals to baseband would simplify the digital processing. This was rejected because it is too expensive to be employed in thousands of BPMs. Figure 6.1 shows a simplified schematic of a very different, hybrid based system. Here, the reference cavity signal is split and combined with the signal from each channel of the position cavity in 180° hybrids. The sum and difference outputs of the hybrids are then rectified using diode detectors so that no analogue LO is required. When the beam is centred in the cavity, the signal level from both hybrid outputs is the same, the output of

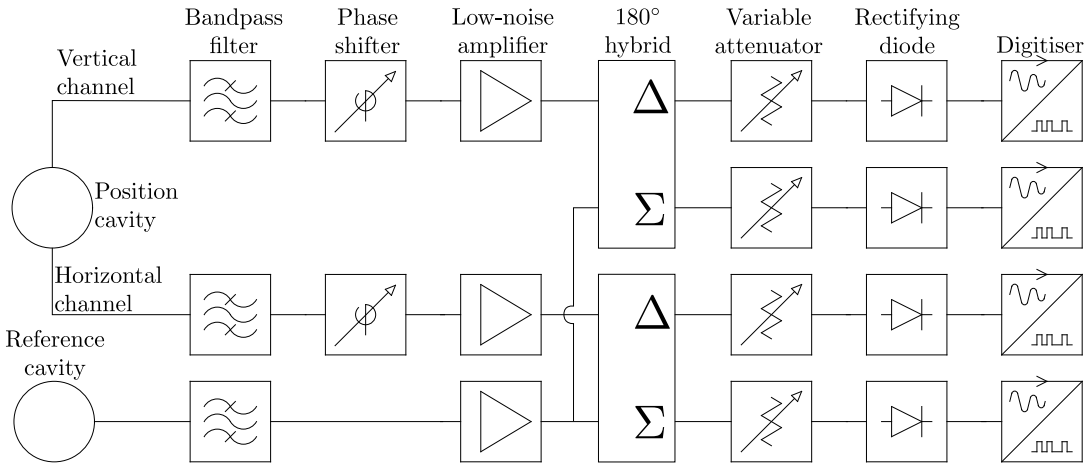

Figure 6.1: Diagram of cavity BPM electronics similar to those at Fermi@Ellettra

the reference cavity. When the beam has some offset, the difference between the signals from the two hybrid channels represents the position signal and the sum of the two can be used to normalise for charge. The phase of the position cavity channels can be adjusted so that a change in beam position only corresponds to a maximum change in interference in the hybrids. Such a scheme has been tested at Fermi@Elettra [58]. It has the advantage of simplicity with no further processing required in the digital domain. However, it performs significantly worse in terms of linearity due to the diode detectors. In [58], a linearity of 1 % is achieved over a dynamic range of only 20 dB, much less than the 66 dB required for CLIC. Variable attenuators can be used to increase this range but they would require more frequent user input or some automation for their control. Furthermore, as seen in Section 3.2.1.4 and Section 5.4.1, changing the attenuation also results in a change in the signal phase and this would alter the interference in the hybrids, introducing further systematic uncertainties.

6.2.3 Digitiser

A digitiser with a fast sampling rate of 2 GS s^{-1} and a low 10 bit resolution, was available for the first beam tests and so was used. The fast sampling rate allowed for a high IF and a for a large bandwidth to be maintained throughout

the electronics. The digitiser was also used with a virtual scope interface which allowed the signals to be viewed and controlled easily. A digitiser window length of 400 samples and a voltage range of ± 1 V was chosen for the down-converter channels. Measurement of the negative polarity, diode rectified signal was usually done over a smaller voltage range and with a positive pedestal value. The noise of the digitiser was measured with the input ports terminated in matched loads and this gave a resolution of 8.4 effective bits. The resolution on a signal at 200 MHz would be smaller than this because of the digitiser clock jitter but this effect was estimated to be less than 0.1 bits.

In order to achieve both the target operational beam position range of ± 100 μm and resolution of 50 nm, a dynamic range of 66 dB is required. This corresponds to at least 12 effective bits. A digitiser with at least 14 bits should therefore be used in future tests. With current technology, this will limit the sampling rate to at most 250 MSs^{-1} which in turn, limits the electronics bandwidth to the corresponding Nyquist frequency of 125 MHz. If the 200 MHz intermediate frequency is to be maintained, this would mean under-sampling. The signal to noise ratio coming from the digitiser clock for a signal of frequency f_{sig} sampled by a digitiser is given by

$$SNR_{\mathrm{sig}} = \frac{1}{(2\pi\sigma_t f_{\mathrm{sig}})^2} \tag{6.1}$$

where σ_t is the RMS clock jitter. It is clear that the signal to noise ratio is degraded for a larger intermediate frequency [40]. Choosing an IF of 70 MHz instead of 200 MHz would improve the signal to noise ratio by 9 dB. The digitiser used in the first beam tests has an RMS clock jitter of $\sigma_t = 1.2$ ps which would lead to a signal to noise ratio of 56 dB at 200 MHz IF. Therefore, for large enough signals and a 14 bit digitiser, the clock noise may be sufficient to dominate the thermal and quantisation noise and so must be taken into consideration.

The disadvantage of a lower IF is the required bandwidth of the front end electronics before the mixer. As discussed in Section 3.1.1, to achieve the best possible phase noise, the output of the mixer that comes from input at the image frequency must be rejected to below the digitiser resolution. This can be achieved

through filtering, the use of an image reject mixer or a combination of both. Because the cavity BPM pick-up has a low quality factor, the rejection at the image frequency will have an effect on the signal shape. This is another argument for having a pick-up with a higher quality factor [35]. Furthermore, narrow band filters are harder to produce at high frequencies. More reproducibility can be achieved if the electronics are band-limited at the lower IF. In general, the electronics bandwidth should be matched to the pick-up in order to extract the pick-up's optimal performance.

It was noted in Section 5.6.3 that, in order to make position measurements that are close together in time, the synchronisation of the position and reference channel signals must be as precise as possible. This may be achieved by improving the multiple bunch timing measurements after digitisation or else, ensuring that the two channels are sampled synchronously at the digitiser. A hardware solution could be similar to the bunch synchronous ADC sample clock phase that is used in [57] to ensure that the BPM signals are always sampled at the peak.

6.2.4 Resolution

If three or more cavity BPMs of the same design were installed close to each other, the resolution of each could be measured. A prediction of the resolution of the current cavity BPM prototype can be made using the measured parameters. A summary of the calculation is shown in Table 6.1 where the measured single bunch sensitivities of the position and reference cavity signals have been used along with the gain as measured for each analogue processing channel. The thermal noise has been estimated using Equation 1.56 and the noise gain has been calculated using Equation 1.59. The electronics bandwidth used was that of the front end bandpass filter to show that even with the widest possible electronics bandwidth, the digitiser noise is larger than the electronics noise and dominates the achieved resolution. To confirm this in future, the noise of the electronics that is not correlated to changes in beam position and charge could be measured by splitting the signal from a single pick-up output port into multiple electronics channels. The resolution has been calculated for the 14 bit digitiser proposed in Section 6.2.3 whose resolution has been estimated at 12 effective bits.

Parameter		Position cavity (Reference cavity)
Measured sensitivity/V nC^{-1}		17 mm^{-1} (65)
Channel gain/dB		-6.0 (-17.7)
Electronics bandwidth/MHz		300
Thermal noise/μV		0.9
Noise gain/dB		17.2
Electronics noise at digitiser/μV		6.6
Bunch charge/nC		0.1
Signal level/V	50 nm offset	4.3×10^{-5} (0.85)
	100 μm offset	0.085 (0.85)
Digitiser resolution/effective bits		12
Digitiser voltage range/V		1
Digitiser noise/μV		490
Resolution/nm		570
Maximum measurable offset/μm		1200

Table 6.1: Calculation of the predicted single bunch resolution from the measured parameters of the prototype cavity BPM.

The predicted resolution of 570 nm is more than ten times the targeted resolution of 50 nm and so more gain is required in the processing electronics. The maximum measurable offset of 1200 μm is also more than ten times the 100 μm required and so the dynamic range is sufficient and above the target of 66 dB. Therefore, as well as removing the 6 dB of fixed attenuation, at least 14 dB of gain must be added to each position cavity channel of the processing electronics without sacrificing dynamic range. The desired performance should then be obtainable and the BPM prototype would be ready for the resolution measurement to be made. The electronics constructed for the first beam tests did not have the required gain because it was known that resolution measurements would not be possible and uncertainties in the beam conditions and pick-up performance were of greater concern.

The calculation in Table 6.1 uses the pick-up sensitivity to single bunch excitation. If a long beam pulse was used for the resolution measurement, a sensitivity of 130 V nC^{-1} mm^{-1}, 8 times larger than the single bunch sensitivity, would be the correct value to use. This would bring the achievable resolution to 74 nm, which is much closer to the desired value. However, the measurements of the average train position would be affected by the intra-train beam motion and this may reduce the correlation of the position readings from the different BPMs.

Operating with a single bunch is therefore preferable for a measurement of the spatial resolution.

6.3 Future Tests

Future beam tests of the current and next iterations of the cavity BPM prototype will include the mover stages, which were not operational during the first beam tests because of delays in procuring the supporting hardware. This will allow a direct comparison between the BPM signal and the beam position without the further beam based measurements that were required to determine the angular response of the correctors. Furthermore, no model at all will be needed to calculate the position change and so the upstream quadrupole triplet can be used to focus the beam in the centre of the position cavity. This will reduce the beam position jitter and hopefully make it easier to steer the beam through the pick-up's 8 mm diameter beam-pipe.

One of the most important measurements needed to assess the performance of the system is the spatial resolution. This will be done using at least three prototype BPMs so that the beam motion that is larger than the resolution and correlated in all the BPMs, can be subtracted. Having the multiple pick-ups on individual mover stages will make it easier to steer the beam through their 8 mm beam-pipe aperture, which will have to be maintained for the length of the three pick-ups with some distance between them to allow for their independent positioning. The individual mover stages will also allow each BPM to be calibrated individually.

6.3.1 Measurement Bandwidth

In order to determine the measurement bandwidth of the BPM, it may be necessary to add a known kick to the beam that varies along the bunch-train. Since any tests would have to be compatible with the main goal of TBTS, which is to test two-beam acceleration and CLIC prototype accelerating structures, and since the accelerating structures have irises of less than 7 mm in diameter, the

most suitable location for a kicker would be downstream of them which would leave about 4 m between the kicker and the BPM test installation.

The displacement Δx due to an electrostatic kicker of length L, installed at a distance d upstream, can be approximated by

$$\Delta x = \frac{2LeV}{Eh}d \tag{6.2}$$

where V is the potential applied to each of the electrodes (with opposite sign), h is the gap between electrodes, E is the beam energy and e is the elementary electric charge. Equation 6.2 is a good approximation when the kicker electrodes are much wider than the beam, the kick angle is small and the edge effects are negligible. Installing a kicker 4 m upstream of the BPM with a stripline electrode of length 300 mm (such as the one used by the Feedback On Nanosecond Timescales group at the Accelerator Test Facility, KEK [59]) with an aperture equal to the 40 mm beam-pipe diameter of CALIFES, the expected maximum displacement of the 200 MeV beam would be 300 nm V^{-1}. Therefore, if it is possible to supply a potential of 100 V to each of the electrodes, the expected displacement would be 30 μm. If this could be supplied within a 150 ns bunch train, it would be possible to test the ability of the BPM to make two measurements within this timescale.

In order to calibrate the kick given to the beam, the timing of the kick could be moved with respect to the beam arrival and the resulting change in the projection of the bunch-train on a screen could be measured [60]. This would allow the time structure of the kick to be probed. Timing signals with a precision of hundreds of picoseconds are available at CTF.

A different method, which would not require a calibration of the kick, would be to modulate the kick given to the beam. If the amplitude of the modulation is independent of the frequency or if the dependence is known, the measurement bandwidth of the BPM could be determined from the amplitude of the measured position modulation. Assuming that the frequency response of the BPM is completely defined by the digital Gaussian filter, the modulation in position as measured by the BPM would be a Gaussian function of frequency centred at 0.

The measurement bandwidth BW_{meas} would then be given by

$$BW_{\text{meas}} = 2\pi f_{\text{c}} \tag{6.3}$$

where f_c is the width of the Gaussian function in the frequency domain. The frequency of the modulation could be scanned from DC to close to the nominal measurement bandwidth of the BPM (≈20 MHz) and a fit made to the measured modulation amplitude to determine f_{c}.

In terms of existing hardware in the CALIFES beamline, it may be possible to use the travelling wave deflecting cavity, which is a 3 GHz structure designed to kick a single bunch vertically for bunch-length and longitudinal charge distribution measurements. This is able to create transverse electric fields in the vertical direction of 1.54 MV m^{-1} with 7 MW of input power [61] although this is far greater than what would be required. The displacement at the BPM Δy would then be given by [62]

$$\Delta y = R_{32}\frac{eV_0}{E} \tag{6.4}$$

where V_0 is the voltage amplitude, E is the beam energy and e is again the elementary electric charge. In the case of CALIFES R_{32}, which is the transfer matrix element from vertical beam angle to position, is not simply the distance between the deflecting cavity and the BPM because the deflecting cavity comes before any focusing elements so some would have to be switched on. The following calculations have therefore been made assuming focusing with the first CALIFES quadrupole triplet only, in which case R_{32} is about 16 m. Scaling the electric field E_y with the root of the input power P gives:

$$\frac{\Delta y}{\sqrt{P}} = R_{32}\frac{E_y eL}{E\sqrt{P}} \tag{6.5}$$

where L is the cavity length. The resulting beam displacement in terms of input power is therefore 9 mm MW$^{-\frac{1}{2}}$ for the 200 mm long cavity. This suggests that an input power of 65 dBm or 3 kW would be enough for a vertical displacement of 30 μm at the BPM. In order for the kick to switch polarity over the length of

a single bunch train, the deflecting cavity would have to be detuned by 3 MHz with respect to the bunch arrival frequency.

Another potential method would be to use the BPM to measure the energy spread along the bunch train. This could be compared with a measurement made using the spectrometer. The current in the bending magnet required to steer the beam to the centre of the BPM pick-up is currently about 2.3 A while, for operation of the spectrometer, it is close to 63 A. A 2 MeV (1 %) spread in the energy along the bunch train would therefore, with a naive approximation, result in a position change within the bunch train of 130 μm in the BPM, assuming that the design energy of 200 MeV results in the design deflection of 22.75°. This method could be made more accurate by performing magnetic field measurements at the two current settings.

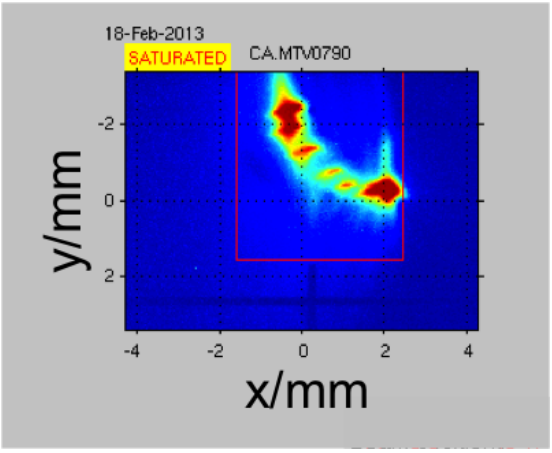

Figure 6.2: Image from a beam profile monitor where the individual bunches in a single train can be seen [4].

Figure 6.2 shows an image from a beam profile monitor where the individual bunches in a single train can be seen with large separations in transverse position. Such an image is usually seen with only a few bunches and high charge and so the cavity BPM would not be able to measure the positions of each bunch. However, if such a beam could be steered through the BPM, some information about the BPM measurement bandwidth could be obtained. The OTR screen downstream of the cavity BPM would be useful for a comparison in this case.

Appendix

A.1 Twiss Parameters and Emittance

The particle beam in a linear collider must be transported from the damping rings, along the main linac and through the beam delivery system, through all the beam line elements, without compromising the beam size and intensity at the interaction point. It is therefore being continually focused using quadrupole magnets. The rest of this section gives a brief overview of the equations describing the beam transport and introduces Twiss parameters that are used to characterise the beam in terms of its size and divergence.

A particle of charge q (negative unit charge in the case of an electron) with momentum p and speed v in a perpendicular magnetic field of strength B will follow a circular trajectory of radius ρ such that

$$\frac{pv}{\rho} = Bqv\,. \tag{A.1}$$

This formula is used to define the magnetic rigidity $B\rho$ given by

$$B\rho = \frac{p}{q}\,. \tag{A.2}$$

Figure A.1 illustrates the trajectory of a charged particle moving off-axis in a perpendicular magnetic field B_y where the magnetic field is zero on axis and increases linearly with the offset. The following equations for the motion of the

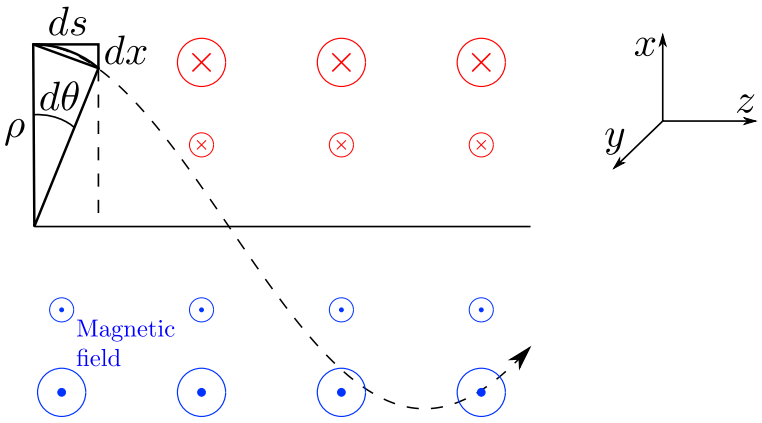

Figure A.1: Trajectory of a charged particle in a quadrupole field.

particle in the horizontal direction x can be taken from the diagram:

$$\frac{dx}{ds} = d\theta \tag{A.3}$$

$$\frac{d\theta}{ds} = \frac{1}{\rho} = -kx \tag{A.4}$$

where k is the magnetic field gradient normalised by the rigidity and is given by

$$k = -\frac{q}{p}\frac{dB_y}{dx}\,. \tag{A.5}$$

Differentiating Equations A.3 with respect to s and substituting in Equation A.4 gives the equation of motion for the particle:

$$\frac{d^2x}{ds^2} = -kx\,. \tag{A.6}$$

Equation A.6 is valid under the assumption that the particle's velocity along the longitudinal axis is much greater than its velocity in the two transverse directions, x and y, and so the trajectory angle is very small.

In the case where k is positive, Equation A.6 has the following general solution:

$$x(s) = A\cos(\sqrt{k}s) + B\sin(\sqrt{k}s) \tag{A.7}$$

$$x'(s) = -A\sqrt{k}\sin(\sqrt{k}s) + B\sqrt{k}\cos(\sqrt{k}s) \tag{A.8}$$

where A and B are constants. Applying initial conditions at $s = 0$, expressions for A and B are obtained:

$$A = x(0) \text{ and} \tag{A.9}$$

$$B = \frac{x'(0)}{\sqrt{k}} \tag{A.10}$$

Equations A.7 and A.8 can be expressed as a matrix equation

$$\begin{pmatrix} x(s) \\ x'(s) \end{pmatrix} = \begin{pmatrix} \cos(\sqrt{k}s) & \frac{1}{\sqrt{k}}\sin(\sqrt{k}s) \\ -\sqrt{k}\sin(\sqrt{k}s) & \cos(\sqrt{k}s) \end{pmatrix} \begin{pmatrix} x(0) \\ x'(0) \end{pmatrix}. \tag{A.11}$$

The matrix in Equation A.11 is known as the transfer matrix and is of determinant one [7]. It can be used for particle tracking as described in Section A.2.

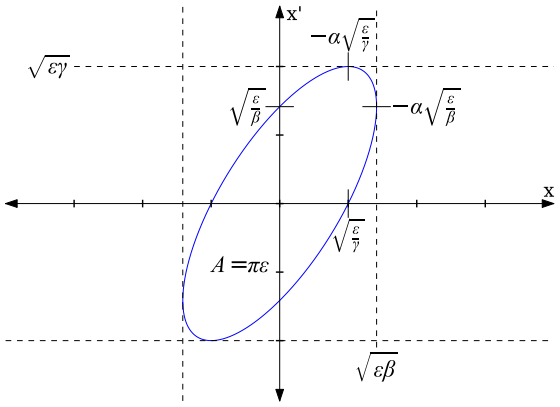

Figure A.2: Phase space ellipse with expressions for the extents and intercepts in terms of the Twiss parameters.

A straight beamline is made up of many elements with different values for k. Therefore, by replacing k with $k(s)$, a more general solution to Equation A.6 can be written as

$$x(s) = \sqrt{\epsilon\beta(s)}\cos(\psi(s) + \phi) \tag{A.12}$$

$$x'(s) = -\sqrt{\frac{\epsilon}{\beta(s)}}[\alpha(s)\cos(\psi(s) + \phi) + \sin(\psi(s) + \phi)] \tag{A.13}$$

where ϕ is a constant phase offset and

$$\alpha(s) = -\frac{\beta'(s)}{2} \text{ and} \tag{A.14}$$

$$\psi(s) = \int_0^s \frac{d\sigma}{\beta(\sigma)}\,. \tag{A.15}$$

Eliminating $\psi(s) + \phi$ from Equations A.12 and A.13 gives an equation for an ellipse in position-angle coordinates:

$$\gamma(s)x^2 + 2\alpha(s)xx' + \beta(s)x'^2 = \epsilon \tag{A.16}$$

where $\beta\gamma = 1 + \alpha^2$. An example of such an ellipse is shown in Figure A.2. ϵ is known as the beam emittance and is equal to the area of the ellipse divided by π and is conserved throughout linear beam transport. It can be defined such that the phase space ellipse represents the 1σ contour of the distribution in position-angle phase space of the particles in a bunch. $\alpha(s)$, $\beta(s)$ and $\gamma(s)$ are known as Twiss parameters and can be propagated along the beamline using the transfer matrices. As can be seen from Figure A.2, the beam size is given by $\sigma_x = \sqrt{\beta(s)\epsilon}$ [7].

The ellipse in Figure A.2 represents a diverging beam since the majority of particles with a positive position offset have a positive trajectory angle while most of the particles with a negative offset have a negative trajectory angle. The beta Twiss parameter, and therefore the beam size, is increasing. A focusing quadrupole magnet would change this around so that most particles with a negative offset have a positive trajectory angle and vice versa and the beam would

be converging [7].

A.2 Particle Tracking

In order to determine a change in position of the beam at the cavity BPM due to the deflection in an upstream corrector magnet, the beam must be tracked from one to the other. When the final quadrupole triplet is switched off, a ballistic model can be used where the change in position Δx is simply the deflection angle $\Delta x'$ multiplied by the longitudinal distance d from the corrector to the BPM:

$$\Delta x = \Delta x' d \,. \tag{A.17}$$

However, when the quadrupole triplet is switched on, a model of the effects of the quadrupole magnets on the beam trajectory must be applied. A brief derivation of the equations that are used is given in Section A.1.

Using Equation A.11 over a finite length, a matrix equation can be obtained relating the beam position and trajectory angle at the exit of a single quadrupole magnet (x_1, x_1') to the position and trajectory angle upon entering (x_0, x_0'). This is

$$\begin{pmatrix} x_1 \\ x_1' \end{pmatrix} = \begin{pmatrix} \cos(\sqrt{k}l) & \frac{1}{\sqrt{k}}\sin(\sqrt{k}l) \\ -\sqrt{k}\sin(\sqrt{k}l) & \cos(\sqrt{k}l) \end{pmatrix} \begin{pmatrix} x_0 \\ x_0' \end{pmatrix} \tag{A.18}$$

where l is the magnetic length of the quadrupole and k is the normalised strength given by Equation A.5. Equation A.18 is applicable in the case where the quadrupole magnetic field gradient is negative. In the case where it is positive, the following equation must be used instead:

$$\begin{pmatrix} x_1 \\ x_1' \end{pmatrix} = \begin{pmatrix} \cosh(\sqrt{|k|}l) & \frac{1}{\sqrt{|k|}}\sinh(\sqrt{|k|}l) \\ \sqrt{|k|}\sinh(\sqrt{|k|}l) & \cosh(\sqrt{|k|}l) \end{pmatrix} \begin{pmatrix} x_0 \\ x_0' \end{pmatrix} . \tag{A.19}$$

The 2×2 matrices in Equations A.18 and A.19 are the transfer matrices of horizontally focusing and defocusing quadrupoles respectively. Elements that

do not affect the beam trajectory can be treated as drift sections. The two-dimensional transfer matrix for a drift section R_{drift} is simply

$$R_{\text{drift}} = \begin{bmatrix} 1 & l \\ 0 & 1 \end{bmatrix} . \tag{A.20}$$

Because the model of each element is linear, a single transfer matrix for a sequence of elements can be found by multiplying the transfer matrices of the individual elements [63].

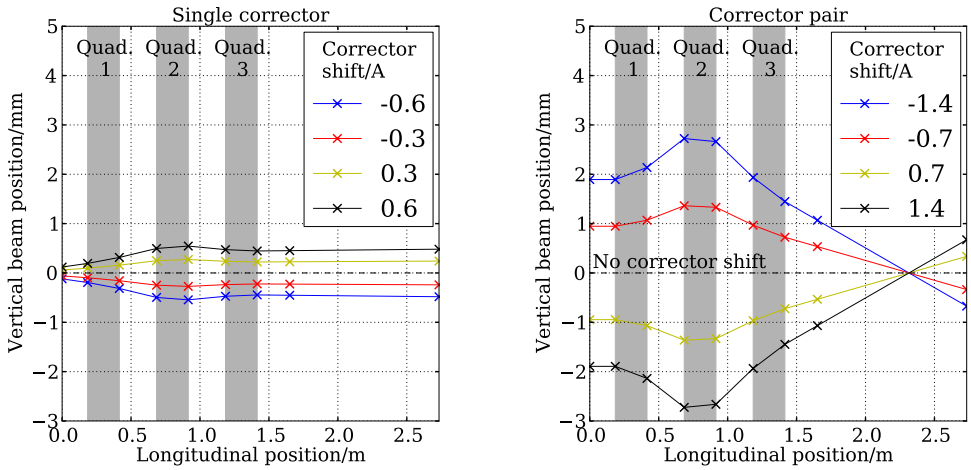

Figure A.3: Tracked beam position between the second inductive BPM in Figure 5.4 and the cavity BPM for changes in a single corrector magnet current (left) and in a corrector magnet pair (right).

Figure A.3 shows the tracking performed in order to determine the beam position at the cavity BPM. The tracking is started at the location of the second inductive BPM in Figure 5.4 since this BPM was originally used to measure the corrector angular response and there is only drift between it the final corrector upstream. The final point on each plot is at the cavity BPM location. The effects of the quadrupole magnets are clear. It can also be seen that changes in the initial position mostly lead to changes in the trajectory angle at the BPM while changes in the initial trajectory angle lead to almost pure changes in position. Therefore, in order to achieve a similar change in the final beam position using the corrector

pair, the change in the current in each corrector magnet must be much larger than when a single corrector is used.

A.3 Position Scan Results

Tables A.1 to A.5 contain the results of the individual scans in the beam position performed in order to measure the cavity BPM sensitivity. They include the parameters of the scan and the results of the fits both to the peak output voltage and peak output power.

Correctors	Steps	Half-range/mm	Charge/nC	Linear fit		Quadratic fit	
				Sensitivity/ V nC^{-1} mm^{-1}	$\frac{\chi^2}{DOF}$	Sensitivity/ V nC^{-1} mm^{-1}	$\frac{\chi^2}{DOF}$
1	7	0.58	0.0605 ± 0.0006	16.7 ± 0.2	5.94	16.72 ± 0.13	5.83
1	7	0.58	0.0609 ± 0.0006	16.7 ± 0.2	8.25	16.85 ± 0.12	8.79
1	9	0.39	0.0622 ± 0.0006	16.1 ± 0.3	12.05	16.45 ± 0.18	7.86
1	7	0.58	0.0648 ± 0.0007	17.1 ± 0.2	5.97	16.99 ± 0.13	6.14
1	7	0.58	0.0810 ± 0.0007	16.5 ± 0.2	4.52	16.81 ± 0.13	5.51
1	7	0.58	0.0826 ± 0.0006	17.35 ± 0.18	19.63	17.10 ± 0.09	18.89
2	6	0.65	0.0641 ± 0.0006	16.17 ± 0.14	4.97	16.53 ± 0.11	8.15
2	6	0.65	0.0670 ± 0.0005	16.49 ± 0.15	3.50	16.91 ± 0.10	8.21
2	6	0.65	0.0641 ± 0.0005	16.85 ± 0.16	10.11	16.79 ± 0.11	10.13
2	6	0.65	0.0618 ± 0.0006	16.47 ± 0.17	2.84	16.73 ± 0.10	4.18
2	6	0.65	0.0613 ± 0.0006	16.7 ± 0.2	1.39	16.88 ± 0.12	1.88
2	6	0.65	0.0600 ± 0.0006	16.3 ± 0.2	3.74	16.62 ± 0.10	5.01

Table A.1: Summary of the position scans with short beam pulses used to measure the position cavity sensitivity in the horizontal direction.

Correctors	Steps	Half-range/mm	Charge/nC	Linear fit		Quadratic fit	
				Sensitivity/ V nC^{-1} mm^{-1}	$\frac{\chi^2}{DOF}$	Sensitivity/ V nC^{-1} mm^{-1}	$\frac{\chi^2}{DOF}$
1	7	0.58	0.0731 ± 0.0006	17.2 ± 0.7	2.06	17.3 ± 0.7	1.82
1	7	0.58	0.0762 ± 0.0006	15.4 ± 0.4	0.95	16.3 ± 0.4	1.03
1	6	0.49	0.0610 ± 0.0006	15.4 ± 0.5	0.41	15.4 ± 0.4	0.39
1	6	0.49	0.0598 ± 0.0005	15.2 ± 0.4	13.02	16.0 ± 0.4	7.24
1	5	0.39	0.0672 ± 0.0006	14.5 ± 0.5	0.72	16.3 ± 0.4	0.80
1	5	0.39	0.0603 ± 0.0007	17.0 ± 0.4	3.97	17.9 ± 0.4	4.71
1	5	0.39	0.0618 ± 0.0007	14.0 ± 0.6	0.37	16.2 ± 0.5	3.03
2	9	0.52	0.0698 ± 0.0006	16.8 ± 0.7	2.67	17.5 ± 0.5	1.46
2	9	0.52	0.0634 ± 0.0005	15.3 ± 0.8	11.44	15.9 ± 0.5	6.02
2	9	0.52	0.0584 ± 0.0005	15.5 ± 0.8	9.15	15.8 ± 0.5	4.55

Table A.2: Summary of the position scans with short beam pulses used to measure the position cavity sensitivity in the vertical direction.

Iteration	X		Y	
	Energy output/J nC^{-2} mm^{-2}	χ^2/DOF	Energy output/J nC^{-2} mm^{-2}	χ^2/DOF
1	$(2.34 \pm 0.04) \times 10^{-8}$	6.51	$(2.6 \pm 0.2) \times 10^{-8}$	2.13
2	$(2.35 \pm 0.04) \times 10^{-8}$	7.72	$(2.36 \pm 0.10) \times 10^{-8}$	1.58
3	$(2.27 \pm 0.05) \times 10^{-8}$	5.70	$(2.09 \pm 0.11) \times 10^{-8}$	0.49
4	$(2.40 \pm 0.04) \times 10^{-8}$	6.22	$(2.23 \pm 0.09) \times 10^{-8}$	7.28
5	$(2.37 \pm 0.03) \times 10^{-8}$	6.63	$(2.35 \pm 0.11) \times 10^{-8}$	0.78
6	$(2.42 \pm 0.03) \times 10^{-8}$	21.49	$(2.80 \pm 0.14) \times 10^{-8}$	4.80
7	$(2.27 \pm 0.03) \times 10^{-8}$	7.25	$(2.33 \pm 0.13) \times 10^{-8}$	2.92
8	$(2.38 \pm 0.03) \times 10^{-8}$	7.83	$(2.70 \pm 0.16) \times 10^{-8}$	1.58
9	$(2.34 \pm 0.03) \times 10^{-8}$	12.65	$(2.24 \pm 0.14) \times 10^{-8}$	7.21
10	$(2.33 \pm 0.03) \times 10^{-8}$	4.46	$(2.22 \pm 0.14) \times 10^{-8}$	4.15
11	$(2.37 \pm 0.03) \times 10^{-8}$	1.51	-	-
12	$(2.30 \pm 0.03) \times 10^{-8}$	4.87	-	-

Table A.3: Summary of the position sensitivity measurements in both transverse directions in terms of signal energy for short beam pulses.

Correctors	Steps	Half-range/mm	Charge/nC	Linear fit		Quadratic fit	
				Sensitivity/ V nC^{-1} mm^{-1}	$\frac{\chi^2}{DOF}$	Sensitivity/ V nC^{-1} mm^{-1}	$\frac{\chi^2}{DOF}$
1	6	0.97	3.476 ± 0.013	134.6 ± 0.7	81.57	130.7 ± 0.4	114.18
1	6	0.97	3.509 ± 0.013	133.9 ± 0.8	65.44	130.3 ± 0.5	77.44
1	6	0.97	3.538 ± 0.013	134.5 ± 0.8	100.11	129.5 ± 0.4	138.73
1	6	0.97	3.492 ± 0.012	136.3 ± 0.7	98.51	133.5 ± 0.5	124.23
1	6	0.97	3.540 ± 0.010	132.2 ± 0.8	102.55	130.7 ± 1.6	110.48
1	6	0.97	3.528 ± 0.011	130.3 ± 0.7	128.53	129.0 ± 1.9	135.95
1	6	0.97	3.421 ± 0.013	134.4 ± 0.7	71.51	130.8 ± 0.4	109.36
2	6	0.65	3.462 ± 0.012	128.3 ± 0.6	5.82	128.4 ± 0.6	6.09
2	6	0.65	3.478 ± 0.009	127.6 ± 0.6	11.33	127.4 ± 0.4	11.84
2	6	0.65	3.498 ± 0.013	130.4 ± 0.6	1.12	131.5 ± 0.6	1.62
2	6	0.65	3.411 ± 0.011	127.5 ± 0.6	5.02	128.6 ± 0.6	1.74
2	6	0.65	3.486 ± 0.011	128.8 ± 0.8	4.17	127.5 ± 0.8	0.84
2	6	0.65	3.525 ± 0.010	127.5 ± 0.8	6.36	128.5 ± 0.7	2.77
2	6	0.65	3.510 ± 0.011	130.5 ± 0.8	7.48	130.2 ± 0.7	5.68

Table A.4: Summary of the position scans with long beam pulses used to measure the position cavity sensitivity in the horizontal direction.

Correctors	Steps	Half-range/mm	Charge/nC	Energy gradient/mJ nC^{-2} mm^{-2}	χ^2/DOF
1	6	0.97	3.470 ± 0.014	$(9.91 \pm 0.07) \times 10^{-3}$	68.62
1	6	0.97	3.504 ± 0.013	$(9.85 \pm 0.09) \times 10^{-3}$	35.46
1	6	0.97	3.532 ± 0.013	$(9.82 \pm 0.07) \times 10^{-3}$	63.16
1	6	0.97	3.486 ± 0.012	$(1.019 \pm 0.008) \times 10^{-2}$	58.58
1	6	0.97	3.540 ± 0.011	$(9.9 \pm 0.3) \times 10^{-3}$	64.00
1	6	0.97	3.527 ± 0.011	$(9.4 \pm 0.3) \times 10^{-3}$	96.02
1	6	0.97	3.417 ± 0.014	$(9.90 \pm 0.08) \times 10^{-3}$	43.61
2	6	0.65	3.456 ± 0.013	$(9.65 \pm 0.10) \times 10^{-3}$	8.74
2	6	0.65	3.472 ± 0.010	$(9.50 \pm 0.07) \times 10^{-3}$	9.28
2	6	0.65	3.492 ± 0.013	$(1.017 \pm 0.011) \times 10^{-2}$	3.00
2	6	0.65	3.405 ± 0.011	$(9.68 \pm 0.09) \times 10^{-3}$	0.77
2	6	0.65	3.479 ± 0.011	$(9.53 \pm 0.11) \times 10^{-3}$	3.80
2	6	0.65	3.519 ± 0.011	$(9.60 \pm 0.12) \times 10^{-3}$	2.01
2	6	0.65	3.504 ± 0.011	$(9.92 \pm 0.11) \times 10^{-3}$	6.67

Table A.5: Summary of the horizontal sensitivity measurements with long beam pulses in terms of the total signal energy.

References

[1] R. Ruber. CTF3 Two-Beam Test Stand Illustrations. http://ctf3-tbts.web.cern.ch/ctf3-tbts/images/, 2011. 10, 16, 27, 148

[2] S. Boogert. personal communication, 2013. 10, 41, 43

[3] J. R. Towler, T. Lefevre, L. Søby, M. Wendt, S. T. Boogert, F. J. Cullinan, and A. Lyapin. Technologies and R&D for a High Resolution Cavity BPM for the CLIC Main Beam. In *IBIC 2013*, number TUPC20, 2013. 16, 75, 153, 154, 161, 208

[4] W. Farabolini. personal communication, 2013. 20, 148, 198, 217

[5] M. Aicheler, P. Burrows, M. Draper, T. Garvey, P. Lebrun, K. Peach, N. Phinney, H. Schmickler, D. Schulte, and N. Toge. A Multi-TeV Linear Collider Based on CLIC Technology: CLIC Conceptual Design Report. Technical Report CERN-2012-007, CERN, Geneva, 2012. 21, 26, 27, 28, 31, 32, 33, 36, 61, 64, 101

[6] R. Ruber, V. Ziemann, T. Ekelf, A. Palaia, W. Farabolini, and R. Corsini. The CTF3 Two-beam Test Stand. *Nuclear Instruments and Methods in Physics Research A*, 729(0):546 – 553, 2013. 22, 37, 147

[7] K. Wille. *The Physics of Particle Accelerators: An Introduction*, pages 44–46,65–81. Oxford University Press, New York, 2000. 28, 220, 221, 222

[8] N Merminga, P. L. Morton, J. T. Seeman, and W. L. Spence. Transverse Phase Space in the Presence of Dispersion. In *PAC 1991*, 1991. 29

[9] D. Schulte. Different Options for Dispersion Free Steering in the CLIC Main Linac. In *PAC 2005*, number RPPP011, 2005. 30

[10] F. Peauger, W. Farabolini, P. Girardot, A. Andersson, G. Riddone, A. Samoshkin, A. Solodko, R. Zennaro, and R. Ruber. Wakefield Monitor Development for CLIC Accelerating Structure. In *LINAC 2010*, number TUP098, 1998. 30

[11] J. Pfingster, J. Snuverink, A. Latina, and D. Schulte. Recent Improvements of the Orbit Controller and Ground Motion Mitigation Techniques for CLIC. In *IPAC 2012*, number TUPPR028, 2012. 31

[12] S. Walston, S. Boogert, C. Chung, P. Fitsos, J. Frisch, J. Gronberg, H. Hayano, Y. Honda, Y. Kolomensky, A. Lyapin, S. Malton, J. May, D. McCormick, R. Meller, D. Miller, T. Orimoto, M. Ross, M. Slater, S. Smith, T. Smith, N. Terunuma, M. Thomson, J. Urakawa, V. Vogel, D. Ward, and G. White. Performance of a High Resolution Cavity Beam Position Monitor System. *Nuclear Instruments and Methods in Physics Research A*, 578:1–22, 2007. 31, 45, 46, 47, 98

[13] R. E. Shafer. Beam position monitoring. In *American Institute of Physics Conference Series*, volume 212, pages 26–58, October 1990. 32, 33, 35, 36

[14] A. Benot-Morrell, L. Søby, A. Faus-Golfe, J. J. Garcia-Garrigós, J. M. Nappa, J. Tassan-Viol, S. Vilalte, and S. Smith. Characterization Tests of a Stripline Beam Position Monitor for the CLIC Drive Beam. In *IPAC 2012*, number MOPPR042, 2012. 35

[15] P. Urschütz, H. H. Braun, R. Corsini, S. Döbert, E. Jensen, and F. Tecker. Efficient Long-Pulse Fully Loaded CTF3 LINAC Operation. In *LINAC 2006*, number MOP002, 2006. 36

[16] R. Corsini. Experimental Verification of the CLIC Two-Beam Scheme, Status and Outlook. In *IPAC 2012*, number TUOBC01, 2012. 36, 37

[17] P. K. Skowroński, J. Barranco, S. Bettoni, R. Corsini, M. Divall Csatari, A. E. Dabrowski, S. Döbert, A. Dubrovskiy, O. Kononenko, M. Olvegaard,

A. Rabiller, F. Tecker, W. Farabolini, T. Persson, R. L. Lillestol, E. Adli, A. Palaia, and R. Ruber. The CLIC Feasibility Demonstration in CTF3. Technical Report CERN-ATS-2011-177, CERN, 2011. 37, 147

[18] D. M. Pozar. *Microwave Engineering*, pages 6,124,284–286. Wiley, New Jersey, 3 edition, 2005. 40

[19] E. L. Ginzton. *Microwave Measurements*, pages 349,347. McGraw-Hill, New York. 42

[20] Z. Li, R. Johnson, S. Smith, T. Naito, and J. Rifkin. Cavity BPM with Dipole-Mode-Selective Coupler. In *PAC 2003*, 2003. 43

[21] P. B. Wilson. Introduction to Wakefields and Wake Potentials. Technical Report SLAC-PUB-4547, SLAC, 1 1989. 44, 46

[22] M. Ross, J. Frisch, D. McCormick, and H. Hayano. RF Cavity BPMs as Beam Angle and Beam Correlation Monitors. In *PAC 2003*, number WPB062, 2003. 47

[23] W. Bruns. GdfidL: A Finite Difference Program with Reduced Memory and CPU Usage. In *PAC 1997*, number 9P118, 1997. 48

[24] K. Ko, A. Candel, L. Ge, A. Kabel, R. Lee, Z. Li, C. Ng, V. Rawat, G. Schussman, and L. Xiao. Advances in Parallel Electromagnetic Codes for Accelerator Science and Development. In *LINAC 2010*, number FR101, 2010. 48, 69, 73

[25] R. Garg. *Analytical and Computational Methods in Electromagnetics*, pages 233–280,355–444. Artech House, Boston, 2008. 51

[26] D. H. Whittum and Y. Kolomensky. Analysis of an Asymmetric Resonant Cavity as a Beam Monitor. *Review of Scientific Instruments*, 70(5):2300–2313, 1999. 51, 92

[27] W. H. Hayt and J. E. Kemmerly. *Engineering Circuit Analysis*, pages 431–450. McGraw-Hill, New York, 2 edition. 52, 71

[28] Fundamentals of RF and Microwave Noise Figure Measurements. Technical Report Application Note 57-1, Agilent Technologies, 2010. 55

[29] F. Ramian. Intermodulation Distortion Measurements on Modern Spectrum Analyzers. Technical Report Application Note 1EF79, Rohde & Schwarz, 2012. 57

[30] A. Lunin. High Resolution RF Cavity BPM Design for Linear Collider. In *8th DITANET Topical Workshop on Beam Position Monitors*, CERN, Switzerland, 1 2012. 58, 59, 91

[31] N. Chritin, H. Schmickler, L. Søby, A. Lunin, N. Solyak, M. Wendt, and V. Yakovlev. A High-Resolution Cavity BPM for the CLIC Test Facility. In *BIW 2010*, number TUPSM32, 2010. 58, 60

[32] N.Y. Joshi. *Design and Analysis Techniques for Cavity Beam Position Monitor Systems for Electron Accelerators*. PhD thesis, Royal Holloway, University of London, 2013. 74

[33] Pico Technology: PC Oscilloscope and Data Acquisition Products. http://www.picotech.com/, 2014. 86

[34] Agilent Technologies. http://www.home.agilent.com/en/pd-1184897-pn-U1065A/acqiris-10-bit-high-speed-cpci-digitizers, 2012. 102, 126

[35] S. Smith. Low vs. High Q Cavity BPMs. In *Post-IBIC Cavity BPM Workshop*, JAI, Oxford, 9 2013. 102, 208, 212

[36] Hittite Microwave Corporation. http://www.hittite.com/, 2012. 103, 109, 110, 112, 128

[37] Mini-circuits. http://www.minicircuits.com/, 2012. 103, 104, 108, 110

[38] Welcome to wfProc's Documentation. http://wfproc.sourceforge.net/, 2012. 104, 107

[39] RLC Electronics. http://www.rlcelectronics.com/, 2012. 107

[40] P. Smith. Little Known Characteristics of Phase Noise. Technical Report Application Note AN-741, Analogue Devices. 127, 211

[41] Y. I. Kim, R. Ainsworth, A. Aryshev, S. T. Boogert, G. Boorman, J. Frisch, A. Heo, Y. Honda, W. H. Hwang, J. Y. Huang, E-S. Kim, S. H. Kim, A. Lyapin, T. Naito, J. May, D. McCormick, R. E. Mellor, S. Molloy, J. Nelson, S. J. Park, Y. J. Park, M. Ross, S. Shin, C. Swinson, T. Smith, N. Terunuma, T. Tauchi, J. Urakawa, and G. R. White. Cavity beam position monitor system for the accelerator test facility 2. *Phys. Rev. ST Accel. Beams*, 15:042801, Apr 2012. 135, 136, 162, 164, 166, 173

[42] G. E. Carlson. *Signal and Linear System Analysis*, pages 178–180, 688–690. Wiley, New York, 2 edition, 1998. 137

[43] S. Smith. Beam Position Monitors for the CLIC Drive Beam. In *8th DITANET Topical Workshop on Beam Position Monitors*, CERN, Switzerland, 1 2012. 138

[44] Diagsim. http://sourceforge.net/projects/diagsim/, 2012. 139

[45] M. Gasior. An Inductive Pick-Up for Beam Position and Current Measurements. In *DIPAC 2003*, number CT01, 2003. 149

[46] W. Farabolini, A. Guillaume, P. Girardot, R. Granelli, F. Harrault, C. Lahonde-Hamdoun, T. Lerch, and F. Orsini. Video Profile Monitor Development for the CTF3 Probe Beam Linac. In *EPAC 2008*, number TUPC024, 2008. 150

[47] C. Simon, S. Chel, P. Contrepois, P. Girardot, M. Luong, N. Baboi, and N. Rouvière. Beam Position Monitors Using a Re-entrant Cavity. In *DIPAC 2007*, number TUPB15, 2007. 151

[48] Newport. http://www.newport.com/, 2013. 153

[49] National Instruments Labview System Design Software. http://www.ni.com/labview/, 2014. 154

[50] Oasis. https://project-oasis.web.cern.ch/project-oasis/, 2013. 154

[51] CMW: Controls Middleware. http://proj-cmw.web.cern.ch/proj-cmw/, 2003. 154

[52] M. Arruat, L. Fernandez, S. Jackson, F. Locci, J-L. Nougaret, M. Peryt, A. Radeva, M. Sobczak, and M. V. Eynden. Front-End Software Architecture. In *ICALEPS 2007*, number WOPA04, 2007. 155

[53] J. Laskar. Frequency Map Analysis and Particle Accelerators. In *PAC 2003*, number WOAB001, 2003. 162, 164, 165

[54] PyMinuit. https://code.google.com/p/pyminuit/, 2013. 163

[55] F. James and M. Roos. Minuit - A System for Function and Minimization and Analysis of the Parameter Errors and Correlations. *Computer Physics Communications*, 10:343–367, 1975. 163

[56] R. Fandos. RF Design of a New Precision Beam Position Monitor. Technical Report EUROTeV-Report-2008-033, EUROTeV, 10 2008. 208

[57] B. Keil, R. Baldinger, R. Ditter, W. Koprek, R. Kramert, F. Marcellini, G. Marinkovic, M. Roggli, M. Rohrer, M. Stadler, and D. M. Treyer. Design of the SwissFEL BPM System. In *IBIC 2013*, number TUPC25, 2013. 209, 212

[58] R. D. Monte, A. O. Borga, P. Craievich, M. Ferianis, and G. Gaio. The Fermi@Elettra Cavity BPM System: Description and Commissioning Results. In *DIPAC 2011*, number MOOC03, 2011. 210

[59] J. Resta-López and P. N. Burrows. Layout and Simulations of the FONT System at ATF2. Technical Report EUROTeV-Report-2008-053, EUROTeV, 8 2008. 215

[60] N. Delarue. FEATHER: A Fast Intra-Pulse Feedback System for the JLC. *arXiv*, 5 2003. 215

[61] W. Farabolini, D. Bogard, A. Brabant, A. Curtoni, P. Girardot, F. Gobin, R. Granelli, F. Harrault, C. Lahonde, T. Lerch, M. Luong, A. Mosnier,

F. Orsini, F. Peauger, and C. Simon. Diagnostics for the CTF3 Probe Beam Linac CALIFES. In *DIPAC 2007*, number TUPB17, 2007. 216

[62] P. Emma, J. Frisch, and P. Krejcik. A Transverse RF Deflecting Structure for Bunch Length and Phase Space Diagnostics. Technical Report LCLS-TN-00-12, SLAC, 8 2000. 216

[63] S. Y. Lee. *Accelerator Physics*, pages 35–48,51. World Scientific, Singapore, 2 edition, 2004. 223