Measuring the Electron Beam Energy in a Magnetic Bunch Compressor

Dissertation

zur Erlangerung des Doktorgrades

des Department Physik

der Universitaet Hamburg

vorgelegt von

Kirsten Hacker
aus Minneapolis, USA

Hamburg

2010

Gutachter der Dissertation:	Joerg Rossbach Bernhard Schmidt
Gutachter der Disputation:	Joerg Rossbach Eckhard Elsen
Datum der Disputation:	16.09.10
Vorsitzende des Pruefungsausschusses:	Caren Hagner
Vorsitzender des Promotionsausschusses:	Jochen Bartels
Dekanin des Fachbereichs Physik:	Heinrich Graener

Abstract

Within this thesis, work was carried out in and around the first bunch compressor chicane of the FLASH (Free-electron LASer in Hamburg) linear accelerator in which two distinct systems were developed for the measurement of an electron beam's position with sub-5 μm precision over a 10 cm range. One of these two systems utilized RF techniques to measure the difference between the arrival-times of two broadband electrical pulses generated by the passage of the electron beam adjacent to a pickup antenna. The other system measured the arrival-times of the pulses from the pickup with an optical technique dependent on the delivery of laser pulses which are synchronized to the RF reference of the machine. The relative advantages and disadvantages of these two techniques are explored and compared to other available approaches to measure the same beam property, including a time-of-flight measurement with two beam arrival-time monitors and a synchrotron light monitor with two photomultiplier tubes.

The electron beam position measurement is required as part of a measurement of the electron beam energy and could be used in an intra-bunch-train beam-based feedback system that would stabilize the amplitude of the accelerating field. By stabilizing the accelerating field amplitude, the arrival-time of the electron beam can be made more stable. By stabilizing the electron beam arrival-time relative to a stable reference, diagnostic, seeding, and beam-manipulation lasers can be synchronized to the beam.

Zusammenfassung

Im Rahmen dieser Doktorarbeit wurden an dem Freien Elektronen Laser FLASH, in Hamburg, zwei unterschiedliche Techniken zur Vermessung der transversalen Elektronenstrahlposition in magnetischen Schikanen mit einer Auflösung von 5 μm über einen 10 cm breiten Meßbereich, entwickelt. Eine diese Technik basiert auf der Bestimmung der Ankunftszeiten zweier kurzer elektrischer Signale, welche beim passieren des Elektronenstrahls an einer Hochfrequenzantenne erzeugt werden, mittels Hochfrequenzelektronik. Die zweite Technik verwendet kurze Laserpulse, die zur Hochfrequenz des Beschleunigers synchronisiert sind, um die Ankunftszeiten der elektrischen Antennensignale mit hoher Präzision zu ermittelt. Die Vor- und Nachteile dieser beiden Methoden werden in dieser Arbeit theoretisch und experimentell untersucht und verglichen mit anderen Methoden, wie zum Beispiel, der Detektion der Flugzeitdifferenzen des Elektronenstrahls durch die magnetische Schikane oder der Positionsbestimmung der Elektronenpakete durch optische Synchrotronstrahlung.

Die Messung der transversalen Elektronenstrahlposition in einer magnetischen Schikane ist ein direktes Maß für die Elektronenstrahlenergie und kann für ein schnelles Regelungssystem zur Stabilisierung der Beschleunigungsgradienten genutzt werden. Durch die Stabilisierung der Beschleunigergradienten kann eine Stabilisierung der Ankunftszeit des Elektronenstrahls relative zur Synchronisationsreferenz der Anlage erzielt werden. Dies verbessert entscheidend Experimente die auf das seeden oder manipulieren des Elektronenstrahls durch externe Laserstrahlen angewiesen sind.

Contents

List of Tables

List of Figures

8.3.6	RF phase measurement drift without temperature control, without disturbances (people in room).
8.3.7	Out-of-loop measurement drift without temperature control and without disturbances (people in room).
9.1.1	A synchrotron light monitor system with CCD screen.
9.1.2	A picture of the beam as imaged with the synchrotron light camera.
9.2.1	Two Photomultipliers used to measure the beam position in the chicane.
10.1.1	Measurements of energy stability in the chicane taken by the coarse and fine HF front-ends of the chicane BPM plotted with energy setpoint values from the upstream accelerating module.
10.2.1	Correlation between the measurements of the beam position in the chicane taken by the chicane BPM (labeled EBPM) and the photomultiplier tube monitor (PMT).
10.2.2	Fine HF front-ends position measurement and photomultiplier tube position measurement in good agreement.
10.2.3	Fine HF front-ends position measurement and photomultiplier tube position measurement in poor agreement.
10.3.1	Fine HF front-ends position measurement and photomultiplier tube position measurement.
10.4.1	Optical (EOM) front-end position measurement and photomultiplier tube position measurement along with a time of flight measurement involving 2 BAMs and a line showing how the setpoint of the gradient changed.
10.4.2	The beam energy was changed by 0.3 % with the accelerator gradient setpoint and the beam energy measured by the chicane BPM changed by a comparable amount.
10.4.3	Optical (EOM) front-end chicane BPM measurement and photomultiplier tube BPM measurement along with a time-of-flight measurement involving 2 BAMs and a line showing how the setpoint of the gradient predicted an energy change of 0.1%.
10.4.4	The beam-arrival time upstream of the chicane measured with both the transversely mounted stripline BPM installed in the chicane and with a button-type pickup BAM installed upstream of the chicane.
10.4.5	Optical (EOM) front-end position measurement, 10.4 GHz front-end measurement, photomultiplier tube position measurement, time-of-flight measurement involving 2 BAMs and the setpoint of the gradient are plotted together over several hours.

1 Introduction

Knowing the exact position of an electron beam under the influence of a magnetic field has been at the heart of many important experiments in the history of physics. From the first cathode ray tube that was placed next to a piece of magnetized metal to the high-energy beams of modern accelerators traveling through lattices of powerful electromagnets, the position of the beam under the influence of a magnetic field gives information about the momentum of the beam. High-precision knowledge of the beam momentum can enable higher precision control of the beam. Control of the beam momentum is critical for the stability of both the wavelength and arrival-time of the light pulses generated by free-electron lasers. The measurement of the beam momentum in a free-electron laser with magnetic bunch compressor chicanes is the topic of this thesis.

The bending radius, r, of an electron with charge e traveling through a magnetic dipole field perpendicular to the beam direction, B, depends on the momentum, p, of the particle.

$$\frac{1}{r} = e \cdot \frac{\mathrm{B}}{p} \tag{1.1}$$

is an equation derived from the Lorentz force law, $\vec{\mathbf{F}} = q(\vec{\mathbf{v}} \times \vec{\mathbf{B}})$, and the relation between force and momentum $F=dp/dt=pv/r$. For a rectangular dipole magnet in which the electron beam enters perpendicularly to one of the magnet's faces, the path length of the electron's trajectory is given by

$$l_{arc} = r \cdot \alpha, \tag{1.2}$$

where α is the bending angle introduced by the magnet. This allows us to write the effective length of the dipole, l_{eff}, in terms of the bending angle,

$$l_{eff} = r \cdot \sin\alpha \,, \tag{1.3}$$

As well as the x offset of the particle from a straight ahead trajectory,

$$x_{offset} = l_{eff} \cdot \frac{1-\cos\alpha}{\sin\alpha} \quad . \tag{1.4}$$

We can use equations 1.3 and 1.4 to describe the situation that one finds in a magnetic bunch compressor chicane where, not one, but four magnetic fields are applied, two of which cause the beam to deviate from its straight-ahead trajectory and two of which bring the beam back to its original trajectory (Fig. 1.1).

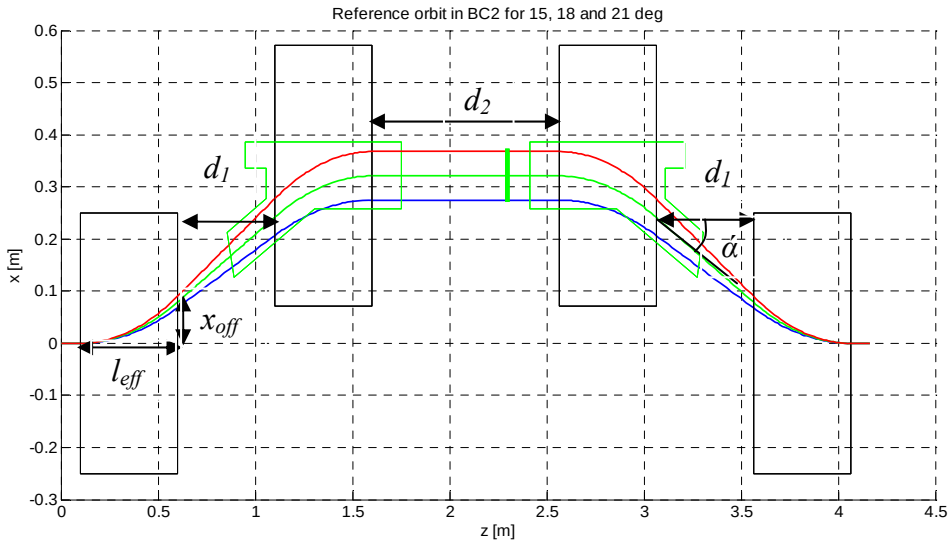

Figure 1.1 Magnetic bunch compressor chicane. High-energy particles travel a shorter path than low-energy particles.

The path-length of an electron traveling through the symmetric chicane shown above can be expressed by

$$l_{BC} = 4r\alpha + 2 \cdot \frac{d_1}{\cos\alpha} + d_2 \tag{1.5}$$

and the x position in the middle of the chicane is

$$x_{BC} = 2 \cdot r \sin\alpha \cdot \frac{1-\cos\alpha}{\sin\alpha} + d_1 \tan\alpha \tag{1.6}$$

where d_1 is the drift space between the first and second dipoles and d_2 is the drift space between the second and third dipoles. If these equations are rewritten in terms of the magnetic field and the momentum of the electrons (Eqs. 1.1 and 1.3), we have

$$l_{BC}(p) = \frac{4p}{eB}\arcsin\left(\frac{l_{eff}eB}{p}\right) + \frac{2d_1}{\sqrt{1-\left(l_{eff}eB/p\right)^2}} + d_2$$

and (1.6)

$$x_{BC}(p) = \frac{2B}{ep}\left(1-\sqrt{1-\left(l_{eff}eB/p\right)^2}\right) + \frac{l_{eff}eBd_1}{p\sqrt{1-(l_{eff}eB/p)^2}}.$$

Because high-momentum particles' trajectories are bent less than those of low-momentum particles, the high-momentum particles will travel a shorter path through the chicane and will arrive at the end of the chicane earlier than the low-momentum particles. We would like to know how a momentum change of a group of particles will affect the path-length through the chicane and the horizontal position in the chicane. To find this, we will Taylor expand Eq.s 1.5 in terms of a small change in momentum $\delta=\Delta p/p$. This is given by

$$\begin{aligned} l_{BC}(p(1+\delta)) &= l_{BC} + p\frac{\partial l_{BC}}{\partial p}\cdot\delta + \frac{p^2}{2}\frac{\partial^2 l_{BC}}{\partial p^2}\cdot\delta^2 + \ldots \\ &= l_{BC} + R_{56}\cdot\delta + R_{566}\cdot\delta^2 + \ldots \end{aligned}$$

(1.7)

$$\begin{aligned} x_{BC}(p(1+\delta)) &= x_{BC} + p\frac{\partial x_{BC}}{\partial p}\cdot\delta + \frac{p^2}{2}\frac{\partial^2 x_{BC}}{\partial p^2}\cdot\delta^2 + \ldots \\ &= x_{BC} + R_{16}\cdot\delta + R_{166}\cdot\delta^2 + \ldots \end{aligned}$$

where R_{56}, R_{566} R_{16} and R_{166} are functions of the magnetic field and the effective length of each magnet. They are named after their locations in a transfer matrix used to calculate beam transport and they are used to predict the arrival-time and x position changes of particles traveling through dipole fields for given momentum changes. For short bunches, like those in FELs, the R_{566} and R_{166} terms are typically small compared to the R_{56} and R_{16} terms. So, for the majority of the calculations in this thesis, only the first-order terms will be used. The first order terms are frequently referred to as momentum compaction ($R_{56}=\acute{\alpha}_c$) and linear dispersion ($R_{16}=D$). Doing the first order derivatives of Eqs. 1.6, one finds

$$R_{56} = \frac{-4l_{eff}}{\sqrt{1-(eBl_{eff}/p)^2}} - \frac{2d_1\left(eBl_{eff}/p\right)^2}{\sqrt[3]{1-(eBl_{eff}/p)^2}} + \frac{4p}{eB}\arcsin\left(\frac{eBl_{eff}}{p}\right)$$

(1.8)

$$R_{16} = \frac{d_1\left(eBl_{eff}/2\right)}{\sqrt[3]{1-(eBl/p)^2}} + \frac{2p}{eB}\left(\frac{1}{\sqrt{1-(eBl_{eff}/p)^2}} - 1\right)$$

which, by substituting cosα wherever the square root of $1-(eBl_{eff}/p)^2$ arises, can be written more simply in terms of the bend angle and the effective dipole length.

$$R_{56} = 2d_1\frac{\tan^2\alpha}{\cos\alpha} + \frac{4l_{eff}}{\sin\alpha}(\tan\alpha - \alpha)$$

$$R_{16} = d_1\frac{\tan\alpha}{\cos^2\alpha} + \frac{2l_{eff}}{\sin\alpha}\left(\frac{1}{\cos\alpha} - 1\right) \qquad (1.9)$$

Both R_{56} and R_{16} increase when the length of the chicane increases and when the bending angle increases. With a rough, small-angle approximation, $R_{56} \approx 2 \cdot R_{16} \cdot \alpha$. Although the results for the single-chicane are the same as those presented here, a different and more complete treatment of dispersion and momentum compaction in different types of chicanes is given in [1].

Because the particles in the chicanes are relativistic, we can use the approximation $\Delta p/p \approx \Delta E/E$ and will use the term 'relative momentum changes' synonymously with 'relative energy changes'. This allows us to write the useful formulas for the change in path length, Δl, and change in horizontal position, Δx, that occur as a result of momentum compaction and dispersion,

$$\Delta l = R_{56}\frac{\Delta E}{E} \qquad \Delta x = R_{16}\frac{\Delta E}{E} \qquad (1.10)$$

Because the momentum compaction and dispersion in a chicane are non-zero, changes in the x position of the beam within the chicane correlate with changes in the energy of the beam. Likewise, the path-length changes in and after the chicane will result in beam arrival-time changes which are correlated with beam energy changes. If the electron beam after a chicane is used to generate a photon beam, as in the case of a Free-electron Laser (FEL), the arrival-time stability of the photon beam will be directly affected by the stability of the beam energy, a quantity which can be measured through the position or arrival-time of the beam in or after a dispersive section.

To detect the position of an electron beam, one could look at the optical transition radiation that is produced as the particles travel through a metal film, but the beam is significantly disturbed by the film and it cannot be used parasitically. Instead, one could look at the synchrotron light that is produced by the beam as its trajectory is changed by the magnetic field. The resolution of the synchrotron light based beam position measurement would then depend on how much light there is and how well it can be detected. If, instead, the position of the beam is detected with a metal antenna in which a current is induced as the beam passes close to it, the resolution would depend on how

well the beam couples to the antenna and how well the properties of the current pulses generated in the antenna can be measured.

In this thesis, the design and characterization of a pickup antenna and pulse measurement system was completed in order to measure the position of an electron beam over a range of 10 cm and with a resolution below 2 μm. Expanding the range, even to the 40 cm required by the XFEL chicanes, should also be possible with this design. Typical beam position measurements achieve few-micron resolution over a few-millimeter range. The monitor system developed within this thesis work represents a new and unique tool in the spectrum of accelerator diagnostics.

Several types of pickups were considered for this task: a pin that detects the ringing in a cavity, an array of closely spaced striplines oriented in the direction of the beam, and a rod in a coaxially shaped channel oriented perpendicularly to the direction of the beam. In the end, the perpendicularly oriented pickup was constructed and utilized with two different types of pulse measurement systems, one of which made use of high-frequency electronic techniques and the other which made use of an optical technique involving Mach-Zehnder electro-optical modulators and the pulses from a master laser oscillator. Both techniques delivered the required resolution, but each had unique advantages and disadvantages. Particular emphasis is placed on the cost and robustness of the high-frequency electronic system compared to the lower potential for systematic error of the optical system.

Understanding the influence of the shape of the beam on chicane beam position measurements is also critical. For this reason, simulations of the beam transport from the start of the accelerator up to the middle of the chicane were undertaken in order to predict the likely beam properties at the Free-electron LASer in Hamburg, FLASH. The impacts of beam shape, charge, position jitter, and accelerating RF properties were also investigated. Measurements were undertaken to verify the predictions of these simulations and benchmark the measurements against those of an existing synchrotron light detection scheme and a beam arrival-time monitor scheme.

2 Free-electron Lasers

The facility at which the experiments were performed is the Free-electron LASer in Hamburg (FLASH). FLASH consists of an RF photo-injector electron source, superconducting RF accelerator sections, two magnetic bunch compressors, and an undulator section (Figure 2.1).

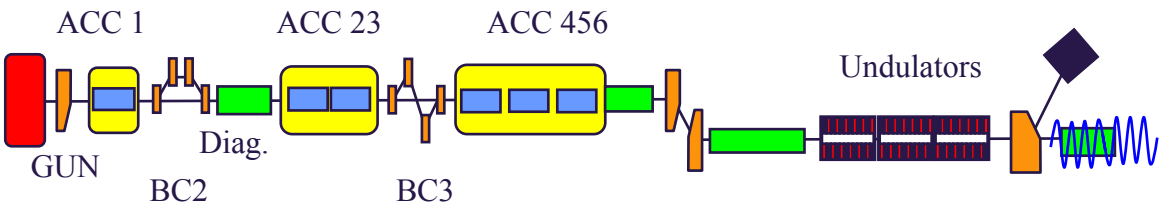

Figure 2.1 Free-electron LASer in Hamburg (FLASH) including RF gun (GUN), superconducting RF accelerator sections (ACC1-6), two bunch compressors (BC2-3) and an undulator section.

The electron beam is produced in the RF photo-injector, accelerated, compressed, accelerated and compressed again, and finally sent through the undulator magnets, producing FEL light pulses with wavelengths between 6 and 40 nm for experimenters [2, 3].

Experimenters use the pulses of light to, among many other sorts of experiments, take pictures of molecules with crystallographic and pump-probe methods. The pulses range in duration from 10 to 50 femtoseconds with the possibility of making them as short as a few hundred attoseconds [4]. If experimenters cannot synchronize their measurement equipment to the beam within these time-scales, they cannot directly make

use of the short pulses in, for example, pump-probe experiments. In pump-probe experiments, one laser, the pump, excites some behavior in a sample and a second laser, the probe, records the behavior of the sample in, for example, a picture called a diffraction pattern. In the THz beamline of FLASH, it is possible to simply measure the arrival-times of the electron bunches relative to the pump laser pulses when they arrive and then use those measurements to make sense of the data [5]. This is like filming a movie with all of the frames taken at random and then later sorting the frames to make a sensible sequence. This has been done with 5 fs resolution over a range of 500 fs with the possibility to deal with beam arrival-time drifts of several picoseconds. For such a measurement, the timing-jitter of the electron beam must only be kept within the dynamic range of the measurement, a requirement which is already fulfilled by the present machine.

Less than a year ago, this measurement of the arrival-time of the electron beam relative to the pump laser was unprecedented. It was previously anticipated that it would not be possible to make such a high-resolution measurement with such a large dynamic range. It was believed that the entire accelerator would need to be actively stabilized with longitudinal intra-bunch-train feedbacks so that the arrival-time jitter and drift of the FEL pulse would be less than 30 fs relative to an optical reference to which the pump laser could be synchronized. This would make it possible for the pump laser to be used to make high-resolution beam arrival-time measurements within a limited, 30 fs, dynamic range. This was the goal of the FLASH optical synchronization system.

The development of the newer, THz beam arrival-time measurements relative to the pump laser does not, however, make the optical synchronization system for FLASH obsolete. When one can control the beam arrival-time with femtosecond precision, one can create defined timing patterns in the bunch train, enabling new sorts of experiments. In addition, THz radiation is not easy to transport and for femtosecond resolution, the length of every hundred meters of optical path must be stable on the sub-micron level. While this may be possible at FLASH, the distances involved at the European XFEL are prohibitive.

For the European XFEL [6] and for sFLASH [7] the optical synchronization system is absolutely critical for the success of the experiments. The European XFEL is a much larger-scale and higher-energy FEL that will be commissioned in 2014. In the XFEL, the THz beam-line is located a kilometer away from where the other experimenters are working and it would be very difficult to transport the THz radiation over that distance to provide the corresponding beam arrival-time measurement. In the case of sFLASH, the machine is the same as FLASH except for the addition of a seeding laser. In sFLASH, a short, seeding laser pulse needs to overlap with the short electron bunch in order to stimulate the FEL process with a defined wavelength. If these two sources are not synchronized, the seeding process will not be effective. This is why the electron bunch arrival-time must be measured and kept under control throughout both the XFEL and sFLASH.

It is possible to stabilize the beam arrival-time at the expense of the energy stability of the beam, but this would be unacceptable because the wavelength of the light generated by an FEL depends on the energy of the electron beam. If the energy of the beam is not precisely controlled, seeding schemes, like sFLASH, will not work. To simultaneously stabilize both the beam energy and arrival-time, the stability of the beam

arrival-time and energy prior to a bunch compressor must be measured and the individual sub-systems must be controlled. The individual sub-systems are described in more detail in the following sections.

2.1 RF Photo-injector

The RF photo-injector [8] generates the electrons for the machine by shooting laser pulses onto a cesium telluride cathode and then accelerating the electrons that are ejected with a strong electric field in an RF cavity. A solenoid around the beam pipe provides additional transverse focusing to counteract the strong space-charge forces that push the beam apart. (Figure 2.1.1).

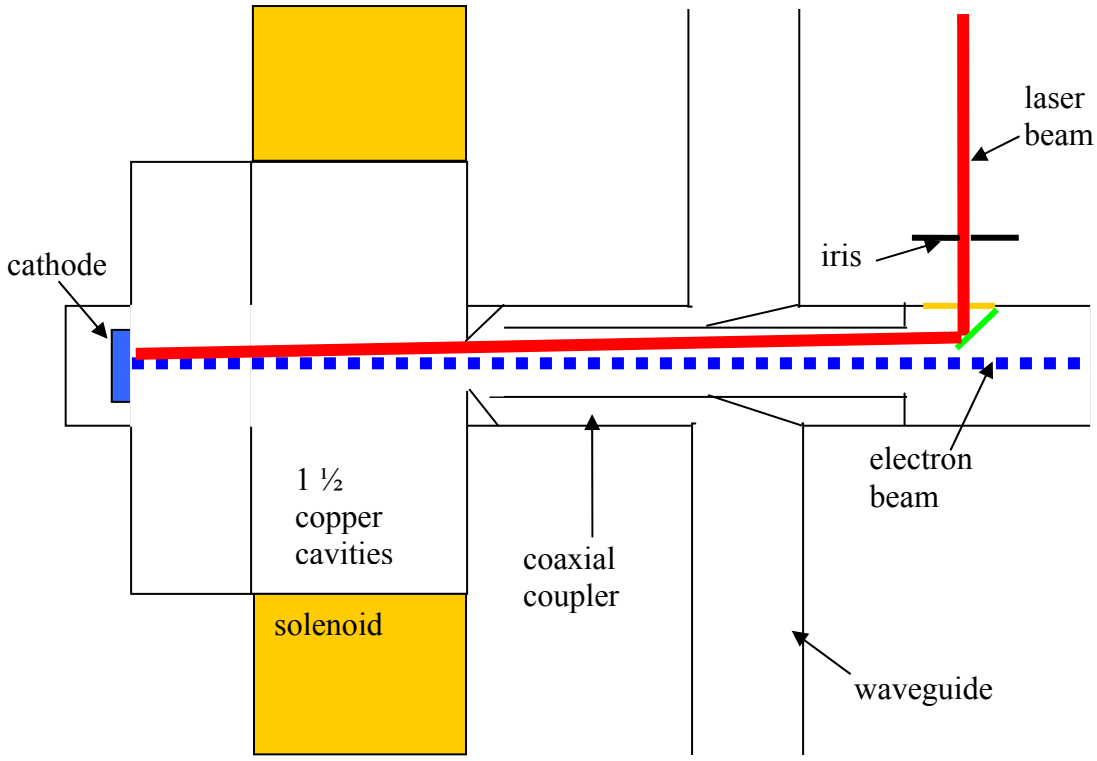

Figure 2.1.1 Simplified cross-section of the RF Gun (not-to-scale). A laser pulse is impinged upon a cesium telluride cathode and the electrons that are ejected are accelerated in an RF cavity; a solenoid provides additional focusing.

In order to produce the high charge-density necessary for the FEL lasing process, the gun is designed to produce as low an emittance as possible. Emittance represents the transverse extent of the beam and when it increases, the charge density decreases. In terms of the horizontal particle distribution of the beam, x, the emittance is defined as

$$\varepsilon_x = \sqrt{\langle x^2 \rangle \langle x'^2 \rangle - \langle xx' \rangle^2} \qquad (2.1.1)$$

and the width of the beam is given in terms of the emittance and the beta function, a function which describes how much the beam is focused at a given phase in the beam transport lattice

$$\sigma_x = \sqrt{\beta \varepsilon_x} \,. \qquad (2.1.2)$$

The same definitions also apply to the vertical particle distribution.

A long list of effects must be correctly balanced in order for the non-relativistic beam to emerge from the photoinjector with a low emittance: non-linear space-charge effects, the gradient at the cathode, RF field distortions, residual magnetic fields, and wakefield kicks name the largest effects [9]. Each of these effects can distort the shape of the bunch, creating and exacerbating asymmetries. The solenoid field counteracts some of the space-charge effects by focusing the beam radially, while the RF field provides some longitudinal focusing. In order to minimize the detrimental effects of wakefields and dispersion, the beam must travel directly through the middle of the beam pipe where the focusing is the most effective and where the distorting effects are the weakest.

Another aspect of producing a low-emittance electron beam in the photoinjector is the quality of the laser pulse. If it is unstable, then the electron beam will be unstable. If it is badly longitudinally shaped, the emittance produced by the photoinjector will be too large for lasing to occur at the end of the machine. The laser pulses originate in an actively mode-locked Nd:YLF laser operating at a wavelength of 1047 nm [10]. These pulses are amplified and then frequency-quadrupled in order to produce a beam charge of around 1 nC. Changes in the amplification of the laser result in changes in the charge produced at the cathode. In order to reduce the effect of the laser's pointing jitter and to make the intensity of the laser spot more uniform, the laser pulse is sent through an iris before it is impinged upon the cathode. The synchronization of the laser pulses and the RF in the cavity to the reference of the machine are critical to keeping the electron beam properties stable and maintaining synchronization to other devices, including the downstream accelerator section.

2.2 Accelerator section

The accelerator sections consist of superconducting niobium cavities which are filled with electromagnetic waves that have a frequency of 1.3 GHz. They form standing waves in the cavity, producing gradients ranging from 12 to 30 MV/m. The waves are produced by a klystron in 800 μs long pulses with a repetition rate of 5 Hz in order to accelerate up to 800 bunches per pulse. Typically, however, bunch trains of only 1-30 bunches have been produced for standard operation.

A klystron consists of a cathode from which electrons are generated and an anode toward which the electrons are accelerated with a voltage drop of many kilovolts (~200kV at FLASH). The electrons then enter a cavity that is filled with gigahertz waves by a modulator. This buncher cavity gives the electrons an energy modulation which is transformed into a density modulation in the subsequent drift section where high-energy

electrons travel faster than low-energy electrons. The electrons are thereby bunched with a periodicity equal to the period of the wave in the buncher cavity. These bunched electrons travel through a second cavity, causing it to ring with the same frequency that was generated by the modulator. The resonant wake fields from this cavity are transported via waveguide to the accelerating structure to accelerate the beam. The phase of the klystron output is strongly influenced by changing the voltage drop of the driving electrons and the amplitude of the klystron output is controlled by the amplitude of the modulator output. At FLASH, a multi-beam klystron from Thomson Tubes Electronics is used to generate the 1.3 GHz for the acceleration modules. A simplified depiction of the basic structure of a klystron is shown in Fig. 2.2.1

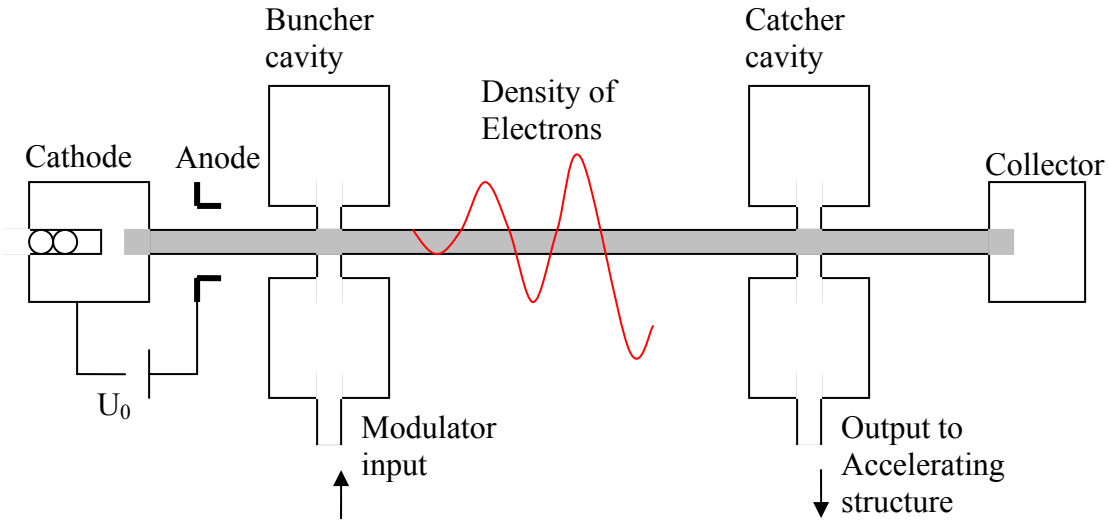

Figure 2.2.1 Basic structure of a klystron. Electrons travel from the cathode to the collector and along the way they are bunched and used to excite waves that are sent to the accelerating structure.

Large amounts of effort have been invested into actively controlling the amplitude and phase of the RF generated by the klystron with digital signal processing feedbacks and feedforward loops on Field Programmable Gate Arrays (FPGAs) [11]. RF pickups inside of the couplers of the modules detect the 1.3 GHz field and the signals are converted to a lower frequency signal that can be sampled with Analog-to-Digital Converters (ADCs). FPGAs execute a real-time algorithm to determine the phase and amplitude of the signals from the ADC outputs and use feedback and adaptive feedforward loops to send signals to actuators that adjust the phase and amplitude of the RF in the cavity. The goal of this is to make it possible that every bunch in every bunch-train experiences the same accelerating field, gains the same amount of energy and behaves in the same manner in the bunch compressors. These efforts are described in more detail in Chapter 3.

2.3 Bunch Compression

When an electron bunch travels through a magnetic bunch compressor chicane, the high-energy electrons travel a shorter path than the low-energy electrons, and, therefore, arrive

at the exit of the chicane earlier than the low-energy electrons. If the electron bunch is given a longitudinal energy dependence (chirp) by accelerating it off of the crest of the RF wave in the cavity, the electrons in the head of the bunch gain less energy than the electrons in the tail of the bunch (Fig. 2.3.1). The energy dependent path through the bunch compressor will make the head arrive later and the tail arrive earlier, thereby shortening the bunch.

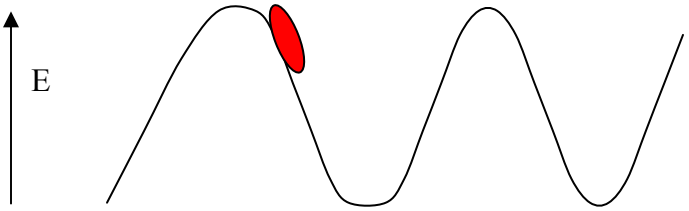

Figure 2.3.1 The phase of the accelerating RF relative to the beam determines the energy chirp of the beam. Here, the head gains less energy than the tail of the bunch.

Due to the curvature of the RF wave, the bunch acquires a non-linear energy chirp. When a bunch with a non-linear energy chirp travels through the chicane, it is compressed inhomogeneously, resulting in a sharp, leading spike with a high charge density followed by a long trailing tail. An RF module operating at the third-harmonic of the 1.3 GHz accelerator frequency will remove the non-linearity of the energy chirp and enable homogeneous compression of the bunch.

In Fig. 2.3.2, plots of the beam distribution for non-linear compression in the first chicane of FLASH were generated by a particle tracking simulation which used a particle distribution from the gun which was generated with ASTRA together with transport matrices for the first accelerator section and for the bunch compressor. The evolution of the transverse distribution is poorly described by this simulation because the lattice is designed for a space-charge-limited beam and space charge effects were not taken into account after the gun. This results in a simulated transverse beam size which is much larger than it would be in reality. The plots, however, serve to demonstrate the gymnastics of longitudinal bunch compression.

In the first pair of plots in Fig. 2.3.2, the energy chirp is shown on the right and the particle distribution in the horizontal and longitudinal planes is shown on the left. The beam is spread out over ~7 mm of longitudinal space and it has a longitudinal energy chirp which is curved due to the curvature of the accelerating RF. In the second pair of plots, the higher energy particles have begun to overtake the lower energy particles and the bunch has become shorter longitudinally and wider in the horizontal plane of the bunch compressor. The slice emittance increase in the tail of the bunch is due to a mismatched lattice; if a quadrupole strength is changed slightly, the emittance increase appears in the head of the bunch and not the tail. This is partially an artifact of the particular simulation used. Space-charge forces were only used up to the exit of the RF gun, but they too significant to ignore in the first four modules of the accelerator section, and this creates a mismatched result in the simulation. With proper matching, the slice emittance of the beam will be smaller and the beam will look like a curved strip in the second pair of plots. In the last pair of plots, the horizontal position spread has returned to the value it had before entering the compressor and the bunch has been compressed into a

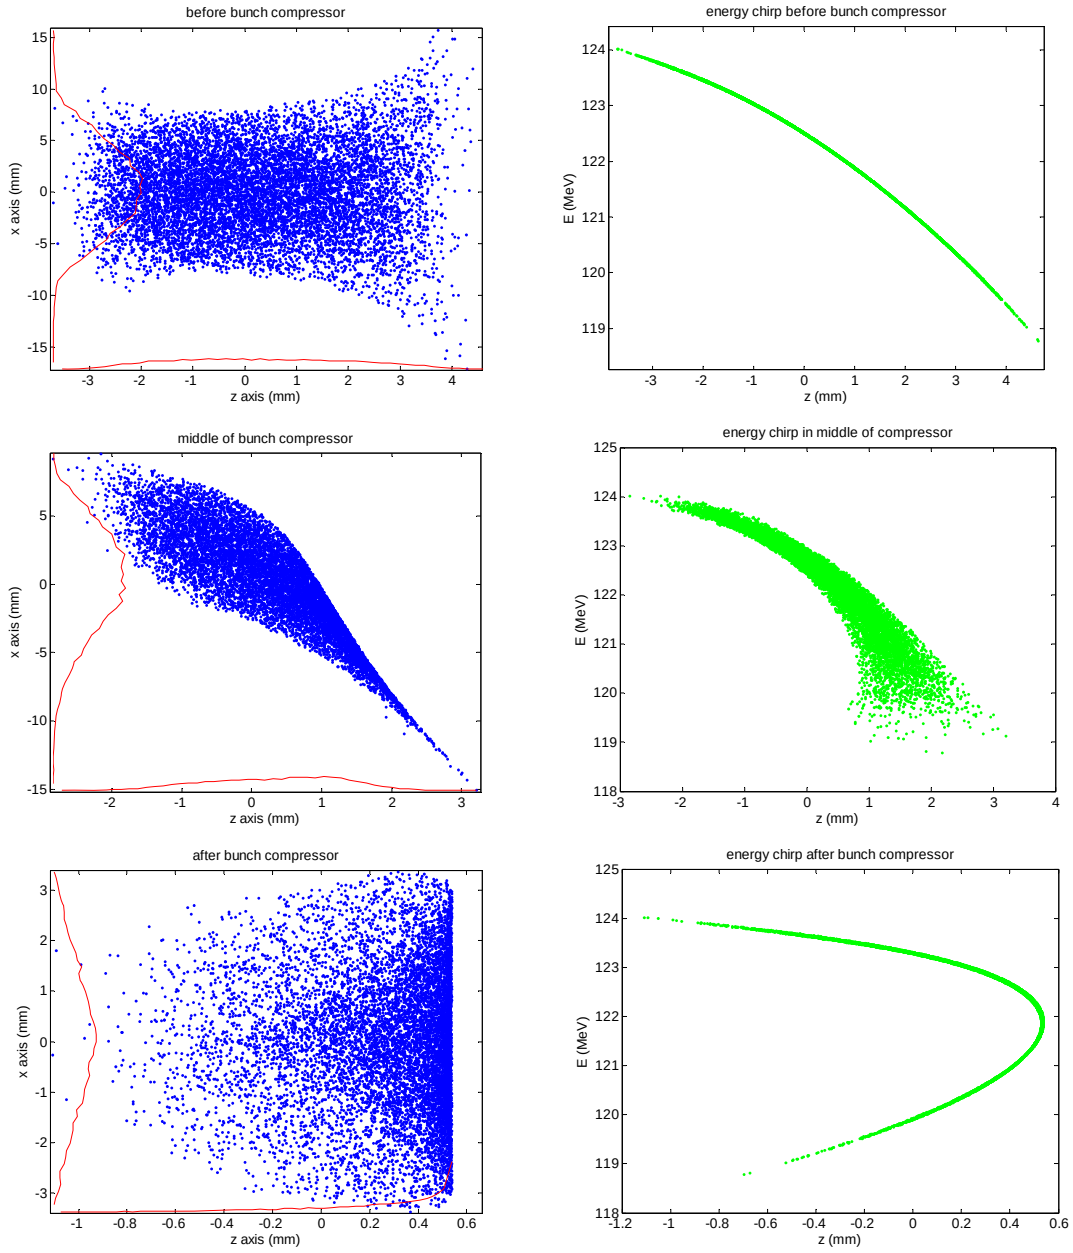

Figure 2.3.2 Bunch profiles in the first bunch compressor for inhomogeneous compression without the third-harmonic cavity. The transverse beam size is larger than in reality because space charge effects were not taken into account after the RF photo-injector.

sharp leading spike of charge distribution followed by a millimeter-long trailing tail. The energy distribution of the bunch has folded over on top of the bunch for this 10 degrees off-crest simulation. When the head of the bunch overtakes the tail of the bunch in the chicane, this is frequently referred to as over-compression. Over-compression is typically employed for beams with a non-linear energy chirp. The end result of this process is that the bunch is longitudinally compressed, the energy chirp increases and the energy spread remains constant.

Bunch compressors are frequently designed to accommodate a range of energies and compression schemes which are defined to first-order by parameters denoted R_{56} and R_{16} after their location in a six-dimensional transfer matrix used to calculate the beam transport [12]. In the introduction, analytic formulas for these parameters were derived for a symmetric, single chicane. Physically, they relate the change in position to the change in energy deviation, δ, according to

$$R_{56} = \frac{\partial z}{\partial \delta} \qquad R_{16} = \frac{\partial x}{\partial \delta} \tag{2.3.2}$$

They are related to one another by

$$R_{56} = \int_0^s \frac{R_{16}(s')}{r(s')} ds' \tag{2.3.3}$$

where the integral is along the reference trajectory, *s,* and *r(s')* is the bending radius of the magnets. From Eq. 2.3.3, we can see that for smaller bending radii, the R_{56} increases and for drift spaces where *r* is infinite, the R_{56} vanishes.

It is useful to be able to make a quick estimation of the energy chirp of an electron bunch after an accelerator section and calculate the resulting bunch length change or *x* position spread change. To do this, let us first write down the energy of an electron subject to an accelerating module with an acceleration voltage of *U* and a phase of $\varphi = k_{rf}\Delta s + \varphi_0$ with $k_{rf} = 2\pi/\lambda_{rf}$, in terms of the wavelength of the accelerating RF, and φ_0 equal to the phase for which the longitudinal position is equal to that of the reference trajectory, $\Delta s = 0$

$$\delta = E_f - E_i = eU\cos\varphi \tag{2.3.5}$$

where E_i is the initial energy of the particle and E_f is the energy after the accelerating module. We can describe the energy chirp produced by the accelerating RF by doing a Taylor expansion of δ about a small longitudinal position change, Δ*s*,

$$\begin{aligned} \delta(s) &= \delta(\Delta s) + (s-\Delta s)\delta'(\Delta s) + \frac{1}{2}(s-\Delta s)^2\delta''(\Delta s)\ldots \\ &= R_{66}\delta_{initial} + R_{65}\Delta s + R_{655}\Delta s^2 . \end{aligned} \tag{2.3.6}$$

The first term describes the initial energy spread over the position change, the second term describes the linear chirp acquired over the position spread and the third term describes the quadratic chirp acquired. The indices of the *R* coefficients describe the coordinates of the values in the beam transport matrix. The first index coordinate equal to six corresponds to energy deviations and the second index coordinate equal to five

corresponds to terms that are linear within Δs. The third index coordinate equal to 5 is for terms that are quadratic within Δ*s*. The coefficients are given by,

$$R_{65} = -\frac{eU_{acc}\sin\varphi}{E_f}k_{rf} \qquad R_{66} = \frac{E_i}{E_f}$$

and

$$R_{655} = -\frac{1}{2}\frac{eU_{acc}\cos\varphi}{E_f}{k_{rf}}^2 \tag{2.3.7}$$

For a quick calculation of the change in beam properties resulting from a chicane, the following formulas make use of the above transfer matrix parameters in order to relate a beam energy chirp given by the difference in the energy of the particles in the head of the bunch and the energy of the particles in the tail of the bunch, δ, to a change in bunch length, σ_z, or a change in position spread in the middle of the chicane, σ_x. To find the beam width in the middle of the chicane, one can use the linear transformation

$$\begin{pmatrix} x \\ \delta \end{pmatrix} \approx \begin{pmatrix} 1 & R_{16} \\ 0 & 1 \end{pmatrix}\begin{pmatrix} x_0 \\ \delta_0 \end{pmatrix},$$

to write $\langle x^2 \rangle = \langle (x_0 + R_{16}\delta_0)^2 \rangle = \langle x_0^2 \rangle + 2R_{16}\langle x_0\delta_0 \rangle + \langle \delta_0^2 \rangle$, where the middle term is equal to zero because we assume that there is no dispersion upstream of the bunch compressor and, therefore, no correlation between x_0 and δ_0. We can also write $\langle x \rangle = \langle x_0 \rangle + R_{16}\langle \delta_0 \rangle = 0$, with an assumption that the beam is centered about zero. Now, with the definition of standard deviation, $\sigma_x = \sqrt{\langle x^2 \rangle - \langle x \rangle^2}$, we can describe the beam width in the middle of the chicane,

$$\sigma_x \approx \sqrt{\sigma_{x0}^2 + R_{16}^2\delta_0^2}\ . \tag{2.3.8}$$

Likewise, to calculate the bunch length after an accelerator section and a bunch compressor, one can use the linear transformation

$$\begin{pmatrix} z \\ \delta \end{pmatrix} \approx \begin{pmatrix} 1 + R_{65}R_{56} & R_{56}R_{66} \\ R_{65} & R_{66} \end{pmatrix}\begin{pmatrix} z_0 \\ \delta_0 \end{pmatrix}$$

to find

$$\sigma_z \approx \sqrt{{\sigma_{z0}}^2(1 + R_{56}R_{65})^2 + \delta_0^2 R_{56}^2 R_{66}^2} \tag{2.3.9}$$

This result can be used to write the compression factor, $C=\sigma_{z0}/\sigma_z$, a term that is used frequently in the following chapter.

The ranges of values that the R_{56} and R_{16} assume for FLASH and XFEL are listed in Table 2.3.1, along with the expected range of beam positions, X, and widths, DX, that a monitor in the middle of the chicane would have to measure.

Chicane	R_{56} *(mm)*	R_{16} *(mm)*	*X (mm) range*	*DX (mm) ±3 sigma*
FLASH BC2	140-228	284-358	0-74	2-10
FLASH BC3	14-84	100-250	0-150	2-6
XFEL BC1	100	500-600	0-400	2-60
XFEL BC2	40	200-300	0-400	2-9

Table 2.3.1 R_{56} and R_{16} values for FLASH and XFEL, the corresponding dynamic apertures of the chicanes (X range) and the position spreads of the beam within the chicanes (DX). The X range starts at zero for all of the bunch compressors, because it is desirable to allow for operation with the compressors off.

Whereas at LCLS (Linac Coherent Light Source at SLAC) [13], another FEL facility, the entire bunch compressor beam pipe was placed on motorized movers in order to accommodate different compression modes, the FLASH and XFEL bunch compressors have wide and flat vacuum chambers that do not move. The reasons for building chicanes that have adjustable properties are twofold: it is sometimes desirable to turn the chicane off for a different mode of machine operation and it may be desirable to have the freedom to independently adjust the bunch length and energy chirp emerging from the chicane.

The limitation of the adjustable range of a bunch compressor is not determined by the feasibility of constructing large aperture diagnostics, which are the subject of this thesis, but instead, by the feasibility of constructing dipole magnets with a high field quality over the entire dynamic aperture. Rectangular bends are used and as the field quality deteriorates at the outer limits of the dynamic aperture, chromaticity and higher order dispersion may have to be corrected due to the quadrupole and sextupole field errors contained therein. With a chicane on movers, the field tolerances on the dipole magnets are reduced and it is easier to add quadrupole magnets to the chicane in order to correct the chromaticity of the beam.

The energy spread of a bunch is unaffected by the bunch compression process and can be used to further compress the bunch in multiple bunch compressor stages until a minimum bunch length is reached. The minimum bunch length for a given energy spread is taken from the second term of Eq. 2.3.9

$$\sigma_{z\min} = R_{56}\frac{E_i\sigma_E}{E_f}$$

It is not, however, always desirable to reach this value in one bunch compressor alone.

Bunch compression is not typically done in only one stage because of non-linear energy spread, wakefield and space-charge issues. At FLASH, the bunch is shortened two times, once at 130 MeV and again at about 460 MeV. When the bunch is highly compressed, the space-charge forces become stronger and force the beam to expand. When the beam is not longitudinally compressed enough, the non-linearity of the energy

spread will be larger and the transverse wakefields will be stronger. Transverse wakefields are the reason that the bunch is partially compressed early in the machine and the issue of space-charge is one of the reasons that the bunch is fully compressed near the end of the machine. A bunch compressor at the end of the machine can also compress some of the timing jitter generated up-stream and can possibly provide a compensation effect for some emittance increases generated in the first compressor [17].

A horizontal emittance increase will occur if a significant energy spread is generated within the chicane. This emittance increase occurs because the energy spread breaks the linear achromaticity of the chicane. In a linear achromat, the particles will have the same transverse position after the achromat that they had before they entered the achromat. When the achromaticity is broken, the off-energy particles no longer follow the same orbit as the on-energy particles and they will emerge at a different transverse position than the on-energy particles. This is depicted in Fig. 2.3.3

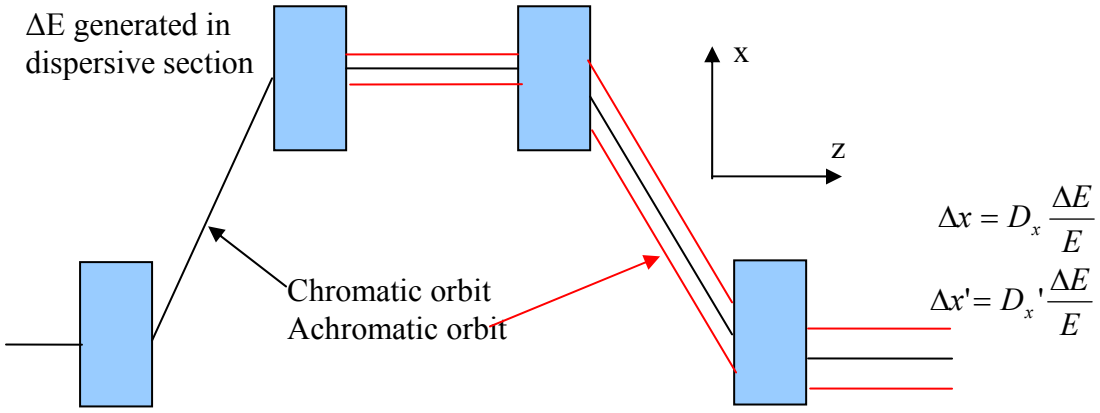

Figure 2.3.3 Energy spread generated within chicane breaks linear achromaticity and results in an increased emittance after the chicane.

Within the chicane dipoles, incoherent and coherent synchrotron radiation (ISR and CSR) generate an energy spread. Collective effects and space charge effects can also generate an energy spread in the chicane. Because these energy spreads are generated in a dispersive section, they will result in emittance growth. Emittance growth will not, of course, result from an energy spread generated in a non-dispersive section, like the accelerator section.

When the path of an electron beam is influenced by a magnetic field, the beam emits synchrotron light. With wavelengths corresponding to the length scales longer than the bunch length, the bunch radiates coherently and for length scales shorter than the bunch length it radiates incoherently. ISR is generated through a random process and cannot, therefore, be corrected. It increases for higher energy beams [14] making it less significant in the first bunch compressor and more significant in the second bunch compressor. Unlike most other emittance increases, this acts solely on the slice emittance. The CSR is much more powerful than the ISR, it is correlated along the bunch and acts strongly on the projected emittance.

The power of the CSR increases in proportion to the bunch length raised to the power -1/3, meaning that the power of the CSR increases for shorter bunches. As the CSR and the electron bunch co-propagate in the bends of the chicane, the CSR can catch

up with and then interact with the electron bunch, giving it an energy spread which is correlated along the bunch. For an rms bunch length, σ_z, dipole length, L_B and dipole bend radius, R, the CSR-induced rms relative energy spread per dipole for a Gaussian bunch under steady-state conditions is

$$\Delta E_{CSR} \approx 0.22 \frac{N r_e L_B}{\gamma R^{2/3} \sigma_z^{4/3}} \qquad (2.3.4)$$

Where N is the number of electrons, r_e is the classical electron radius and gamma is the Lorentz energy factor [15]. Since the bunch length is shorter after each bend, the local energy spread generated in the last bend of the chicane will be the largest.

Aside from the increase in emittance, bunch compressors can also cause something called a micro-bunching instability when a density modulation created by the impedance of geometric wakefields, longitudinal space-charge, or CSR get caught in a feedback loop in which these energy modulations are coupled into density modulations via the momentum compaction of the chicane. For example, the CSR creates an energy modulation of the bunch and the dispersion chops the beam up into slices with lengths that are comparable to the coherent synchrotron radiation wavelengths. These micro-bunches then interact with one another, experiencing resonant oscillatory motion that can break the macro-beam apart. Although it was initially suspected that CSR would be the primary driver of the microbunching instability, the longitudinal space-charge effects in the injector are now the primary focus of concern. The microbunching instability can be avoided by increasing the residual energy spread of the bunch with a laser heater that imposes a periodic energy modulation over the bunch and then smears it out longitudinally via dispersion [16].

At FLASH, a single chicane is used as the first bunch compressor and a symmetric double-chicane is used as the second bunch compressor (Fig. 2.3.4). The microbunching instability is more likely in the second chicane, due to the added bending of the double-chicane design and the higher energy of the beam [13]. In general, the gain of the instability increases with the inverse of the characteristic wavelength of the modulation squared. It is cut off for wavelengths that are shorter than $R_{56}*C*\delta$. For the first bunch compressor, R_{56}, C and δ are larger than for the second bunch compressor and so the cut-off occurs at longer wavelengths, making the microbunching less effective. With careful balancing of CSR and other effects, this can be avoided and the second chicane can be used to compensate for CSR-based emittance growth generated in the upstream bunch compressor [17].

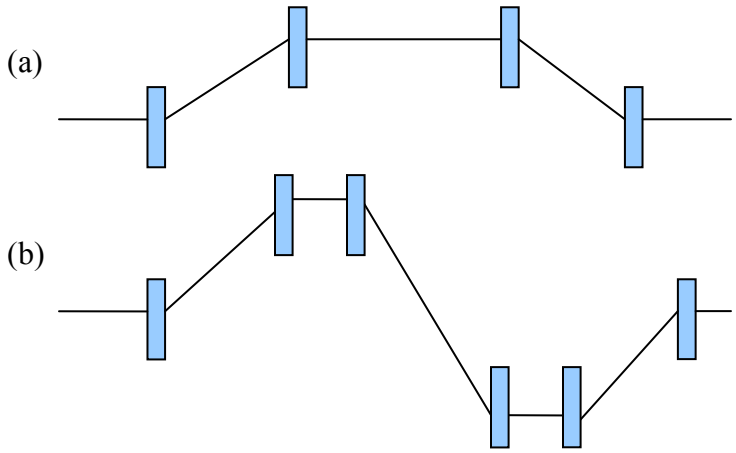

Figure 2.3.4 Single chicane (a) and symmetric double-chicane (b).

The higher order dispersion terms, R_{566} and $R_{166,}$ are opposite for bunch compressors of the single and double-symmetric types, adding another opportunity for canceling out destructive effects. These cancellation effects only become possible with the third-harmonic module in operation to linearize the compression process.

2.4 Undulator Section

The electron bunch with a high current-density is sent through a series of undulator magnets in order to create a short and high-energy pulse of light. Undulator magnets with a period of λ_u and a magnetic field of B_0 bend the path of the electrons back and forth many times and cause them to radiate synchrotron light with a fundamental wavelength of [18]

$$\lambda_l = \frac{\lambda_u}{2\gamma^2}\left(1+\frac{K^2}{2}\right) \text{ with } K = \frac{eB_0\lambda_u}{2\pi m_e c}$$

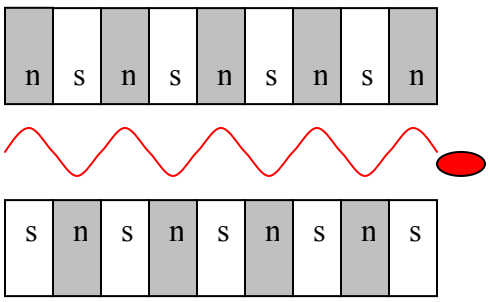

Figure 2.4.1 Undulator magnet and electron bunch producing synchrotron radiation.

Based on the presence of the Lorentz factor in the above formula, it is clear that if the beam energy changes, then the wavelength of the fundamental mode of the light will change.

One of the keys to making the light generated by the undulators coherent, as in a laser, is to maintain a sustained interaction between the light pulse and the electron bunch over a distance known as the gain length. After several gain lengths, the sustained interaction creates an energy transfer from the light pulse to the electron bunch, causing the electron bunch to break up into microbunches with a periodicity equal to the wavelength of the synchrotron light. These microbunches begin to radiate coherently; the light generated by each microbunch adds coherently to the light generated by the other microbunches. This increases the intensity of the coherent light pulse and increases the energy transfer from the electron bunch to the light pulse, furthering the microbunching process. When this sustained interaction is maintained over a sufficient number (~20) of gain lengths and high intensity synchrotron light is generated, the machine is said to be 'in saturation' or 'lasing'.

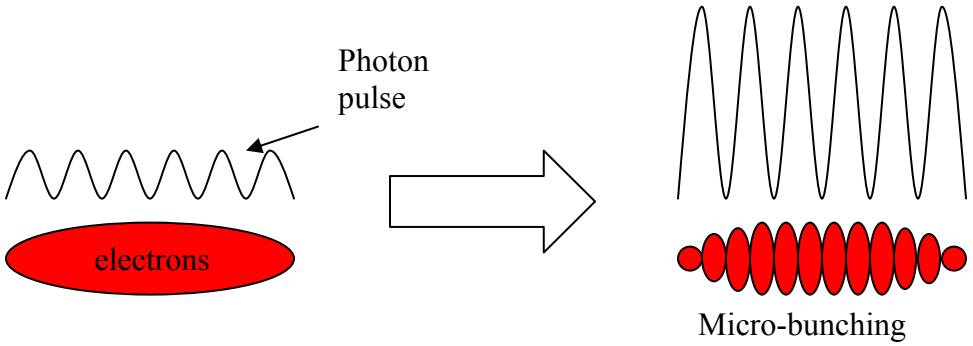

Figure 2.4.2 Interaction of electron beam and photon pulse.

The process described above is known as SASE, Self Amplified Spontaneous Emission, wherein the micro-bunching structure develops spontaneously from shot-noise and grows more distinct as saturation is achieved. An alternative to SASE is to use a seed laser to initiate the micro-bunching process at a desired frequency. This makes the light generated in the undulators much more monochromatic and intense. Such a seeding project, called sFLASH, is underway for commissioning during the coming year [7]. A complete introduction to VUV and X-ray FEL techniques is given in [18].

3 Beam Arrival-time Stabilization

As we learned in the previous chapter, beam energy changes upstream of a bunch compressor become arrival-time changes after the bunch compressor and it is desirable for the electron beam to have a stable arrival-time relative to a reference signal so that seed laser, diagnostic laser, and pump laser pulses arrive synchronously with the beam. A feedback to control the arrival-time of the FEL beam can be made no more complicated than a single monitor, like a beam arrival-time monitor after the chicane, that tells a single klystron how to set the energy of the beam. While such a feedback can produce a dramatic improvement in the arrival-time jitter measured at one point in the machine, it can have the flaw that it feeds back on arrival-time jitter that is generated somewhere other than in the module it is controlling. It also constitutes a single point-of-failure, an unfortunate design flaw for a system that requires a high level of robustness. A feedback architecture that prevents these conditions is described here, along with a description of the relative contributions various machine sections bring to beam arrival-time jitter.

3.1 Baseline Control

Proportional gain feedbacks and adaptive feedforward loops that utilize measurements of the cavity fields in order to stabilize the cavity fields have been the workhorses of FLASH beam energy stabilization since its inception [19], while beam-based feedbacks, feedbacks that utilize measurements of beam parameters to control the cavity fields, have only been tested briefly [20]. Additional improvements in the feedback architecture, involving an enhanced low frequency gain profile and reference injection to reduce drifts, have also only recently been tested [11, 21]. The relative merits of these systems will be described below along with the limitations of what is currently available.

Feedforward is a term describing an element in a control system that delivers commands in a pre-defined way, without responding to how the system reacts. A fixed setpoint table that takes into account various calibrations done at an earlier time is an example of feedforward. After the effects of Lorentz-force detuning, cavity detuning, and imbalances in the actuator chain have been calibrated away over weeks of studies in the absence of beam production, a fixed feedforward (setpoint) table can be determined. When the beam is then added to the system, a slope on the RF pulse arises due to something called beam loading. Beam loading occurs when an electron bunch enters the accelerating cavity. The beam takes energy out of the accelerating field and this energy must be replaced by increasing the klystron's output. If, in one bunch train, each bunch took a certain amount of energy out of the cavity, the same thing is likely to happen in a subsequent bunch-train, as long as the beam charge or orbit is not significantly changed. A fixed feedforward table may be appropriate for one set of beam parameters, but as soon as the machine operator changes the setup of the machine, the feedforward table will have to be manually tuned to compensate for changes in beam loading.

In the absence of an expert to tune the feedforward table, an adaptive feedforward algorithm using "Iterative Learning Control" can automatically change the feedforward table in order to counteract the changes in beam loading. The control decisions of the adaptive feedforward algorithm do not take place within the bunch-train, but after averaging over multiple bunch-trains. The adaptive feedforward can not only remove slopes from the bunch train, it can also, in principle, remove ripples. If ripples are periodic and appear in bunch-train after bunch-train with the same phase, they can be removed through the feedforward. An adaptive feedforward control can identify patterns in the cavity signals during one klystron pulse and attempt to remove those patterns by applying a pattern of equal and opposite amplitude in a subsequent klystron pulse. This control option has not, however, been used for day-to-day beam operation due to the incompleteness of its implementation. The incompletely debugged failure modes of the controller have caused the superconducting cavities to quench. New versions of the controller are under development [11].

Feedback is different from feedforward in that it sets control parameters based on the reaction of the system to the control parameters. It utilizes measurements taken at the beginning of the pulse in order to change the settings of the klystron within the pulse. The current system can implement Proportional, Integral and Differential (PID) feedback control, but for typical operation, only the proportional feedback is used. Feedback is limited by measurement resolution and latency, how long it takes for a signal to be measured, interpreted, and converted into a control parameter. With Field Programable Gate Arrays (FPGAs), Analog-to-Digital Converters (ADCs), and Digital-to-Analog Converters (DACs) able to process signals at more than 100 MHz, or every 10 ns, and signal transport times that can be kept below 100 ns, bunch-to-bunch feedback within the FLASH bunch spacing of 1 μs and even the XFEL bunch spacing of 200 ns becomes a goal within reach. The latency of present systems is currently limited to ~3 μs, but in certain locations where cable lengths can be minimized, faster performance could be realized with future hardware.

While the delivery of control decisions to the klystron between bunches is possible, making large changes in the field of the accelerating structure at that rate is physically limited due to the large quality-factor, or small bandwidth of the cavity. Large

changes in the amplitude of the klystron produce only small changes in the cavity gradient. When the corrections demanded by the feedback loop become small enough, then the large quality-factor of the cavity is no longer a problem and the cavity phase and amplitude can, in principle, reach a stability determined solely by the resolution and drift of the monitoring system used in the feedback.

In Fig. 3.1.1, the current digital processing architecture is depicted [11]. The cavity field probe signals 1 through n are sent to the field detectors wherein they are down mixed from 1.3 GHz to 54 MHz and sampled with an 81 MHz ADC. The digitized signals from the ADC are sent into a digital phase and amplitude (I & Q) detection algorithm. This calculation requires a multiplication by a calibration constant. The phase and amplitude information from each individual cavity are then added together in a vector sum. The vector sum is compared to the setpoint values generated from a feed-forward table. A correction to the cavity fields is then calculated from a gain setting and the difference of the measured signals from the setpoint. This correction is multiplied by calibration factors appropriate for each cavity and sent to a DAC which generates a 250 kHz signal. The 250 kHz signal goes to a vector modulator (labeled VM in the diagram) which shifts the phase and amplitude of the 1.3 GHz that is sent to the klystron.

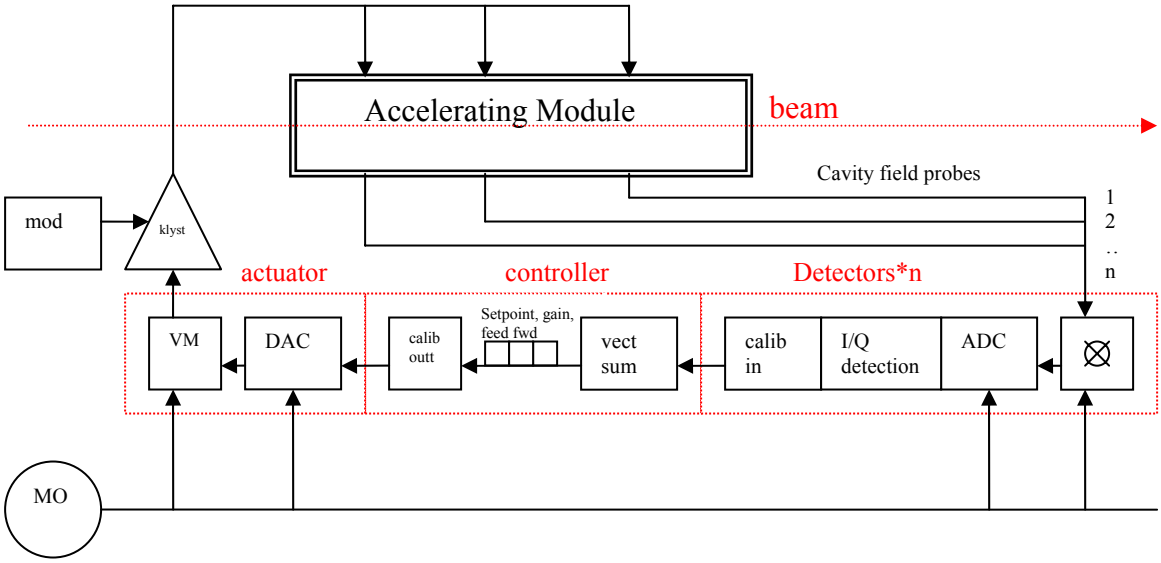

Figure 3.1.1 System for controlling the cavity fields of the accelerating module. The phase and amplitude of the fields are detected from cavity pickups, the difference from the setpoint is calculated and a correction is sent to the klystron [11].

The boxes labled: I/Q detection, calibration in, vector sum, setpoint, gain, feedforward, calibration out comprise the routines that take place on an FPGA. This routine is depicted in more detail below in Fig. 3.1.2. In Fig. 3.1.2, the phases and amplitudes detected from each cavity are added together in a vector sum and then subtracted from the setpoints generated by the feed-forward table. The feed-forward table is generated from a setpoint given by a user and an older feed-forward table [11].

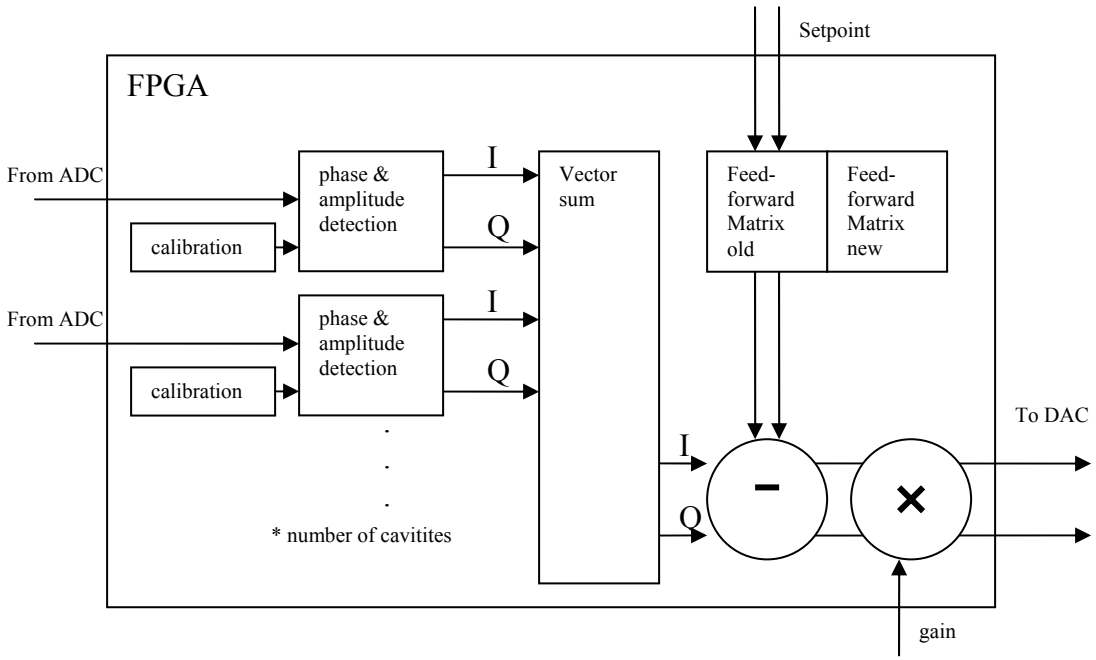

Figure 3.1.2 Simplified block diagram of the cavity regulation routines on the FPGA.

The limitations of the current digital processing system are primarily due to cavity amplitude and phase measurement resolution and drifts; the field measurements inside the cavities are not completely accurate. When an inaccurate cavity measurement is used to regulate the cavity parameters, the measurement drifts and fluctuations can turn into cavity drifts and fluctuations which, in turn, become beam drifts and fluctuations.

The drift problem of the cavity amplitude and phase measurements has recently been approached through reference injection and reference tracking techniques [21]. Reference tracking refers to the use of a reference signal to measure the out-of-loop performance of a regulation loop. Reference injection refers to the use of an out-of-loop measurement to actively improve the drift performance of the loop. In the system under consideration, the MO signal was the reference used to measure the out-of-loop performance of the vector sum of the phases and amplitudes of all of the cavities. This sum drifts due to the influence of temperature changes on the RF cables and circuits that make the cavity field measurements. The reference injection and tracking method reduced the impacts of temperature changes in the electronics racks from several degrees in phase and 1e-2 in amplitude to below 2e-4 in amplitude and 0.008 degrees in phase. The amplitude drift was verified with a beam-based measurement, but the phase measurement is only relative to the 1.3 GHz reference delivered on an RF cable, not relative to the beam. While reference injection can remove the 1e-3/degC drift of the field detectors, the phase of the reference signal on every meter of RF cable will still drift with temperature with a coefficient of 1e-4 /degC. After 10 meters of cable, the cable drift is the same as the field detector drift. With more expensive cables, this phase drift can be reduced by a factor of 5. With active drift compensation, using RF reflectrometry, the phase drift on the long RF cables could be made as small as 0.005 degrees. Cable drift compensation will be discussed in Chapters 7 and 8 in the context of beam arrival-time measurements.

The problem of cavity field measurement resolution resulting in bunch-to-bunch energy fluctuations has recently been approached with a new version of the feedback control algorithm that incorporates a system identification scheme in a so-called Multiple Inputs Multiple Outputs (MIMO) framework [11]. The advantage of the MIMO feedback is that it increases the gain for low frequency fluctuations without increasing sensitivity to high-frequency fluctuations and it allows for easier incorporation of new information into the control algorithm. In contrast, the existing proportional gain feedback responds to all frequencies up to the bandwidth of the ADC and is incompatible with beam based feedback. When measured with the reference tracking (out-of-loop vector sum) setup, the MIMO feedback with iterative learning control, a new version of the feedforward, successfully reduced the pulse-to-pulse amplitude stability of the first accelerator section from 2e-4 to below 5e-5 and it reduced the phase stability from 0.008 down to 0.003, but it developed resonances and instabilities over the course of an hour [11]. With notch filters applied to the resonance frequencies, the system should be more stable. The feedback is limited primarily by the resolution of the cavity measurement front-ends. An improvement in the resolution of the cavity front-ends by a factor of 5 to meet 1e-5 stability is conceivable in the near future given the implementation of new 16-bit front-ends [22].

If the amplitude stability measured by the reference tracking system translated directly into beam energy stability, then 5e-5 energy stability would have been measured with the MIMO controller, but this was not the case. The best pulse-to pulse beam energy stability that was achieved with this controller was only 1.3e-4, not a major improvement over the best-case 2e-4 beam energy stability produced by the proportional gain controller. This just serves to reinforce that stabilizing cavity field measurements does not always stabilize the beam energy.

A summary of the module controller benchmarks described above is given below in Table 3.1.1. Only the amplitude stability measurements have been verified with beam-based measurements. The jitter and drift of the phase refers to the jitter of the module RF phase relative to the 1.3 GHz reference signal phase and not relative to the beam phase.

	Current System		*Reference Injection*		*System Identification*	
	drift	*jitter*	*drift*	*jitter*	*drift*	*jitter*
Phase(deg)	1 to 3	0.008	0.008	0.008	1 to 3	0.003
Amplitude	1e-2	2e-4	2e-4	2e-4	1e-2	1.3e-4

Table 3.1.1 Module controller performance benchmarks: as it now stands, for the system with the addition of reference injection and for the system with the addition of the system identification algorithm.

It should be noted that the numbers from Table 3.1.1 represent the best measurements of the jitter performance; the typical performance is 4e-4 in amplitude and 0.07 degrees in phase as measured with beam-based devices [23]. These beam-based measurements used the beam image on a screen in the bunch compressor to measure the energy jitter and a bunch compression monitor pyrodetector to measure the phase jitter. While the energy jitter measured with this method is primarily due to the amplitude jitter of the first accelerator section, with a smaller contribution from injector phase jitter, the phase jitter

measurement is limited by the injector phase jitter. This means that 0.07 is a rather pessimistic value for the phase jitter of the first accelerator section, since it represents more the jitter between the phase of the gun laser and gun RF. The true value of the first accelerator section jitter could be anywhere between 0.01 and 0.07, but based on principles of the regulation algorithm, the amplitude jitter should always be ~1.8 times the phase jitter [19], so for future calculations, we will assume a phase jitter of 1e-4 or 0.02 degrees.

These numbers are also only applicable for short bunch trains and 1 nC bunch charge. When long bunch trains with 3 nC were used in the 9 mA experiment for ILC research, due to the large beam loading effect, the energy jitter of the bunches at the end of the bunch train was ten times worse than the energy jitter of the bunches at the beginning of the bunch train [24].

While the combination of a new down-conversion front-end [22], reference injection [21] and system identification [11] could conceivably stabilize the cavity amplitude jitter and drift to within 1e-5 for short bunch trains in a well tuned machine, it is unclear what the system performance would be under more typical circumstances. Using the present machine as an example, the best results are different from the typical results for short bunch trains by about a factor of 2. For long bunch trains the difference can be a factor of 10-30

A beam-based feedback system is being developed to complement this cavity-based system because beam-based measurements can often provide a more accurate measure of how much energy the beam gained in the cavity than a measurement of how much energy the cavity lost when the beam traversed it. The reason for this is that the beam-based measurement can often be accomplished in one location, with one pickup and one set of electronics, while the cavity measurement requires a sum of signals from many different pickups in an RF vector sum. The beam-based measurements can also be drift-free relative to an optical reference which is used to synchronize various laser systems through optical cross-correlation. It is not possible to use an RF reference to synchronize the beam to lasers with femtosecond precision.

The beam position monitor developed in this thesis is ideally suited for use in a beam-based feedback system because it provides beam energy measurements with <1e-5 resolution for every bunch in the train as well as a measurement of the beam arrival-time in the chicane. A pair of beam arrival-time monitors installed up and downstream of the chicane provides lower resolution information about the energy changes of the beam and higher resolution information about the arrival-time of the beam at the chicane. Because of the systematic errors that both the chicane BPM and the arrival-time monitor suffer from when the beam shape changes, an ideal solution is to use both measurements to cross-check one another. No other existing monitors have the resolution to make such a high-resolution cross-check.

It is envisioned to use beam-based feedback to provide a small correction to the work that the cavity feedback is already doing. Under typical operating conditions, the proportional gain feedback is on and the adaptive feedforward is off. But in all tests of beam-based-feedback to-date, the proportional gain module feedback was always off and the adaptive feedforward was on [20]; the beam-based corrections were too large to be implemented without the adaptive feed-forward and the system was unstable with the proportional gain feedback on. This sort of beam-based feedback architecture, however,

creates a single point-of-failure system that is not robust enough for long-term operation, especially since the high resolution monitors have such a limited dynamic measurement range. Ideally, multiple systems should be used simultaneously, so that if one measurement is out-of-range, another system can step in. An architecture that takes beam-based and reference tracking information into account in the cavity controller is depicted below in Fig. 3.1.3.

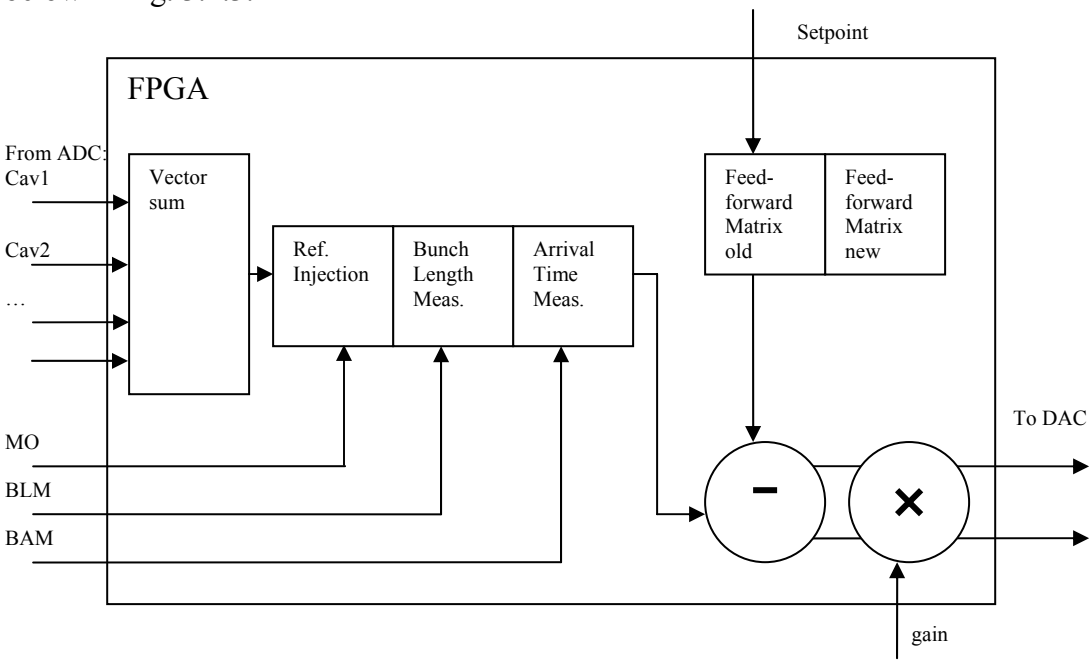

Figure 3.1.3 A desired FPGA algorithm structure incorporating reference injection and beam-based information.

So far, there has only been a description of what the module controller can do with respect to amplitude and phase stability, but we have not given a description of what it must do in terms of beam arrival time stability. This is determined by the relation of the accelerator RF parameters to the bunch compressor parameters and this is described in the following section.

3.2 Arrival-time Changes after a Bunch Compressor

The equation 3.2.1 gives a good representation of how an incoming arrival-time jitter, $\Sigma_{t0,}$ is altered by transport through an accelerating module followed by a bunch compressor (Fig. 3.2.1).

$$\Sigma_t^2 = \frac{1}{C^2}\Sigma_{t0}^2 + \left(\frac{R_{56}}{c_0}\frac{\sigma_A}{A}\right)^2 + \left(\frac{C-1}{C}\right)^2\left(\frac{\sigma_\varphi}{c_0 k_{rf}}\right)^2 \qquad (3.2.1)$$

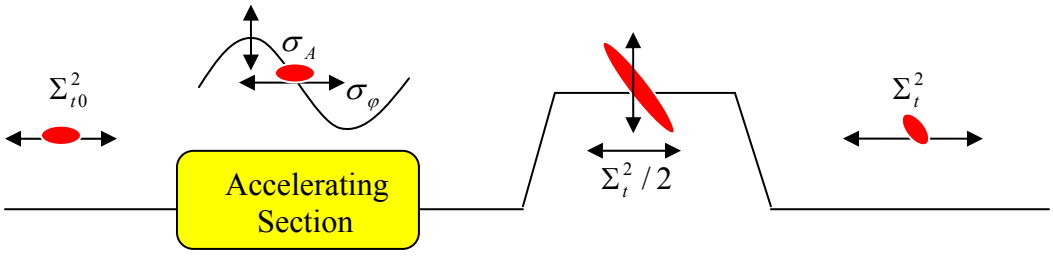

Figure 3.2.1 Transformation of arrival-time jitter with an accelerator section followed by a bunch compressor.

In Eq. 3.2.1, R_{56} is the longitudinal dispersion constant of the chicane, C is the compression factor of the bunch (defined in Sect. 2.3), A is the amplitude of the upstream accelerating voltage, φ is the phase of the accelerating gradient, and Σ_{t0} is the arrival-time jitter upstream of the accelerator section. It makes several approximations: the bunch is short relative to the wavelength of the RF, the initial energy is small compared to the energy after the accelerator section, the incoming energy chirp is small compared to the energy chirp gained in the first accelerator section and the jitter is statistically uncorrelated. It was first published in [25] and a derivation is written in the Appendix A.

The first thing to notice about the equation, reading it from left to right, is that the incoming arrival-time jitter, Σ_{t0}, is compressed in the chicane. The second thing to notice is that at FLASH, the amplitude stability, σ_A/A, of the klystron will be more critical than the phase stability, $\sigma_\varphi/c_0 k_{rf}$, for most typical values of σ_A/A and $\sigma_\varphi/c_0 k_{rf}$. This becomes more apparent when typical values from the first bunch compressor of FLASH (BC2) are inserted into the equation, giving:

$$\Sigma_{\text{after BC2}}^2 \approx (0.2\text{ps}*0.1\text{ps/ps}_{\text{injector}})^2 + (0.04*5.5\text{ps/\%}_{\text{gradient}})^2 + (0.01*2\text{ps/deg}_{\text{phase}})^2$$

With a compression factor of 10 in the first bunch compressor, the 200 fs injector jitter would be compressed to 20 fs arrival-time jitter after the bunch compressor. The 0.04% gradient stability of the first accelerator section would limit the arrival-time jitter downstream of BC2 to 220 fs. The 0.01 deg phase stability of the first accelerator section would limit the arrival-time jitter downstream of the first bunch compressor (BC2) to 20 fs. Given no additional contributions to arrival-time jitter from the second accelerator section amplitude or phase and further compression of the injector jitter by a factor of 2 in the second bunch compressor, the arrival-time stability at the end of the machine would be limited by the first accelerator section to 220 fs. An improvement by a factor of 10 in the gradient stability would make the arrival-time jitter contributions of the first accelerator section amplitude and phase approximately equal. The arrival-time stability at the end of the machine would then be 30 fs rms. This calculation has, however, ignored the second accelerator section.

Because of the assumptions of Eq. 3.2.1, one cannot simply use the equation recursively for the second accelerator section. For the second accelerator section, the incoming energy chirp is not small compared to the outgoing chirp and the incoming energy is not small compared to the outgoing energy. Eq. 3.2.2 describes the arrival-time jitter after the second bunch compressor for a beam that is on-crest in the second

accelerator section and it was first published in [26]; the derivation of Eq. 3.2.2 is written in the Appendix A.

$$\Sigma_t^2 = \frac{1}{C^2}\Sigma_{t0}^2 + \left(\left[R_{1,56} + \frac{E_1}{E_2}R_{2,56}\right]\frac{\sigma_{A_1}}{c_0 A_1}\right)^2 + \left(R_{2,56}\frac{E_2 - E_1}{E_2}\frac{\sigma_{A_2}}{c_0 A_2}\right)^2 + \left(\frac{C-1}{C}\frac{\sigma_{\varphi 1}}{c_0 k_{rf}}\right)^2$$

Where the $R_{1,56}$ is from the first bunch compressor and $R_{2,56}$ is from the second bunch compressor. The RF amplitude, A, and beam energy, E, are given with indices corresponding to whether they refer to the first or second accelerator section. $C=C_1*C_2$ is the compression factor for the combination of both bunch compressors, and $\sigma_{\varphi 1}$ is the phase stability of the first accelerator section. When the beam is on-crest in the second accelerator section, judging from Eq. 3.2.2, an arrival-time monitor after the second bunch compressor will measure primarily arrival-time jitter caused by the amplitude fluctuations of the first accelerator section for the following reasons:

- Injector jitter will be compressed by both bunch compressors (C_1=10, C_2=2)
- Arrival-time jitter induced by the first accelerator section amplitude jitter RF will not be compressed in the second chicane.
- The R_{56} of the second chicane is a factor of 5 smaller than that of the first chicane, so the amplitude jitter of the second accelerator section will make a smaller contribution to the total arrival-time jitter than the amplitude jitter of the first accelerator section.

These are the reasons that when a single arrival-time monitor after the second chicane was used to feed back on the amplitude jitter of the first accelerator section, it was able to stabilize the beam arrival-time to within 30 fs [20]. This is possible when the second accelerator section is operated on-crest, but not when it is operated off-crest. When the beam is off-crest in the second accelerator section, the arrival-time jitter from the first accelerator section is compressed in the second chicane making the first and second accelerator section arrival-time jitter contributions more equal after the second bunch compressor.

Given off-crest operation in the second accelerator section and a larger R_{56} in the second bunch compressor, the amplitude stability of the second accelerator section amplitude becomes almost as important to the timing stability of the beam as that of the first accelerator section. This can partially be seen by using Eq. 3.2.1 recursively for the second bunch compressor, but because the ratio of the incoming energy chirp to the outgoing energy chirp and the ratio of the incoming energy to the outgoing energy are not small, as assumed in the derivation of the equation, the prediction of the equation will be wrong by up to 30%, depending on the machine configuration.

In general, because a beam that arrives earlier or later on the falling slope of the RF wave will gain a lesser or greater amount of energy, whenever the accelerator section upstream of a bunch compressor is operated off-crest, it is not advisable to use the energy measurement from in the chicane or the arrival-time measurement from after the bunch compressor to directly feed back on the upstream accelerator section without first taking into account the effect of incoming arrival-time jitter on the quantity that is measured. Although, for large compression factors, the incoming arrival-time jitter may be compressed enough in the chicane that it can be ignored, this is not always the case. It is

not possible to disentangle which energy change was caused by incoming arrival-time jitter and which energy change was caused by accelerating gradient and phase jitter unless the arrival-time jitter generated upstream of the accelerator section has been measured.

There are two different strategies to deal with this problem. One could use an accelerator section upstream of each chicane in order to stabilize the arrival-time after each chicane, regardless of how much the beam energy jitter is increased. One would then need to stabilize the beam energy using an accelerator section located after the chicanes. Alternatively, one could use a single accelerator section gradient setpoint to simultaneously stabilize the beam energy and arrival-time after each chicane. The end-result of both schemes would, in principle, be the same.

For the sake of machine stability, the author believes that the latter option is better: a feedback on the first accelerator section should not respond to changes in the injector jitter and a feedback on the second accelerator section should not respond to changes in the first accelerator section. To simultaneously stabilize beam energy and arrival-time after a bunch compressor, one could execute a combination of the following:

- use measurements of the arrival-time jitter upstream of a bunch compressor to keep the energy/arrival-time feedback from responding to the energy/arrival-time jitter that it creates.
- stabilize the arrival-time jitter upstream of the bunch compressor before correcting the energy/arrival-time jitter downstream.

3.3 Beam-Based Feedback Strategy

Two schematics of the synchronization sensitive components in the machine are shown in Fig. 3.3.1 [27]. An optimal feedback setup is depicted in Fig. 3.3.1(a.) and a more quickly realizable architecture is depicted in Fig. 3.3.1(b.). The injector laser (Laser), the injector RF (Gun), the super-conducting accelerator sections (ACC1-7) and the third-harmonic module (third-) are depicted in block diagram format with arrows connecting various optical cross-correlators (OCC), chicane beam position monitors (EBPM), bunch length monitors (EO-1D, THz-1D), and beam arrival-time monitors (BAM) to the digital processing boards (μTCA or SIMCON-DSP) with which they would be connected in a beam-based feedback system that controls, on a bunch-to-bunch basis, the amplitude and phase (A, ϕ) of a normal-conducting RF cavity (3GHz NRF cavity), the 1.3 GHz super conducting acceleration cavities (1.3GHz SRF ACC1(2&3)) and the super conducting third- harmonic linearization cavity (3.9GHz SRF).

The reason that the system shown in Fig 3.3.1(b.) will be built before that shown in Fig 3.3.1(a.) is that the μTCA crate system, shown in blue in the figures below, along with the corresponding ADC and FPGA boards will not be available in 2010. VME is the crate system that has been used at FLASH since its inception but will be phased out as μTCA crates become viable. In the currently available VME crate infrastructure, the beam arrival-time is calculated on an in-house built Analog Carrier Board (labeled ACB in the figure) that contains ADCs, delay-chips and FPGAs. Beam arrival-time information from this board is delivered to the cavity controller via an optical Gigalink. The cavity controllers reside on VME based SIMCON-DSP boards that each have ADCs, DSPs, DACs, and an FPGA.

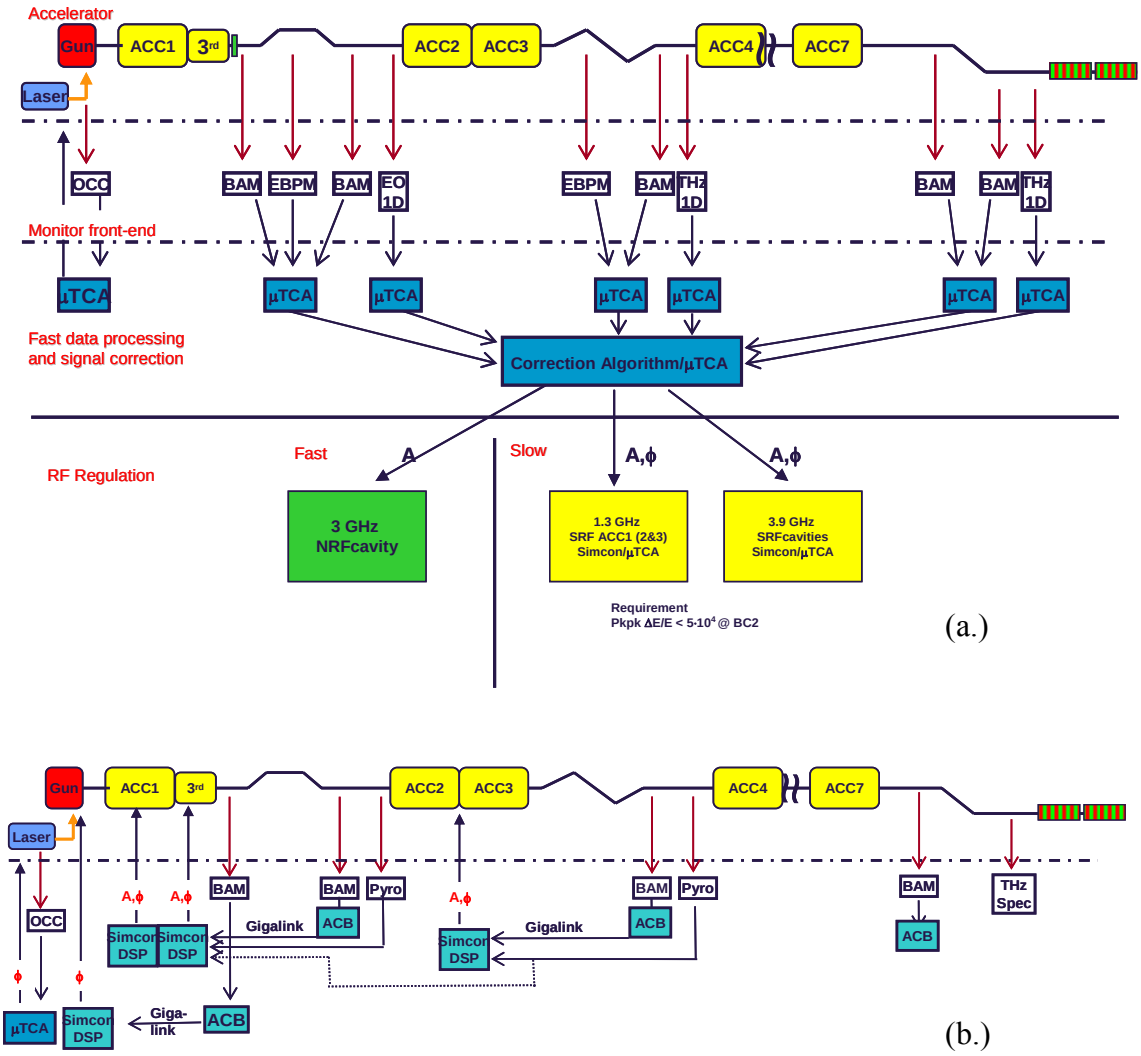

Figure 3.3.1 Layout of synchronization sensitive components at FLASH, along with desired feedback loops in an optimal configuration (a.) and a more expedient configuration (b.). Schematics taken from in-house presentation of Holger Schlarb.

In Fig. 3.3.1(a.), the arrival-times of the photo-injector laser pulses relative to the Master Laser Oscillator (MLO) pulses are measured in an optical cross-correlator. The arrival-times of the photo-injector pulses can then be adjusted with a vector modulator that controls an electro-optical modulator in the actively mode locked injector laser cavity. This can be done at a speed of 27 MHz with the use of a µTCA crate with an ADC, a DAC, and an FPGA installed within. All of the RF cavities' phases and amplitudes would then be influenced by a central correction algorithm operating on a single crate that collects the beam energy, arrival-time and compression information from monitors throughout the machine and delivers corrections to the cavities' controllers. The

super conducting cavities' controllers would be sent commands from the central controller on an intra-bunch-train basis with a moderate bandwidth, while the normal-conducting RF cavity amplitude could be sent commands on a fast, bunch-to-bunch basis. Each super-conducting cavity would have its own, independent, fast, vector-sum controller feedback loop using multiple cavity pickups as the diagnostic references. All of this would be accomplished with a newly developed μTCA crate system; this is the crate system that will replace the existing VME infrastructure in the coming years. The advantages of the system from Fig. 3.3.1(a.) include

- the ability to use a high-level system identification algorithm to tune and stabilize the entire machine,
- built in cross-checks and redundant measurements that use different techniques to measure the same quantities,
- robustness afforded by distributed cavity controllers that do not require the central controller to operate.

The disadvantage is that with a centralized decision making process, the latency of the signal transport increases due to the multiple digital processors and long cables to and from the central decision making crate. The latency problems can be offset with the use of a normal conducting cavity with a low quality factor in which the accelerating field can be changed quickly. The alternative, more expedient architecture using existing hardware and distributed control loops, shown in Fig 3.3.1(b.), requires fewer digital processors and cable lengths that are shorter, however, the high quality-factor of the super-conducting cavities used as actuators limits the speed with which the energy of each bunch can be adjusted. This means that the first 10 or more bunches of the train are un-stabilized.

In Fig 3.3.1(b.), the photo-injector laser phase is synchronized to the Master Laser Oscillator (MLO) in the same manner depicted in Fig 3.3.1(a.), but that is where the similarity ends. The injector RF phase is stabilized with a cavity controller that incorporates beam-based feedback from a beam arrival-time monitor downstream of the first accelerator section, but upstream of the first chicane. The first accelerator section amplitude is stabilized with a cavity controller that incorporated beam-based feedback from an arrival-time monitor located directly after the first chicane. The first accelerator section phase is stabilized with a bunch-length monitor located after the second chicane. The second accelerator section amplitude and phase are stabilized in a similar fashion with an arrival-time monitor and a bunch length monitor after the second chicane.

The system shown in Fig 3.3.1(b.) has the following disadvantages

- If an upstream feedback loop fails to deliver acceptable beam stability, the downstream loop will start to feed back on noise that is generated some where other than in the cavity it is controlling.
- The downstream feedback loop has to be slower than the upstream loop in order to avoid instabilities.
- It is not able to make use of cross-checks, monitors that measure the same quantities in different ways.
- The energy changes due to arrival-time jitter from upstream of a bunch compressor are not subtracted from the energy changes measured after the bunch compressor.

- The lack of a normal-conducting cavity limits the speed with which the beam energy can be changed. As a result, the first 20 bunches in the train will have a different energy than the stabilized bunches that follow

Because the scheme depicted in 3.3.1(a.) will take 2-3 years to be realized, the scheme shown in 3.3.1(b.), to be commissioned in the coming months, will be the focus of the following sections.

3.4 Injector Jitter

The injector needs to be stabilized on two fronts: the laser timing jitter and the cavity RF phase jitter. The present infrastructure synchronizes both devices to an RF reference and the delivery of this reference is subject to drifts and noise. The infrastructure described in the previous section synchronizes both devices to an optical timing reference for which drifts have been actively compensated.

The injector laser is actively mode locked and has an electro-optical device within the laser cavity that regulates the arrival-time of the laser pulses at the cathode. These electro optical modulators (EOMs) are driven with the 1.3 GHz reference from the master oscillator. Temperature changes and noise picked up by the cable and amplifier that bring the master oscillator signal to the EOM have, of course, an impact on the phase stability and drift of the laser. Temperature changes and noise picked up on the cables involved in cavity field regulation will likewise have an impact on the phase stability of the field in the cavity.

The jitter of the RF phase relative to the phase of the laser was most accurately measured in [28]. In this measurement, the phase of the laser was changed with a vector modulator that acted on the master oscillator signal feeding the EOM such that the beam arrived on the falling slope of the cavity RF signal. At this phase, the changes in the RF phase relative to the laser phase produced a change in beam charge that could be measured with a downstream toroid. By scanning the phase of the laser relative to the RF phase, a calibration of the beam charge dependence on the phase relationship between the RF and the laser could be determined. Multiplying this calibration by the charge fluctuations measured at the toroid gave a measurement of the phase jitter between the laser and the cavity RF. This jitter was larger than 0.5 degrees from pulse to pulse, a quantity that requires significant improvement. Nevertheless, it is of little use to the machine if the laser and cavity are locked together if they are drifting or jittering relative to a downstream reference.

Relative to the optical timing reference, the injector laser timing jitter can be measured by optically cross-correlating the injector laser pulse with a pulse from the optical reference (Fig. 3.4.1) [29]. This measurement can then be used to feed back on the phase of the RF signal sent to the EOM in the injector laser cavity and thereby stabilize the injector laser timing relative to the optical reference. This can be done at a speed of 100 kHz, thereby counteracting much of the noise produced in the amplification of the reference RF signal. The RF module phase jitter can then be stabilized by measuring the arrival-time of the beam relative to the optical reference and then feeding back on the phase setpoint of the RF module. It does not, however, make sense to stabilize the cavity phase in this manner without first stabilizing or measuring the laser jitter because any beam-based feedback would respond to both RF module phase jitter and laser jitter.

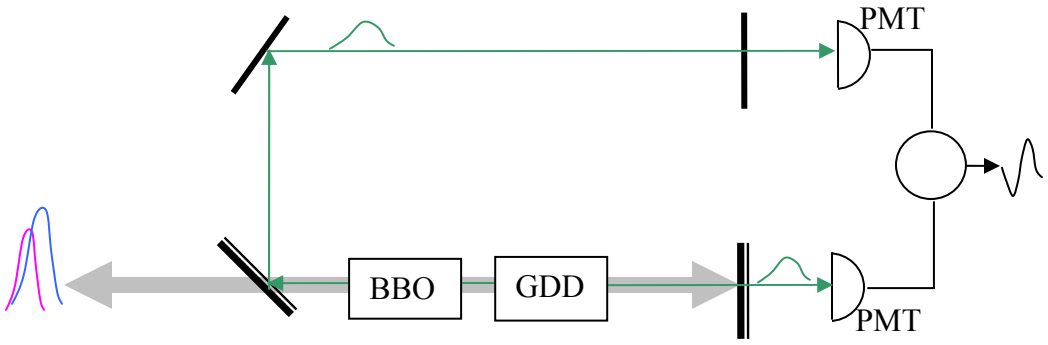

Two laser pulses: 800 nm 1550 nm	Dichroic mirror: Reflects 527.7 nm Transmits 800 and 1550nm	BBO crystal: second harmonic light generated when pulses overlap.	GDD: Generate a group delay a dispersive medium	Dichroic mirror: Transmits 527.7 nm Reflects 800 and 1550nm

Figure 3.4.1 Measurement of the arrival-time of the injector laser pulse relative to a timing laser reference (MLO) using a two-color single-crystal balanced optical cross-correlator. From S. Schulz, *Synchronization of Injector Laser and Master Laser Oscillator* PAC'10.

In Fig. 3.4.1, the arrival-time of the injector laser pulse is measured relative to the optical reference laser in an optical cross-correlator. In a two-color optical cross-correlator, two laser pulses with different colors (800 nm and 1550 nm) are sent through a dichroic mirror that reflects the sum frequency (527.7 nm) and transmits the higher frequencies. The input laser pulses are then sent through a BBO crystal. When the pulses overlap in the crystal, new pulses with the sum frequency are generated and emitted in both forward and backward directions. With the aid of dichroic mirrors and a group delay generated in a dispersive medium, the pulses generated in the crystal each travel to a photomultiplier tube (PMT). The incoming laser pulses return from whence they came. In this balanced detection arrangement, the measurement of the relative arrival-times of the two input pulses is insensitive to laser noise and is background and drift free.

3.5 Third-harmonic Module Jitter

The third-harmonic module is comprised of 4 cell cavities which are filled with 3.9 GHz. It is used to linearize the energy chirp of the beam. In the first bunch compressor, the energy of an electron at position z in the bunch is

$$E_1 = E_0 + V_1 \cos(k_{rf} z + \Phi_1) + V_3 \cos(3k_{rf} z + \Phi_3) \qquad (3.5.1)$$

where (V_1, Φ_1, V_3, Φ_3) are the RF amplitudes and phases of the first accelerator section and the third- harmonic module. As in Eq. 1.7, the path-length through the chicane as a function of the energy of a given particle is written by

$$L(E) = L_0 + R_{56}\frac{E - E_0}{E_0} + R_{566}\left(\frac{E - E_0}{E_0}\right)^2 \qquad (3.5.2)$$

where E_0 is the energy of the bunch center ($z=0$). The first and second order chirps of the beam energy as a function of the compression factor C are [30]

$$E' = \frac{E_0}{R_{56}}\left(1 - \frac{1}{C}\right) \qquad E'' = -2\frac{R_{566}}{R_{56}}\frac{(E')^2}{E_0}. \qquad (3.5.3)$$

Taken all together, we have three equations with four free parameters, (V_1, Φ_1, V_3, Φ_3), in Eq. 3.5.1. To optimize this system, the second-order energy chirp will be compensated by an appropriate phase and amplitude setting in the third-harmonic module.

Under the current plan for the operating parameters, the gradient of the third-harmonic cavity is only a 9th of the gradient of the downstream accelerator sections, so the amplitude jitter contribution will be only a 9th of that of the downstream accelerator sections. This means that it doesn't make sense to stabilize the gradient of the third-harmonic module until the down-stream accelerator sections' gradient stabilities are improved by a factor of nine. Because this has yet to be accomplished, the third-harmonic module is a lesser worry. Schemes that attempt to compensate for the third-order energy chirp call for large amplitudes in the third- harmonic module and this would begin to have an impact on the arrival-time stability.

The jitter contribution of the phase of the third-harmonic module depends much more dramatically on the setpoint of the module. Some settings make the module the dominant contributor to bunch length jitter, while others make it the weakest contributor. It can be argued that the ideal setting for timing jitter and peak current considerations is one for which the phase jitter contributions of the 3.9 and 1.3 GHz modules are approximately equal [31].

3.6 First Accelerator Section Jitter

The first accelerator section needs to be stabilized on two fronts: the phase and the amplitude. The phase in this section determines primarily the bunch length, while the amplitude strongly affects the beam arrival-time and the energy.

Following the argument given in Sect. 3.1 about arrival-time stability after a chicane, the desired energy stability in the first accelerator section is approximately 0.004%, a factor-of-ten improvement over the current 0.04% stability. A monitor for a feedback system should be at least two times better than the stability that it hopes to achieve, and so the monitor for an energy feedback should resolve 0.002% energy changes. By this logic, with an R_{56} of 180 mm (550 ps) and an R_{16} of 345 mm, an arrival-time measurement of the beam's time-of-flight through the first chicane of FLASH should resolve 10 femtoseconds and a position measurement in the chicane should resolve 7 μm (21 fs). It follows that the resolution required by the position measurement is lower than that required by a time-of-flight arrival-time measurement in proportion to the ratio between the R_{16} and R_{56} terms. In the case of the first bunch compressor of FLASH, this ratio is 2:1, in favor of the position measurement, whereas for the XFEL the

ratio is 6:1. Future FLASH configurations also call for a reduction of the R_{56} in the first chicane by a factor of two, increasing the resolution requirements of BAMs used in a time-of-flight energy measurement, while not affecting the BPM resolution requirements.

As argued in Sect. 3.2, while we know that the 200 fs injector timing jitter is compressed in the chicanes, it is not compressed enough to make it negligible, especially if the compression factor of the first chicane is reduced from 10 to 5, as is planned for upcoming operation with the third- harmonic module. The injector jitter must, therefore, be measured and incorporated in any energy/arrival-time feedback scheme.

While the arrival-time is stabilized using the accelerating amplitude as an actuator in a feedback, the bunch length is stabilized using the accelerating phase as an actuator in a feedback. The bunch length stability is less of an issue of the total bunch length, but more of an issue of the lasing bunch length, namely the tiny slice of the bunch that lases in the undulator section. The length of this is ~15 μm and no existing monitor can truly resolve it, but one, in particular, can, in practice, stabilize it. An array of pyrodetectors located after a diffraction screen can be arranged with various filters in order to produce single-shot spectra of all of the bunches in the bunch train [32]. Previous spectrometers required multiple shots as a delay stage moved over several centimeters. The single-shot measurement detects fewer frequencies than the scanning measurement and cannot reconstruct the longitudinal profile as well as the scanning measurement, but this is offset by the advantage of having an unchanging beam emittance, position, and profile over the course of the spectrum measurement. Moreover, it is not mathematically possible to accurately reconstruct the asymmetrical longitudinal profile of the beam with either method. This is also not important because, when the beam is lasing, certain frequency components measured with the single-shot spectrometer become stronger, and if they are maintained at a constant level with a feedback on the accelerating phase, then the beam has been shown to lase at a more constant level [20].

There is a complication that has not yet been mentioned, namely the the energy jitter that is caused by accelerator phase jitter. When the R_{56} is small in Eq. 3.2.1, the contribution of the amplitude jitter of the accelerating section to the overall timing jitter becomes smaller relative to the contribution of the phase jitter to the timing jitter. Using information from the bunch compression monitor, it is then necessary to disentangle which energy changes are caused by accelerating amplitude jitter and which are caused by accelerating phase jitter. To do this, one must consider the transformation

$$\begin{pmatrix} \Delta t \\ \sigma_z \end{pmatrix} = \begin{pmatrix} R_{16} & \frac{C-1}{C}\left(\frac{1}{ck_{rf}}\right) \\ \approx 0 & M \end{pmatrix} \begin{pmatrix} \frac{\Delta A}{A} \\ \sigma_\varphi \end{pmatrix} . \tag{3.6.1}$$

The first line of this transformation should be familiar from Eq. 3.2.1. and the second line might be derived along the lines of Eq. 2.3.9, but in practice, the quantity M must be measured. To use this transformation to generate feedback commands for the upstream accelerating section, the inverse of the matrix must be found and used in the following way,

$$\begin{pmatrix} R_{16} & \frac{C-1}{C}\left(\frac{1}{ck_{rf}}\right) \\ 0 & M \end{pmatrix}^{-1} \begin{pmatrix} \Delta t \\ \sigma_z \end{pmatrix} = \begin{pmatrix} \frac{\Delta A}{A} \\ \sigma_\varphi \end{pmatrix}. \tag{3.6.2}$$

3.7 Second Accelerator Section Jitter

When operated on-crest, the amplitude jitter of the second accelerator section affects the arrival-time stability in a way that is similar to that of the first accelerator section, only the effect is smaller because the R_{56} of the second bunch compressor at FLASH has typically been five times smaller than that of the first bunch compressor. This means that the conversion of energy jitter into timing jitter has been five times less dramatic. This also reduces the required sensitivity of the second bunch compressor monitors by about a factor of 5 compared to that of the monitors in the first bunch compressor chicane. In future machine configurations, however, the compression factors of both compressors will be more equal. In these future machine configurations, the R_{56}s will be reduced, thereby increasing the resolution requirements of beam arrival-time measurements for both the first and second accelerator section energy jitter measurements.

In future compression schemes, the second accelerating section will be operated off-crest with a larger R_{56} in the following bunch compressor, making the trouble with arrival-time jitter upstream of the chicanes the same as it is in the case of the first accelerator section. As in the case of the first accelerator section, if the arrival-time jitter upstream of the second chicane is not small enough to ignore, an arrival-time monitor upstream of the chicane must be used in any energy/arrival-time feedback scheme.

3.8 Outlook

While two beam arrival-time monitors (BAMs) placed one after the other in a drift section have been shown to produce 9 fs jitter relative to one another over few-minute time-scales, if the longitudinal profile of the beam changes slightly, the arrival-time measurement will change by up to ~100 femtoseconds [20]. While the bunch head is only ~25 fs long at the end of the machine [32], the picosecond-long beam tails that result from non-linear compression [33] (i.e. without the third- harmonic module) cause the arrival-time measurement to be more sensitive to changes in the beam profile that it would be without a long tail which changes in length when the RF phase changes. Changes in the beam profile limit the stability of the measurement under stable beam conditions to 30 fs (rms) over 7 hours [20]. But, the beam profile is not the only source of measurement error. Although the BAM equipment is in a temperature stabilized enclosure, the arrival-time measurement will drift by ~2 femtoseconds when the temperature of the fibers changes by 0.1 degree. A beam energy measurement done by using BAMs to measure the time-of-flight through the chicane will be especially sensitive to these temperature changes because it is carried out with two measurements separated by >10 meters.

The chicane BPM is less sensitive to temperature changes because it is carried out within a single temperature stabilized enclosure, but it is possibly more sensitive to beam profile changes than the BAM. Such sensitivity to beam profile changes can be a good thing if it is related to the stability of the peak current, but it can be a bad thing if it is unrelated. Simulations of these effects are presented in Chapter 6, but when more BAMs are commissioned, a true comparison can be made. If the chicane BPM is used to measure beam energy changes, the incoming orbit jitter must be measured using BPMs

from before and after the chicane. It is important that these BPMs have resolution that is comparable to that of the chicane BPM and that there are not any quadrupoles between these BPMs and the first or last dipole of the chicane. BPMs with <5 μm resolution were installed before and after the first bunch compressor for this very purpose. Initial attempts to benchmark the available beam energy measurements are made in Chapter 10 of this thesis, but most measurement devices of interest were only commissioned for a couple weeks during the course of these studies.

In general, the measurements that deliver the highest resolution have a limited dynamic range and require a mechanical delay line to accommodate larger changes. The designs of these monitors will be presented in Chapters 7, 8, and 9. Experience with these monitors gained in the course of producing the measurements in Chapter 10 showed that they have an intrinsically high failure rate when the beam conditions are unstable. If any beam parameter changes significantly, the monitor can be pushed out of range and require mechanical movement and recalibration, all of which can take several seconds to complete. Such a beam-based monitor can only be used to add small corrections to the more robust cavity field measurement based regulation.

4 Beam Shape in the Bunch Compressor

Because the beam position monitors of interest for this thesis are located in the middle sections of the chicanes, where the dispersion is at a maximum, it is important to understand the likeliest shapes of the electron bunches at these locations for various machine configurations. The impacts of various beam shapes on the chicane BPM performance will be described in chapter 6.

The transport of the beam has been simulated with a code called ASTRA [34] and with the use of transfer matrices [12]. ASTRA was developed for the space-charge dominated beams that one finds in the injector and will, in the following simulations, be used to generate a charge distribution at the exit of the RF gun. This charge distribution will then be transported to the middle of the chicane by multiplying together transfer matrices with Matlab. Simulation results will be compared to measurements of the beam done with the beam image from the synchrotron light monitor screen.

4.1 Perfect Alignment

Ideally, the beam is centered on the cathode, travels straight through the RF gun and solenoid field, and straight through the accelerator section, thereby minimizing the effects of wakefields and dispersion. The beam is transversely round in the cavities and when it reaches the middle of the chicane, if it has been accelerated on the slope of the accelerating wave, it is elongated in the *x*-*y* plane. There is, however, a curvature to the beam due to the curvature of the accelerator RF. It results in inhomogeneous compression, consisting of a sharp leading spike and a long trailing tail. A third-harmonic module has recently been installed upstream of the bunch compressor in order to remove this curvature by sending the beam through a set of cavities that operate at thrice the frequency of the accelerating cavities. By selecting the phase and amplitude of the third-harmonic module appropriately, the bunch will be compressed homogeneously.

The beam shape in the second bunch compressor is also impacted by the addition of the third-harmonic module. The slice emittance of the tail is small, but due to over-compression, transverse tails are generated at the spike of the bunch [32]. This, in addition to the long, longitudinal tail, causes the projected emittance to be significantly larger than the slice emittance. With the third-harmonic module, the projected emittance will be reduced, hopefully enabling a larger portion of the beam to lase

With a non-linear energy chirp, the portion of the beam that lases is but a fraction of the sharp, leading spike. While the duration of the spike has been measured with a transverse deflecting cavity (a sort of streak camera) and is <60 fs FWHM [32], measurements of the bunch spectrum suggest that the fraction of the bunch that is responsible for the lasing process is closer to 25 fs in duration [33]. With the addition of the third-harmonic module, the beam will acquire a linear energy chirp. With a linear energy chirp, more of the beam will acquire the charge density and emittance characteristics necessary for lasing, but this is only true if the projected transverse emittance from the injector is sufficiently small, a requirement that can be met with a perfectly aligned injector.

4.2 Mis-aligned Injector

The effect of a mis-aligned injector can be best observed on a beam that is accelerated on-crest. This is due to the fact that the minimal energy spread minimizes the effects of dispersion downstream of the injector. It is often observed on an OTR screen in the middle of the first bunch compressor that when the beam is on-crest, the shape on the screen is not round, but rather like a c or a boomerang. The head and the tail of the bunch are offset in the y-plane, as shown in Fig. 4.2.1.

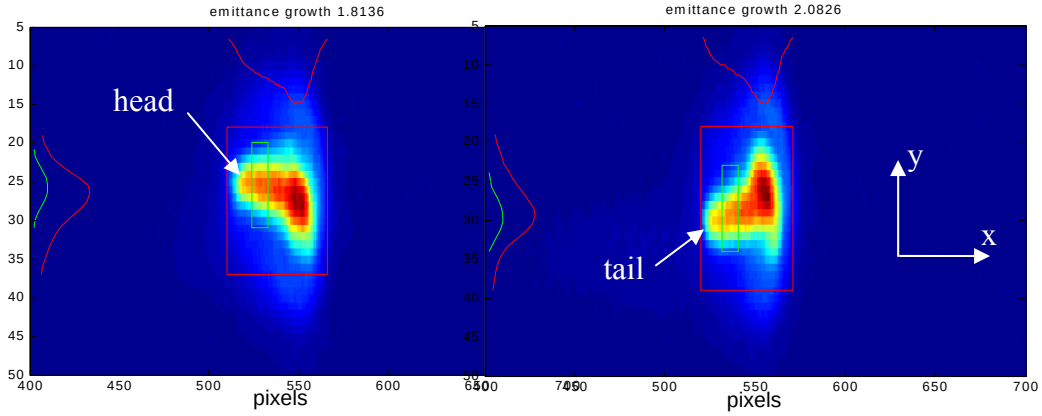

Figure 4.2.1 The shape of the beam in the bunch compressor for on-crest operation, as viewed on the synchrotron light monitor in the first bunch compressor. As the phase of the upstream accelerator section is moved slightly off-crest relative to the longitudinal center of the bunch, the head of the bunch is visible sticking out above (left). When the phase is changed slightly in the opposite direction, the tail is seen sticking out below (right). This head-tail separation of 400 μm accounts for a projected emittance growth of 1.8.

There has been some debate over the cause of this shape, but not much interest in removing it. Because only a slice of the bunch lases in the present machine, without the third-harmonic module, the shape of the rest of the bunch has been irrelevant. This situation, while true now, will not be true when the third-harmonic module is commissioned and the projected emittance becomes almost as important as the slice emittance. Some wondered if the c-shape was the direct result of often simulated but never measured coupler kicks in the first accelerator section. Others believed it was from oft-simulated but never measured wakefields, due to the mirror in the injector. The contribution of linear dispersion induced downstream of the RF gun but before the first accelerator section would be too small to create this effect all by itself, as would the effect of a tilted first accelerator section. To put these effects in a frame of reference, the strength of each effect will be quantified below.

Coupler kicks are caused by the forces that the beam experiences when it passes the RF couplers, ports from which the accelerating cavities are filled up with RF waves or from which higher-order modes are removed. The kick arises because the coupler introduces a field asymmetry into the structure. There are two HOM couplers per module, one power coupler, and eight cavities per module. The HOM couplers are oriented in such a way that the kicks of pairs of upstream and downstream couplers should partially cancel one another out, making the power coupler the coupler of concern (Fig. 4.2.2).

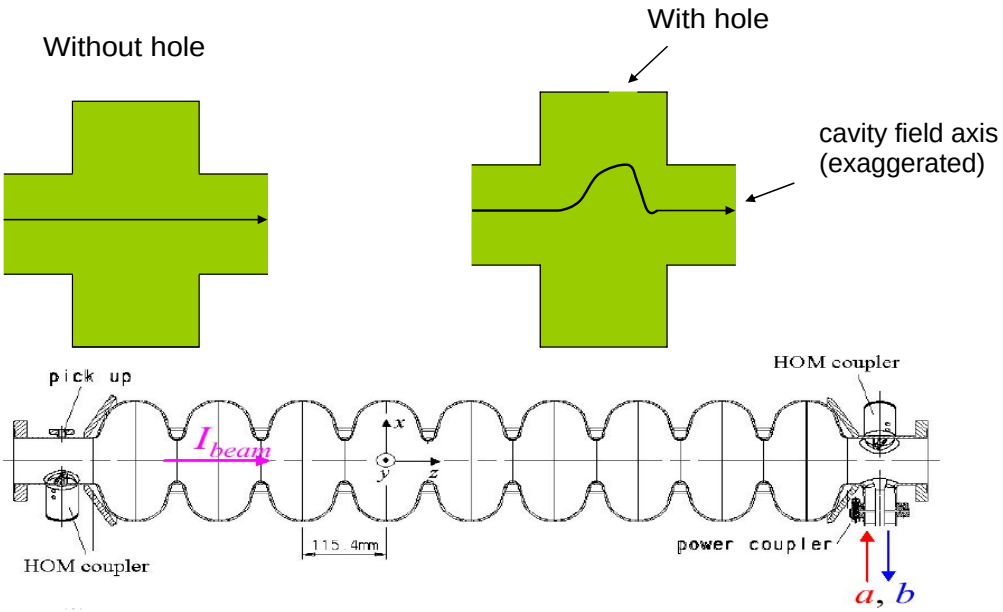

Figure 4.2.2 Coupler kick concept. A hole in a cavity changes the axis of the cavity field. There are two sets of couplers per accelerating module. They are oriented in such a way that the kicks of pairs of HOM couplers should partially cancel one another. The power coupler produces the strongest kick and it scales with the gradient in the cavity.

Estimates for the strength of the coupler kicks in a steady state situation have been made using software that calculates the fields in cavities of arbitrary shape [35]. The coupler

geometry and results from this simulation are shown in Fig. 4.2.3. One can use the voltages V_x and V_y along with the accelerating voltage $V_{||}$ in order to predict the strength of the kicks for different gradients. The voltage is given in complex notation in order to take into account the dependence of the kick on the phase of the RF. The main effect is a time-varying dipole kick for a situation in which the forward power and reflected power have reached a steady state (reflected=0). The predictions made in [35] were incorporated into a beam transport simulation in order to check if it is possible to measure the kicks with existing diagnostics. In this simulation, the cumulative effect from each coupler of 20-30 μrad tilt was so weak that it could not be measured with existing diagnostics. The separation between the head and tail due to the coupler kicks alone is only 17 μm after the first accelerator section.

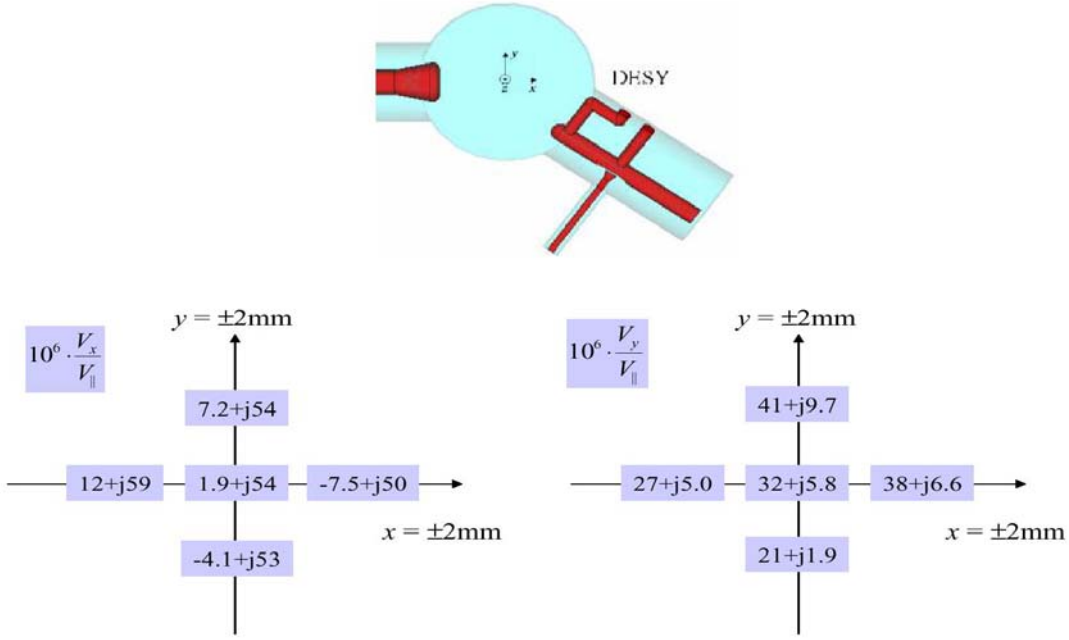

Figure 4.2.3 Coupler geometry with pickups (top) and voltage of kick (bottom). V_x and V_y are proportional to the accelerating voltage $V_{||}$. The voltage is given in complex notation in order to take into account the dependence of the kick on the phase of the RF. The main effect is a time-varying dipole kick for a situation in which the forward power and reflected power have reached a steady state (reflected=0). From M. Dohlus, *Field Asymmetries and Kicks* [35].

A beam that travels diagonally through an accelerator section will acquire a tilt due to the difference in the orientation of the field axis relative to the beam path. The vectors contributing to this tilt are shown in Fig. 4.2.4 and are defined below in terms of V_{cav}, the accelerating voltage of the cavity, $V_{||}$, the portion of the accelerating voltage in the beam direction, and V_T, the portion of the accelerating voltage acting on the beam in the plane of the module's tilt:

$$V_{T0} = V_{cav} \sin\theta$$

$$V_T = V_{T0}\cos(k_{rf}\cdot z+\phi)$$
$$= V_{T0}\left(\cos\phi - k_{rf} z\sin\phi - \tfrac{1}{2}\left(k_{rf} z\right)^2\cos\phi...\right)$$
$$\Delta y' = V_T(z)/E$$

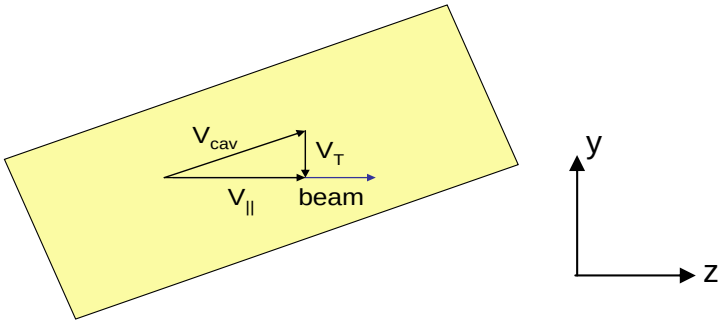

Figure 4.2.4 Voltages acting on a beam as it travels in *z* through a cavity tilted in the *y*-*z* plane.

where z is the beam direction in the internal coordinates of the beam, y is the direction of the tilt, θ is the angle of the tilt, ϕ is the phase of the cavity RF field, and k_{rf} is the wavenumber of the cavity RF field. Using this description of the cavity, we can estimate that for a 1 mrad cavity tilt at 25 MeV where the beam length is 3 degrees of the cavity wave at a phase of 10 degrees off-crest, the transverse accelerating voltage is 25 keV. This would produce a 227 V potential for one sigma of the beam dimension, corresponding to bunch tilt of ~10μrad. This is even smaller than the effect of coupler kicks.

Wakefields tend to magnify any particle offsets, driving the head and tail even further apart and they are best described by the, so called, wake function which is the Fourier transform of the coupling impedance. The wake function is an integral of the electric field of a particle over an accelerator segment of length L. For longitudinal wakefields [36],

$$w_{\|} = \int_L E_{\|}(\mathrm{s})\mathrm{ds}\ ,$$

where $E_{\|}$ is the electric field parallel to the beam direction. For transverse wakefields per unit-transverse-offset, Δu, one can see the dependence of the wake on the offset of the beam from the center of the chamber

$$w_{\perp} = \frac{1}{\Delta u}\int_L \left[E + c(\beta\times B)\right]_{\perp} ds$$

Convolving the wake functions with the longitudinal charge distribution, *q(s)*, will give the wake potential

$$V(s') = \frac{1}{Q}\int_{s'}^{\infty} w(s-s')q(s)ds$$

Integrating over all slices *ds'* gives the total energy change of the bunch due to the wake.

$$L(s') = \frac{1}{Q}\int_{-\infty}^{\infty} q(s')ds'V(s')$$

This energy change amounts to a transverse kick or a longitudinal energy loss. In general, the transverse wakes are stronger for long bunches and the longitudinal wakes are stronger for short bunches. The primary wakefield effect that one would see in the injector comes from the geometric wakes of the accelerating cavities [37]:

$$w_{\parallel}(s) = -344\exp(-\sqrt{s/s_0})[\text{V/pC/module}],$$
$$\text{w}_{\perp}(s) = 1000\left(1-\left(1+\sqrt{s/s_1}\right)\exp(-\sqrt{s/s_1}\right)[\text{V/(pC}\cdot\text{m}\cdot\text{module)}]$$

where s_0=1.74e-3 m and s_1=0.93e-3 m. For a beam that travels directly down the center of the acceleration module, these effects are negligible, but when large orbit deviations are present, the effect can become significant. Nevertheless, without any other effects, geometric wakes in the first accelerator section could only produce a tenth of the tilt that it is observed on the screen in the bunch compressor for a large 4 mm orbit bump through the first accelerator section.

When the beam has an energy chirp, the effect that usually influences the beam tilt more than any other is that of dispersion. When a beam with an energy chirp travels through a dipole field, like that of a corrector, an offset quadrupole magnet or an offset focusing field of an accelerating structure, the lower energy particles in the head of the bunch are deflected more than the higher energy particles in the tail of the bunch and therefore acquire an offset relative to the lower energy particles. The head and the tail of the bunch get further apart in that deflecting plane (Fig. 4.2.5).

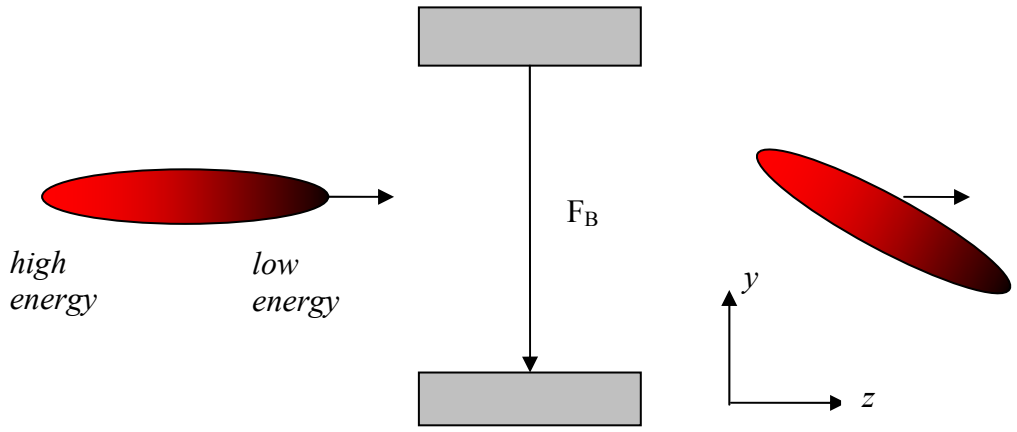

Figure 4.2.5 The effect of a dipole magnet on a beam with an energy chirp.

Upstream of the first accelerator section, where there is a ~1% energy chirp and there are no quadrupole magnets, only a couple of correctors, the contribution to the beam tilt from dispersion for a 2 mm offset out of the gun would only be a 20 μm vertical separation between the head and tail of the beam. For a beam of 2 mm length, this corresponds to a tilt of 170 μrad. This is about as strong as the wakefield effects, but even in combination with the wakefields and other effects, it is still too small to create the >400 μm of head-tail separation seen in BC2 at the phase advance for which the tilt is maximized.

After some investigation, simulating and measuring the effects listed above, it was concluded that the c-shape seen on the OTR screen in the dispersive section of the bunch compressor for on-crest operation is due to a badly (several mm) mis-aligned solenoid. The first hint of this was that in ASTRA beam transport simulations, no combination of wakefield, coupler kick, and dispersive effects was strong enough to cause the shape seen on the screen, but if the solenoid was given a 1 cm offset from the beam axis, the beam shape that would be clearly visible on the screen in the dispersive section of the first bunch compressor would be that of a C. The solenoid alignment was not previously suspected as the cause of the c-shape because a lot of trouble is generally taken to align the solenoid with sub-mm precision [38]. Measurements confirming the conclusion that the solenoid alignment was the culprit involved changing the position of the iris on the cathode. This changed the position of the electron beam within both the RF module and the solenoid. As the electron beam position relative to the module and solenoid was changed, changes in the maximum separation of the head and tail of the beam were observed on the bunch compressor screen. The maximum separation of the head and tail was determined for each iris position by scanning the current of a quadrupole in order to find the phase advance for which the head-tail separation was maximized. When the charge was changed between 0.5 nC and 3 nC, the c-shape did not change appreciably, ruling out the influence of wakefields. If the gradient of the first accelerator section was changed, the c-shape was unaffected, ruling out the influence of the coupler kicks. The c-shape was present both before and after a major injector upgrade. In fact, it was more pronounced after the upgrade. This suggests that, the screen in the bunch compressor should be used as a diagnostic to verify proper solenoid alignment.

With transport through even a perfectly aligned accelerator section and lattice, any offsets gained in the injector are magnified. This is clear from the simulation shown in Fig. 4.2.6, in which the offset of a beam is tracked with ASTRA from the cathode through the first accelerator section. An offset of 100 μm in the horizontal plane (X) at the cathode becomes 140 μm in Y and 40 μm in X at the exit of the first accelerator section. As the beam travels through the accelerating module, the focusing effect of the module is evident through the curvature of the beam's path. But, before the beam enters the accelerating module, the offset particles are over-focused in the injector. This is because the injector lattice is optimized for a space-charge dominated beam while the centroid orbit is not affected by the space charge effect. A mismatched beam line is a group of focusing fields that are either too weak or too strong for the beam that is being transported. The mismatch in the injector magnifies any position chirps that are generated in the injector.

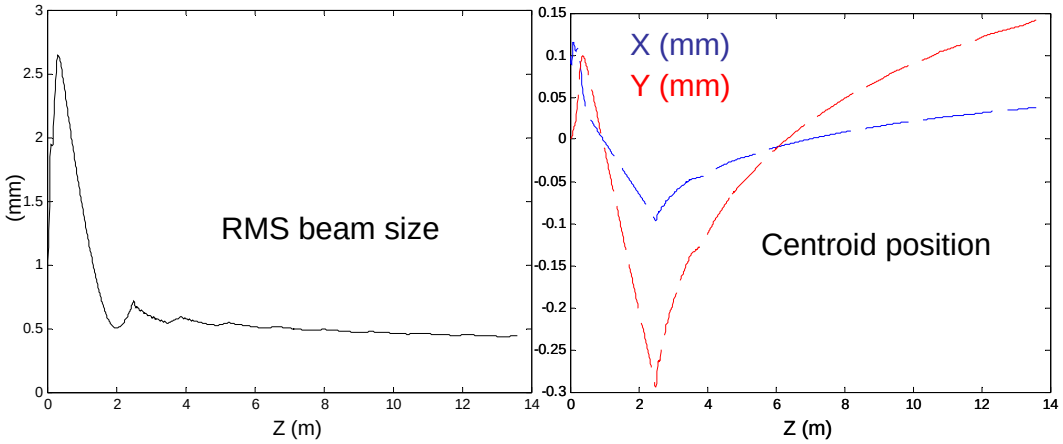

Figure 4.2.6 A beam offset in the injector is magnified as it travels through the first accelerator section. After the RF gun, the RMS beam size does not change significantly.

This mismatch cannot, of course, be avoided in the injector because the extra focusing strength is needed in order to counteract the space-charge forces. The only way to avoid the exacerbation of asymmetries that it promotes is align the injector with sub-mm precision.

4.3 Downstream of the Injector

When the beam is accelerated off-crest in the first accelerator section, a much larger head-tail separation is observed on the OTR or synchrontron light monitor in the bunch compressor than for the on-crest case described in the previous section. This off-crest tilt is the result of the effects of a mis-aligned injector described in the previous section, plus the effect of dispersion in and after the accelerator section. The tilt produced by a mis-aligned injector is, however, insignificant compared to the tilt created by dispersion from beam offsets in dipole fields downstream of the first accelerator section.

To measure the strength of these effects, the beam was placed off-crest in the first accelerator section and then given a closed orbit bump with several correctors. The tilt of the beam was maximized by adjusting the phase advance at the screen and then measured by analyzing the images taken with the synchrotron light monitor screen. The measurements were done together with C. Gerth and E. Pratt [39]. In Fig. 4.3.1, the closed orbit bumps are depicted on the left and the resulting tilts are depicted on the right.

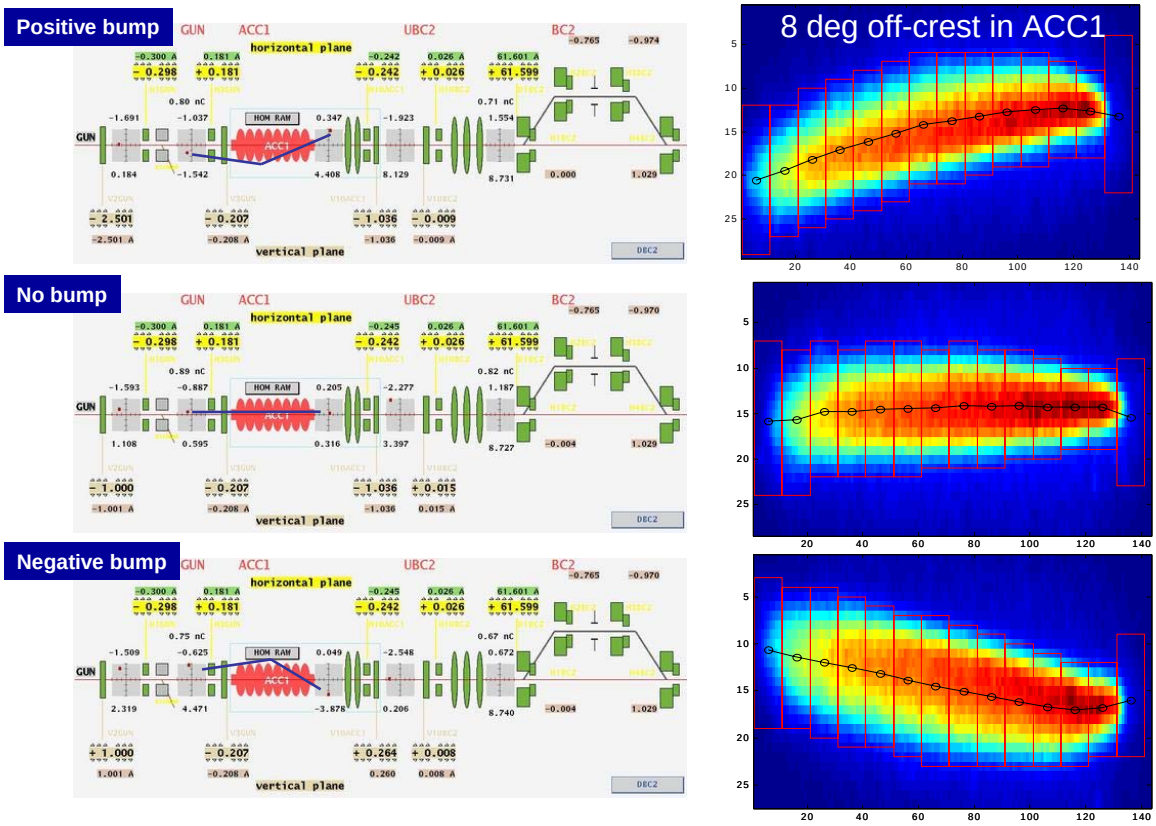

Figure 4.3.1 The tilt of the beam for various closed orbit bumps through the first accelerator section. Screen shots taken from a synchrotron light monitor located after the 3rd dipole of the bunch compressor. Bunch is streaked out longitudinally, but the image is in the x-y plane. Plot by C. Gerth.

The images seen on the monitor are projections of the beam streaked out in a longitudinal direction. The x-axis of the monitor corresponds to a combination of the z-axis and x-axis of the beam. The y-axis of the monitor corresponds to the y-axis of the beam. Each beam projection is composed of beam slices that are tilted in the x-y plane and streaked and tilted in the x-z plane (Fig. 4.3.2). This means that if one makes a fit to the entire beam image, there will be an error associated with the x-y tilt of the beam, when the quantity that is sought is the x-z tilt of the beam. This error can be avoided by fitting only to the central portion of the beam and ignoring the head and tail sections. This is shown through the difference between the red line which is fitted to the entire beam and the blue line which is a fit to the central portion of the beam. One can see that the central portion is not affected by the x-y tilt of the individual slices, while the end portions are. In practice this is done by cropping the image of the beam, so that the head and tail of the beam are cut off. The colors of the pixels are assigned weights according to the intensity of the detected synchrotron light. A line is then fitted to the weighted distribution.

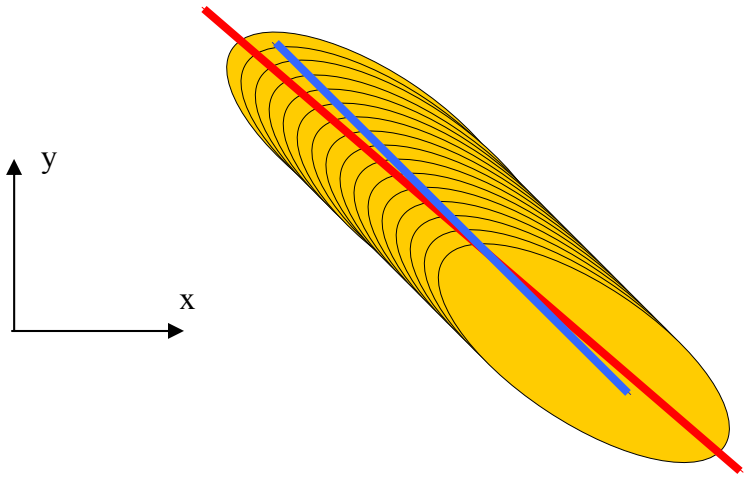

Figure 4.3.2 Images on the screen contain projections of the beam streaked out in a longitudinal direction. The red line is a fit to the centroids of the slices of the entire beam. The blue line is a fit to the centroids of the slices from the central portion of the beam.

The orbit bumps, wakefields, and coupler kicks were then simulated with a combination of ASTRA in the injector and the transport matrices constructed in MATLAB for the propagation of the simulated beam through the first accelerator section and first half of the first bunch compressor. The simulation results for the energy spread of the beam coming out of the injector, the horizontal beam path through the accelerating module, and the resulting tilt in the first bunch compressor are shown in Fig. 4.3.3 for simulations with and without the energy dependent deflection of the offset trajectory through the quadrupole magnets (chromatic effect). The simulation was done for beams generated with and without the small influence of the energy chirp from the gun. The Rosenzweig-Serafini model for beam transport through an accelerator section was used without a high-energy approximation [40]. In the plot of the transport through the first accelerator section of a beam with an incoming angle and offset in the horizontal plane, the angle, goes to zero and the offset is also focused. It is clear from the comparison between the simulations for which the chromatic effects of the quadrupole magnets were on and off, that dispersion induced by the quadrupole field offsets is, by far, the largest contributor to the tilt of the beam seen in the first bunch compressor. When the chromatic effect of the quadrupole is applied to an offset beam, the resulting tilt it 3.38 degrees in the middle of the bunch compressor. When the chromatic effect is “off”, the tilt is only 0.2 degrees.

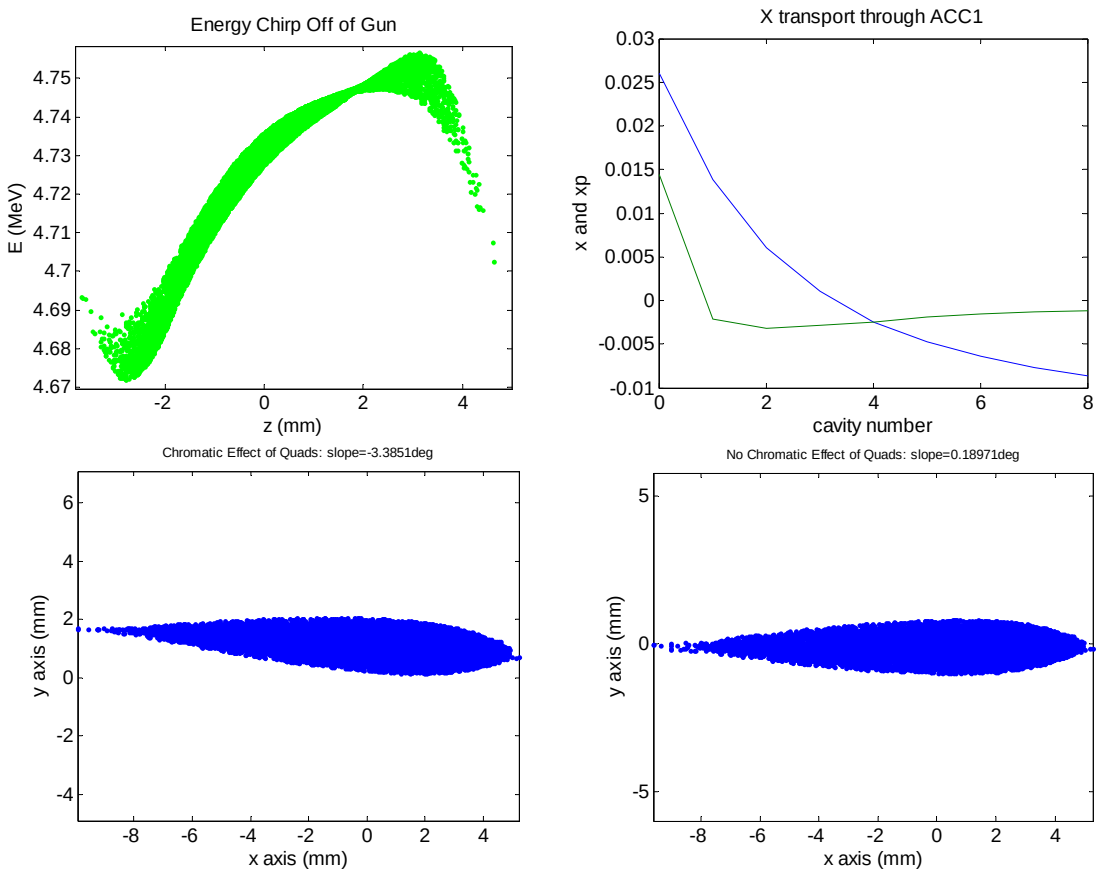

Figure 4.3.3 Simulation results for energy spread of the beam coming out of the injector (top-left), the horizontal beam path through the accelerating module (top-right), and the resulting tilt in the first bunch compressor with (bottom-left) and without (bottom-right) the chromatic effect of the quadrupoles. The beam was 8 degrees off-crest in the first accelerator section.

Table 4.3.1 summarizes the contributions of the various effects to the tilt simulated in the bunch compressor.

	ACC1 entrance: No offset No angle (degrees)	*ACC1 entrance: 3 mm offset 15 mrad (degrees)*	*ACC1 entrance: gun energy dist. added <-same orbit (degrees)*
Dispersion after inj.	0.00	2.38	2.09
+Wakefields (1nC)	0.00	2.34	2.12
+Coupler kicks	0.04	2.46	2.19

Table 4.3.1 Contributions to x-z beam tilt for an off-crest bunch, observed in x-y beam images taken in the first bunch compressor. In the first two cases, a flat energy distribution coming out of the gun was used and in the last case, the energy chirp that comes from the gun was used.

In the first two cases of Table 4.3.1, a flat energy distribution coming out of the gun was used and in the last case, the energy chirp that comes from the gun was used. Dispersion contributes most to the tilt, while the impact of wakefields was barely measurable by changing the charge of the beam, especially given a tilted module in which the same tilts could be generated by different paths through the module. Coupler kicks simulated with equations from [38] are so small that they are not possible to measure with either an on- or off-crest beam. When the beam is on-crest, there is no dispersion contribution after the injector and the primary effect that should be visible on the beam are the coupler kicks and injector mis-alignement effects. The coupler kicks have such a small impact that it would not be possible to measure them with the screen in the dispersive section of the chicane. The beam tilts predicted with the simulation matched the measured results with high accuracy when offsets at the exit of the gun were taken into account (Fig.4.3.4).

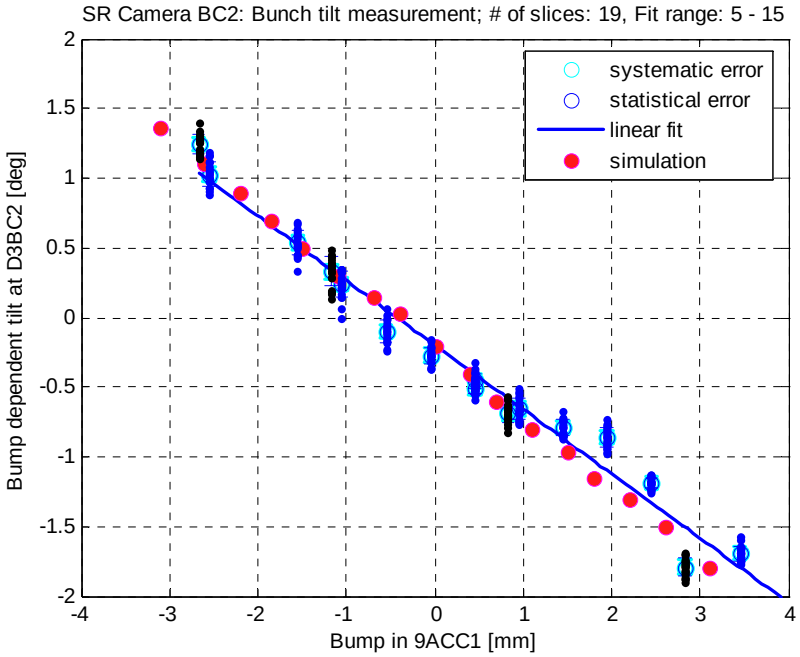

Figure 4.3.4 Beam tilts in the middle of the first bunch compressor measured and simulated. Each tilt corresponds to a defined orbit bump through the first accelerator section. The individual measurement points for two different measurements are represented as blue and black solid dots.

The beam is not only tilted in the x-z plane, but also in the y-z plane. After the second bunch compressor, the beam is streaked vertically in the transverse deflecting cavity. On the streaked beam image, one can see that beam is tilted, resulting in a slice emittance that is a factor of two smaller than the projected emittance. Measurements conducted with the transverse deflecting cavity indicate that the head of the bunch is separated from the tail by ~200 μm. In [32], sources for this tilt originating downstream of the first bunch compressor were simulated and only a third of the tilt measured with the transverse deflecting cavity could be accounted for when a non-tilted distribution in the first bunch compressor was used. This simulation propagated a particle distribution from the exit of the first bunch compressor through the transverse deflecting cavity. This

leads one to suspect that as in the x-z tilt case, the tilt in the vertical plane is also generated primarily upstream of the first bunch compressor.

The monitor developed in this thesis can be made highly sensitive to the tilt of the beam and that is why it was given so much attention here. The response of the measurements of the beam position in the chicane to beam tilts is described in Chapter 6.

5 Beam Pickups

Several types of beam position monitors have been used in accelerators and could be appropriate for use in the dispersive section of a bunch compressor. Of these, the button pickup, the cavity monitor and various forms of stripline pickups will be described with respect to their relative merits for this application. The primary focus will be on the design that was actually built and installed, the transversely mounted stripline. CST Microwave Studio simulations of some of the pickups were conducted and matched the measured performance of the monitors with high accuracy.

5.1 Button Pickups

Before the commencement of this thesis, an array of button pickups was the first method attempted to meet the challenge of measuring the beam position over a 10 cm aperture. The measurement principle relies on determining the normalized beam position, $X_{normalized}$, as a function of the voltages of the pickup outputs,

$$X_{normalized} = C\frac{V_{left} - V_{right}}{V_{left} + V_{right}}, \qquad (5.1.1)$$

where V_{left} and V_{right} are the voltages from the buttons on the right and left sides of the beam and C is a calibration constant. It was a relatively easy design to quickly construct and install with components that were already on-hand and since there was a hurry to get even a low-resolution measurement up and running, the button pickups were installed with an eye to arranging them as closely together as possible, but without detailed simulations of their interaction with the beam. As it turned out, due to the size of the

vacuum feedthroughs, the buttons were not close enough together in order to measure a significant voltage difference signal between two buttons. Due to the poor performance of the system and limited potential for improvement, the concept was abandoned. The reasons for this failure and the limitations of the method, along with the principles of button pickup operation, are described in the following paragraphs. These principles are also useful in describing how other types of pickups work.

In a round chamber, button pickups work well because the electric field lines from the electron beam spread out radially in the transverse plane and terminate over the entire surface of the chamber (Figure 5.1.1). This means that in a round chamber, the button pickup will always experience the image current of the beam as it passes and will always generate a signal. In the case of the wide, flat chamber in the middle of the bunch compressor, the electric field lines tend to concentrate in regions directly above and below the location of the beam and it is possible that if the button pickup is located at a distance from the beam, it will produce no signal at all (Figure 5.1.1).

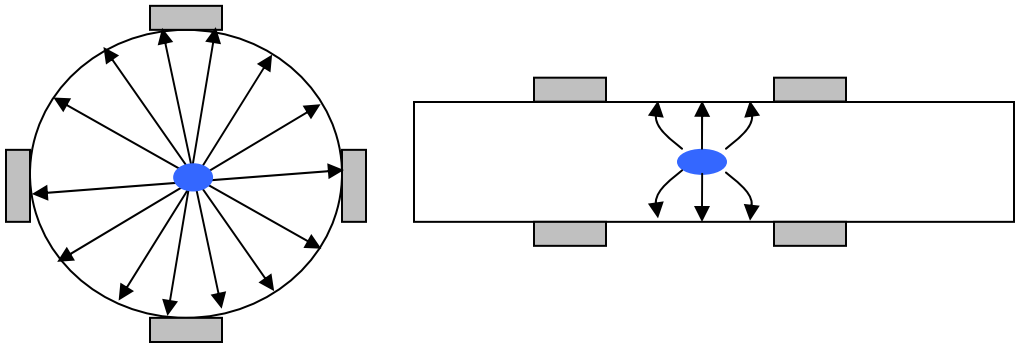

Figure 5.1.1 A cross-section of button pickups in a round vacuum chamber (left) and button pickups in a flat vacuum chamber (right). If the electric field lines from the beam do not terminate on the pickup, no signal will be produced.

As a beam travels through a vacuum chamber, an image current, mirroring the beam, travels along the walls of the vacuum chamber. In the absence of wakefields, when the beam velocity is approaching the speed-of-light, the electric field lines are longitudinally concentrated above and below the beam (Fig. 5.1.2).

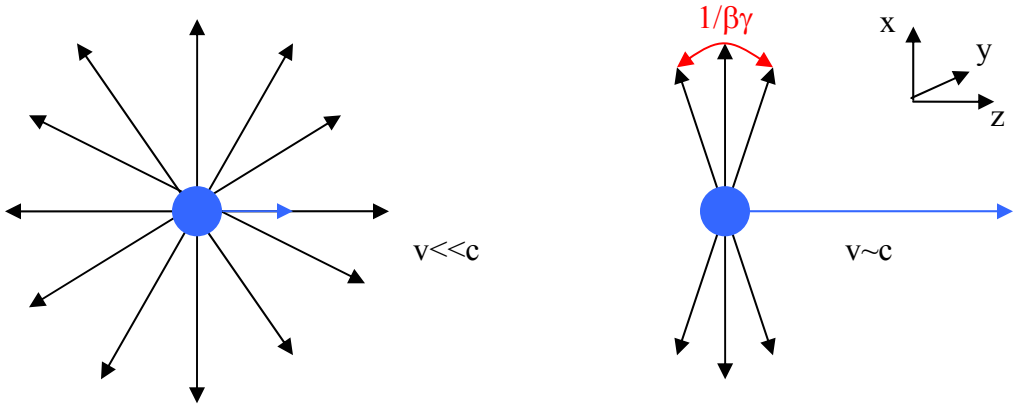

Figure 5.1.2 The electric field lines of charged particle beam moving at much less than the speed of light (left) and close to the speed of light (right).

One can numerically calculate the distribution of the image currents in the plane perpendicular to the beam's direction of motion with a 2-d magneto- or electrostatic simulation of the vacuum chamber cross-section, or, for a simple geometry, like a circle, the distribution of the image current density can be found analytically from the static 2-D version of Ampere's law,

$$\vec{j}_{beam} = \vec{\nabla} \times \vec{H} \ , \tag{5.1.4}$$

where the curl of the magnetic field, in units of Amperes per meter, is equal to j_{beam}, the free current density of the beam. The displacement current term is zero because the electric field of the beam is perpendicular to the beam direction.

Since an image current flows on the surface of the vacuum chamber to cancel out the magnetic field tangential to the metal surface, the magnetic field at a distance r from the electron beam with current I_{beam} is given by the Biot-Savart law in which a vector describing the direction of the current flow of the beam is crossed with a vector pointing perpendicularly out towards the vacuum chamber walls

$$H_{\varphi}(r) = \int_0^r \frac{I_{beam} d\mathrm{l} \times \hat{\mathrm{r}}}{4\pi r'^2} = \frac{I_{beam}}{2\pi r} \ . \tag{5.1.5}$$

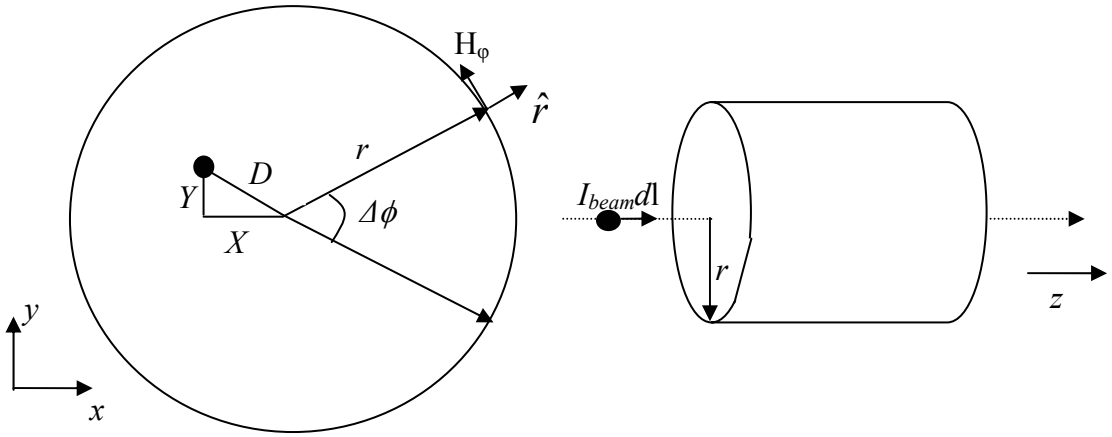

Figure 5.1.3 Coordinate system for a circular vacuum chamber.

For a circular vacuum chamber (Fig 5.1.3) with a cross-section of radius r and with the beam in the center of the chamber, the image line current density is then

$$J_{image} = \frac{-I_{beam}}{2\pi r} \ , \tag{5.1.6}$$

and for a beam that is displaced from the center of the chamber by X and Y at a position given by (D, θ), the image current contained within an angular spread of φ can be calculated from either Laplace's equations in two dimensions, yielding [43]

$$J_{image}(\varphi) = \frac{-I_{beam}}{2\pi r}\left[1 + 2\sum_{n=1}^{\infty}\left(\frac{D}{r}\right)^{n}\cos\left[n(\varphi-\theta)\right]\right] \tag{5.1.7}$$

or from Biot-Savart, to get

$$J_{image}(\varphi) = \frac{-I_{beam}}{2\pi r}\frac{r^2-(X^2+Y^2)}{r^2+X^2+Y^2-2r(X\cos\varphi+Y\sin\varphi)}. \tag{5.1.8}$$

A perfect pickup electrode spanning an arc $\Delta\phi$ integrates a fraction of J_{image}. Using normalized beam displacements $x = X/r$ and $y=Y/r$, we get the image current integrated by the pickup

$$I_{pickup}(\Delta\phi) = I_{beam}\frac{2}{\pi}\arctan\left(\frac{\left((1+x)^2+y^2\right)\tan(\phi/4)-2y}{1-x^2-y^2}\right). \tag{5.1.9}$$

The same exercise can be done for a wide, rectangular chamber. Let, $2h$ be the height of the chamber and zero be the vertical position of the beam centered between the top and bottom surfaces located at $+h$ and $-h$. This means that a sum of image currents of $+I_{beam}$ at $y=\pm 2h, \pm 4h, \pm 6h$ yields the magnetic field seen by a particle at location (x,y) [41],

$$H_y(x,y) = \frac{I_{beam}}{2\pi}\left[\sum_{n=1}^{\infty}\frac{1}{\sqrt{x^2+(nh-y)^2}}+\sum_{n=1}^{\infty}\frac{1}{\sqrt{x^2+(nh+y)^2}}\right], \tag{5.1.10}$$

which, for $y = h$, becomes

$$H_y(x,y=h) = \frac{I_{beam}}{2\pi}\frac{1}{\sqrt{x^2+h^2}}. \tag{5.1.11}$$

This means that the image current can be written,

$$J_{image}(x) = \frac{-I_{beam}}{2\pi\sqrt{x^2+h^2}}, \tag{5.1.12}$$

and integrating x over a pickup of width w with a horizontal offset of X gives,

$$I_{pickup} = \frac{I_{beam}}{2\pi}\left[\ln(X+d+\sqrt{(X+w)^2+h^2})-\ln(X+\sqrt{X^2+h^2})\right]. \tag{5.1.13}$$

We can now use Equations 5.1.10 and 5.1.14 to compare the relative sensitivities of pickups installed in flat-chamber and round-chamber configurations. Three different

configurations of pickups are shown below in Fig. 5.1.4. The unsuccessful but expedient design which was installed in the flat chambers of the first and second bunch compressors is shown on the left. It had a button diameter of 8 mm and a distance between buttons of 55 mm. A flat-chamber design that would have had performance comparable to existing button pickups installed in round chambers is shown in the middle with 20 mm diameter pickups and a separation between the pickups of ~21 mm. Typical pickup dimensions in a 35 mm diameter round chamber are shown on the right.

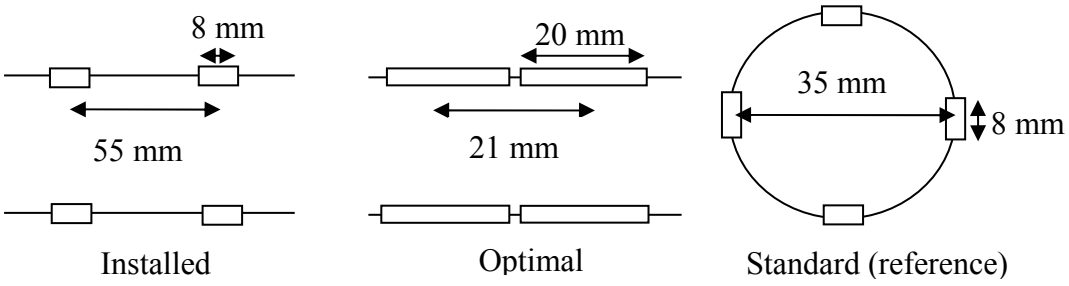

Figure 5.1.4 Dimensions of pickups in three different configurations: the unsuccessful, but expedient design which was installed in the tunnel (left), a flat-chamber design that would have had performance comparable to existing button pickups installed in round chambers (middle), standard pickup dimensions in a round chamber (right).

In Fig. 5.1.5, the sensitivities of the three different pickup configurations shown previously to changes in the beam position are shown below. The sensitivity of the expedient configuration that was installed in the bunch compressors is shown in green. This configuration has roughly half of the sensitivity of the same 8 mm diameter buttons installed in a typical 35 mm diameter round chamber. 20 mm diameter button pickups installed with a spacing of 21 mm would have delivered resolution that is comparable to the resolution of 8 mm button pickups in a 35 mm round chamber. Both 8 mm and 17 mm buttons in a round chamber, when coupled with typical DESY front-end electronics, get 15-20 μm resolution over the central, linear portion of their ~10 mm dynamic range at 1nC. This means that it is possible to install large, closely spaced buttons that have a sensitivity that is comparable to existing button installations that get 15-20 μm resolution with existing electronics. It is not, however, possible to find a design with buttons that can get the required sub-5 μm single-bunch resolution unless a new electronics concept is invented.

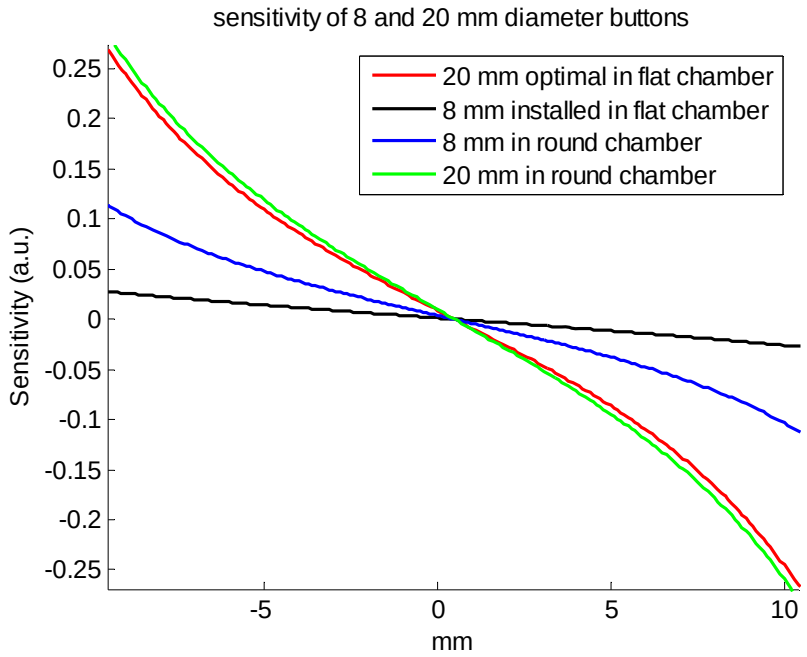

Figure 5.1.5 Sensitivities of button pickups in flat-chamber and round-chamber configurations. The sensitivity of the button configuration which was installed in the bunch compressor flat chamber is plotted in green, the sensitivity of an optimal flat-chamber button pickup configuration is plotted in red and the sensitivity of a standard button pickup configuration in a round chamber is shown in red.

Fig. 5.1.5 also implies that the flat-chamber button configuration that was initially installed might have worked with a resolution of about three times the 15-20 μm resolution produced by the round chamber button pickups at 1 nC, a surprising result, given that, under test, the measured difference between the amplitudes of the voltages of the pickup outputs was below the signal-to-noise-ratio of the pickup output voltage for a significant range of beam positions close to the middle of the two buttons. The lesson from this is that it is not always possible to make scaling assumptions about button monitors because below a certain threshold, noise dominates.

To calculate this signal-to-noise threshold, we need the beam current, the transfer impedance of the pickup, and the noise floor of the measurement. The time dependence of the current of the short Gaussian electron bunches of an FEL can be approximated as

$$I_{beam}(t) = \frac{1}{\sqrt{2\pi}} \frac{Q}{\sigma_t} e^{-\frac{1}{2}(t/\sigma_t)^2} \tag{5.1.15}$$

and the frequency response as

$$I_{beam}(\omega) = Q e^{-\frac{1}{2}(\omega\sigma)^2} \tag{5.1.16}$$

where σ_t is the rms duration of the beam. A Gaussian is not an accurate representation of a bunch for high frequencies because the tails of a Gaussian go out to infinity, while a real bunch has a finite extent. The result of this discrepancy is that any edges in the distribution will cause the beam current spectral density to fall off with $1/\omega$ at high frequencies.

To describe the transfer impedance, it is helpful to first write an equivalent circuit for the pickup. Each segment of a button pickup, or any pickup, can be described with some equations from transmission line theory [44]. A small section of transmission line can be modeled as a series inductance with a shunt capacitance as shown in figure 5.1.6.

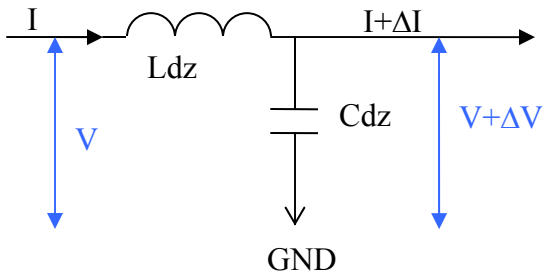

Figure 5.1.6 Electrical circuit representation of a small section of transmission line.

The current going through the inductor is proportional to the voltage drop across the inductor:

$$-\frac{\partial V}{\partial z} = L\frac{\partial I}{\partial t} \tag{5.1.17}$$

and the current exiting the transmission line segment is the difference between the input current and the capacitance current:

$$-\frac{\partial I}{\partial z} = C\frac{\partial V}{\partial t} \tag{5.1.18}$$

where L and C are the inductance and capacitance per unit length. The voltage and current can then be written down in terms of forward and backward propagating waves:

$$V = V_+ \ (t - \frac{z}{v_{phase}}) + V_-(t + \frac{z}{v_{phase}})$$
$$I = \frac{V_+}{Z_0} \ (t - \frac{z}{v_{phase}}) - \frac{V_-}{Z_0}(t + \frac{z}{v_{phase}}) \tag{5.1.20}$$

where v_{phase} is the phase velocity of the waves, given by

$$v_{phase} = \frac{1}{\sqrt{LC}} \quad . \tag{5.1.21}$$

Finally, the characteristic impedance of the transmission line is defined by the ratio between the wave voltage and the wave current and can be written in terms of the line inductance, L, and capacitance, C, as

$$Z_0 = \sqrt{\frac{L}{C}} \quad , \tag{5.1.22}$$

where, for a coaxial line, the inductance and capacitance per unit length, L and C, can be written in terms of the radius of the inner conductor to the radius of the outer conductor

$$C = \frac{2\pi\varepsilon_r\varepsilon_0}{\ln(r_{out}/r_{in})} \qquad \text{and} \qquad L = \frac{\mu_r\mu_0}{2\pi}\ln(r_{out}/r_{in}) \tag{5.1.23}$$

This solution of equations 5.1.17 and 5.1.18 is very general and does not predict any distortion of the pulse as it travels down the transmission line when ε_r and μ_r are constants. It is instructive to substitute Eqs. 5.1.23 into Eq. 5.1.21 to see that the phase velocity of a pulse traveling in an in-vacuum coaxial line is equal to the speed-of-light. While this is not very important for button pickups, it is very useful for the stripline pickups that will be described in the following sections. For the button pickups described here, Eqs. 5.1.22 and 5.1.23 are absolutely necessary when designing a pickup geometry that keeps Z_0 constant, (Fig 5.1.7). Issues related to the resistance of the transmission line have not been covered.

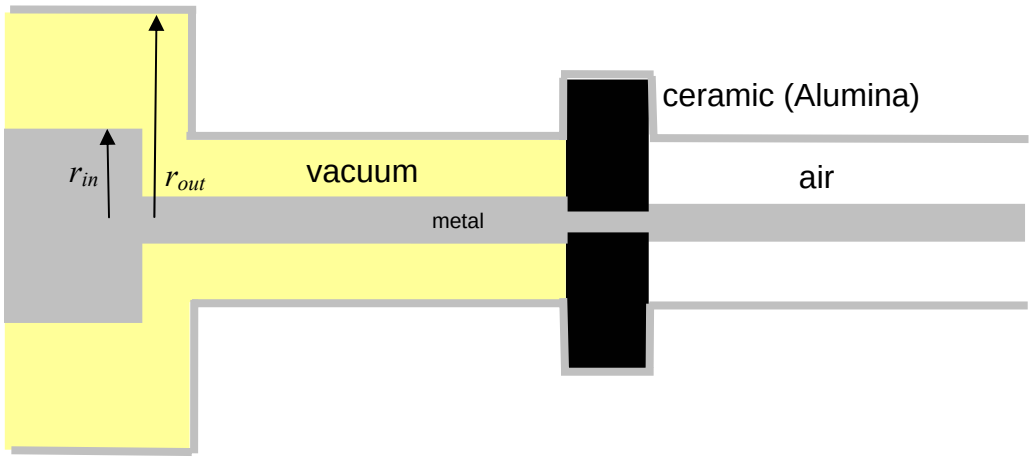

Figure 5.1.7 Button geometry that keeps the impedance constant will keep the ratio between the inner and outer conductor constant.

A button pickup can be modeled by the equivalent circuit shown in Fig. 5.1.3 in which the beam current, I_{beam}, in parallel with the transfer impedance, Z_t makes a voltage

source that is experienced by the high-pass filter composed by capacitance, C_{pickup}, and R_l, the load impedance.

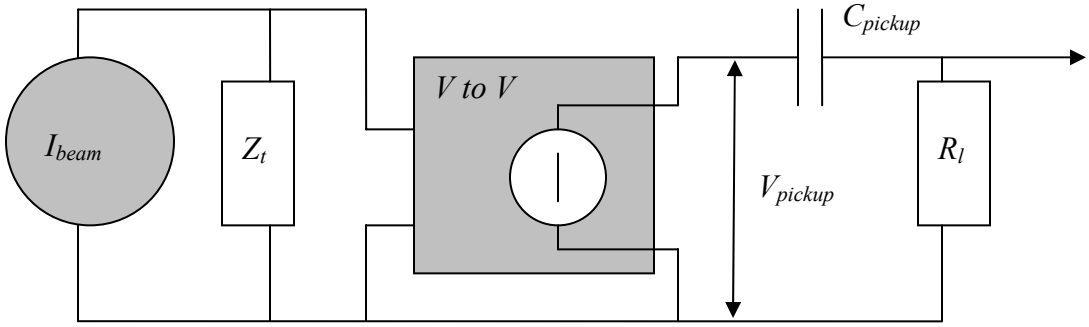

Figure 5.1.8 Button pickup equivalent circuit. I_{beam} is the current of the beam. Z_t is the transfer impedance, C_{pickup} is the electrode capacitance and R_l is the load impedance.

The transfer impedance Z_t is the quantity that we are looking for, and to get it, we start by writing down the voltage induced at the button with no load

$$V_{pickup}(t) = \frac{1}{C_{pickup}} \int_{-\infty}^{t} j_{pickup}\, dt' \tag{5.1.24}$$

If the bunch is longer than the button radius and the beam is in the middle of the pipe, we use an approximation depicted in Fig. 5.1.9

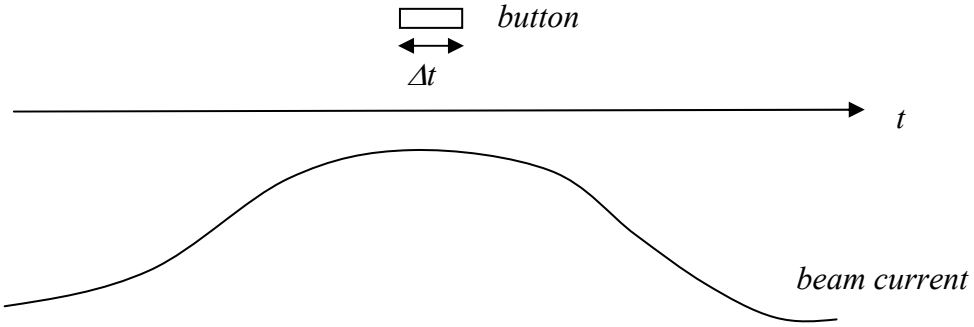

Figure 5.1.9 Approximation that bunch is longer than button radius allows for integration over the beam current in steps of Δt. It is valid for frequencies with wavelengths longer than the button radius.

along with Eq. 5.1.10 to get

$$V_{pickup} \cong \frac{I_{pickup}}{C_{pickup}} \frac{d}{c_0} = \frac{I_{beam}}{C_{pickup}} \frac{d^2}{8rc_0}, \tag{5.1.25}$$

where d is the diameter of the button and r is the distance from the beam. This is, of course, not the case in an FEL, where each bunch is less than a millimeter in length in

some locations. Nevertheless, it is convenient to ignore the time dependence at this stage as the equation is true at least for wavelengths longer than the button size and it gives us a useful rule: if the vacuum chamber diameter quadruples, the button size only needs to double in order to conserve the amplitude of the signal.

Now, we consider the load, R_l, determined by the cable impedance and the termination of the cable. In the frequency domain, the voltage at the load is given by [45],

$$V(\omega) = V_{pickup} \frac{R_l}{R_l + 1/i\omega C_{pickup}}, \tag{5.1.26}$$

where V_{pickup} is given by Eq. 5.1.19. Then, according to Ohm's law, the transfer impedance is given by

$$Z(\omega) = \frac{d^2}{8rc} \frac{1}{C_{pickup}} \frac{R_l}{R_l + 1/i\omega C_{pickup}} \tag{5.1.27}$$

Taking the absolute value of Z, we see the behavior of a high-pass filter,

$$|Z(\omega)| = \frac{d^2}{8rc} \frac{\omega R_l}{\sqrt{1 + (\omega R_l C_{pickup})^2}}, \tag{5.1.28}$$

We also see that the transfer impedance is more sensitive to the button diameter than to the capacitance of the pickup, which is generally less than 10 pF and can be calculated from Eq. 5.5.23(a) multiplied by the length of the button. This is the reason that button pickup BPM front-ends that use wavelengths which are longer than the size of the button use buttons with a diameter that is maximized at the expense of impedance matching; i.e. the ratio between the inner and outer conductor is not the same as it is in the transmission line. This mismatch, however causes a low-pass filter-like response which will be described in the following paragraphs.

The time constant of the high-pass filter is given by $\tau = R_l C_{pickup}$ and can be measured with a network analyzer, but first, the time constant needs to be given as a function of frequency. To calculate the time constant of a button pickup for a forward traveling sinusoidal wave

$$V_+(t,z) = V_0 \cos(\omega(t - \frac{t}{v_{phase}})) = \mathrm{Re}(\mathrm{A}e^{i\omega t}) \tag{5.1.29}$$

where A is a phasor with

$$|A| = V_0 \qquad \text{and} \qquad \arg(A) = -\omega \frac{z}{v_{phase}} \tag{5.1.30}$$

The derivative of the phase with respect to the frequency gives the time delay through a transmission line of length z:

$$-\frac{\partial}{\partial\omega}\arg(A)=\frac{z}{v_{phase}}=\tau \tag{5.1.31}$$

The behavior of this phasor as a function of frequency can be measured with a network analyzer. Most network analyzers measure a circuit with one or two ports. At each port there is an incoming wave A generated by the network analyzer and an outgoing wave B (Fig. 5.1.10).

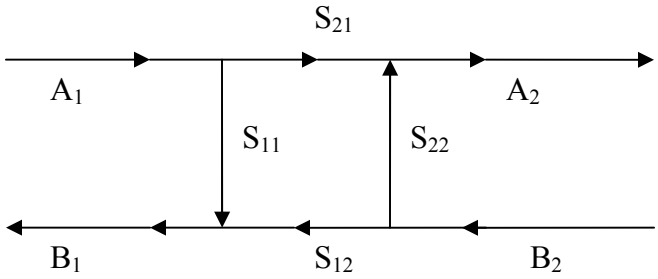

Figure 5.1.10 Definition of two-port S parameters in terms of incoming wave A and outgoing wave B.

The network analyzer measures the ratio of the outgoing wave, A, to the incoming wave, B, by delivering complex "S" parameters according to:

$$\begin{aligned} B_1 &= S_{11}A_1 + S_{12}A_2 \\ B_2 &= S_{21}A_1 + S_{22}A_2 \end{aligned} \tag{5.1.32}$$

The diagonal elements S_{11} and S_{22} are reflection coefficients and S_{12} and S_{21} are transmission coefficients. These "S" parameters are also the terms produced by electromagnetic simulation software, such as CST. A typical configuration for a test of a button monitor might involve a stretched wire running through the middle of the vacuum chamber. The network analyzer would send its incoming signal down the wire, representing the electron beam, the outgoing parameters would be measured from the pickup output. From this measurement, one should be able to see if the ratio of the outgoing to the incoming waves is large enough that there will be enough voltage for the front-end electronics to function for a given beam current at the frequency of interest.

Button pickups are used frequently in the FLASH linac for both position measurements and for arrival-time measurements. Position measurement (BPM) pickups usually require a large voltage at frequencies below a GHz, whereas, for arrival-time measurement (BAM) pickups for an optical setup that will be described in chapter 8, a large bandwidth that stretches up past 40 GHz without any notches in the spectrum is desirable. These properties can be adjusted by tuning the impedance of the pickup design.

In the spectra shown in Fig. 5.1.11, the BPM button delivers more power at lower frequencies than the BAM button, while the BAM button delivers a steeper signal slope.

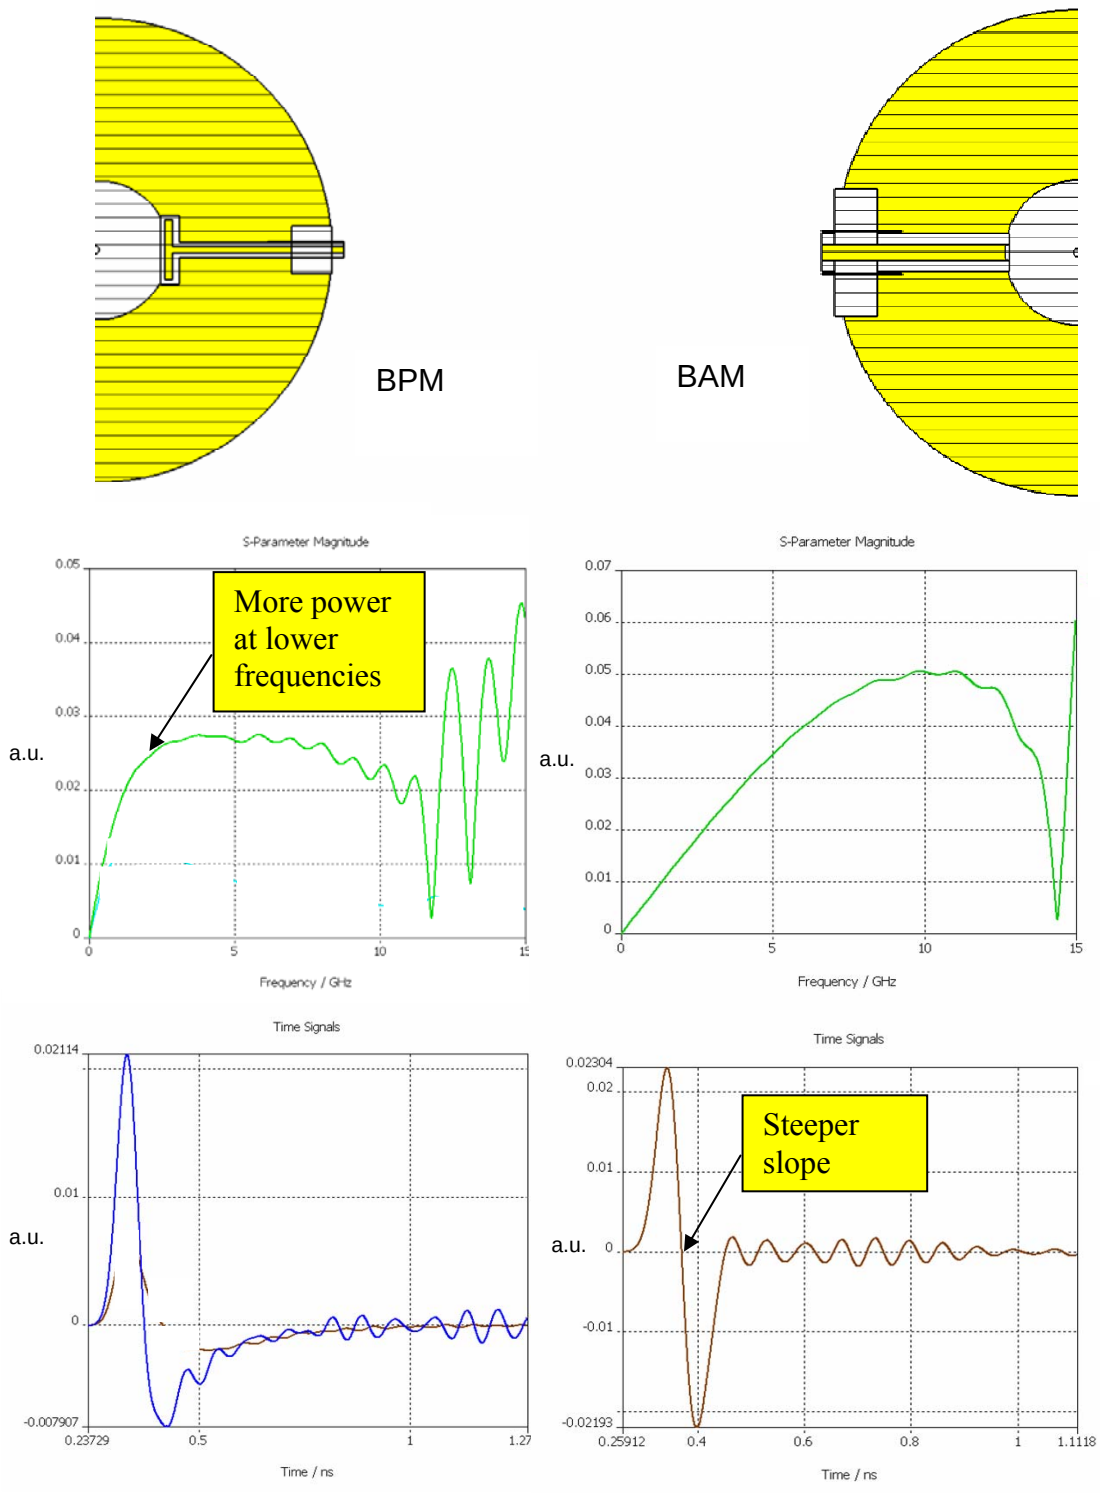

Figure 5.1.11 Comparison of frequency and time domain simulations of two pickups. Above left is the cross section of a button pickup for a BPM. Above right is a button pickup that is used in a Beam Arrival-time Monitor (BAM).

So far, the button has been treated as a simple high-pass filter without accounting for the various notches that appear in the spectrum due to resonances within the pickup and cavity. Some of these resonances include:

- A resonant cavity between the button and the body of the vacuum pipe; this resonator is usually tuned to frequencies greater than 10 GHz and is excited by short bunches.
- The impedance variation on the transmission line from vacuum to air will produce impedance mismatching, generating reflections and standing waves.

To deal with unwanted resonances, most BPM front-ends include a low-pass filter to remove the higher frequency resonances. This is not possible when the pickup is used in an arrival-time measurement (BAM) where a broadband pickup output is desired.

The BAM pickup that was originally developed for a 1.3 GHz RF front-end was a ring supported by two SMA connector sized feedthroughs (Fig. 5.1.12(a.)). This effectively delivered a signal with a voltage that was high enough for the 1.3 GHz processing electronics. When a front-end with a higher resolution was used, the shape of the signal produced by the pickup was more clearly resolved and it was noticed that a bump in the signal, at the sampling location desired for the arrival-time measurement, moved from side-to-side when the beam position changed in the horizontal plane. Since the monitor was supposed to measure beam arrival-time and not position, this was undesirable and a new pickup was constructed with button like pickups combined with external cables (Fig. 5.1.12(b.)).

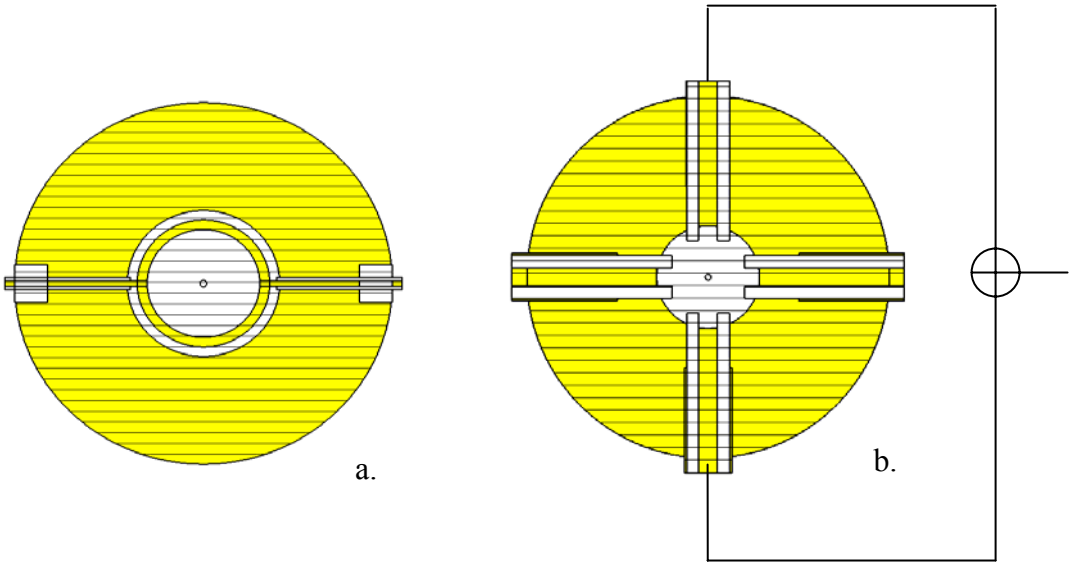

Figure 5.1.12 Cross sections of the old (a.) and new (b.) beam arrival-time pickups.

The beam position dependent bump observed in the time domain signal from the old, ring pickup was the result of a notch in the frequency spectrum at ~ 5 GHz, with a corresponding wavelength of ~ 6cm. This length scale was close to the length of half of the circumference of the ring. The output of the ring pickup is shown below in Fig. 5.1.13 as given by CST simulation and measurements with an oscilloscope done with and without an RF limiter. The use of the combiner to reduce the beam position dependence is shown below in Fig. 5.1.14.

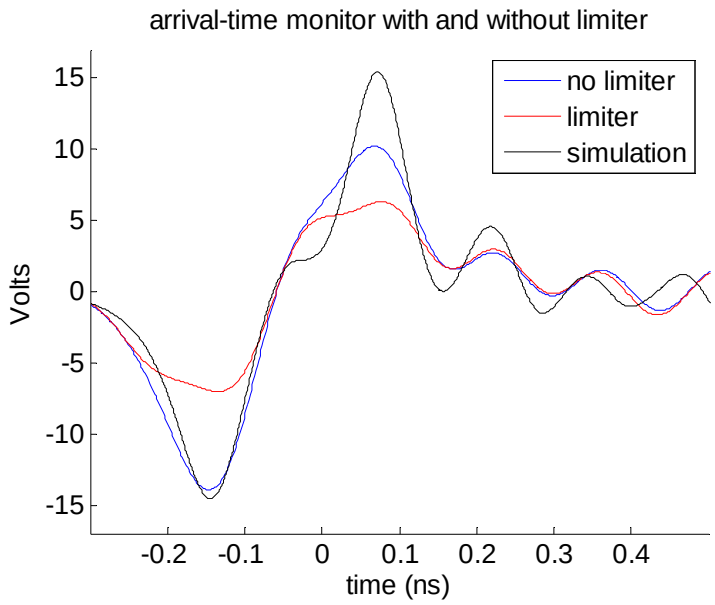

Figure 5.1.13 Ring pickup output with (red) and without (blue) limiter. Simulation shown in black. Undesired bump in signal changes position when beam position changes.

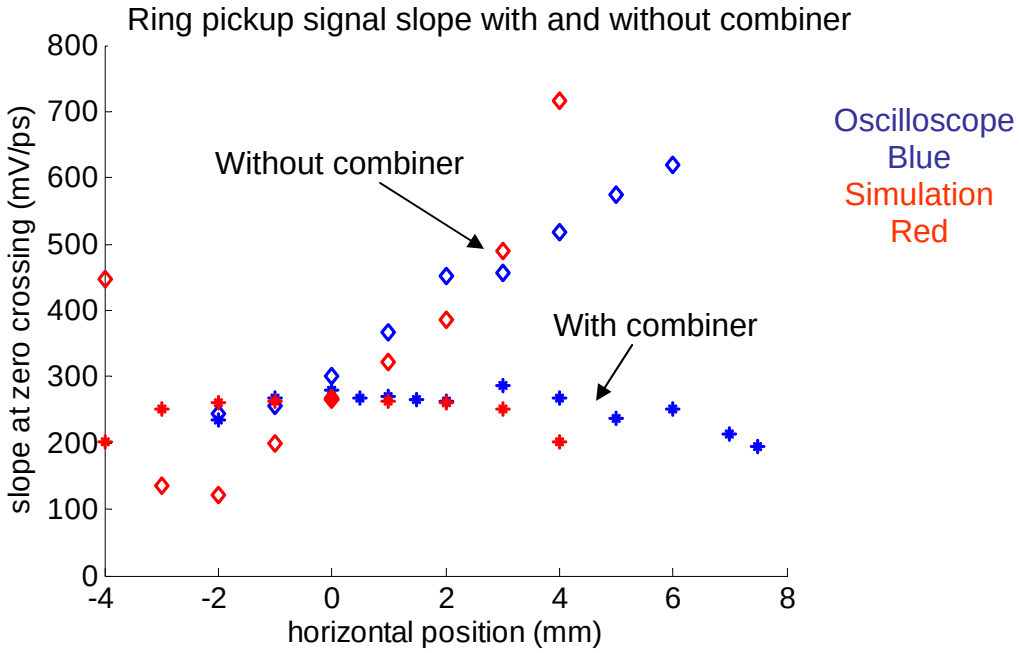

Figure 5.1.14 Position dependence of pickup output slope with and without combiner.

The new, button-type pickup was designed to maximize the bandwidth of the output signal without sacrificing too much of the amplitude. Another concern is the length of time that the signal from the pickup rings. The XFEL bunch spacing is only 200 ns and the ringing from the first bunch must be gone by the time the second bunch comes.

While the ringing is gone within the FLASH bunch spacing of 1 μs, it is not gone within 200 us. The limitation of the design has not, however, yet been completely evaluated because the in-house combiner that has been used so far has a 6 dB insertion loss and creates standing waves on the cable between the pickup and the combiner. A better, commercial combiner will reduce this effect and reduce the ringing observed. The remaining ringing will likely be caused by the cavity created by the channels opening into the beam pipe and the sensitivity of type N size pickups (Fig. 5.1.15(a.)) to resonances with wavelengths on length scales corresponding to the diameter of the coaxial channel and the distance between the feedthrough ceramic and the beam pipe. An SMA feedthrough option was simulated (Fig. 5.1.15(b.)) and it does not suffer from all of these resonances, but it would take longer to manufacture and have a higher risk of breakage at the ceramic in the feedthrough.

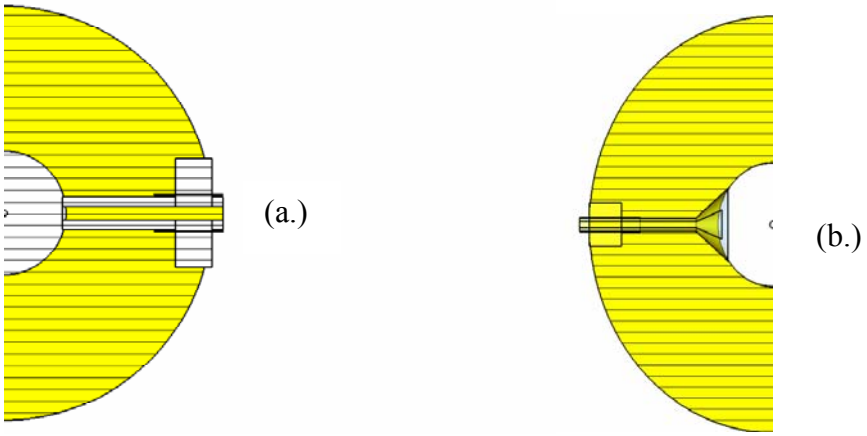

Figure 5.1.15 BAM pickup designs. Design (a.) provides an expedient production and installation process and is the design that has been installed at FLASH for BAM applications. Design (b.) would be very complicated to realize but it would have better performance than design (a.).

It was decided that the type-N design (Fig. 5.1.15a.) would be produced for FLASH because of expediency of manufacture, installation, and low risk of breakage. At ELETTRA, a similar design was installed, but it used an SMA sized pin and channel instead of a type-N sized pin and channel. Out of all of the designs, the SMA sized pin delivers the best ratio of the amplitude of the signal to the signal slope. This is ideal for avoiding problems related to AM (amplitude modulation) to PM (phase modulation) conversion problems, but because the amplitude of the signal it generates is so much smaller than that of any other pickup, the resolution of any low-charge BAM scheme would be severely limited by this pickup. That is why it was not selected. It would, however, be an ideal pickup for high-charge applications (>1nC).

An SMA design with a tapered button (Fig. 5.1.15b.) would be preferable for the XFEL due to the shorter distances between electron bunches and the necessity to limit the amount of time that the pickup rings. It is, however, much more complicated to build. For low-charge applications, where 100 pC beams are expected, one would prefer a pickup that is 10 times more sensitive than either of the pickups in Fig. 5.1.15. In order to accomplish this, the diameter of the beam pipe needs to be smaller so that the pickup is closer to the beam.

5.2 Cavity BPM

Cavity BPMs (Fig. 5.2.1) have been built that produce beam position measurements with sub-micron precision and the same sort of monitor could, in principle, be scaled up to accommodate larger apertures. Alternatively, a small-scale cavity BPM could be put on movers so that it slides from side-to-side along the flat chamber of the dispersive section of a chicane. Both concepts, however, present significant challenges given the 20 cm of longitudinal space allocated for the chicane BPM.

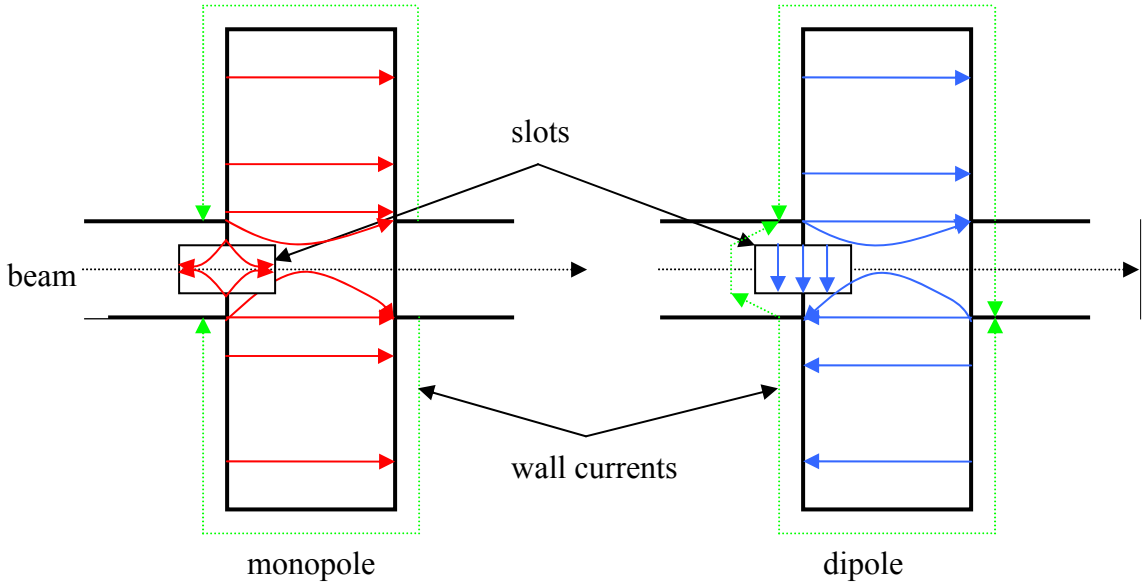

Figure 5.2.1 Side-view cavity BPM. The monopole modes are most strongly excited, but they do not contain information about the beam position. That is why a slot that selectively couples out the dipole mode and not the monopole mode is desired. A waveguide is attached to the slot.

In order to make a high-resolution cavity monitor, the dipole mode must be coupled out without picking up the monopole mode or higher order modes. These requirements put strict limits on manufacturing tolerances and installation alignment. While sub-micron resolution has been demonstrated at multiple labs, the higher the resolution that is required, the smaller the dynamic range of the monitor generally becomes.

One major impediment to simply scaling up an existing design is that the vacuum chamber in the middle of the FLASH bunch compressor is flat and the standard designs of high-resolution cavity monitors have round vacuum chambers. This means that a good design cannot simply be scaled up. Assuming that a quality large scale cavity design for a flat chamber could be made, one must consider several other factors. A larger cavity rings for longer than a smaller cavity and when the bunch spacing is very close, as it is for FLASH and XFEL bunch trains, the ringing from one bunch could overlap with that of the following bunch. Slight asymmetries and deformations in the cavity production

impact the quality of the measurement and must be carefully investigated with respect to what is possible in manufacturing. Finally, space considerations can be prohibitive for a large cavity installation.

Likewise, 20 cm longitudinal space considerations limit the possibility of taking a smaller, existing cavity monitor and putting it on movers, so that it slides from side-to-side along the flat chamber. Long stretches of bellows would be required in order to limit the discontinuities in this wakefield sensitive area.

5.3 Stripline Pickups

Longitudinally oriented stripline pickups are used frequently throughout the FLASH linac because they deliver a better resolution than button pickups. When optimized, they have achieved 5 μm resolution at FLASH. This is due to the larger voltages that they produce at the frequencies for which the front-end electronics are designed. Both button and stripline front-ends rely on measuring the differences in amplitudes between pickups on opposite sides of the vacuum chamber. Striplines, however, cannot be used everywhere because they take up a lot of longitudinal space; they can require 100-1000 mm compared to the button's 20 mm.

Stripline pickups consist of a channel carved out of the vacuum pipe in which a metal rod or strip is suspended and terminated on one end with the characteristic impedance of the stripline (Fig. 5.3.1). Because the width of the stripline is only a fraction of the beam pipe circumference, the stripline will not carry all of the image current. The fraction of the image current not carried by the stripline will travel across the upstream gap as a displacement current which will give rise to a voltage pulse on the upstream end of the stripline. Because the stripline is terminated on both ends with the same impedance, the pulse will split into two equal pulses and travel to each end of the pickup. The fraction of the image current continues to travel downstream until it encounters the downstream gap where another pulse is created that is equal in amplitude to the upstream pulse but opposite in polarity. This pulse also splits into two pluses traveling to each end of the stripline. This is depicted in the lower of the two pictures in Fig. 5.3.1. If the velocity of the beam and the phase velocity of the pulse are equal, the pulse that was created at the upstream gap will arrive at the downstream gap at the same time that the image current induces the downstream voltage pulse. Since the pulses have opposite polarity, they will cancel one another out and no energy will be dissipated in the downstream termination.

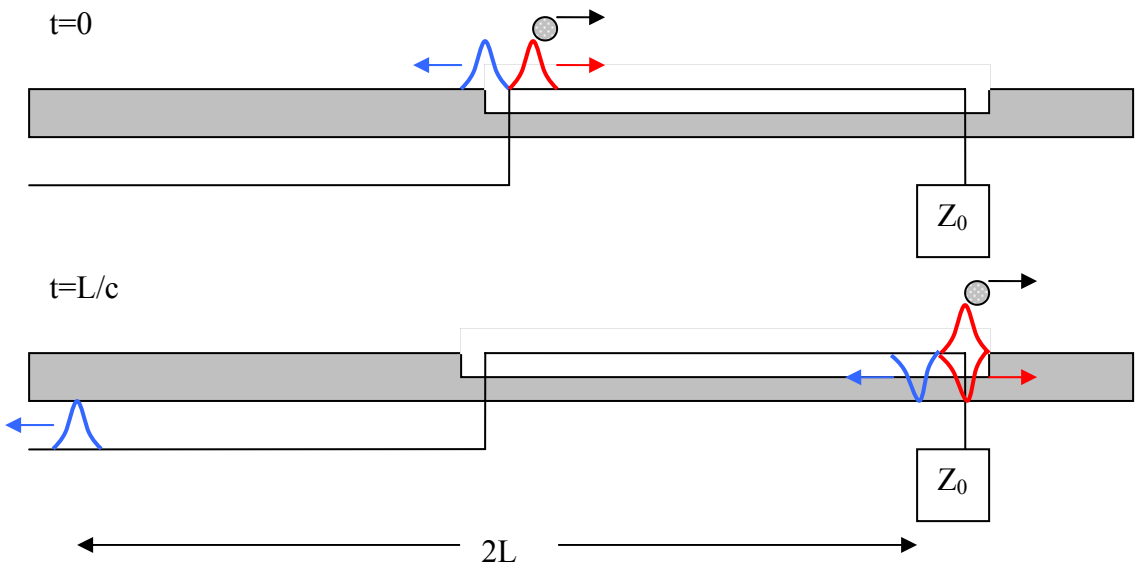

Figure 5.3.1 Longitudinally oriented stripline BPM principle of operation. The output consists of a bipolar signal with peaks separated by twice the length of the stripline.

In order to calculate the frequency response of this type of pickup, let us model the bunch as a Dirac impulse,

$$I_{beam} = q \cdot c \cdot \delta(z) \qquad (5.4.1)$$

This is an appropriate approximation for the short bunches of an FEL. We can use equation 5.1.10 for a stripline as well as for a button to give a function of the change in image current per change in beam position. Let us call that function ΔI and use equations 5.4.1 and 5.1.10 to write

$$\Delta V(z) = \frac{Z_0 \Delta I}{2}[\delta(z) - \delta(z - 2L)] \qquad (5.4.2)$$

which can be Fourier transformed with

$$\Delta V(k) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \Delta V(z) e^{-ikz} dz \qquad (5.4.3)$$

to get

$$\Delta V(k) = \frac{Z_0 \Delta I}{2\pi} \cdot e^{i\pi/2} \cdot e^{-ikL} \sin(kL) \qquad (5.4.4)$$

This value has a maximum at frequencies where the length is an odd multiple of quarter wavelengths:

$$kL = \frac{\pi}{2}(2n-1) \qquad \text{or} \qquad f_0 = \frac{c}{4L}(2n-1) \qquad (5.4.5)$$

where n =1,2,3,.. . For n=1 and a stripline length of 10 cm, the central frequency would then be 750 MHz with a bandwidth of 1.9 GHz given by the 3 dB points

$$f_{low} = \frac{1}{2} f_0, \qquad f_{high} = 3 \cdot f_0 \qquad (5.4.6)$$

and the time resolution is given by the inverse of the bandwidth (BW) times π,

$$\tau = \frac{1}{\pi \cdot BW},$$

yielding a signal width of 170 ps for the 10 cm long stripline case.

The beam position resolution achieved with FLASH 30 cm long striplines has been around 5 μm using a front-end method that combines the pickup outputs with a delay between them and sends them through a single filter and amplifier, thereby removing the impact of filter and amplifier drifts. LCLS gets 2-3 μm resolution with 10 cm long striplines for beam charges ranging from 0.2 to 8 nC. They send each pickup signal through its own filter and amplifier and calibrate away filter and amplifier drifts by injecting a reference signal prior to the arrival of each bunch.

5.4 Array of Striplines

One idea that was popular in the initial planning stages of the chicane BPM (2005) was to take an array of ~10 cm long striplines and mount them above and below the flat vacuum chamber of the bunch compressor. 1.3 GHz would be filtered out from the pickup signal and the down-conversion scheme used by the FLASH LLRF system would be used to sample the amplitudes of the signals. The beam centroid would be determined by processing the multiple channels of data with the LLRF system FPGA, thereby providing bunch-to-bunch beam position measurements that could be incorporated into an intra-train beam-based feedback. The difficulties with this technique come from drifts and noise of the filters, amplifiers and down-conversion electronics. As described in Sect 3.1, the individual LLRF amplitude measurements suffer from drifts of 2e-3/deg C and rms jitter of 5e-4. This would not be sufficient for the <5e-5 resolution desired for the chicane BPM. While a reference injection scheme has recently reduced these drifts to within the resolution of the measurement and a new down-conversion front-end has recently been developed with 1e-4 resolution, this would still not be sufficient to meet the desired BPM resolution. Nevertheless, using stripline pickups with digital down-conversion and reference injection has been successfully used by LCLS for their stripline BPM system.

The LCLS stripline BPM 140 MHz front-end addresses the problems of filter and amplifier drift with the injection of a reference signal to calibrate the monitor prior to the arrival of each bunch. It is designed to operate with 10 cm long stripline pickups with diameters of 25 mm in a 35 mm diameter beam pipe [42]. It gets 2-3 μm resolution for an

0.2 nC beam charge and can handle beam currents of up to 8 nC with the application of built in attenuators. An array of 10 cm long stripline pickups shown in Fig. 5.4.1 might have been able to achieve the required sub-5 um resolution in the chicane with an adaptation of the LCLS front-end design, even though the pickups in a flat chamber configuration have less sensitivity than the pickups had in their LCLS installation (Fig. 5.4.1 & Fig. 5.4.2).

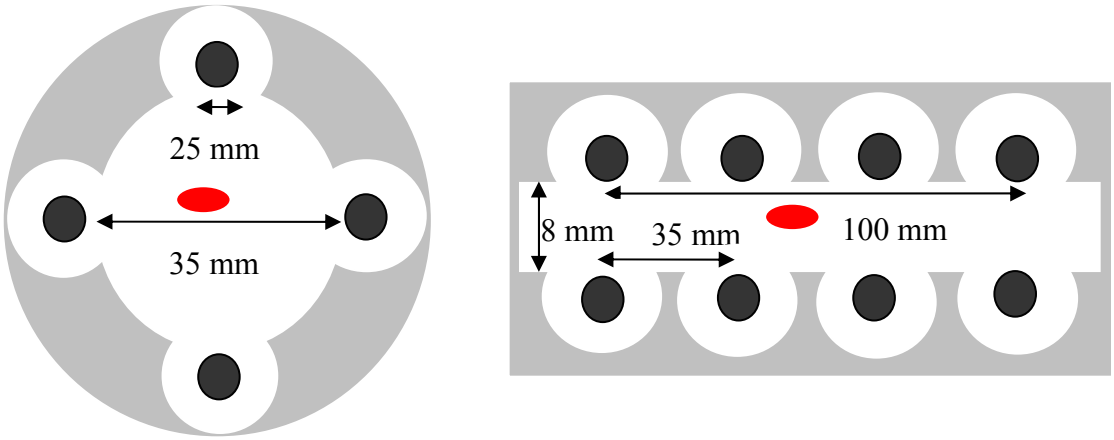

Figure 5.4.1 Cross-section of striplines in a round chamber and in a flat chamber (not-to-scale).

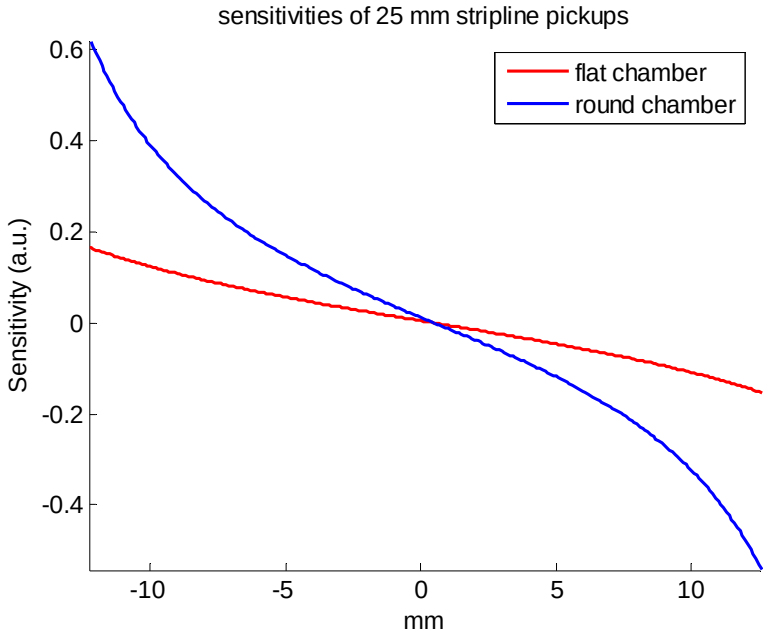

Figure 5.4.2 Simulation of sensitivity of 25 mm diameter striplines in a round chamber configuration and in a flat chamber configuration. The flat chamber has reduced sensitivity.

At the time of the design decision in 2005, this multi-channel reference injection solution that was commissioned in 2008 by LCLS seemed more complicated to implement than

the two-channel solution afforded by a transversely mounted stripline design. In-house support was also available for the front-end that was to accompany it.

5.5 Transversely Mounted Stripline

A transversely mounted stripline BPM pickup is depicted in yellow above a wide rectangular vacuum chamber (Fig.5.5.1). The pickup rests in a coaxial shaped channel which is open to the vacuum chamber below. The pickup is tapered to an SMA sized vacuum feedthrough. The beam path is depicted in green below the pickup.

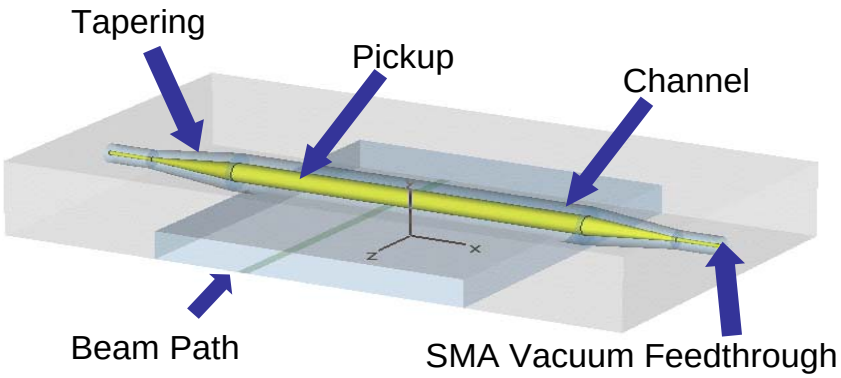

Figure 5.5.1 3-D transparent representation of the upper-half of the chicane BPM pickup (not-to-scale). The beam (green) travels under the pickup (yellow).

If a standard stripline pickup is rotated so that it is perpendicular to the beam direction, the current pulses induced as the beam passes beneath it will travel to each end of the pickup (Fig. 5.5.2). This is an idea that was proposed in early 2005 at DESY by Manfred Wendt. The average difference in the arrival-times of the current pulses, multiplied by the speed-of-light in a coaxial cable, gives the position of the beam:

$$beam\_position = \tfrac{c}{2} \cdot \left[(arrival\_left) - (arrival\_right) \right]. \qquad (5.5.1)$$

Alternatively, the average of the arrival-times of the current pulses multiplied by the speed-of-light gives the arrival-time of the beam:

$$beam\_arrival = \tfrac{1}{2} \cdot \left[(arrival\_left) + (arrival\_right) \right]. \qquad (5.5.2)$$

This requires a measurement of the pickup signals' phases and it is distinct from the typical stripline BPMs that measure differences in signals' amplitudes. The concept requires that the speed of the signals on the pickup are equal to *c* and is based on the basic transmission line equations derived in section 5.

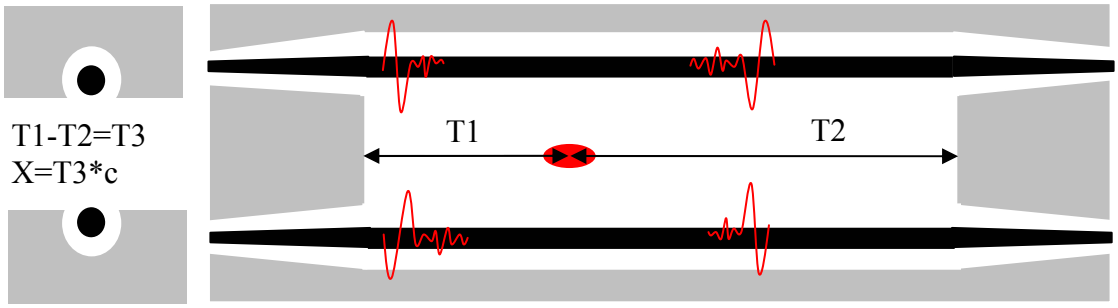

Figure 5.5.2 Cross-section of a transversely mounted stripline pickup with tapering to vacuum feedthroughs. The beam passes beneath or above the pickup and causes current pulses to travel to either end of the pickup. The arrival-times (T1 and T2) of the pulses are measured and the position of the beam (X) can be determined.

The pickup shown in Fig. 5.5.2 was chosen because of its ability to deliver the high-bandwidth pulses that are desirable for optical and high frequency phase measurement techniques. Using the design principles established in the previous sections, the questions that arose during the design process were primarily practical in nature. Existing striplines at DESY consisted of hollow, rounded rods suspended at a diagonal from SMA feedthroughs (Fig. 5.5.3). In principle, a tapered design would offer an improvement in impedance matching over the existing design. The extent of the improvement needed to be simulated and verified. The tapering required more welding stages and had more potential for failure due to the newness of the design. If the older design had been sufficient, then it could have been constructed with existing parts and know-how.

Figure 5.5.3 Stripline feedthrough cross-sections (not-to-scale). An existing DESY design for quadrupole mounted striplines (left). A new design with tapering to the feedthroughs (right). The existing design had been required due to space constraints in other installations.

The figures of merit for measuring the performance of the monitor include the steep slope of the signal at the zero-crossing, low amplitude and the absence of distortions in the signal that occur when the position of the beam is changed. The older and newer designs

were compared with the aid of CST software and it appeared that the performance of the tapered design (Fig. 5.5.4 (middle)) would offer a significant improvement over the existing design (Fig. 5.5.4 (top)).

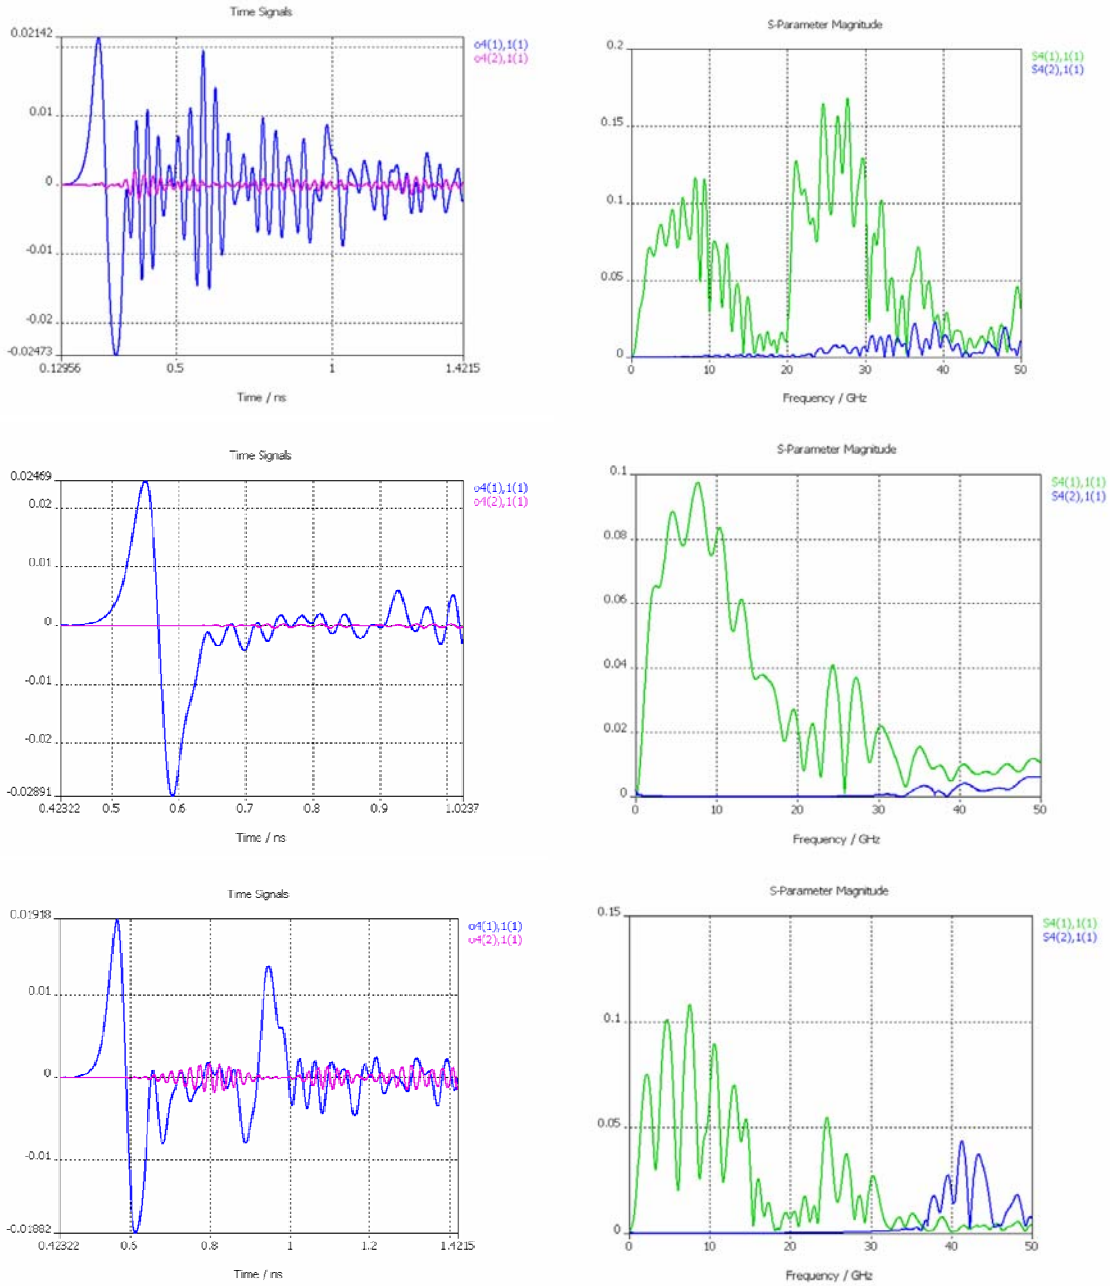

Figure 5.5.4 Comparison of time and frequency domain simulations of three different stripline designs: older stripline design (top), tapered without ceramic support (middle), tapered with ceramic support (bottom). Slope of time-domain signal of middle design is 35 % steeper than top design. Slope of bottom design is only 5% steeper than top design, but the amplitude is 20% smaller. The different colors represent the monopole and dipole modes at the output of the pickup.

None of the designs suffered from any distortion of the signal shape as the beam position was changed, but there were big differences in the steepness of the signal slope and the amplitudes of the signals. It appeared that the tapered design would have a slope at the zero-crossing that was 35% steeper than the existing design. When the tapered design was completed, however, it was apparent that unless the pickup antenna could be made light and hollow, there would need to be a ring made of ceramic Alumina to support the antenna and hold it in a stable position, reducing the risk that the feedthrough ceramic would crack and cause a vacuum leak. Although the diameter of the ceramic was designed to minimize the impact of the impedance mismatch that it creates, when this ceramic ring was added to the simulations, the comparison between the existing design and new, tapered design was much less dramatic (Fig. 5.5.4 (bottom)). While Vespel would offer 30% better performance than Alumina, it is not allowed in vacuum installations because it outgases under the influence of radiation.

The performance of the monitor predicted by the CST simulation cannot be measured up to 50 GHz. For a bandwidth of below 8 GHz, the simulation is in partial agreement with the performance measured with an 8 GHz oscilloscope (Fig. 5.5.5). The most noticeable difference between the two curves shown in Fig. 5.5.5 is the additional ringing that one sees on the oscilloscope signal. This is an artifact of the oscilloscope caused by the interaction of frequencies higher than 8 GHz with capacitive elements in the oscilloscope.

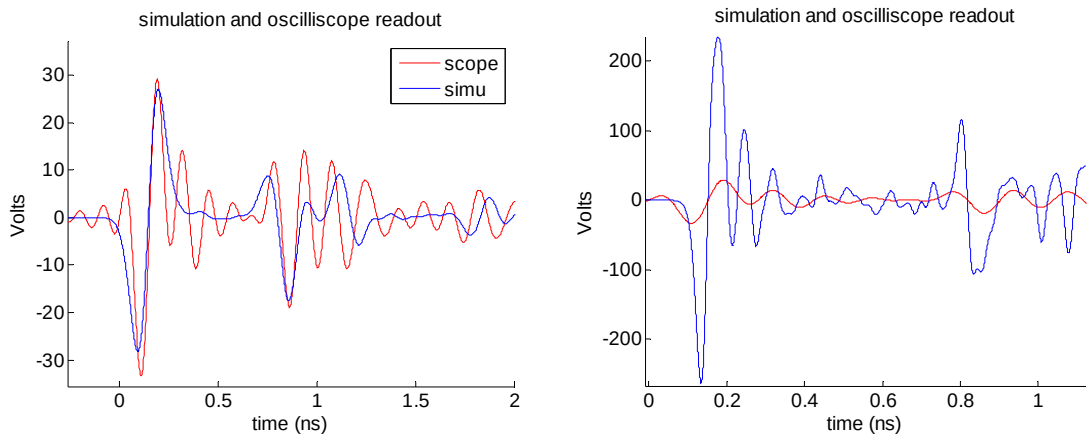

Figure 5.5.5 The simulated (blue) and measured (red) performance of the pickup below 8 GHz (left) and below 50GHz (right). Measurements were done with an 8 GHz oscilloscope.

There is a concern about the length of time that the signal from the pickup rings. The XFEL bunch spacing is only 200 ns and the ringing from the first bunch must be <0.01% of the peak voltage by the time the second bunch comes. When the signal was measured in FLASH with a higher bandwidth setup (the optical front-end to be described in chapter 7), the ringing is gone before the next bunch comes (1 μs), but it looks like there is still a significant amount of ringing at 200 ns after the bunch transient. It was suspected that a significant portion of this ringing comes from the splitter/combiner that was used in the

distribution of the signal, but when it was replaced with a splitter/combiner with a much lower insertion loss, the ringing was the same, suggesting that the ringing is primarily generated in the pickup. For the construction of the XFEL pickup, it is recommended that extra time be allotted for the development of a hollow pickup antenna that can be suspended between the two vacuum feed-throughs, without the need for the ceramic support rings.

When measured with the ~10 GHz optical front-end setup, a slope of 1.5 V/ps is measured. This is less than the 6 V/ps predicted in the 50 GHz simulation and this is due to the bandwidth limitations of the RF components that were used in the distribution of the signals. Nevertheless, at this stage, it is not desirable to further increase the slope of the signal because of the limitation that high signal slopes impose on the dynamic range of the measurement. Bypassing these limitations in order to achieve the highest resolution possible will be discussed in the context of the front-end measurement setup in Chapter 7.

Good agreement was also observed between a simulation of the pickup's frequency domain response and a measurement with a network analyzer. Although in the network analyzer plot (Fig. 5.5.6 left), the blue curve is lower than the green curve, the poor performance is due to oxidation on the contact to the feedthrough. When the contact was cleaned, much better agreement with the simulation (Fig. 5.5.6 right) was achieved.

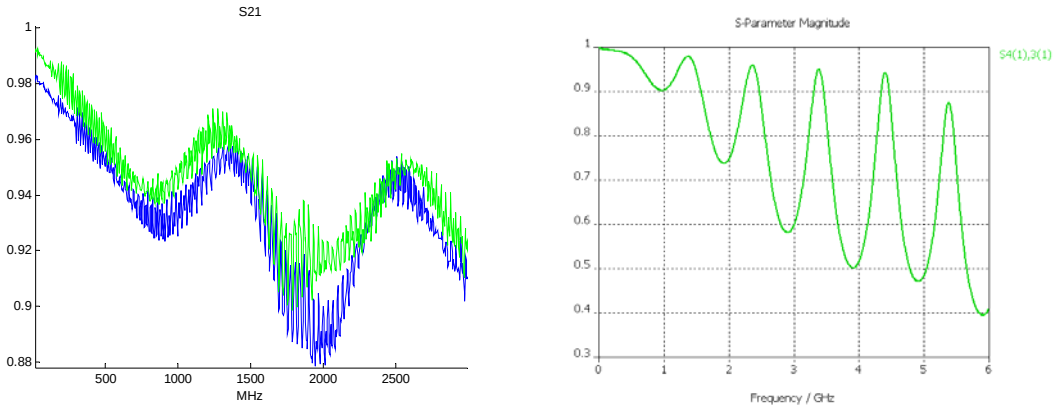

Figure 5.5.6 The pickup network analyzer measurement of two of the stripline outputs (left) and the simulation of a single pickup output (right). The network analyzer measurement shows poor agreement with the simulation at low frequencies due to oxidation on the pickup contacts. The ringing in the measured signals is due to ceramic supports which were not included in the simulation shown here.

Agreement with simulation only really serves to engender a sense of confidence in the simulation. The only test that really matters is to change the position of the beam over the full range and observe the corresponding changes in the zero-crossings of the signals.

The zero-crossings of the pickup signals were tracked as the beam was moved across the full range of the bunch compressor vacuum chamber and the resulting beam position measurements (Eq. 5.5.1) were plotted for both on and off-crest beams as a function of a change in beam energy (Fig 5.5.7). There are no anomalous distortions in the signal shape that disturb the position measurement over the full length of the pickup and the plot looks the same if the beam is on or off-crest. This result implies that the

pickup functions appropriately over the full range of the bunch compressor vacuum chamber.

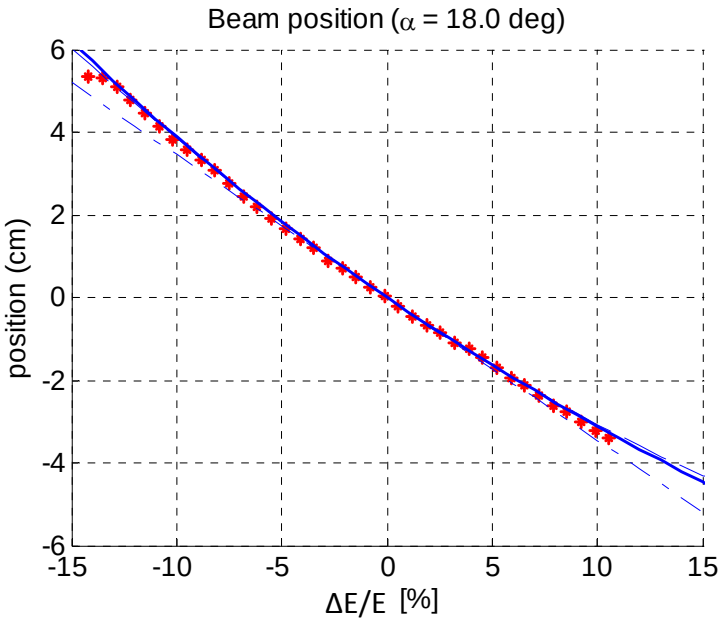

Figure 5.5.7 Beam position across the full range of the vacuum chamber as a function of the beam energy change. The red stars are the beam positions as measured with the pickup and an 8GHz oscilloscope. The solid blue line represents the expected position for various energy deviations as calculated with first, second and third order dispersion. The broken blue line is calculated with first order dispersion alone.

At the upper left edge of the plot in Fig. 5.5.7, the beam was scraping on the edge of the beam pipe. At the lower right end of the plot, the gradient of the cavity could not be increased any more. The higher order dispersion is also plotted; it is the curved line along which the measured positions lie.

Because the oscilloscope samples the signal many times, using an oscilloscope to measure the arrival-times of the pulses ignores the effect of amplitude changes of the signal and only measures the zero-crossing of the signal. When one must rely on only one sample point per zero-crossing measurement, one must sample the signal close to the zero-crossing and have a calibration measurement of the slope of the signal close to the zero crossing. Of course, as the signal amplitude decreases and the slope of the signal decreases, the resolution of the one sample-point measurement will decrease. Changes of the amplitude of the signal will make changes in the slope of the signal, and consequently, after any change in the amplitude of the signal, the measurement will need to be quickly re-calibrated or there will be an error in the measurement. The amplitude of the signal changes when the beam width changes and when the charge or y position changes. The 3 ps error produced by sampling the beam pickup signal 100 ps away from the zero-crossing for a 3% change in the charge of the beam is shown below in Fig. 5.5.8.

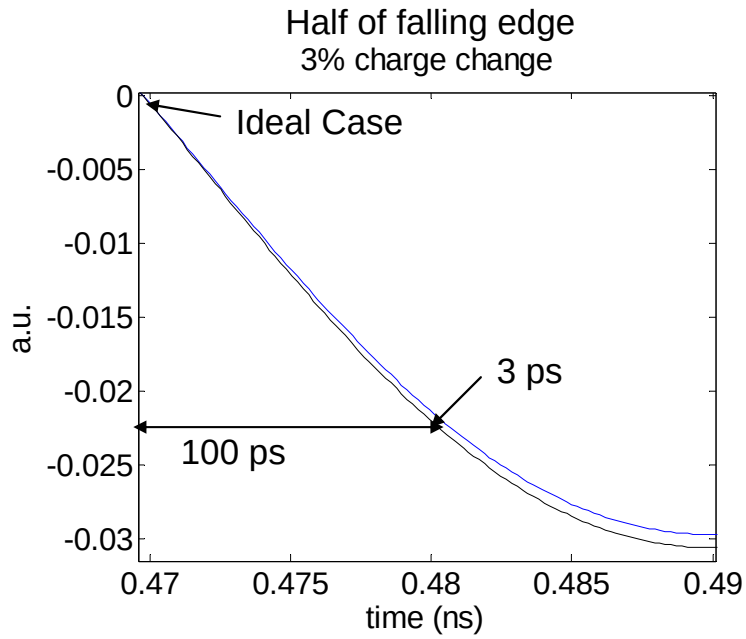

Figure 5.5.8 Impact of charge change on single sample-point which resides 100 ps away from the zero-crossing of the pickup signal. Ideally, the signal would be sampled at the zero-crossing, but a 3 ps error is incurred by sampling 100 ps away from the zero-crossing when the beam charge changes by 3%.

The influence of these pickup signal amplitude changes on the accuracy of the measurement can be removed through a routine that automatically re-calibrates the monitor on a regular basis or through a calibration constant that is updated based on a measurement of the phase, y position, and charge. The influence of these amplitude changes on the resolution of the measurement cannot be removed, but the impact on the resolution is typically small and is only of significance when the charge of the beam is dramatically changed. The dependence of the resolution on the beam charge is linear and if, for example, a lower charge of 0.2 nC were to be used, the resolution would be less than a quarter of what it would be for a 1 nC beam. The effect of changing the beam charge by large amounts can be compensated by adding or removing attenuators on the measurement front-end.

When the beam is vertically centered in the vacuum chamber, the BPM resolution will not be dramatically affected if the vertical position of the beam jitters by a few hundred microns, (Fig. 5.5.9), but for any static position changes, the measurement will need to be re-calibrated. The y-position sensitivity appears to provide the option that for low charge levels, the beam could be steered close to the pickup in order to improve the resolution of the measurement. If the beam gets within a couple of millimeters of the pickup, however, the calibration will become unstable due to the high sensitivity to vertical position changes.

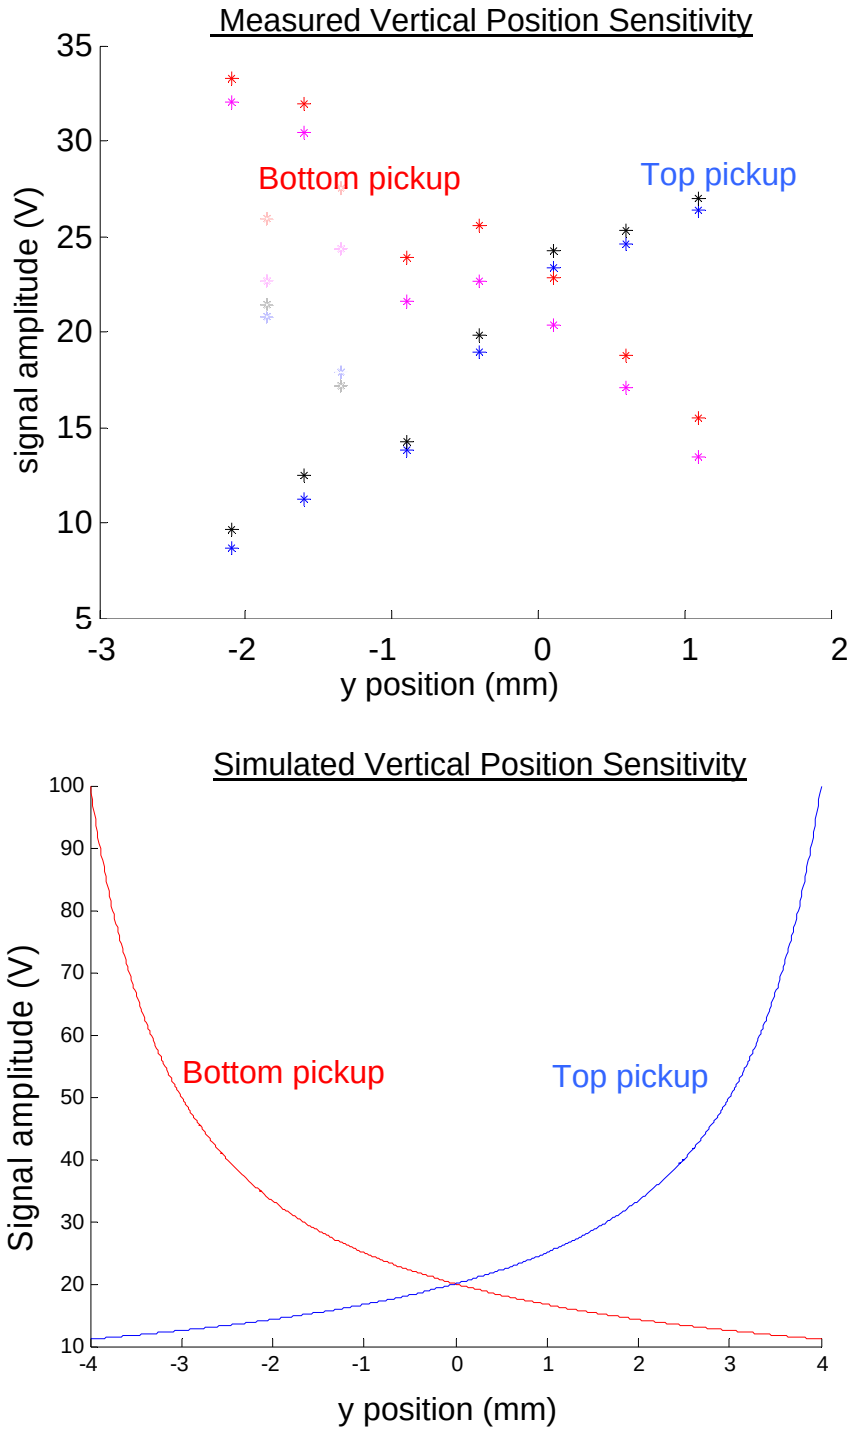

Figure 5.5.9 Measurement (top) and simulation (bottom) of chicane BPM pickup signal amplitude response to changes in y position.

If the top and bottom signals of the monitor are combined with equal-length cables, a cancellation of this dependence will occur. While it is possible to reduce the dependence, experience with signal combination suggests that it is not possible to entirely remove it.

While the impact of the beam charge and y-position is fairly easy to predict, the impact of the horizontal position spread of the beam is not as obvious. A measurement of the influence of the phase of the upstream accelerator section on the slope of the pickup signal is shown in Fig 5.5.10. The amplitude changes of the pickup signal due to RF phase changes are weak and will not impact the resolution of the monitor by more than a few percent.

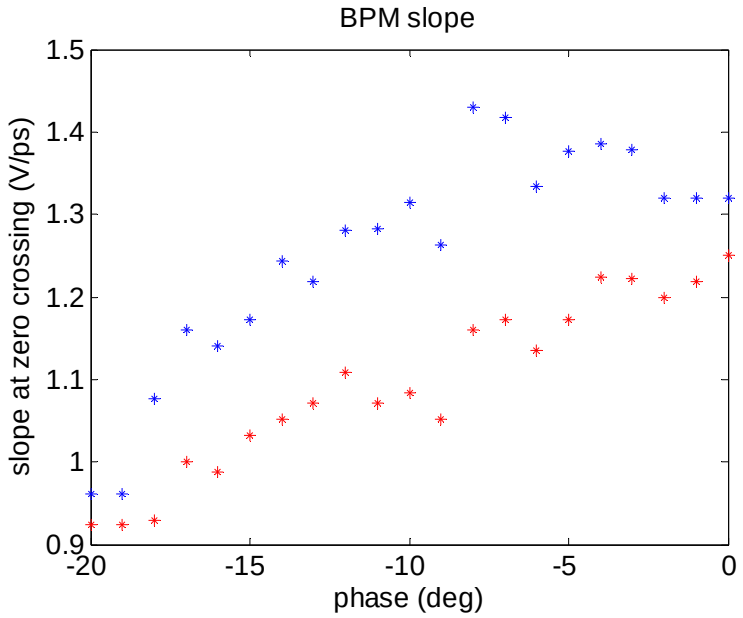

Figure 5.5.10 Dependence of the slope of pickup signal on the phase of the upstream accelerator section. The upstream phase determines the energy spread of the beam and, therefore, the beam width in the chicane. Right (blue) and left (red) pickup output slopes plotted as stars.

The dependence of the beam position on the phase of the upstream accelerator section shows less than perfect agreement between the expected position and the measured position for deviations from on-crest phase which are larger than 15 degrees. For these off-crest phases, the beam is wider, but this alone should not be detected in the beam position measurement. When the beam is wider, the position measurement is more sensitive to tilts of the beam. In Fig. 5.5.11 one sees the change in position measured as the phase of the accelerating RF is changed. The predicted change of beam position due to energy change is shown as the solid line, while the measured positions are shown as stars.

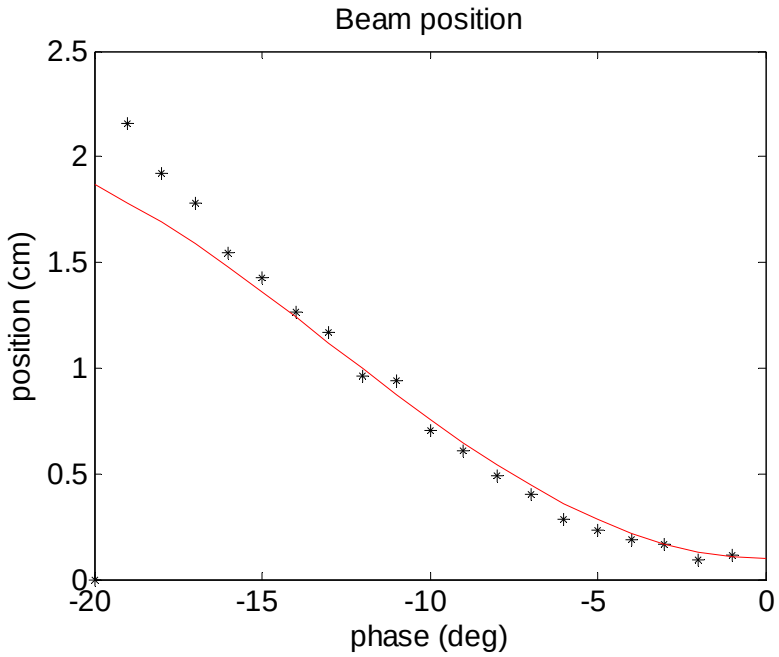

Figure 5.5.11 Change in beam position as a function of RF phase. The beam position change predicted by the change in beam energy is shown as the solid line, while the measured positions are shown as stars. There is poor agreement with the predicted energy change for large phases (wide beams). This is most likely due to a wide and tilted beam.

Wide and tilted beams will be treated in the following chapter.

6 Impacts of beam shape and orientation

A standard BPM measures the beam position by comparing the amplitudes of signals from two pickups. If the transverse size of the beam is small compared to the distance from the beam to the pickup, the position of the beam's center-of-mass is measured and the particular transverse distribution can be neglected. This is not the case for the chicane BPM pickup. The measurements of the transversely mounted stripline pickup use the arrival-times of the pulses at the ends of the pickup and due to the large dispersion in the chicane and energy spreads of up to 1%, the transverse beam-size may influence the measurement. If the beam is tilted in the x-y plane or has a longitudinally asymmetric charge distribution, this will also affect the signal produced by the transversely mounted stripline pickup. First, however, the way in which a wide beam couples to the pickup will be described.

6.1 Pickup Signals from a Wide Beam

In a Green's function approach to the description of the pickup signals generated by a wide beam traveling under a transversely mounted stripline pickup, we first take the voltage output of the stripline resulting from a pencil-like beam passing beneath the middle of the stripline. This voltage can be taken from the CST simulations presented in the previous chapter and will be given the name $U_0(t)$. We can use it to determine the voltage output for various charge distributions by summing together the results from various arrangements of pencil-like beams. There is a Green's function, G, associated with the differential operator L from the linear differential equation

$$L(x_0,t)U_0(t) = \lambda(x_0,t) \tag{6.1.1}$$

where λ is a known, homogeneous, linear charge distribution from a pencil-like beam and where the inverse of the differential operator is defined in terms of the Green's function by

$$L^{-1}(x_0,t)\lambda(x_0,t) = \int G(x_0,t-t')\lambda(x_0,t')dt'. \tag{6.1.2}$$

The solution to Eq. 6.1.1 can then be written in terms of the Green's function

$$U_0(t) = \int_{-\infty}^{\infty} G(x_0,t')\lambda(x_0,t-t')dt' \tag{6.1.3}$$

Let x_0 be the center of the stripline and when the pencil-like beam position is altered by Δx, Δx is much less than x_0 so that the signal at the exit of the pickup will not change in shape but only be delayed by $\Delta x/c_0$. This critical assumption was verified with both CST simulations and a mockup of the pickup and electron beam which could be moved with a micrometer. We will also assume that the pulse shape will not change if the beam arrival-time changes. In terms of the Green's function, these assumptions are written for the left (+) and right (-) pickup outputs as

$$G(x_0 + \Delta x,t) \cong G_{\pm}(x_0,(t-\Delta t) \pm \Delta x / c_0), \tag{6.1.4}$$

meaning that the function is invariant under translation in space and time and can therefore be used as a convolution operator. Let us define the charge distribution of the thin pencil beam as Gaussian in x and t,

$$\lambda(x,t) = \frac{Q}{2\pi\sqrt{\sigma_x\sigma_t}} \cdot e^{\frac{(t-t_0)^2}{2\sigma_t^2}} \cdot e^{\frac{(x-x_0)^2}{2\sigma_x^2}} \tag{6.1.5}$$

Which is normalized to the bunch charge Q according to

$$\int \lambda(t,x)dtdx = Q. \tag{6.1.6}$$

Integrating in over slices of the beam defined by the pencil-like beams (Eq. 6.1.3), we get the pickup output voltage that would result from a wide beam,

$$U(t) = \int_0^{2x_0} \int_{-\infty}^{\infty} G(x',t-t')\lambda(x',t')dt'dx', \tag{6.1.7}$$

Substituting in Eq. 6.1.5, this becomes

$$U_{\pm}(t) \cong \int_{0}^{2x_0} \int_{-\infty}^{\infty} G_{\pm}((t-t') \pm x'/c)\lambda(x',t')dt'dx' . \qquad (6.1.8)$$

If the bandwidth of the beam is much larger than that of the pickup, or, likewise, $\sigma_t << \Delta t$, we can write the result for the pickup output for a wide beam

$$U_{\pm}(t) \cong \int_{0}^{2x_0} G_{\pm}(t \pm x'/c_0)\lambda(x')dx' . \qquad (6.1.9)$$

For a beam that is tilted in the x-y plane an additional function, *y(x),* must be incorporated into the equation

$$U_{\pm}(t) \cong \int_{x_1}^{x_2} G_{\pm}(t \pm x'/c_0)\lambda(x',t)y(x')dx' \qquad (6.1.10)$$

y(x) weights the individual slices of the beam in terms of how their y position influences the amplitude that the pencil-like beam would produce on the pickup.

For a beam that is tilted in the x-z plane, the pulses traveling to the left on the pickup and the pulses traveling to the right on the pickup need to be given separate treatments. If the (+) side of the beam arrives earlier than the (-) side of the beam, the arrival times of the pulses traveling to the (+) side will be condensed with respect to one another and the arrival times of the pulses traveling to the (-) side of the pickup will be spread out with respect to one another. This is described by

$$U_{+}(t) \cong \int_{x_1}^{x_2} G_{+}(t + x'/c_0 + t_{tilt})\lambda(x',t)dx' \qquad (6.1.11)$$

and

$$U_{-}(t) \cong \int_{x_1}^{x_2} G_{-}(t - x'/c_0 - t_{tilt})\lambda(x',t)dx' . \qquad (6.1.12)$$

Further complications from an asymmetric charge distribution can also easily be accommodated with this method.

The goal of this is to determine how the measurement is affected by wide, tilted beams and how the measurement differs from the actual, center-of-mass beam position and arrival time. The position and arrival-time of the beam are measured according to

$$x_{meas} = \frac{c_0}{2}(t_+ - t_-) \quad \text{and} \quad t_{meas} = \frac{1}{2}(t_+ - t_-) \qquad (6.1.13)$$

with $U_{\pm}(t_{\pm}) = 0$ at the zero crossings of the signals. The real center-of-mass position and arrival-time of the beam are given by

$$x_{beam} = \frac{1}{Q}\int\int \lambda(x,t)\cdot x dt dx\,, \quad t_{beam} = \frac{1}{Q}\int\int \lambda(x,t)\cdot t dt dx \qquad (6.1.14)$$

Using either this Green's function method or the numerical simulations from CST, one can predict that the measured beam arrival-time will be different from the real, center-of-mass arrival-time when the beam width changes. Using the Green's function method alone, one can predict that the measured beam position changes compared to the center-of-mass beam position when a wide, asymmetrical charge distribution is tilted in the *x*-*y* or *x*-*z* planes. This approach has the advantage that Green's functions for 2-D and 3-D transient cases can be found through multiplication of 1-D cases. The solution takes the form of the superposition (sum) of several integrals. The magnitude of these tilt and width effects will be detailed below.

6.2 Beam Width Changes

When the phase of the upstream accelerator section is changed, the width of the beam will change. The charge density changes when the beam width changes. Since the coupling of the beam to the pickup occurs at locations of changing charge density, different charge densities will produce signals with different amplitudes. More rapidly changing charge density produces a higher amplitude signal than a slowly changing charge density. This is depicted in Fig. 6.2.1 for an elliptical beam shape and for a flat beam shape; both are types of beams which can be generated with different injector laser parameters.

Figure 6.2.1 Coupling of the beam to the pickup for an elliptical beam (left) and for a flat beam (right) is shown on the top in black. The charge distribution of the beams is shown on the bottom in red. The beam is directly under the pickup.

It is easier to imagine why the pictures above look as they do if one imagines the beams divided up into slices. Each slice generates a pulse traveling to the left and a pulse traveling to the right. When a slice has the same charge and vertical position as its neighboring slices, the pulses that it generates will be canceled-out through destructive interference with the pulses generated by the neighboring slices (Fig. 6.2.2).

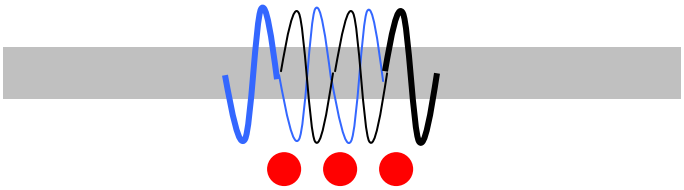

Figure 6.2.2 Cancellation of signals on the pickup through destructive interference for neighboring pencil-like beams. When the beam is no longer under the pickup, the only signals remaining are the ones shown in bold.

While the signals on the pickup look very different directly above the beam for the two cases shown in Fig 6.2.1, by the time that the pulses have been transported to the outputs of the pickup, they have very similar properties. This is due to the dispersion of the pulse as it travels along the pickup and the filtering effects of the impedance mismatches in the pickup and vacuum feedthrough. The higher frequency components of the spectrum of the pulse will be more strongly suppressed than the lower frequency components such that after transport to the output of the pickup, the length of a shorter pulse has increased by more than the length of a longer pulse. It is still clear from the picture, however, that, even for an elliptical beam, when the beam is directly under the pickup, there is a space between the zero crossings of the signals on the pickups. That space is proportional to the width of the beam and, consequently, to the energy-spread of the beam. If an incoming beam arrival-time measurement is available from before the chicane, the difference between the incoming arrival-time measurement and the arrival-time measured with the BPM pickup in the chicane will give a measurement of the beam energy spread according to:

$$arrival_{incoming} - arrival_{BPM} = R_{16} * \Delta E/E.$$

This is because incoming arrival is measured with button pickups and a non-dispersed beam and the BPM arrival-time in the chicane measures the arrival of the locations where the charge distribution is changing for a beam which is stretched out transversely. The arrival-time measured with the BPM is given by the average of the arrival-time measured by both outputs of the pickup. The arrival-time measured by the BPM will then be later for narrow beams and it will be earlier for wide beams. A simulation of this principle for different beam widths is shown in the plot of Fig. 6.2.3.

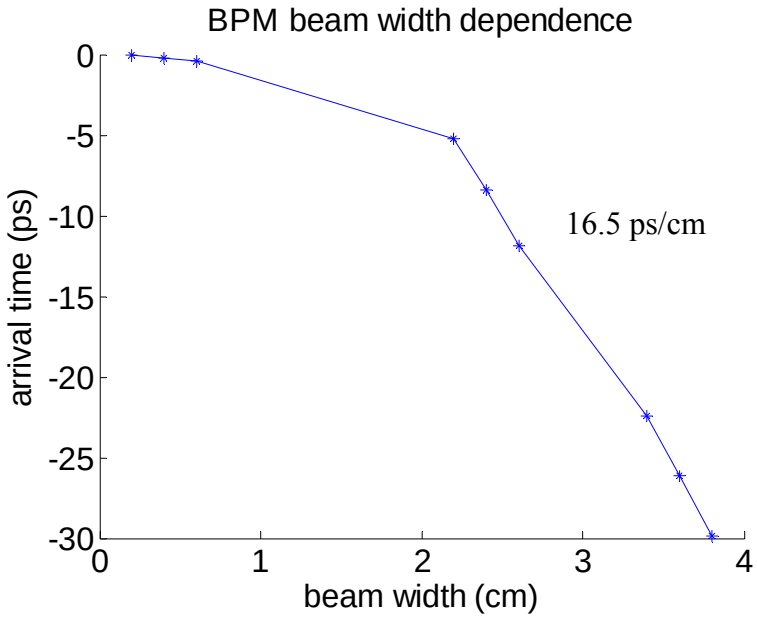

Figure 6.2.3 Sensitivity of the chicane BPM arrival-time measurement to changes in the width of the beam.

Based on this plot generated with CST simulations of the average pulse arrival-times at the stripline outputs for various widths of the elliptical beam shown on the left-hand side of Fig. 6.2.1, the dependence of the beam arrival-time as measured with the chicane BPM on the width of the beam is 1.65 ps/mm for beams that are more than 20 mm wide, where the width is given by $\pm 3\sigma_x$. This is consistent with the 3.3 ps/mm conversion factor of a pulse traveling at the speed of light. When the beam is less than a centimeter wide, this sensitivity drops to <0.1 ps/mm. This is due to the fact that when the length or height of the beam is comparable to the width, there is no longer a significant sensitivity of the arrival-time to the width of the beam. The FWHM bandwidth of the simulation used for this plot was 20 GHz and this corresponds to a beam length of ~15 mm. In the figure above, when the beam width is greater than 15 mm, the sensitivity of the arrival-time measurement to the beam width starts to become significant. The actual length of the beam is much shorter than 15 mm, so the limitation of using the chicane BPM to measure the beam energy spread will be due primarily to the limited bandwidth of the pickup itself. The energy spread of the beam is related to the bunch length after the chicane, and changes in several length scales of the bunch are measured with high precision are measured with high precision with a pyrodetector based single shot spectrometer or bunch length monitor.

Whereas at FLASH, the beam is only about a centimeter wide (6 σ_x), in the XFEL, the beam may be as wide as 6 cm. With wider beams comes not only a higher sensitivity of the chicane BPM arrival-time measurement to the energy spread of the beam, but also a higher sensitivity to tilts of the beam.

6.3 Tilted in x-y plane

A beam tilted in the x-y plane is depicted beneath a bar representing the pickup in Fig. 6.3.1. The lines represent the waves traveling on the pickup as a result of the beam transient. Where the beam is closer to the pickup, the signal amplitude is larger, where it is farther away, the coupled signal is smaller.

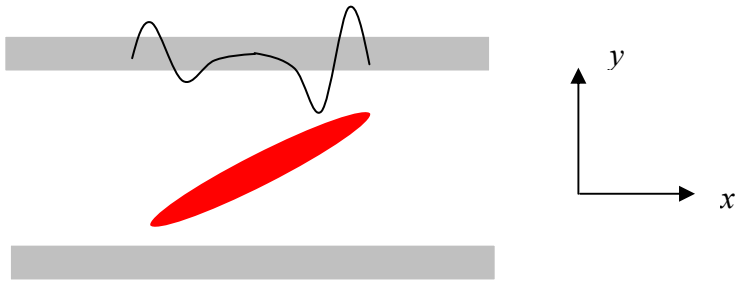

Figure 6.3.1 The beam tilted in x-y plane relative to the pickup. The side of the beam that is closer to the pickup produces a larger amplitude signal than the side that is further away.

One must first know the sensitivity of the monitor to changes in the y position of the beam before making an estimate of the sensitivity of the monitor to changes in y tilt. This was done with the CST simulation and is shown in Fig. 6.3.2. Over a few millimeter range, it was verified with oscilloscope measurements.

Using an estimation of the macro-beam as a collection of pencil-like beams that can be arranged in the x-y plane, the signals produced by each beam slice can be summed together as described in the Greens function integration of Sect 6.1. This provides an estimate of what the pickup output would look like for a charge distribution which is tilted in the x-y plane. In the first step of this integration process, the y-position sensitivity data from Fig. 6.3.2 was used to calculate the amplitudes of the signals that would be induced on the pickup for each slice of a 1 cm wide ($\pm 3\sigma_x$) beam which is tilted by 5 degrees in the x-y plane. From Fig. 6.3.2, it is clear that the further the beam is away from y=0, the stronger the effect of an x-y tilt will become.

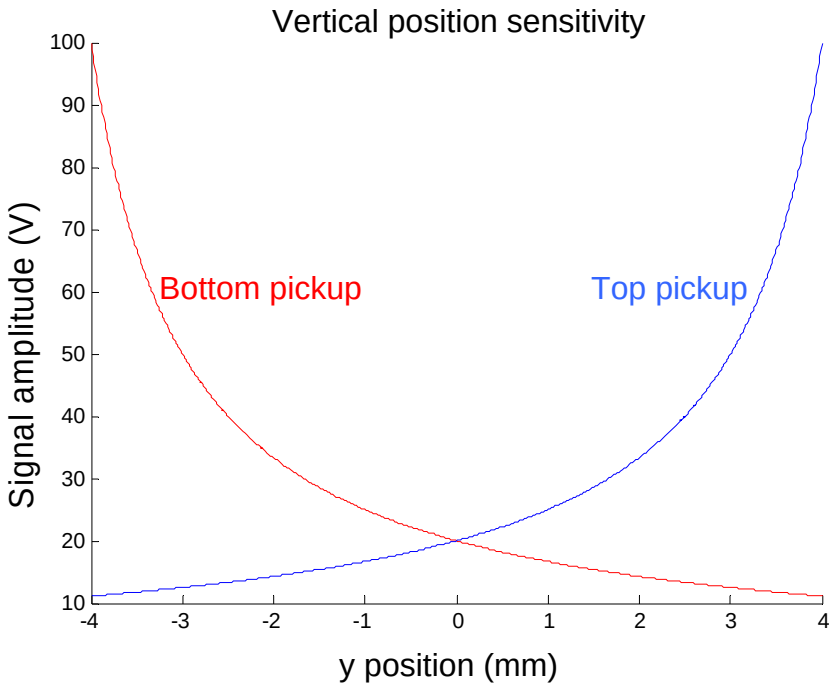

Figure 6.3.2 CST simulation of the sensitivity of the chicane BPM signal amplitude to changes in y position of a pencil-like beam with a charge of 1 nC. For y=0, the beam is in the middle of the vacuum chamber.

In Fig. 6.3.3, the position of the beam slice is shown on the x-axis and the signal amplitude induced on the pickup by the beam slice is shown on the y-axis. The scaling of the amplitude varies according to the number of slices into which the beam is divided and the amount of charge contained within each slice. In Fig. 6.3.3, the pulse amplitudes that slices of a tilted beam with a flat charge distribution would produce on the pickup are shown. The amplitudes from Fig. 6.3.3 can be added together for pulses traveling to the left and to the right in order to produce Fig. 6.3.4. This gives an estimate of the shape of the signals produced by a flat beam that is tilted in the x-y plane compared to the same beam that is not tilted. The beam length was 1 mm, the beam width was 10 mm and the tilt was 5 degrees. The difference in the arrival-time of the tilted signal's zero-crossing at the exit of the pickup compared to the non-tilted case is 1 ps. If the charge distribution is not flat, but Gaussian, with a length of 1 mm, a width of 4 mm (FWHM) and a tilt of 5 degrees, the difference between the tilted and non-tilted cases is 0.3 ps (200 μm). These differences due to x-y tilt constitute errors in the measurement of the beam centroid.

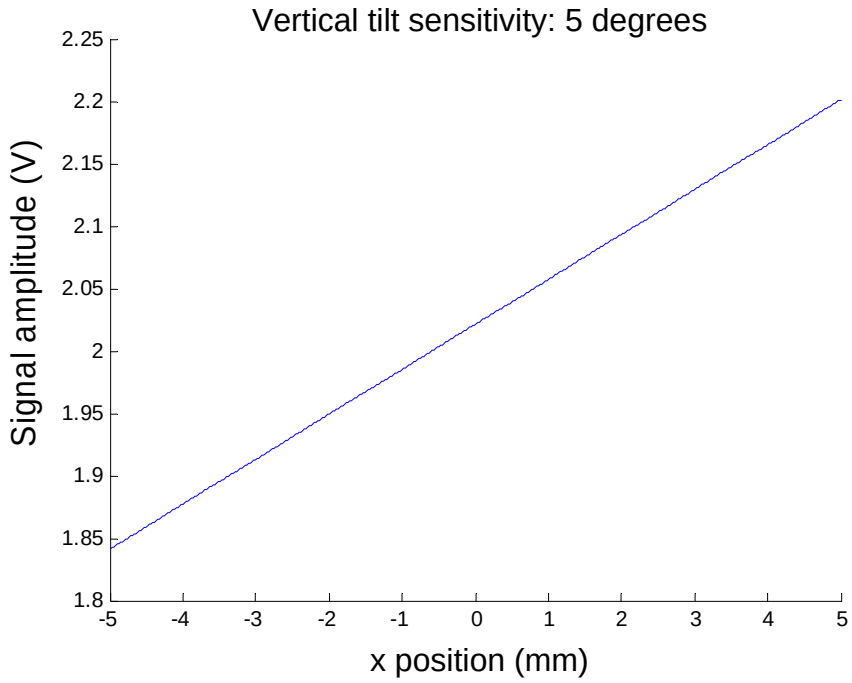

Figure 6.3.3 The amplitude of the signal induced on the pickup by a slice of a beam with a flat charge distribution that is tilted in the x-y plane by 5 degrees as a function of the x position of the slice within the beam.

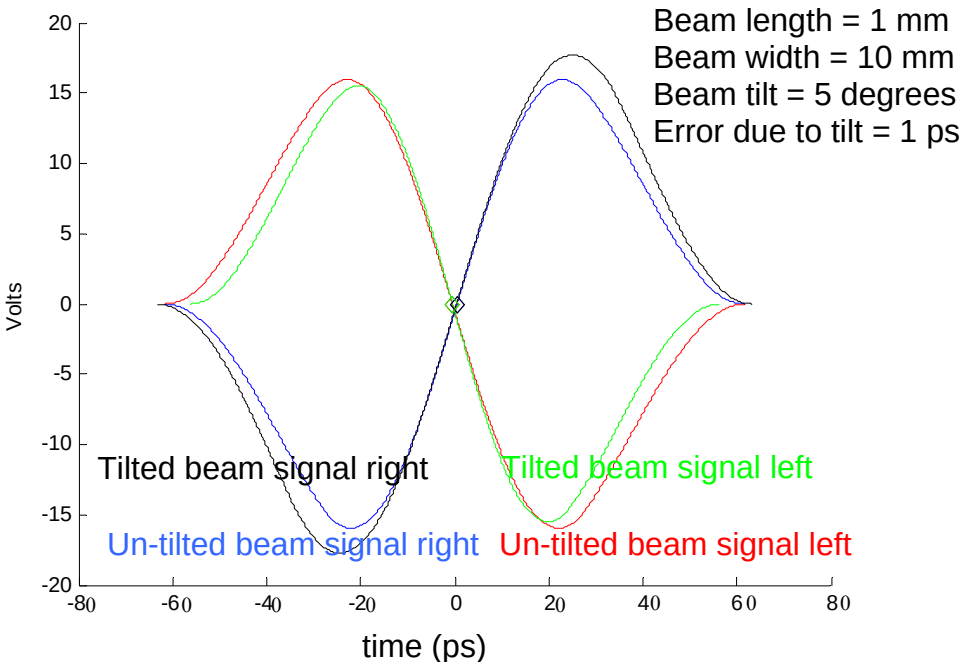

Figure 6.3.4 Pickup outputs for tilted/un-tilted beams with flat charge distribution. Difference between arrival-times of zero-crossings of tilted and un-tilted beams gives measurement error.

Such an x-y tilt was measured with the beam image on the synchrotron light monitor screen and the resulting changes in the beam position measured with the beam pickup and an oscilloscope were recorded (Fig. 6.3.5). The position of the beam was held constant and the tilt of the beam was changed and measured on a screen. The sum of the pickup signals' arrival-times is constant while the difference of the pickup signals' arrival-times has a dependence on the tilt of the beam. The measurement provides an opportunity to cross-check the simulation presented in Fig. 6.2.4. In the simulation, 5 degrees of tilt of a 4 mm (FWHM) wide Gaussian beam produces 200 μm of position measurement error and 5 degrees of tilt in the measurement would produce 250 μm of position measurement error. Since the beam shape is more complicated than a simple Gaussian (Ch. 4), this can be considered good agreement.

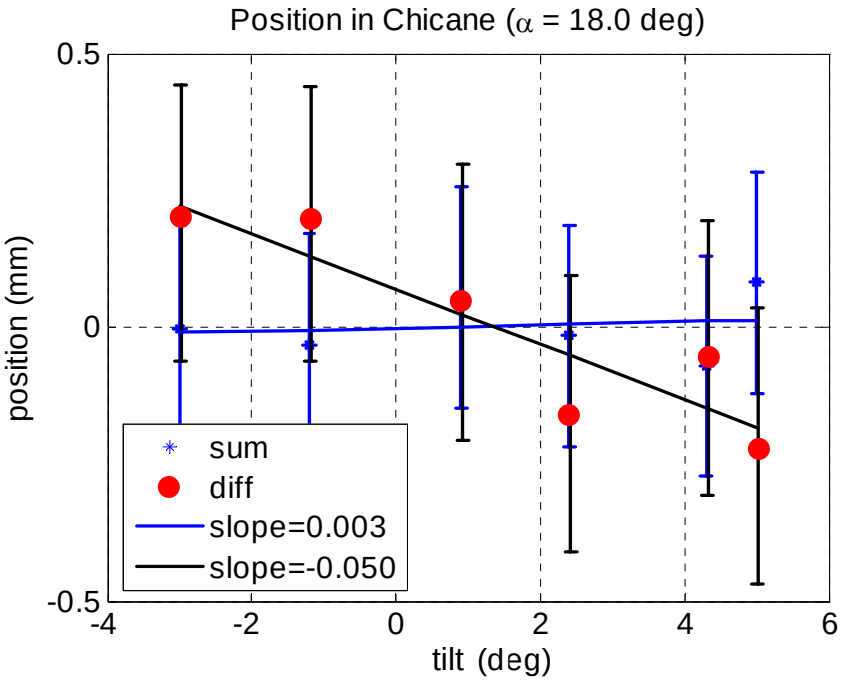

Figure 6.3.5 Impact of x-y beam tilt on beam position measurement. The position of the beam was held constant and the tilt of the beam was changed and measured on a screen. The sum of the pickup signals' arrival-times is constant while the difference of the pickup signals' arrival-times has a dependence on the tilt of the beam. The error due to a beam tilt of one degree will cause 50 um of measurement error.

Although the resolution of the oscilloscope measurement was approximately equal to the energy stability of the beam (5e-4), averaging over 30 pulses reduced the measurement error to a few microns. The tilts were created by making orbit bumps around the first accelerator section as described in the chapter on the beam shape in the bunch compressor (Ch. 4). Using data from the plots in Fig. 6.3.3 and Fig. 4.3.6, a 4 mm orbit bump in the first accelerator section causes 2 degrees of beam tilt in the chicane and causes a ~100 μm position error to be measured by the BPM. This means that, since the tilt effect caused by dispersion downstream of the first accelerator section is a linear

effect, the typical 100 μm orbit jitter downstream of the first accelerator section could cause enough beam tilt jitter to create a 2.5 μm error in the position of the beam as measured by the chicane BPM. Given a target resolution of ~1 μm, this tilt jitter effect could eventually limit the measurement's resolution.

Unless the measurements or signals from the top pickup and the bottom pickup are combined, this sort of tilt will generate an error and lead a user to believe that the centroid of the beam has shifted when in fact the tilt of the beam has changed. While combining signals from different pickups can dramatically reduce this effect, it cannot remove it entirely and it could still become visible for very large horizontal position spreads, large vertical offsets or large beam tilts.

6.4 Tilted in y-z plane

A beam tilted in the y-z plane is depicted beneath a bar representing the pickup in Fig. 6.4.1.

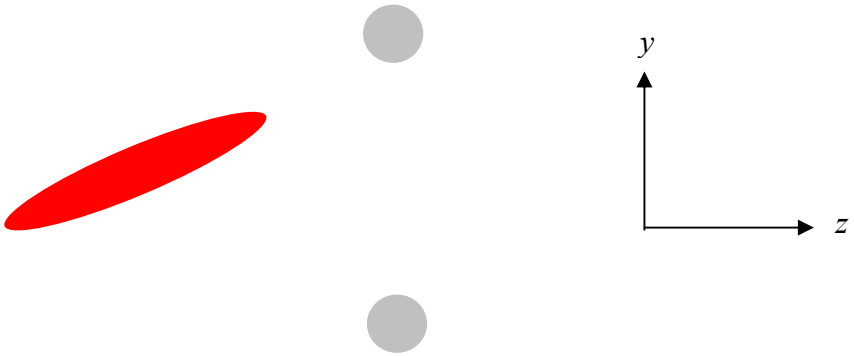

Figure 6.4.1 The beam tilted in y-z plane relative to the pickups above and below the beam.

This effect will only make the bunch seem shorter than it really is. This will not have a measurable impact on the measurement of the beam position, but it could have an effect on the beam arrival-time measurement.

6.5 Tilted in x-z plane

The beam is tilted in the x-z plane as shown in the picture of the particle distribution in the middle of the chicane (Fig 6.5.1). The distribution is more complicated than just a flat distribution or a Gaussian distribution tilted in the x-z plane, but for the sake of simplicity, in the following Green's function calculations, such a simplified beam will be used.

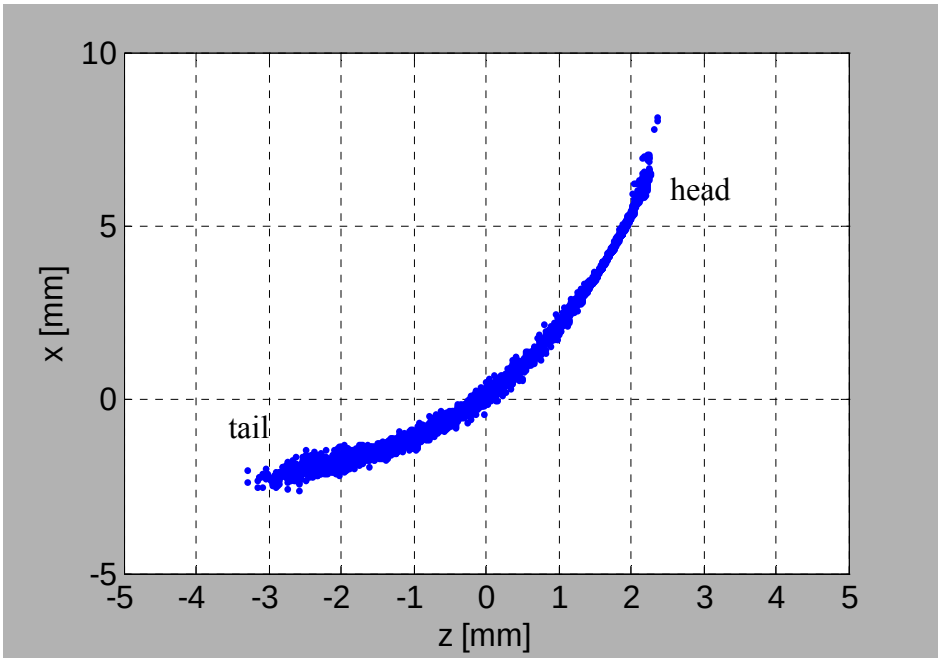

Figure 6.5.1 Particle tracking simulation of a nicely matched beam at the location of the BC2 BPM.

A beam tilted in the x-z plane is depicted alongside a bar representing the pickup in Fig. 6.5.2. One might naively think that since it appears that the head of the bunch would arrive 14 ps earlier than the tail, there would then be a 4 mm error in the measurement of the center of the beam, but this would be mistaken.

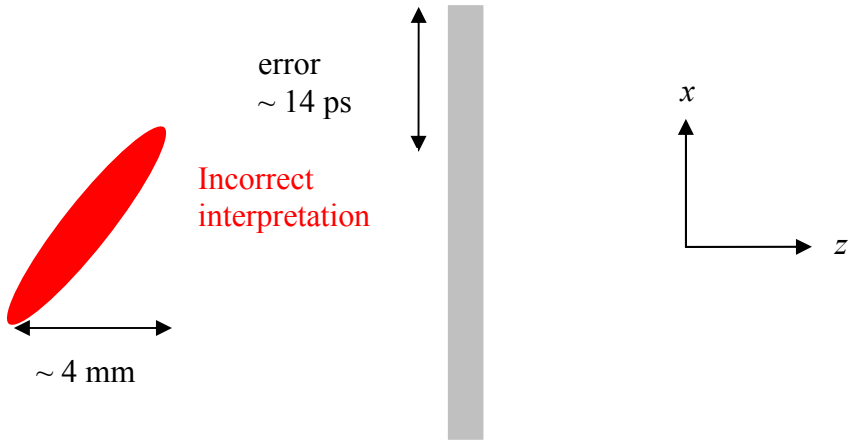

Figure 6.5.2 The beam tilted in x-z plane relative to the pickup. A naïve estimate of the effect of this tilt is sketched alongside the bunch.

While it is true that the bunch is tilted in this plane and that the separation between the head and tail in the x-z plane is about 4 mm, the interpretation shown in Fig. 6.5.2 ignores the Doppler-like effect of the tilted beam coupling to the pickup.

The Doppler effect refers to the change in the frequency of a signal that is measured by an observer which is moving relative to a source. While the frequency of the signal changes, the phase does not. If the beam is tilted, it appears that the beam is traveling more quickly in the reference frame of one end of the pickup, while in the reference frame of the other end of the pickup, it appears that the beam is traveling more slowly. If this still doesn't seem plausible, imagine the beam divided into many small slices. Each slice will couple to the pickup and a pulse will travel to the left and to the right. If you detect the arrival-times of each slice on one side of the pickup, the time elapsed between the slices will be larger on one end of the pickup than the time elapsed between the slices on the other end of the pickup. If all of these slices are added together on each end of the pickup, one finds that the periods of the signals on the opposite ends of the pickup are different but the arrival-time of the zero-crossing is the same as it would have been for a non-tilted beam. Pictorially, this is represented in Fig. 6.5.3.

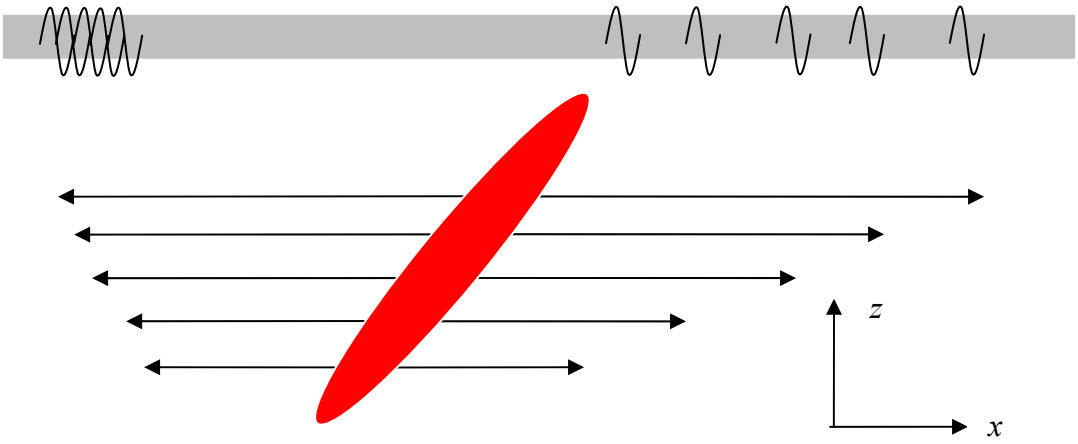

Figure 6.5.3 Illustration of the spacing of the wavelets produced by beam slices as they are transported on the pickup for a tilted beam.

In the numerical simulation of this effect (Eqs. 6.1.6, 6.1.7), the beam charge is broken up into many, small slices. Each slice induces a wavelet on the pickup. The wavelets are added together to generate a transient for the whole beam. When the zero crossings of the transients for the tilted beam case are compared to the non-tilted case it is apparent that, although the amplitudes of the beam transients are affected by the beam's tilt, the phases of the pulses are not; both measurements return the same value for the beam position and there would be no systematic error from this effect. The story changes, however, when the charge distribution is asymmetric.

6.6 Asymmetric Charge Distribution Tilted

When the charge distribution is asymmetric and the beam is tilted in the x-z plane, a systematic error is generated that cannot be removed by any available methods. For an asymmetric Gaussian charge distribution shown in Fig 6.6.1, the width is 4.5 mm, the length is 4 mm and the centroid is offset horizontally from the center by 3.3 ps.

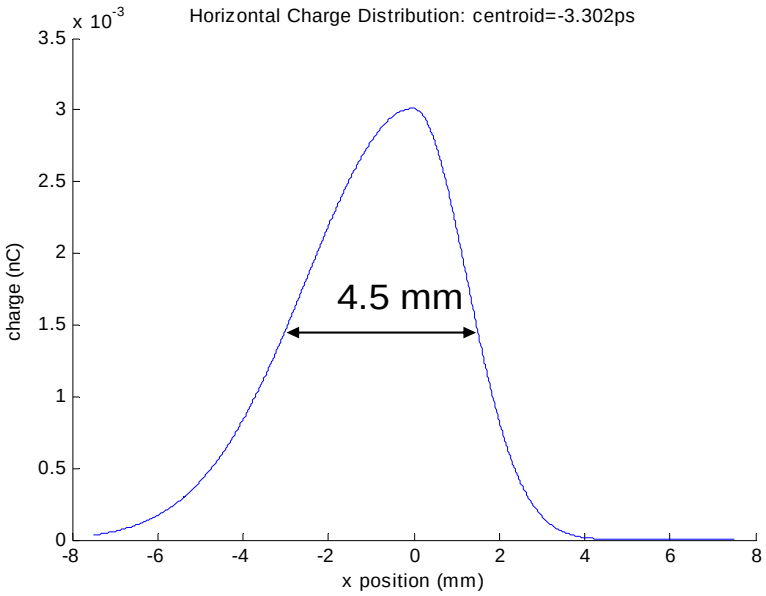

Figure 6.6.1 Asymmetric horizontal charge distribution with centroid offset from center by 3.3 picoseconds.

When a bunch with an asymmetric Gaussian charge distribution, like the one shown in Fig. 6.6.1, is tilted by ~45 degrees like the beam shown in Fig. 6.4.2, a 580 fs systematic error compared to a non-tilted beam results (Fig. 6.6.2). This is called an error because if one is trying to measure the beam centroid, one does not want to measure the tilt as well.

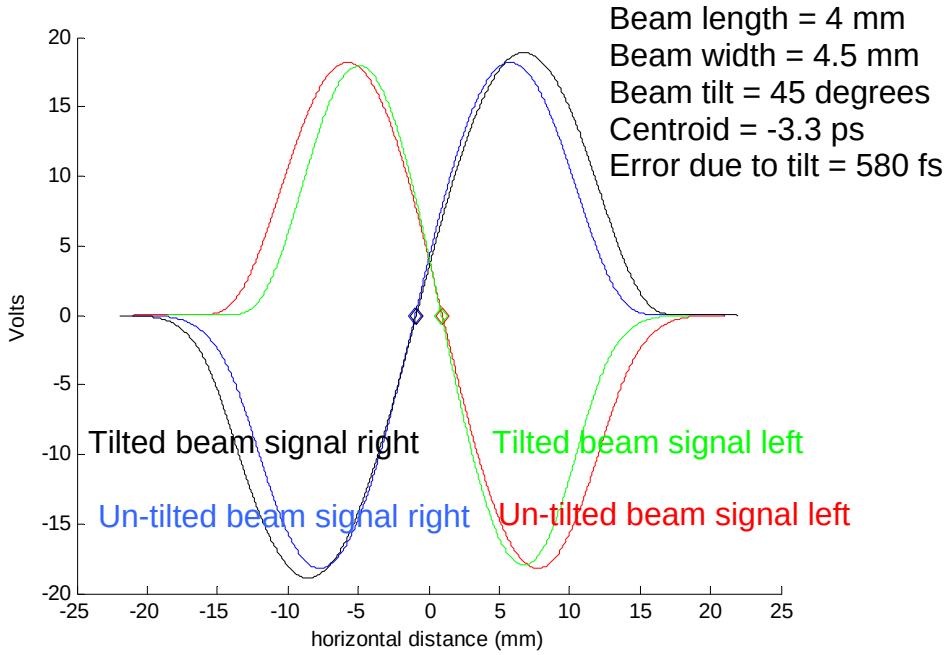

Figure 6.6.2 Error resulting from x-z tilted asymmetric charge distribution. An error of half of a picosecond is seen for the tilted asymmetric, tilted distribution (green, black) compared to the asymmetric, un-tilted distribution (red, blue).

In this simulation, the zero-crossings of the asymmetric charge distribution accurately give the position of the horizontally offset centroid position of the distribution, but the zero-crossings of the tilted and asymmetric charge distribution differ by 580 fs.

6.7 Wakefields

There is a valid concern that, in addition to the image charge of the beam as it passes beneath the pickup, the chicane BPM stripline measures fields originating from the beam's image charge from earlier in the chicane. Due to the large vacuum chamber width, the fields from previous positions before and in the bending magnets are detected by the pick-up. This would manifest as an averaging over previous beam positions.

While averaging over previous positions could reduce the slope of the signal produced by the pickup, thereby reducing the resolution of any zero-crossing sampling scheme, this would be evidenced by a gradual reduction of the amplitude and slope of the signal as the beam is moved from the inside of the vacuum chamber towards the outside. In oscilloscope measurements of the amplitude of the signal over the full range of the pickup, the amplitude and signal slope remained constant over the full range of the pickup (Fig. 6.7.1). This implies that this averaging problem will be difficult to detect unless one's measurement integrates over hundreds of nanoseconds instead of over a picosecond, as in the zero-crossing sampling scheme which will be described in the following chapter.

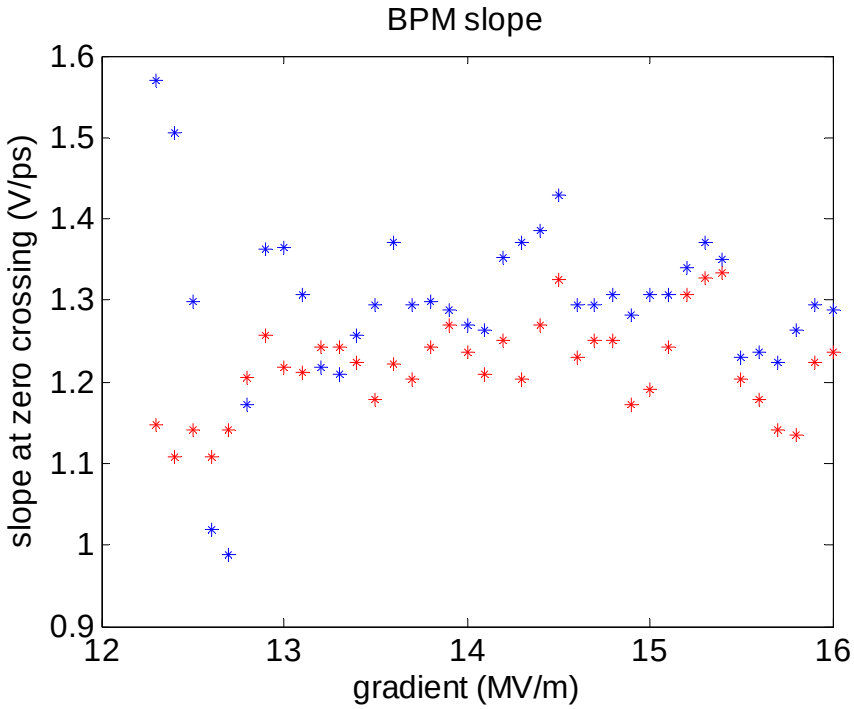

Figure 6.7.1 Slope at the zero-crossing of pickup signal over full dynamic range of monitor. For low gradients, the beam was beginning to scrape on the edge of the vacuum chamber.

The problem that could arise due to this averaging effect is that the average of the positions might change when the second dipole current changes or when the beam position or angle in the beam pipe is altered. While a true simulation of this sort of effect would require high frequency analysis of the entire chicane vacuum chamber, one can be relatively certain that for measurements that integrate over a picosecond, this averaging/wake effect should be negligible. For measurement techniques that integrate over hundreds of nanoseconds, modes and CSR wakes in the chamber could very well impact the accuracy of the measurement.

6.8 Summary

The conclusion that one can take from these studies of the beam tilt is that the beam position monitor measures a combination of the centroid of the beam, the tilt of the beam, and the path of the beam. If the beam is tilted in the x-y plane, the effect of the tilt can be mostly removed by combining the signals from the top and bottom pickups. If the beam has an asymmetric charge distribution and is tilted in the x-z plane, there is a systematic error of up to several hundred femtoseconds that cannot be removed by any means. In this case, one measures a property of the beam that is determined by both the centroid and tilt of the beam. The effect of the path of the beam appears to be so small that it is difficult to measure. A table summarizing the strongest effects is shown below

	X-Y		**X-Z**	
Shape	flat	Gaussian	symmetrical	asymmetrical
Length (mm)	1	1	4	4
Width (mm)	10	4	4.5	4.5
Tilt (deg)	5	5	45	45
Error (fs)	1000	300	0	580
Error (um)	660	200	0	400

Table 6.8.1 Effects of beam tilts on beam position measured with transversely mounted stripline BPM. Several hundred microns of measurement error can be expected from typical beam distributions and tilts.

7 Chicane BPM Front-ends

Using the transversely mounted stripline pickup, the beam position and arrival-time is determined by measuring the arrival-times of the pulses coming from opposite ends of the pickup. So far, we have ignored the fact that these pulse arrival-time measurements must be made relative to a reference signal. In the equation for the beam position, if the same reference is used for both left and right pickup outputs, the influence of the phase stability of the reference cancels out:

$$beam\_position = \tfrac{c}{2} \cdot \left[\left(arrival\_left - reference \right) - \left(arrival\_right - reference \right) \right]$$

where c is the speed-of-light and *reference* is the phase of the reference signal against which the arrival-times of the pulses are measured. If, however, the arrival-time of the beam is measured from the same pickup signals that are used to perform the beam position measurement, the phase of the reference does not cancel out:

$$beam\_arrival = \tfrac{1}{2} \cdot \left[\left(arrival\_left - reference \right) + \left(arrival\_right - reference \right) \right]$$

From this we can conclude that while a front-end for a chicane beam position monitor and for a beam arrival-time monitor must both be able to measure the arrival-times of the beam transient pulses emerging from the pickup with femtosecond precision in order to meet the resolution requirements described in previous chapters, the beam position measurement has much looser tolerances on the stability of the reference signal. Both types of measurements must, however, measure the amplitudes of the pickup signals

around their zero-crossings in order to avoid measuring changes in the charge or vertical position of the beam.

These problems can be approached with either RF or optical methods. The RF front-end and the optical front-end for the chicane BPM are compared and contrasted with respect to their limitations, cost and performance. In general, given certain modifications, they can deliver comparable performance, but the optical measurement has a much lower potential for making systematic errors. The optical measurement is considerably more expensive than the RF measurement and requires a complicated infrastructure to implement.

7.1 RF Front-end Concept

The RF front-end of the chicane BPM contains circuits that operate at 10.4 GHz and 1.3 GHz, delivering two distinct measurements of the beam position. The lower frequency and lower resolution measurement gives the information required to set the position of a mechanical phase shifter for the higher frequency and higher resolution measurement. Both measurements utilize the same down-mixing-to-baseband principle; they take the outputs from the pickups, filter out a certain frequency from the spectrum and mix that with the same frequency generated by the machine reference signal from the nearby master oscillator. A simplified schematic illustrating the down-mixing concept for the higher frequency is depicted in Fig. 7.1.1.

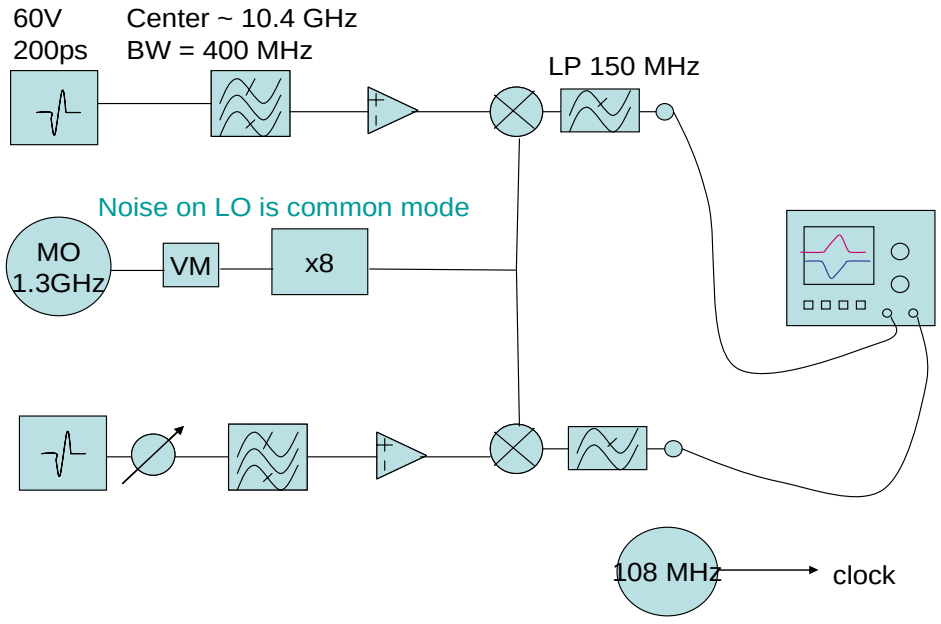

Figure 7.1.1 Down-mixing scheme to measure the relative phases of two pulses.

In Fig. 7.1.1, the output signals from the pickup are depicted on the left as bipolar pulses that are about 200 ps long and 60 V in amplitude. These parameters vary depending on the beam charge and the cables used to deliver the signals. The signals come from each side of the pickup and the difference between their arrival-times is

proportional to the beam position. They are each sent through a 4-pole band-pass filter with a center frequency of 10.4 GHz and a 200 MHz FWHM bandwidth. The filter has four poles so that the group-delay of the signal in the filter is flat in the pass-band. The bandwidth is large so that the group delay doesn't respond too dramatically to temperature changes; the smaller the filter bandwidth is, the more sensitive it is to temperature changes. Even though, the pulse from the pickup is more than a hundred volts and less than 100 ps long, after the 30 meter cable, the bandpass filter, and the mixer only a few mV remain. To compensate for this loss, the filtered signal is amplified and then mixed with what is known in mixer terminology as the Local Oscillator (LO). The LO is generated from the Master Oscillator (MO) reference frequency of 1.3 GHz by multiplying it by 8 in a Hittite frequency multiplier to make 10.4 GHz. The frequency multiplier also provides some amplification to the LO signal through an active component. The output from the Minicircuits mixers is then low-pass filtered with a cut-off frequency of 30 MHz. This serves 2 purposes: it removes some high frequency noise and it broadens the signal so that it is easier to sample with an ADC. The low-pass filtered signal is amplified in order to make it match the $\pm$1 V range of the Struck ADC which is clocked by 108 MHz delivered from the Master Oscillator.

So-far, this description has avoided mention of the various phase shifters shown in the diagram. There is a motorized trombone phase shifter on the lower phase measurement arm that must move to account for changes in the *difference* of the arrival-times of the pulses at the mixers. There is an electrical phase shifter called a vector modulator (labled VM) that shifts the phase of the 1.3 GHz signal from the MO in order to account for any changes in the *sum* of the arrival-times of the pulses from the pickup. Lastly, there is a vector modulator phase shifter that can shift the 108 MHz in order to adjust the sampling time of the ADC.

The weakness of the scheme shown in Fig. 7.1.1 for measuring the *sum* of the arrival-times of the pulses is that all of the noise on the LO will be part of the measurement. The phase of the MO signal drifts on the 30 meter cable with a temperature coefficient of 3 ps/deg C (Table 8.1.1). While the drifts of the cable could be compensated with a reflectrometry scheme to within 100 fs [46], noise picked up on the long cable would limit the measurement resolution to more than 20 fs (rms). This was measured by comparing the resolution of the phase measurement for a short cable to that of a long cable.

The strength of the scheme shown in Fig. 7.1.1 for measuring the *difference* between the arrival-times of the incoming pickup pulses comes from the fact that the LO is common to both arrival-time measurements; any LO noise measured by one arm of the setup will also be measured by the other arm. If one subtracts the one arm's measurement from the other arm's, as in a beam position measurement, the LO phase noise will cancel out. The measurement of the difference between the arrival-times of two pulses will only suffer from inaccuracies if the filters, cables, and mixers in the two different phase measurement arms drift relative to one another. These thermal drift effects are counteracted by active temperature stabilization within the chassis, a system described in a later sub-section, but first, a theoretical investigation of the above circuit will be presented.

If the inputs of a mixer are sinusoidal voltage waves, v, with amplitude, A, frequency, f, and phase, θ,

$$v_i(t) = A_i \sin(2\pi f_i t + \theta_i),$$

the output of a mixer is the product of these signals. According to the trigonometric identity,

$$\sin(A)\cdot\sin(B) = \frac{1}{2}\left[\cos(A-B) - \cos(A+B)\right]$$

we can write the output of a mixer as

$$v_1(t)v_2(t) = \frac{KA_1A_2}{2}\left[\cos(2\pi t(f_1 - f_2) + \theta_1 - \theta_2) - \cos(2\pi t(f_1 + f_2) + \theta_1 + \theta_2)\right]$$

Where K is a constant of the mixer. One can see that the output of a mixer is a superposition of the sum and difference of the input frequencies. The sum and difference of the phases will also govern the output of the mixer.

If one estimates the output of the filters shown in Fig. 7.1.1 as a sine wave with a frequency equal to the LO, the difference frequency will be zero, a DC signal, and the sum frequency will be 20.8 GHz. If one low-pass filters the output of the mixer in order to remove the sum frequency, changes in the phase relationship between the two mixer input signals will produce a change in the DC voltage measured at the output of the mixer. When the signals have a fixed phase difference of $\pm\pi$, changes in the DC output of the mixer are proportional to changes in the amplitude of either input signal. When they have a phase relationship of $\pm\pi/2$, the sensitivity of the mixer output to changes in the amplitude of the input signals will be minimized. This is the ideal phase for measuring differences between the phases of the input signals.

If one approximates the output of the filters as a pulse with the shape of a single cycle of a sinusoid, the output of the mixer will not be a DC voltage, but a pulse. This concept, along with the phase relationships for measurements of a signal's phase or amplitude are illustrated in the drawing below in Fig. 7.1.2.

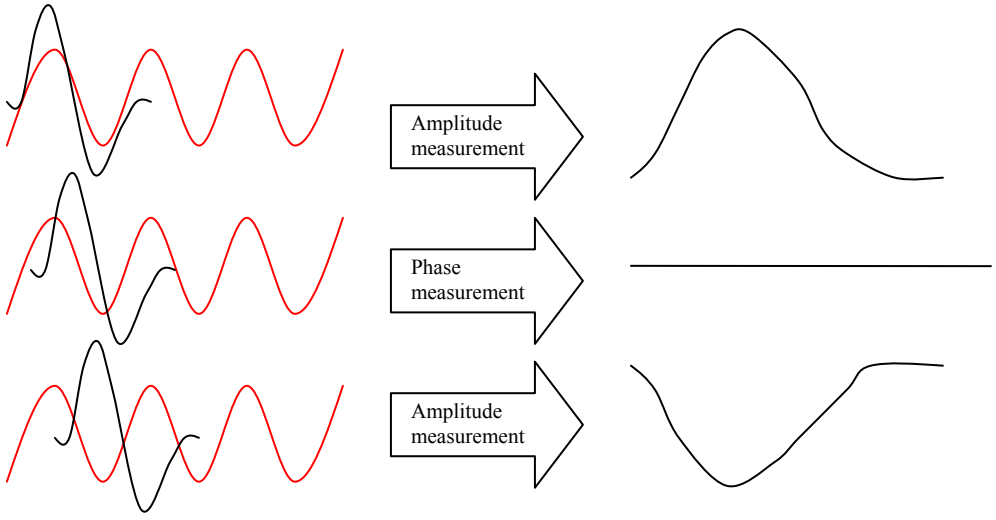

Figure 7.1.2 The input and output of a mixer and how an appropriate phase relationship facilitates the measurement of the phase of the input signal.

In Fig 7.1.2, the signal from the output of the filter is depicted as a single cycle of a sinus. The filter output is not just one frequency; it has a bandwidth determined by the design of the filter. The more poles and the more bandwidth that are used in the filter, the flatter and broader the group delay of the pass band becomes and the wider the mixer output pulse becomes. The wider the output pulse becomes, the easier it is to sample it with an ADC. Given a filter with 4 poles and a bandwidth of 150 MHz at the 3dB attenuation point, the output of the mixer has a pulsed characteristic shown in Fig. 7.1.3 for five different phase relationships centered about the phase for which the mixer output pulse amplitude is a minimum. Each line plotted represents a 1 degree change from the adjacent line. Fig 7.1.3 was produced with a simulation of the circuit shown in Fig 7.1.1 using ORCAD software and it is an accurate representation of what is visible on an oscilloscope when the mixer output is measured for different LO phases.

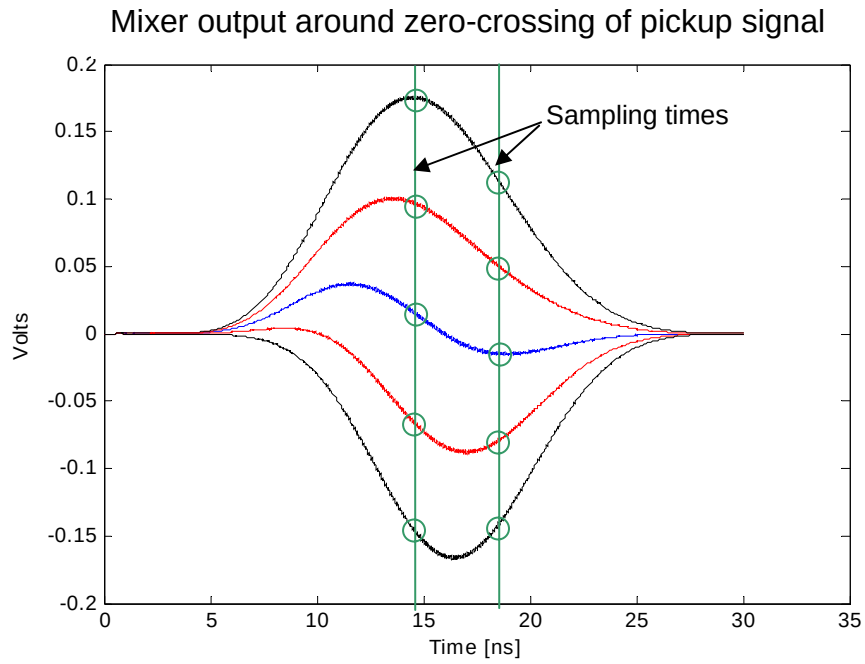

Figure 7.1.3 Simulated behavior of the mixer output around the phase for which the mixer output is minimized. There is an asymmetry in the behavior, but regardless of which sampling time you choose, the amplitude changes in a sinusoidal fashion as a function of the LO phase.

The asymmetrical behavior seen in Fig. 7.1.3 arises due to the bandwidth of the filter and it is seen on an oscilloscope measurement as well as in simulation. The higher frequency components beat with the lower frequency components to cause the asymmetrical signal shape for LO phases that minimize the amplitude of the signal. This is easier to visualize when the pulsed situation from Fig. 7.1.3 is simulated in continuous wave format (Fig. 7.1.4). In Fig. 7.1.4, two different frequencies with the same phases but different amplitudes are added together to make a single RF signal. The RF signal is mixed with the LO. As the phase of the LO is changed, a beating behavior becomes apparent between the lower and higher frequency components that make up the RF signal. The relative

strengths of the peaks seen in Fig. 7.1.4 change as the LO phase is changed. This is the same effect that is seen in the asymmetric behavior of the mixer output in the pulsed signal simulations of 7.1.3.

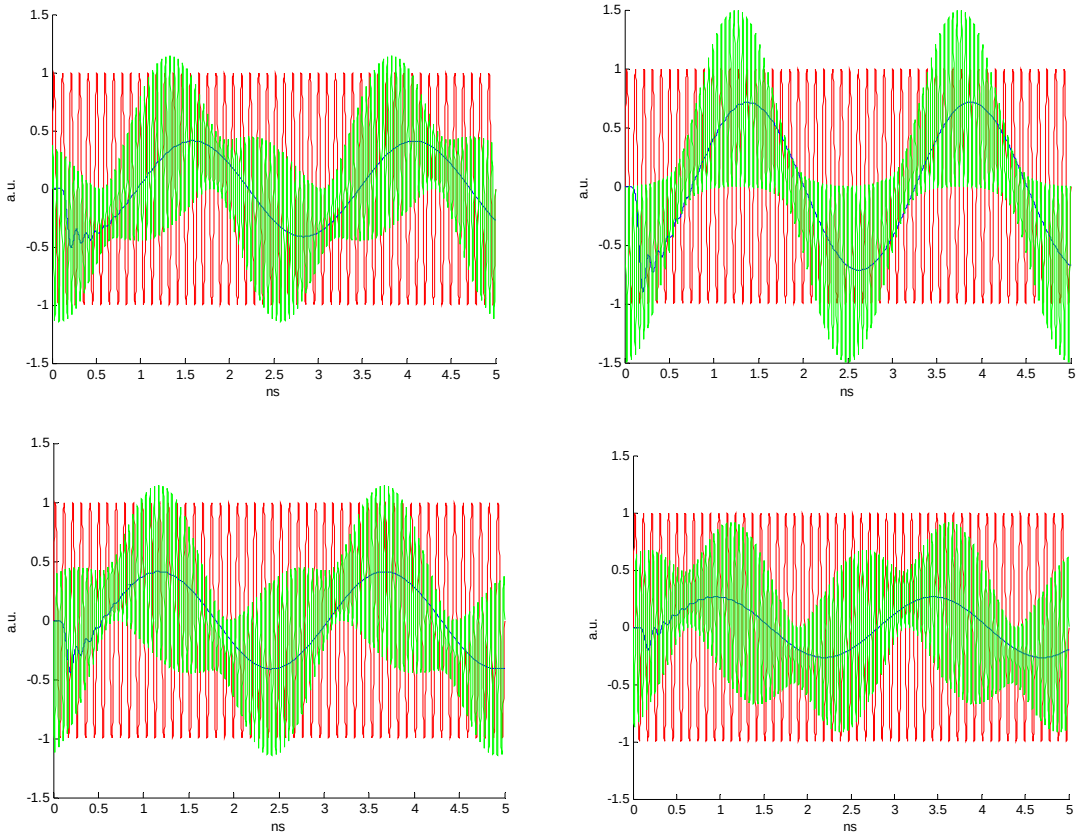

Figure 7.1.4 RF signal composed by sum of two different frequencies mixed is with the LO (red); the output of this mixer is shown in green. As the phase of the LO is changed by a few degrees about the zero crossing of the RF signal, shape of the output of the mixer changes in a manner that is similar to that which is observed in the pulsed case from Fig. 7.1.3. The blue line is the result of low-pass filtering the green signal.

If the amplitude of one of the signals shown in either Fig. 7.1.3 or 7.1.4 is sampled at a given time for an array of LO phases, the amplitude measured at this sample point will change in a sinusoidal pattern as is shown in Fig. 7.1.5. Even though there is an asymmetrical pulse emerging from the output of the mixer, the dependency of the output pulse amplitude on the phase of the LO is still sinusoidal, regardless of the sample point. This is not, however, what was measured (Fig. 7.1.6). There is a bump in the measured mixer output signal that changes its position within the signal when the position of the beam changes. This is not a desired effect because the measurement of the beam's position is given by the difference in the phases of the two signals shown in the plot (red and blue). The difference in the phases is measured by adjusting the various delay lines until the ADC sampling time falls about the zero-crossing on the falling or rising slopes of both signals.

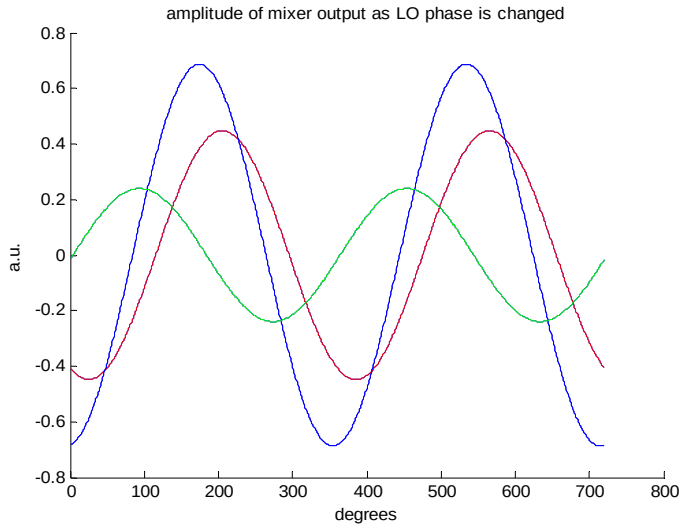

Figure 7.1.5 Three different sample points of the mixer output when the phase of the LO is changed in simulation. The RF signal is composed of two signals with different frequencies but the same phases.

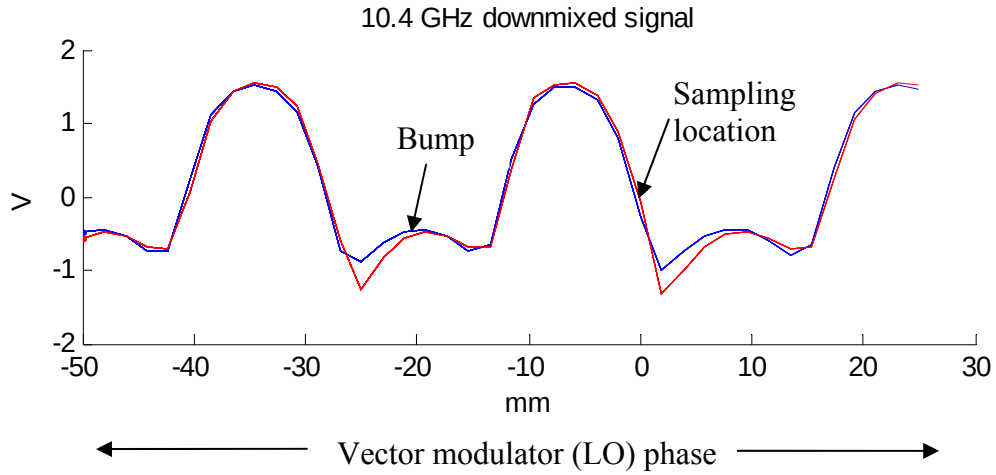

Figure 7.1.6 Distorted sinusoidal pattern that is measured at the output of the mixers when the phase of the LO is changed. The location of the bump in the distorted signal changes when the beam position is changed. The red signal is from the left side of the pickup and the blue is from right side of the pickup. A change in the difference of these signals' phases is proportional to a change in the beam position. The phase of one of these signals is most accurately measured by sampling a point on a steeply rising or falling edge.

The bump could present a problem if it pops up on the slope which has been selected for the signal phase measurement. In principle, the bump can always be avoided by selecting the sampling point that is unaffected by it, but this adds undesirable complexity to the

algorithm required to select the phase measurement sample-point. If the bump cannot be easily removed, its origin should, at least, not remain a mystery.

Perhaps it is naïve to assume that the only pulse of significance to this measurement is that of the initial beam transient. There are, after all, rather significant reflections in the pickup. If the time elapsed between the initial pulse and reflected pulse is shorter than the duration of the mixer output pulse, then the mixer output pulse amplitude is composed by a combination of both the initial and reflected pulses. Since the reflection in the pickup occurs at ~300 ps after the initial beam transient pulse, it would make sense that the phase dependencies of both the incident and reflected pickup pulses are contained within the 100 ns (FWHM) pulse that emerges from the mixer. This, however, could not create the bump seen in the measured mixer output signal. In a simulation where the RF signal is composed by two signals with different phases, when the LO phase is scanned, the mixer output behavior is still sinusoidal, without any distortions.

The only mechanism that could create this bump phenomenon is if there are two LO signals with different rates of phase change. One LO, for example, has a phase that changes at twice the rate of the other LO. This could happen if a harmonic of the RF and LO is in the mix. A continuous wave simulation of this concept is shown below in Fig. 7.1.7. In the simulation, the RF signal is composed by two waves that have the same frequencies, but one signal has twice the phase of the other. When the amplitude of the mixer output is sampled at a single point in time, a distorted sinusoidal pattern will be measured at this sample point as a function of the phase of the LO. In Fig. 7.1.7, the LO phase dependence of the mixer output is plotted for three different sample times.

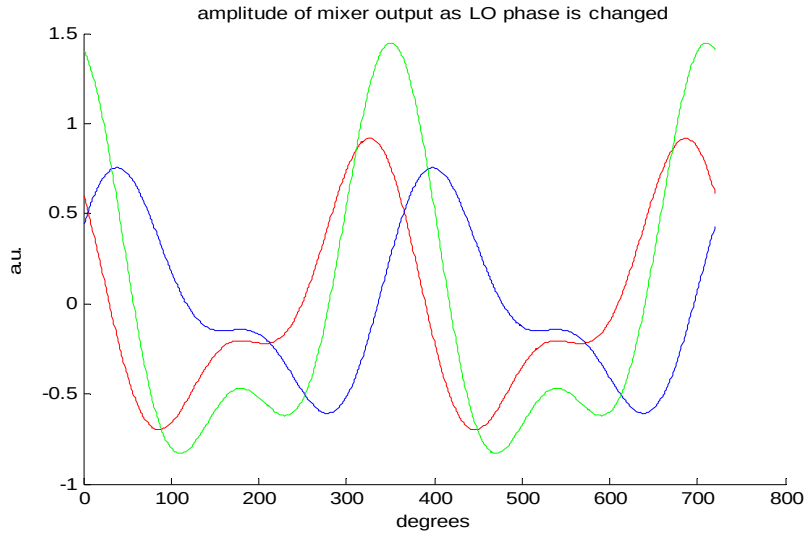

Figure 7.1.7 Mixer output when the RF signal is composed of two signals with the same frequencies but where one signal has twice the phase of the other. The mixer output amplitude for three different sample points are plotted.

Harmonics do emerge from both the LO frequency multiplier and from the band-pass filter output. The frequency multiplier data-sheet warns of harmonic content and band-pass filters typically open at harmonics of the pass-band. When the beam position is

changed, the phase of the harmonic coming from the band-pass filter will change at twice the rate of the fundamental. This is consistent with the bump's beam position dependence.

The harmonic could be removed with a low-pass filter with a cut-off starting around 15 GHz. While such filters exist, they are fairly specialized and one had not been purchased in time for the experiments presented in this thesis. Consequently, in subsequent sections the distortion is simply ignored and a sampling point is found that does not reside on or near the signal distortions. Frequent calibrations, measurements of the signal slope about the zero crossing, also negate any errors this effect could occasion.

Bumps and harmonics aside, the question of the accuracy of the calibration should be approached from a theoretical standpoint. Because the bandwidths of the filters used in the scheme are up to 400 MHz wide, reflections within the pickup are mixed with the LO along with the initial beam transient. The down-mixing technique then measures the sum of the reflections' phases together with the phase of the initial beam transient pulse. If the amplitudes of the reflections are very small compared to the amplitude of the initial beam transient pulse, then this will have no effect at all, but since the amplitude of the first reflection in the pickup is about a third of the initial beam transient pulse, it cannot be ignored. When the beam position changes, the beam transient pulse on one side of the pickup will arrive earlier, but the reflection of that pulse will arrive later. The change of the sum of the phases of the initial and reflected pulses will, therefore, be smaller than the change of the phase of the initial pulse alone. This means that when the monitor is calibrated by scanning the LO with the vector modulator, the calibration factor will not be accurate for measurements of the beam position or arrival-time. For the measurement of the beam position to be accurate, the monitor must be calibrated by scanning the position of the beam, either by changing the accelerating gradient setpoint or the chicane dipole current.

In the measurements presented in the following sections, the calibration of the monitor was done by scanning the LO with the vector modulator phase and not by scanning the beam position with the accelerating gradient. For most measurements with this calibration method, the errors due to the reflected pulses were not apparent because the beam was close to the horizontal center of the pickup and even though the calibrations of each side were incorrect, they were incorrect in the exact same proportions and so the beam position measurement was still correct. When the machine configuration changed, causing the beam to sit at the far end of the pickup and the reflections to move in opposite directions temporally, the calibration done by scanning the LO was two times too large on one end of the pickup and two times too small on the other end of the pickup. As described in the chapter on pickup design, if the pickup were made lightweight and hollow and the ceramic support rings were removed from the assembly, the reflections and their attendant problems would vanish.

7.2 RF Front-end Execution

In Fig. 7.2.1, there is a drawing of the RF circuits contained in the RF front-end chassis. The high resolution 10.4 GHz down-mixing circuit described in Fig. 7.1.1 is drawn alongside a similar 1.3 GHz down-mixing circuit that delivers a lower resolution measurement but does not require any mechanical phase shifters.

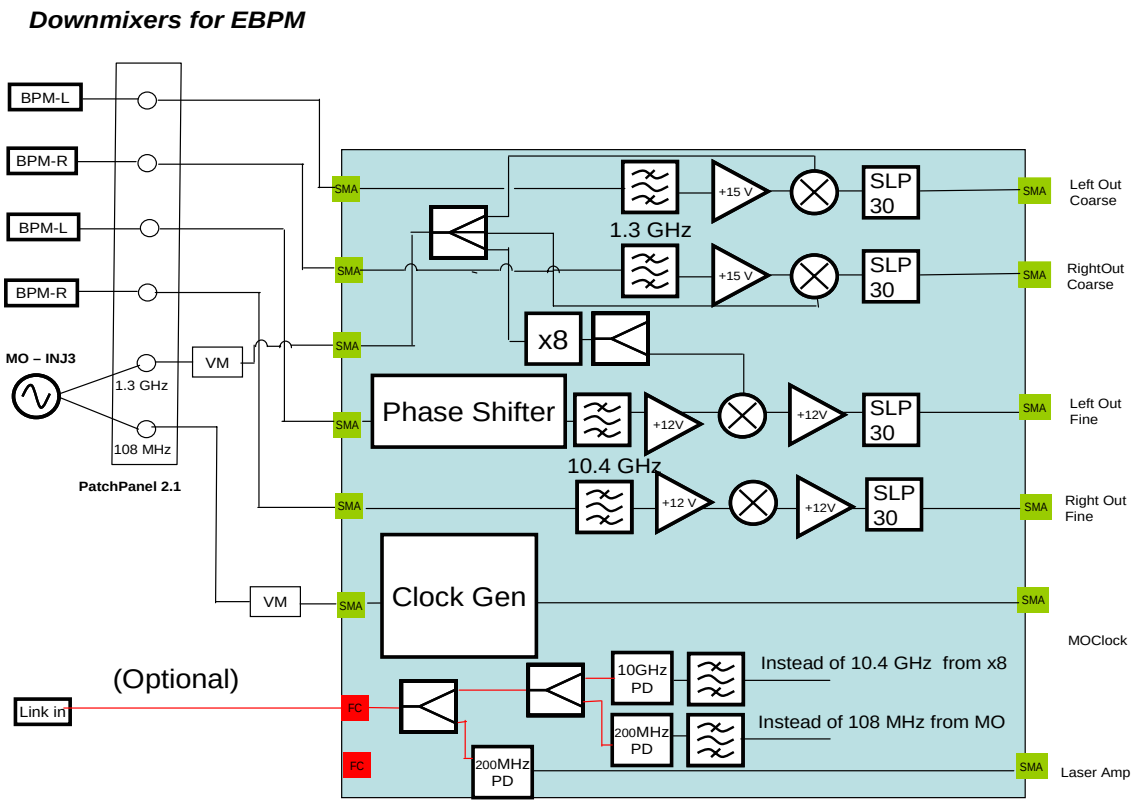

Figure 7.2.1 Chicane BPM RF front-end schematic. The reference frequency of 1.3 GHz is used as an LO for a mixing scheme with signals from the BPM pickup. The reference frequency of 108 MHz is used as the clock for an ADC that samples the outputs of this schematic. An optional scheme for generating the reference frequencies from a master-laser signal is sketched in the bottom -left corner.

The figure above contains a symbol labeled MO-INJ3. This refers to the signal from the master RF oscillator in the injector racks. The MO signals were delivered over a few-meters-long cable to the patch panel, labeled “PatchPanel2.1”, on the left side of the drawing. Signals from the BPM pickup were delivered over 30 meter long cables to the same patch panel. From the patch panel, the signals are sent over 2 meter long RF cables to the chassis depicted by the blue square. Within the chassis, the 108 MHz signal from the MO is sent through a clock generation circuit. The clock generation circuit turns the sinusoidal 108 MHz signal into a square wave that is appropriate for the clock input of the ADC that will be used to sample the mixer outputs. The 1.3 GHz signal from the MO enters the chassis and is split into three signals: two of the signals are mixed with 1.3 GHz waves coming from a filtered pickup signal and one of the signals is multiplied by 8 in a Hittite frequency multiplier. After the frequency multiplier, the 10.4 GHz signals are split and then mixed with 10.4 GHz waves coming from a filtered pickup signal. In one of the two arms of the 10.4 GHz down-mixing circuit, there is a box labeled “phase shifter”. This mechanical, motorized trombone phase shifter from the company ATM is used to synchronize the arrival-times of the signals coming from the pickup. If the phase shifter is adjusted appropriately, the signals from the right and left sides of the pickup

will arrive at the mixers at exactly the same time. The other type of phase shifter is the vector modulator. There are two of them and they are labeled "VM" in the drawing. One shifts the phase of the 108 MHz clock and is used to adjust the sampling time of the ADC and the other will shift the common LO for all of the different phase measurement circuits. When the beam arrival-time is changed, the phase of the 1.3 GHz signal coming from the MO will need to be adjusted in order to maintain the optimal LO phase.

An optional feature that was not built into the tested design is the generation of the LO signals from the optical reference produced by the master laser oscillator. This is shown in the lower portion of Fig. 7.2.1 and again in more detail in Fig. 7.2.2. This scheme would be advantageous if the front-end were relied upon to generate beam arrival-time measurements as well as beam position measurements. In a beam position measurement, the drift of the LO is irrelevant because the position depends on the difference between the arrival-times of the signals. The beam arrival-time, however, is measured by the average of the arrival-times of the signals and hence noise or drifts of the incoming reference signal would affect the stability of the measurement. Because the phase of the optical reference delivered on a length stabilized fiber can be made more stable than that of the 1.3 GHz RF reference delivered on an un-stabilized RF cable, the use of an optical reference could remove the impact of drifts from the delivery of the LO signal and the use of an optical delay line in lieu of the noisy vector modulator.

One method is to use a 10 GHz photo-detector operated in saturation and filter out the desired frequency component from the frequency comb. 10.4 GHz and 1.3 GHz would be filtered out, for example, and then a non-saturated photo-detector would be used in order to actively stabilize the amplitude of the laser signal with DSP feedback on the laser diode driver responsible for the amplification of the laser signal (Fig. 7.2.2). This last step is necessary because the phase of the signal produced by a photodetector will change by 20 fs whenever the laser amplitude changes by 0.1% [47]. Thermal stability of the photodetectors is also important due to the 340fs/degC thermal drift coefficient [48].

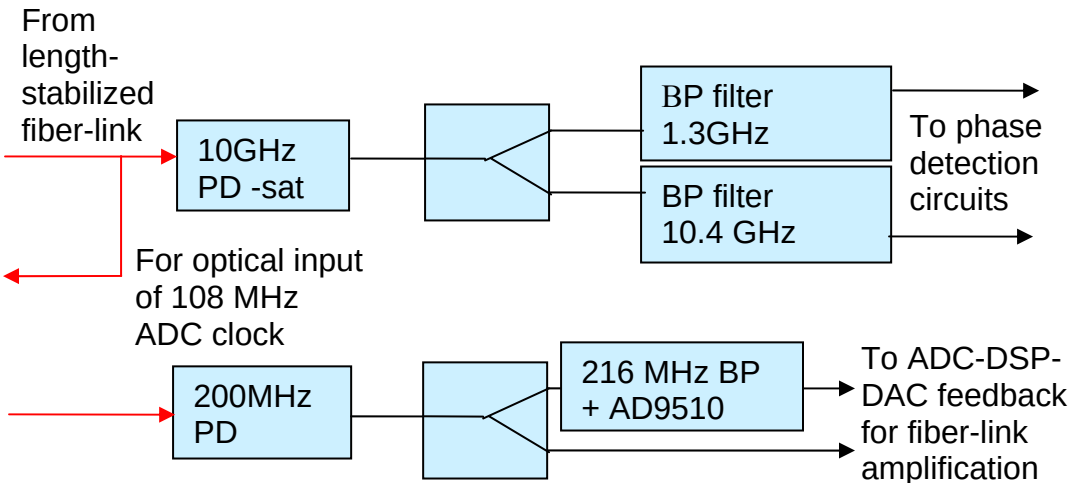

Figure 7.2.2 Scheme for delivering a stable reference signal to the phase-detection circuits using an optical signal from a length-stabilized fiber link.

While this option was not built and tested in this thesis work, these sorts of schemes are under active study for other applications [49].

An alternative way to improve the drift stability of the RF arrival-time measurement is to use RF cable reflectrometry. With the active phase stabilization method from [46], a signal of a few GHz is sent from a Master Oscillator (MO) over a coaxial cable to an end location at which part of the signal is reflected and sent back to the source. The phase of the signal being generated and the phase of the signal returning are then be compared by RF phase detection. The length of the cable is then adjusted until the returning phase matches the sent phase, thereby removing the effects of cable drifts. Sub-100 fs drift performance was achieved with this method. Although the phase detection of the returning pulse can be made with ~10 fs accuracy, the reflection of the pulse at the end of the cable is problematic due to the temperature dependence of the mismatch that produces the reflection. Any drift of the mismatch will be "corrected" by the feedback loop, but this does not accurately represent the drift of the cable and it, therefore, adds an error to the cable length stabilization. Another error comes from the drift of the finite directivity of the coupler. Individual couplers have to be measured to determine their drift sensitivity, but typical high power RF couplers drift by -40dB.

This RF reflectrometry method was not employed for the chicane BPM beam RF front-end because beam arrival-time measurements relative to a pulsed optical reference had already demonstrated ~10 times better accuracy and precision than beam arrival-time measurements relative to an RF reference could ever hope to achieve. What was built and tested is shown below in Fig. 7.2.4. The emphasis will be placed on the robustness of the beam position measurements and details about beam arrival-time measurements are reserved for the following chapter.

The RF front-end shown in Fig. 7.2.4 was constructed in two layers within a 3 rack-units high chassis. The lower-layer contains circuits that are not sensitive to temperature changes and the upper-layer contains circuits which are sensitive to temperature changes. The temperature of the upper-layer was actively stabilized with two Wavelength-HTC temperature controllers which determined the heating and cooling action of two Peltier elements. Because the circuit element that was most sensitive to temperature changes was the band-pass filter, and because the stability of the high-resolution measurement was a key goal of this setup, the Peltier elements were installed as close as possible to the 10.4 GHz band-pass filters. Although the Peltier elements could have been directly mounted on the filters, in order to prevent thermal gradients over the RF circuit, each Peltier element was in direct thermal contact with a 10x300x50 mm metal bar which was mounted to the aluminum plate and not in direct thermal contact with the filter itself or with the 4 mm thick aluminum plate to which the circuit elements were fixed. This intermediate medium distributed the heating and cooling action of the Peltier over a larger surface, reducing the possibility of noise in the temperature controller making its way into the RF circuit stability.

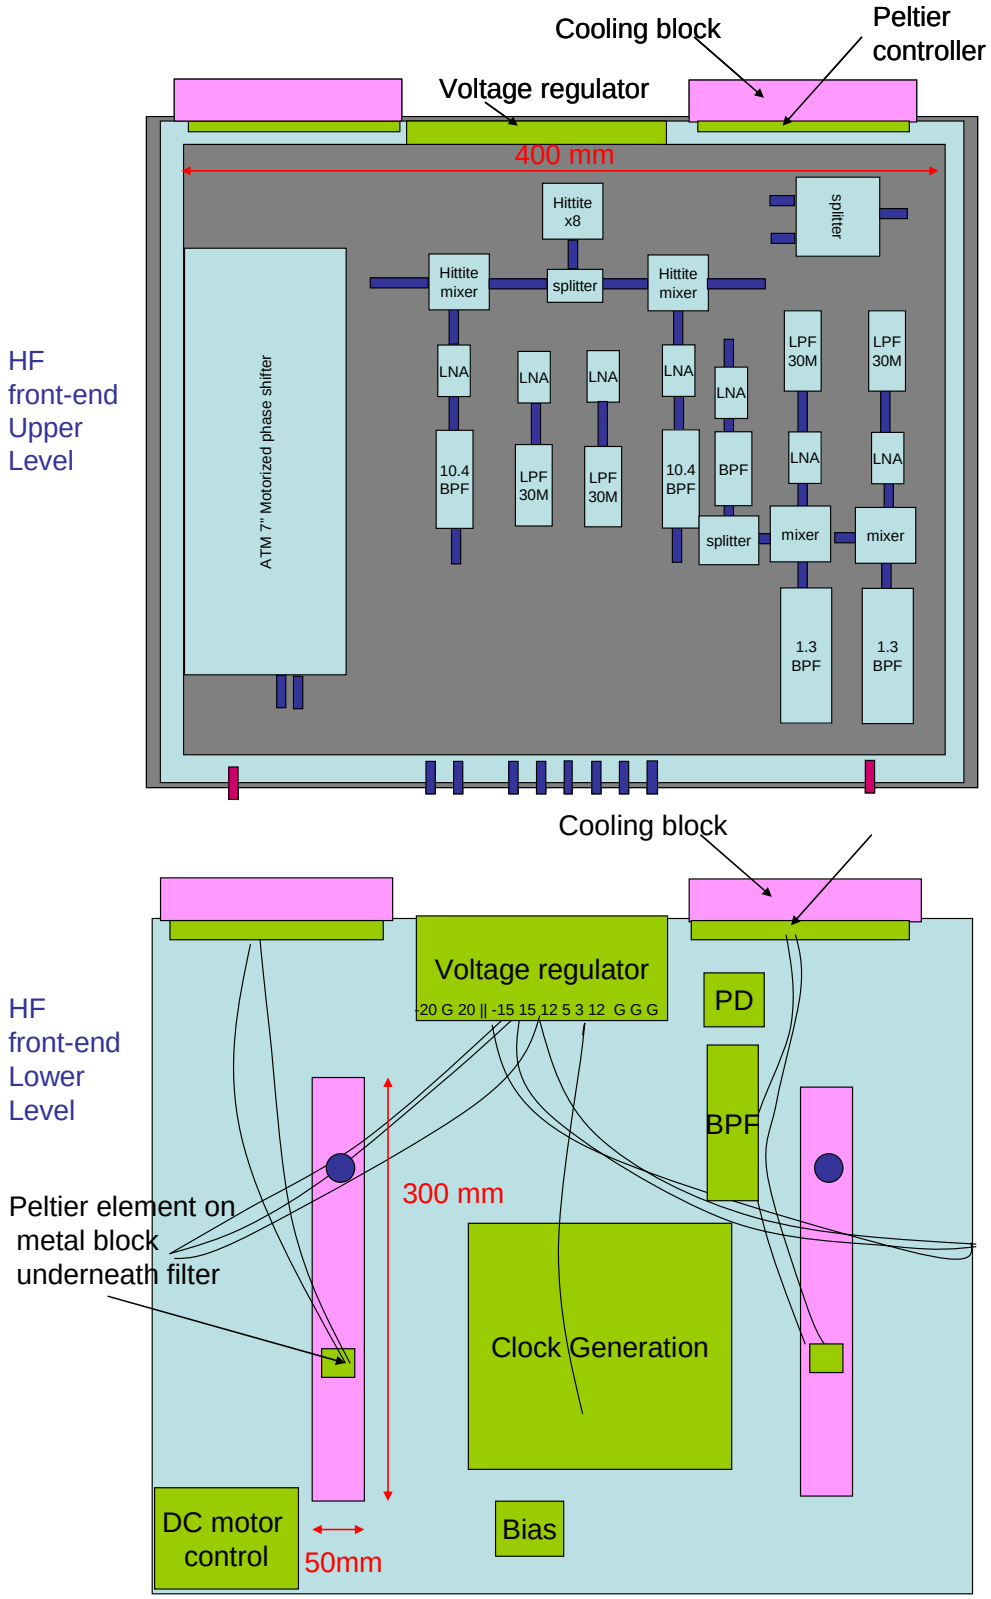

Figure 7.2.4 The upper-level and lower-level of the HF front end chassis.

A Peltier element acts as a heat pump. When an object mounted to the top surface of a Peltier is too cool, the top surface of the Peltier will become warmer and the bottom surface will become cooler. When the object mounted to the top surface of the Peltier is too warm, the reverse is true. When the Peltiers cool the aluminum plate, the heat from the bottom side of the Peltier travels down the metal post mounted to the bottom surface of the Peltier, towards the floor of the chassis. The heat from the post is distributed over the floor of the chassis which is then cooled with an external rack mounted fan that blows air up onto the bottom of the chassis. The two temperature controllers are mounted to the back plane of the 3 RU chassis and their cooling blocks are mounted outside of the chassis. This was a successful arrangement and could stabilize the temperature of the chassis to <0.003 degrees C peak-to-peak with the lid open in an air-conditioned room with a thermal stability of 0.2 degrees C.

If a lid is added to the top of the chassis, the temperature control system is much more unstable and difficult to manage, but it can resist larger, temporary temperature changes, if the controller is properly set-up. When the gain of the controller is too high and the setpoint of the temperature controller is too low, the Peltier will continuously pump heat to the bottom of the chassis, thereby heating up the air within the chassis, warming up the circuit and causing the Peltier to pump even more heat to the bottom of the chassis. If the setpoint is higher than room temperature, the system is more stable even if the controller gain is high. Hence, the strategy for commissioning the temperature control system with a closed lid and a high controller gain was to use a setpoint that was several degrees above room temperature and then slowly, over the course of days, reduce it until it was closer to room temperature. This produced the most efficient operation in the air-conditioned room, but given large, permanent changes in the room temperature of a degree or more, the temperature controller of the chassis could become unstable and need to be re-commissioned.

Peltier elements were used instead of a heating mat because they can cool as well as heat and have a faster reaction time to changes in the control voltage. While it is possible to stabilize the temperature of a chassis by simply making it warmer than everything in the room, this technique was not chosen because if too many warm devices are installed in an air-conditioned room, the temperature stability of the room will be more difficult to manage.

Not everything in the chassis is so sensitive to temperature changes. These less sensitive elements are installed on the lower-level of the chassis. Aside from a voltage regulator and a circuit to control the DC motor for the mechanical phase shifter, the lower level contains a circuit that takes the sinusoidal 108 MHz reference signal and generates a 108 MHz NIM-level square-wave clock signal for the ADC. This was done with an AD9510 clock divider evaluation board. The clock divider generated a TTL-level square wave. The offset voltage of this signal was adjustable with an RF bias-voltage circuit element and a potentiometer. Lowering the bias by a couple of Volts made the TTL signal appropriate for the NIM level clock input of the Struck 100 MHz 14-bit ADC.

The RF front-end chassis was installed in a rack that was adjacent to a patch panel containing signals sent from the master oscillator (MO) over ~5 meter long cables and signals from the pickups in the first bunch compressor sent over ~30 meter long RF cables. The patch panel signals were connected to the chassis with meter-long cables.

The resolution of the 10.4 GHz front-end was evaluated by splitting a signal from the beam pickup and measuring the difference between the arrival-times of the pulses. If the front-end measurement had perfect resolution, the difference measured would be a constant value, regardless of horizontal beam position changes. The difference between the split signals is plotted below in Fig. 7.2.5.

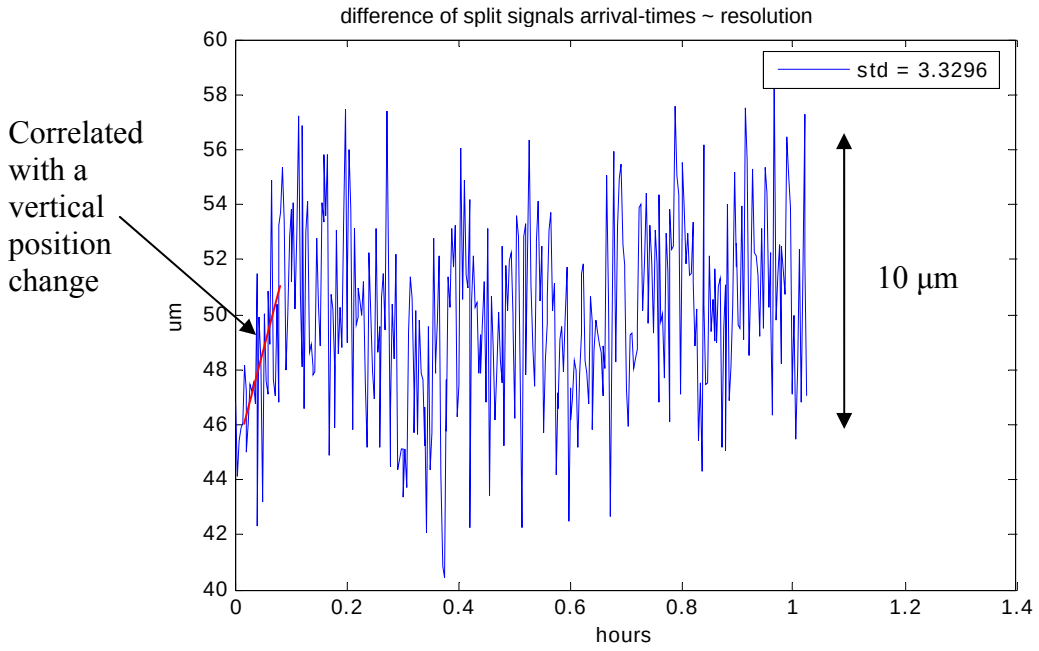

Figure 7.2.5 Resolution of the 10.4 GHz BPM front-end chassis. This was evaluated by splitting a signal from the beam pickup directly in front of the chassis and measuring the difference between the arrival-times of the pulses. If the front-end had perfect resolution, the standard deviation of the difference between the arrival-times of the two identical signals would be equal to zero.

The rms jitter of the difference of the split signals over the course of an hour is 3.3 μm. Based on this, one could say that the RF front-end can achieve 3.3 μm resolution if the cable lengths from the pickup to the chassis are minimized.

At the beginning of the measurement there was a change in the split-signal difference measurement that was correlated with a change in the vertical position of the beam. This was marked with a red line in Fig. 7.2.5. It occurred because the monitor was not periodically re-calibrated during this measurement. Given periodic active re-calibrations using a quick scan of the vector modulator, this sort of correlation is not observed. Alternatively, a calibration constant based on a measurement of the vertical position of the beam using BPMs upstream and downstream of the chicane can be determined and this can be multiplied by the measurements of the beam arrival-time made with the chicane BPM in order to, without active re-calibration, passively remove the influence of vertical position changes on the measurements of the horizontal position of the beam performed with the chicane BPM.

For the measurement shown in Fig. 7.2.5, the pickup signal was split at the end of the 30 meter cable, so one can only claim that the resolution of the beam position measurements produced by this RF front-end would be ~3 μm if it had been installed in the tunnel, with 1-2 meter long cables connecting the pickup to the chassis. If the front-end is installed out of the tunnel, with 30 meter long cables connecting the pickup outputs to the chassis, only ~6 μm resolution can be claimed, based on the fast jitter shown in Fig. 7.2.6.

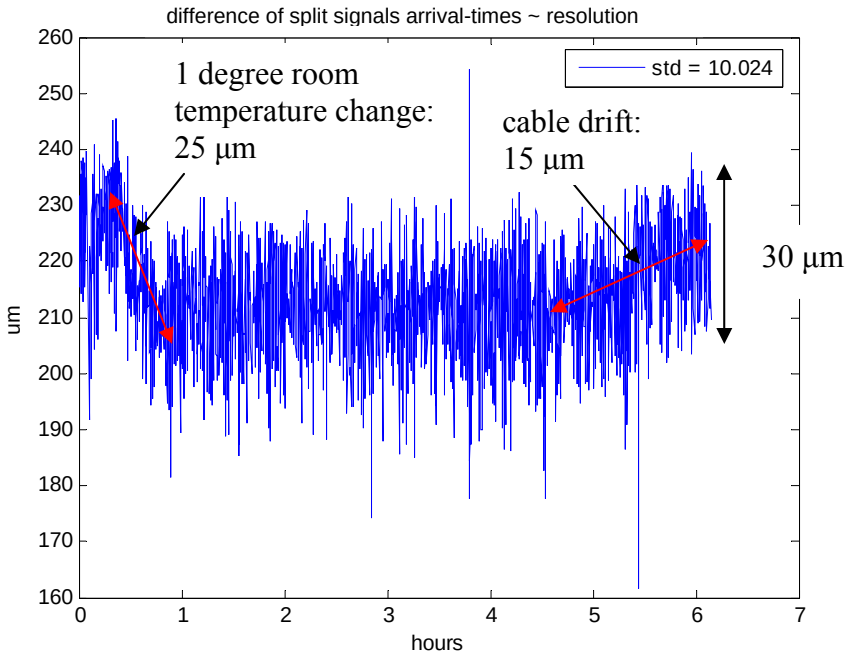

Figure 7.2.6 Resolution of the 10.4 GHz front-end with long cables. This was evaluated by splitting a signal from the beam pickup directly after the pickup, sending the signal over 30 meter long cables to the front-end chassis and measuring the difference between the arrival-times of the pulses. The influence of the temperature dependent drift of the long cables is frequently smaller than the influence of the temperature stability of the front-end chassis. The temperature control of the chassis was not operating during this measurement.

From Fig. 7.2.6, it can be seen that the influence of the temperature dependent drift of the 30 meter long cables can be smaller than the influence of the temperature stability of the front-end chassis. To be fair, the temperature control of the chassis was not working properly during this measurement and a large drift of two cables that are of the same type and run parallel to one another would not be expected, since the majority of the drifts would be common to both cables.

It is not, of course, fair to make claims about the stability of a system without showing data that has been taken over several days; to that end, three days worth of data is plotted below in Fig. 7.2.7.

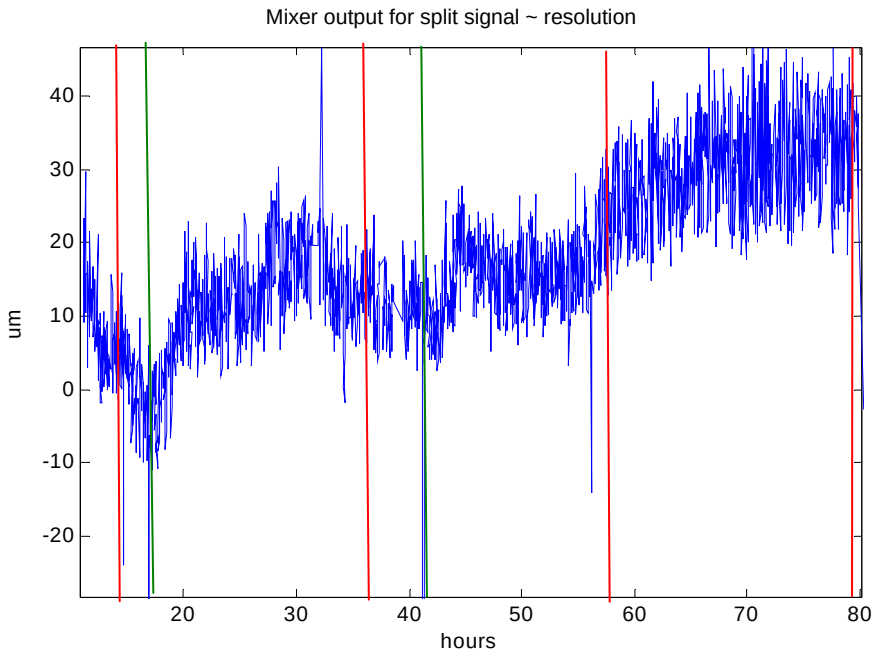

Figure 7.2.7 Three days long measurement of the difference between the split signals. Temperature disturbances (1 degree) in the rack are marked in green and mornings are marked in red. The temperature of the hall in which the cables reside slowly increased by one degree over the course of the three days.

In Fig. 7.2.7, the changes in the temperature of the room in which the chassis is located have a larger impact on the stability of the measurement than the day-night temperature changes that were experienced by the 30 meter long cables. During this measurement, the large temperature changes in the room on the first day were due to the opening and closing of doors. The smaller temperature change observed on the second day is more typical of the changes that the measurement must withstand. While day-night temperature changes are not easy to see, what is more apparent is a slow trend upward as the tunnel temperature warms by one degree over the course of three days. This is a drift that could not be eliminated unless the chassis were moved closer to the pickup, an option that, while not too challenging to implement, was not implemented.

The measurements of the resolution can have no bearing on reality unless they are accompanied by a measurement of the monitor's response to changes of the beam position. In Fig. 7.2.8, the position of the beam was changed by altering the energy of the beam with the first accelerator section amplitude. The mechanical phase shifter position was held constant and the vector modulator kept the measurement centered about the zero-crossing of one signal. As the sampling position of the other signal moves further and further away from the zero-crossing, the measurement of the beam position becomes non-linear and moves out of the range for which the calibration was valid. The beam was off-crest during the scan.

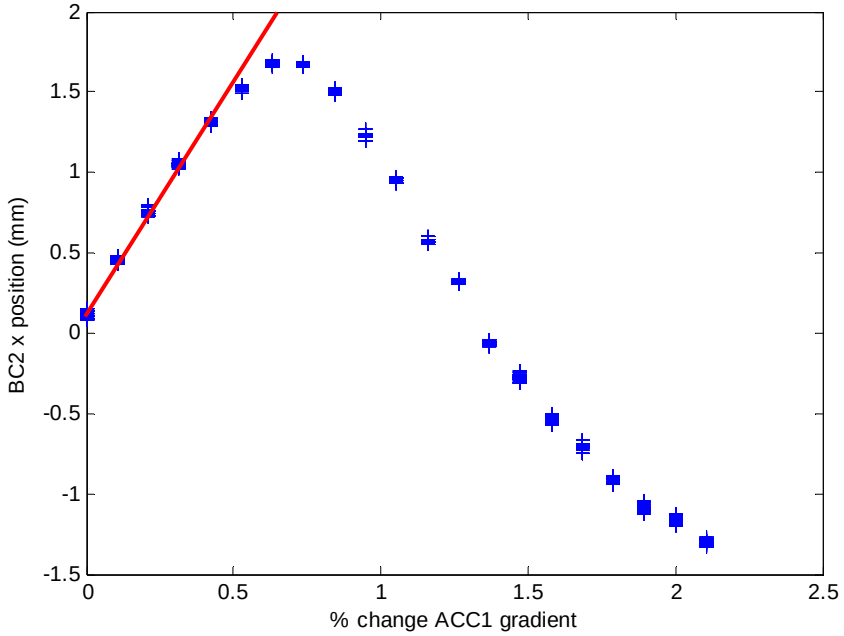

Figure 7.2.8 Scanning the gradient of the first accelerating module and measuring the change in the position of the beam with the chicane BPM. The position measurement is only in-range for the first few data points.

Zooming in on the first 0.5% of the scan shows the 1 mm linear range of the monitor (Fig. 7.2.9). One percent energy change times the R_{16} of the chicane gives 3.5 mm of expected position change. The position change measured by the monitor in the linear range was 3.5 ± 0.1 mm.

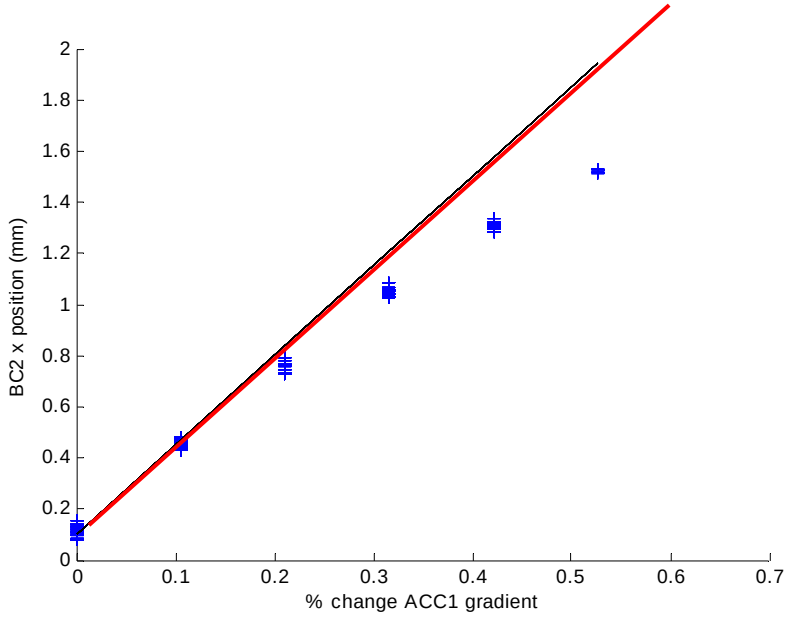

Figure 7.2.9 Beam position change corresponding to a small energy change.

When good agreement with expectations is observed at one measurement location, it is tempting to assume that such agreement will aslso be observed at other locations, but this is not the case for this RF front-end. Due to reflections within the pickup which are one third to one half the amplitude of the initial beam transient, the calibrations produced by scanning the vector modulator are wrong by a factor of ~two. For some beam positions, the errors in the calibrations exactly cancel out, giving a correct measurement of the beam position, but for other beam positions, the measurements done with these calibrations can be wrong by a factor of ~two. If the monitor is calibrated with a beam-based reference instead of the LO reference, these errors do not occur, but it is unfortunate because it is more convenient to calibrate a monitor with a parasitic reference that does not disturb the operation of the machine. Scanning the position of the beam in the chicane in order to calibrate the monitor disturbs the machine operation.

Beam-based calibrations aside, the beam arrival-time change corresponding to a beam energy change can also be measured with the monitor by adding the change of the vector modulator to the average of the pulse arrival-times measured (Fig. 7.2.10). The arrival-time change that would be expected for a 1% energy change is equal to 3.0 ps. The arrival-time change measured for this energy change was 2.8±0.4 ps. The error bars of this measurement are large, because the beam arrival-time jitter is large.

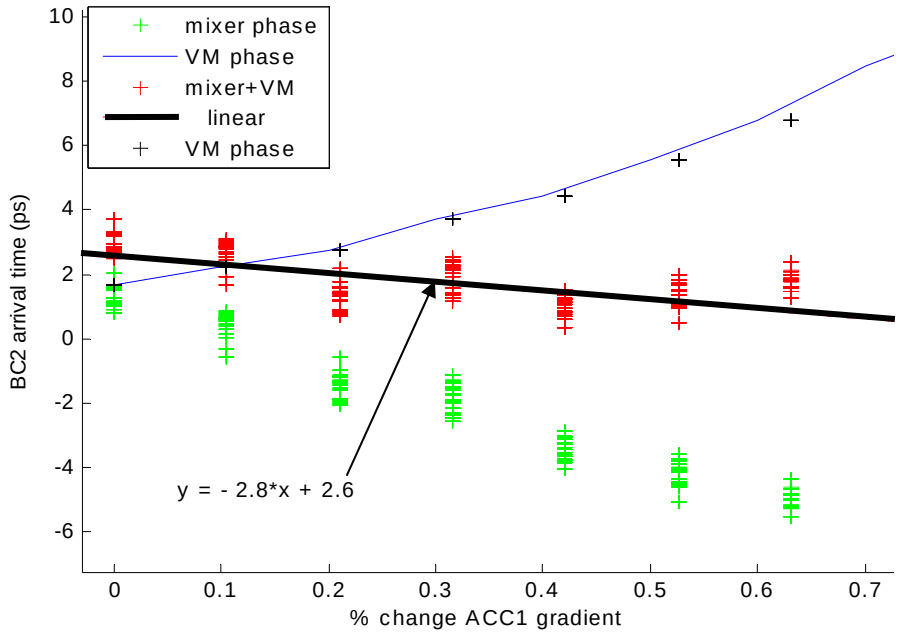

Figure 7.2.10 Beam arrival-time change corresponding to a small energy change.

Using the chicane BPM RF front-end the arrival-time changes from the injector can also be measured (Fig. 7.2.11). Fitting a line to the middle portion of the left-hand plot and to this first half of the right-hand plot gives a net arrival-time change of 2.3±0.2 ps/degree phase change. For one degree of phase shift, the sum of the laser and RF phase shifts should produce 2.125 picoseconds of arrival-time change. This is within the error bars of the measurement.

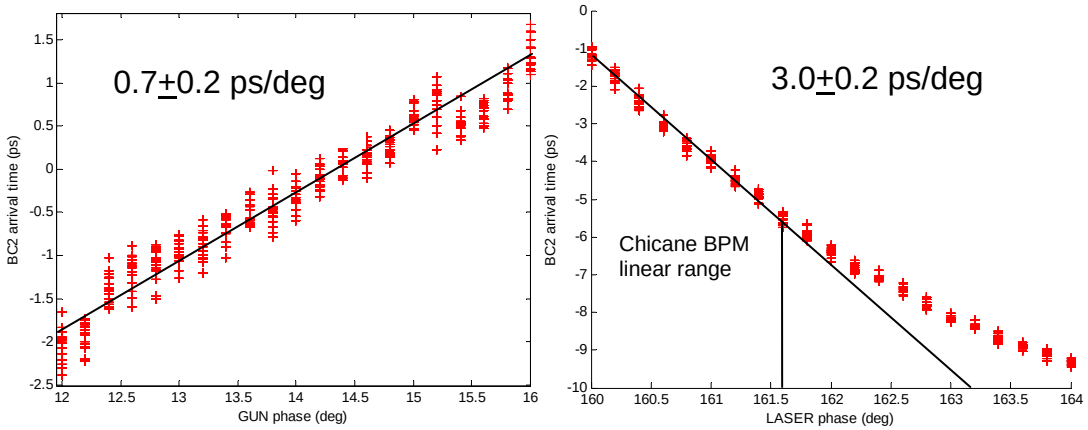

Figure 7.2.11 Measurements of the beam arrival-time changes resulting from scans of the RF (GUN) and laser phases in the photo-injector. For one degree of phase shift, the sum of the laser and RF phase shifts should produce 2.125 picoseconds of arrival-time change. Fitting a line to the middle portion of the left-hand plot and to this first half of the right-hand plot gives a net arrival-time change of 2.3±0.2 ps/degree.

The main concern about the performance of this RF scheme as it is presently built is that the mechanical phase shifter from ATM does not have the position repeatability that one would desire in a measurement that could serve as a reference for measurements of the energy and arrival-time of the beam. The problems associated with the potentiometer-based position readback of the ATM phase shifter are illustrated below in Fig. 7.2.12. The position of the mechanical phase shifter is read by measuring the voltage of a potentiometer mounted to a gear in the motor assembly. The correct scaling of the voltage measured across the potentiometer can be determined by comparing the change in the mixer output produced by changes in the mechanical phase shifter position to changes in the mixer output produced by changes in the phase setting of the vector modulator. If the calibration factor for the potentiometer voltage is correct, then the slope of the signal measured by scanning the phase of the vector modulator will be equal to the slope of the signal measured by scanning the phase with the mechanical phase shifter. Backlash of the motor is seen as a discontinuity between data points measured after a forward movement of the stage (red) compared with the data taken after a backward movement of the stage (blue). The fact that the period of the signal measured with the scan of the mechanical phase shifter is not constant, as it is for scans of the vector modulator phase, indicates that the voltage read-back of the potentiometer is not a reliable indicator of the true position of the stage in the mechanical phase shifter.

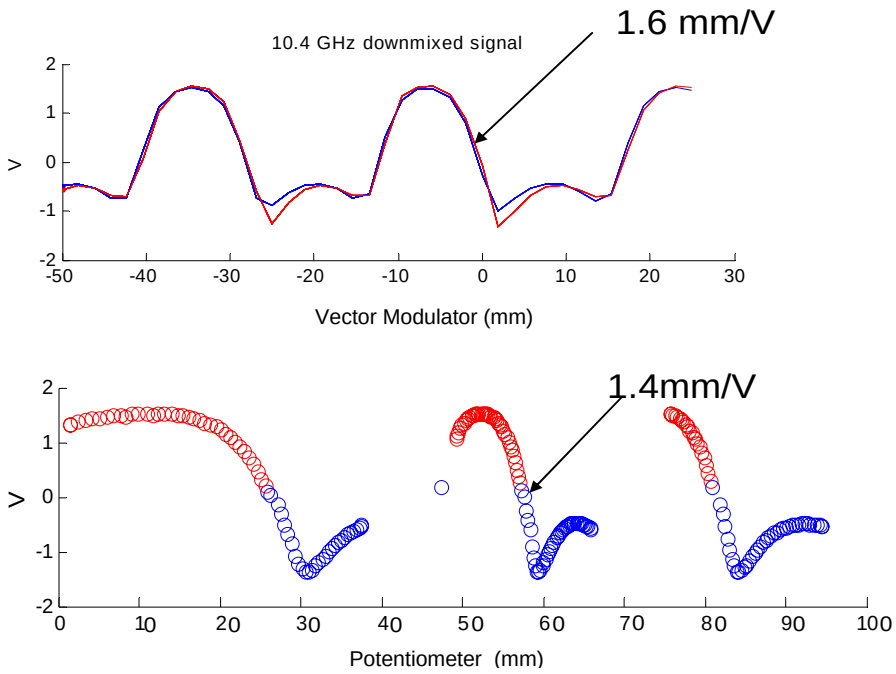

Figure 7.2.12 Fiducializing the mechanical phase shifter potentiometer with the vector modulator. If the calibration factor for the potentiometer voltage is correct, the signals will look exactly the same and the slope of the signal measured by scanning the phase of the vector modulator will be equal to the slope of the signal measured by scanning the phase with the mechanical phase shifter. Backlash of the motor is seen as a discontinuity between data points measured after a forward movement of the stage (red) compared with the data taken after a backward movement of the stage (blue).

Further evidence of the problems with the mechanical phase shifter are seen when a scan of the beam energy is done with a feedback that uses the mechanical phase shifter to center the sampling position about the zero-crossing of the signal (Fig. 7.2.13). The beam position should change in a linear fashion as the beam energy is changed over a small range, but the beam position as a function of beam energy is curved and does not follow the setpoint of the accelerator section. Despite the problems with the mechanical phase shifter, whenever the beam position changes by more than a millimeter, it is necessary to use it to keep the system sampling the zero-crossing of the signal. In the long-term measurements which will be presented in Chapter 10, the mechanical phase-shifter feedback was active.

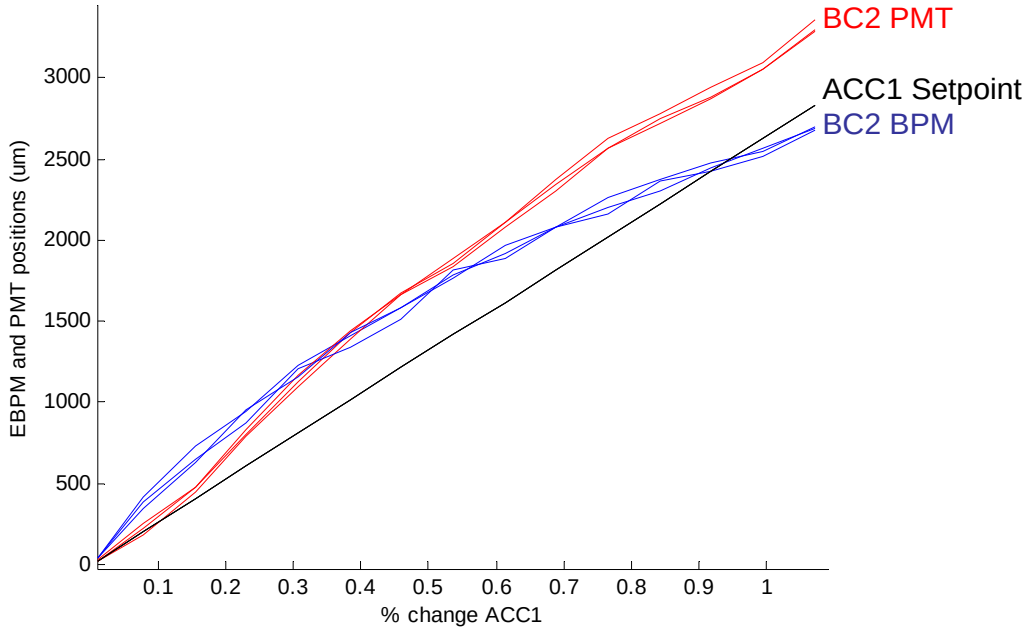

Figure 7.2.13 Curvature of the BC2 BPM measurement results from the problems with the mechanical phase shifter. BC2 PMT is another chicane beam position monitoring system which will be introduced in Chapter 9.

Due to the unreliability of the mechanical phase shifter position measurement, whenever the mechanical phase shifter moves, the accuracy of the measurement is reduced. This is not a concern for the alternative to this scheme, the optical front-end, because a high-resolution linear encoder was mounted to the mechanical delay stage of the optical measurement. If a high-resolution linear encoder were mounted to the mechanical phase shifter of the RF front-end, the concerns about the accuracy of the measurement could be removed. Alternatively, a ~12 GHz vector modulator that is in development through a collaboration between DESY and PSI might be available in the future. An alternative that eliminates the need for any phase shifter at all is to down-mix, not to base-band, but to some intermediate frequency. The phase of the intermediate frequency could be measured through digital down conversion. This would be recommended if the front-end is developed in conjunction with a larger project that also demands low ADC clock-jitter, but it is too complicated to implement as a novelty system.

The main concern about the long-term supportability of this down-mixing-to-baseband scheme is its uniqueness. While the RF front-end is cheaper and easier to set-up than the optical front-end, it is the only front-end of its kind at DESY. Since the optical front-end uses a format that is repeated throughout FLASH in the form of the beam-arrival monitor, it has a higher likelihood of being supported long-term. If it were not for the robustness and dynamic range afforded by the lower resolution 1.3 GHz phase measurement contained together with the higher resolution 10.4 GHz phase measurement in the RF front-end chassis, the chicane BPM RF front-end would probably not be supported long-term. Because it has no moving parts, the 1.3 GHz phase measurement is more stable, reliable and easy to commission than any of the other, higher resolution measurements. The same cannot be said of the 10.4 GHz front-end. While the 10.4 GHz front-end is more robust than the optical front-end, it essentially delivers the same information as the optical front-end and, like the optical front-end, it requires the precise adjustment of a mechanical stage in order to deliver any measurement at all. The 1.3 GHz front-end, in contrast, must be set up once and re-calibrated periodically, but if the beam moves by several centimeters, it will still deliver a measurement and the operator would not have to wait for a minute for the high-resolution measurement to scan a mechanical stage in order to find the correct sampling position. This can be seen by comparing the dynamic range of the 1.3 GHz phase measurement to that of the 10.4 GHz phase measurement in Fig. 7.2.14. There, you see less than 5 mm of dynamic range in the fine measurement before a mechanical phase shifter must be adjusted. The dynamic range of the coarse measurement is 80 mm with no need for mechanical elements.

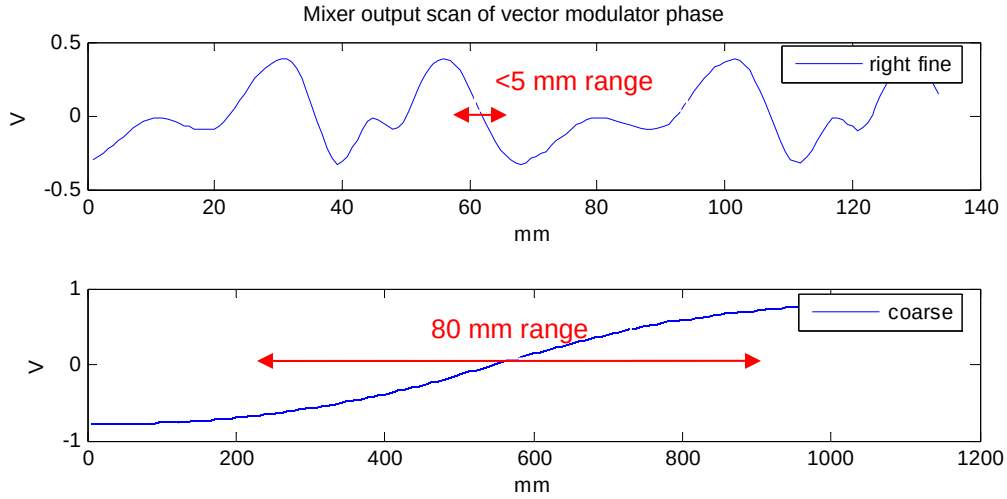

Figure 7.2.14 Comparison of mixer outputs for 10.4 GHz phase measurement (top) and 1.3 GHz phase measurement (bottom). The 1.3 GHz measurement has a much larger dynamic range than the 10.4 GHz measurement.

A 1.3 GHz front-end measurement of the beam position change resulting from a change in beam energy is shown below in Fig. 7.2.15.

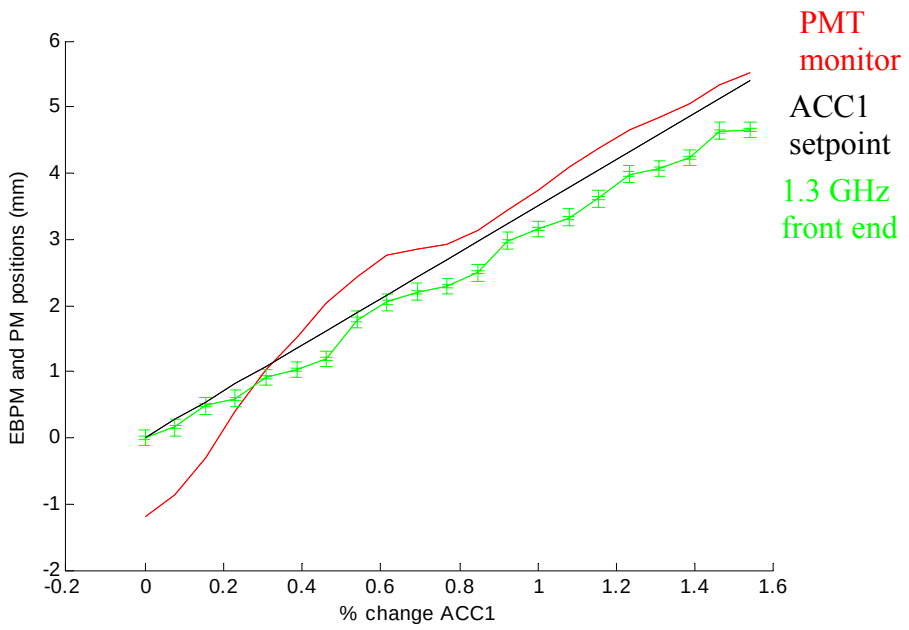

Figure 7.2.15 1.3 GHz front end beam position measurement as a function of beam energy.

The 1.3 GHz front-end can also measure changes in the beam arrival-time (Fig. 7.1.16).

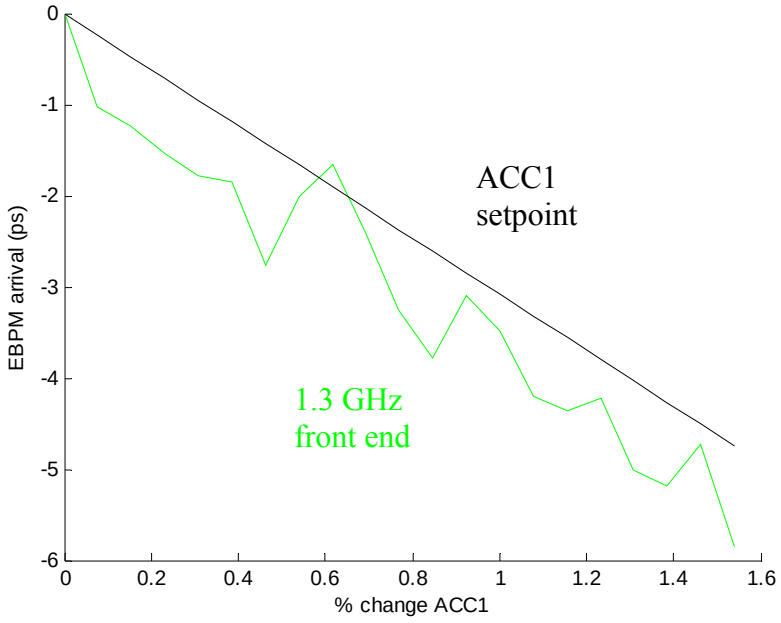

Figure 7.2.16 1.3 GHz front end beam arrival-time measurement as a function of beam energy. Arrival-time jitter is ~400 fs (rms) during this measurement.

Instead of delivering the pickup signals over 30 meter long cables to the RF front-end outside of the tunnel, if the RF front-end is installed in the tunnel, the drifts of the cables will no longer have an impact on the resolution of the measurement and the amplitude of the signals from the pickups will be less attenuated, possibly enabling the

removal of some amplifiers from the front-end circuits. The outputs of the mixers can be delivered over the long cables to the VME crates which are typically installed outside of the tunnel. While the phase of a signal tends to drift on a long RF cable, the amplitude does not. Amplitude drifts would not, therefore, be a problem for the transport of the mixer output signals. Noise picked up on the cable could still, however, be a problem.

Since the RF front-end chassis of the chicane BPM is being used primarily as a beam position monitor and not as a beam arrival-time monitor, drifts of the phase of the reference signal on a long cable into the tunnel do not present a problem. Of course, if the 1.3 GHz reference signal had been delivered to the in-tunnel chassis with a length-stabilized RF cable (Fig. 7.2.3) or with a length-stabilized optical fiber (Fig. 7.2.2), the in-tunnel RF front-end could be used for precise beam arrival-time measurements as well as beam position measurements. This, in conjunction with an improved pickup (reduced reflections), could make the arrival-time resolution and stability of the RF front-end competitive with the optical front-end that will be described in the next section. Under the best circumstances, the RF arrival-time measurement would not be able to easily achieve the < 6 fs few-minute resolution < 30 fs few-hour stability demonstrated by the optical system, but 10-20 fs resolution and comparable stability for this RF front-end would be possible. Given that the 10.4 GHz RF front-end is much more robust than the optical front-end and need not rely on a complicated optical synchronization infrastructure, it is an interesting option to consider in situations that require a quickly built and inexpensive system that doesn't require high-precision synchronization of various laser systems. A cost comparison of the considerably cheaper RF front-end and the optical front-end will be made in the last section of this chapter.

To summarize, the beam position resolution achieved with the 1.3 GHz down mixing scheme was 25 μm. The beam position resolution that could be achieved with the 10.4 GHz down mixing scheme was 6 μm with a 30 meter cable and 4 um with a 2 meter cable delivering the pickup signals. This was measured by splitting the signal from one pickup output and measuring the difference between the outputs of the two arms of the measurement. These numbers ignore the drifts in the measurement which can be ~15 μm over a day and 50 μm over a week. These drifts are due to cable length changes and would be eliminated if the thermally stabilized chassis is installed in the tunnel. Due to reflections in the pickup, the monitor must be calibrated by scanning the position of the beam and not by scanning the LO reference phase. Corrections to the calibration factors need to be calculated based on measurements of the vertical beam position and of the beam charge.

7.3 Optical front-end concept

The fundamental principle of the optical front-end is the same as that of the RF front-end in that for both types of front-ends, a low-resolution phase measurement is used to set the position of a phase shifter that keeps a high-resolution phase measurement in range. As in the RF front-end, four different phase measurements are done in this optical front-end, two of which have a lower resolution and two of which have a higher resolution. The resolution of each measurement is adjusted through the manual application of attenuators, limiters or filters to the signals from the pickups. If a large attenuator or a low-pass filter is used, the dynamic range of the measurement increases while the resolution decreases.

If a limiter or nothing at all is used, the highest resolution is achieved. The maximum resolution that has been achieved to date was <2 μm with 6dB attenuation and a 3 nC beam. With 6dB attenuation and a 1 nC beam, the best resolution was <4 μm. These measurements were conducted with an in-tunnel installation of the optical front-end. With an out-of-tunnel installation of the front-end, the best resolution achieved was 4 μm with a 1 nC beam and a non-split signal. The reduction in resolution for the out-of-tunnel measurement can be attributed to signal attenuation over the 30 meter long cable.

This resolution was achieved through the use of a signal phase measurement technique that takes short pulses from an optical reference and uses them to sample the zero-crossing of the beam transient signal from the pickup. This sampling is enabled through a compact, 5x10x50 mm, device called a Mach-Zehnder Electro-Optical Modulator (MZ-EOM). In Fig. 7.3.1, the EOM is depicted as a rectangular box with red lines, representing a laser beam, going through it and an electrical signal coming in from the top. The polarization of the laser light must be adjusted so that it is in alignment with the polarization axis of the EOM. The incoming laser beam is split and it travels through two Lithium-Niobate crystals which are under the influence of the electric field from the electrical signal. When the crystal is under the influence of an electric field it becomes birefringent and causes a phase shift of the light that is transmitted. Each crystal experiences the opposite polarity of the electric field and will shift the phase of the light in opposite directions. When the laser pulses are recombined, constructive or destructive interference between the pulses will result in a laser pulse amplitude that changes when the amplitude of the electrical signal changes. A less symbolic drawing of the EOM would depict the laser pulse and RF wave co-propagating along a stripline, such that the group velocity of the laser pulse is equal to the phase velocity of the RF wave. Usually this relationship cannot be maintained above a certain frequency. The EOMs are presently available with bandwidths of 10, 20, and 40 GHz. The 40 GHz EOM is ~3 times as expensive as the 10 GHz EOM and until other limiters of the bandwidth are addressed, it will not be used.

In Fig. 7.3.1, the electrical signal (black) shown above the EOM has a wavelength that is long compared to the repetition period of the laser pulses (red) shown below the EOM. Looking at the amplitude of the electrical signal and the amplitude of the laser pulses, it is clear that the amplitude of the electrical wave modulates the amplitude of the laser pulses.

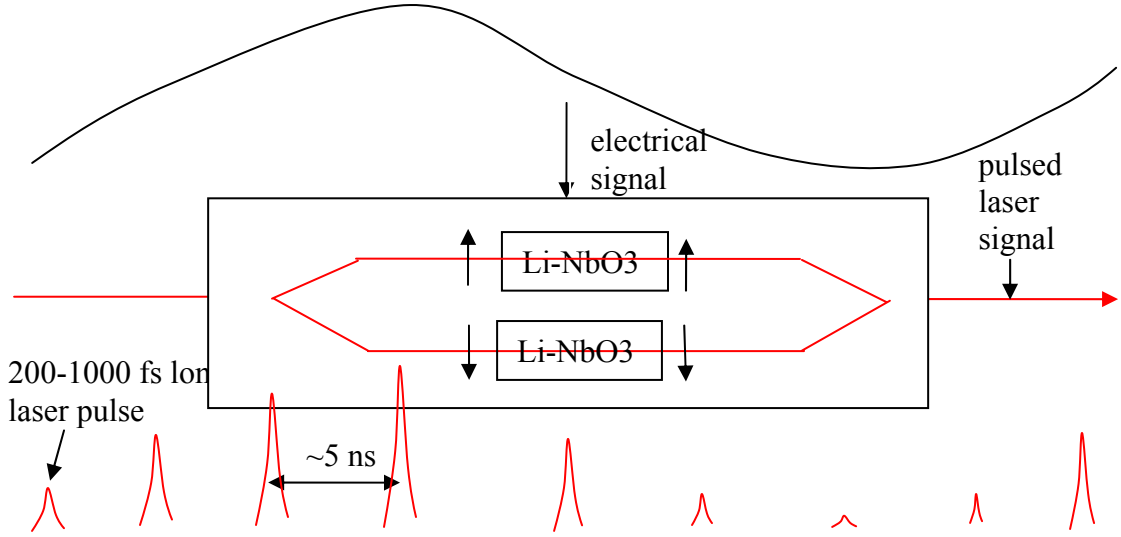

Figure 7.3.1 Mach Zehnder Electro-Optical Modulator (EOM) used to sample the amplitude of an electrical signal. The incoming laser pulse is split and travels through two Lithium-Niobate (Li-NbO3) crystals under the influence of the electric field from the electrical signal. The electric field through each crystal is of opposite polarity.

This is not, however, the way in which the EOM was actually used in the optical front-end of the BPM. The electrical pulse in Fig. 7.3.1 is very long compared to the repetition period of the laser for purely illustrative purposes. In the BPM optical front-end implementation, the electrical pulse was long compared to the duration of the laser pulse, but short compared to the period of the laser pulse train. This is illustrated in the following figure (Fig. 7.3.2).

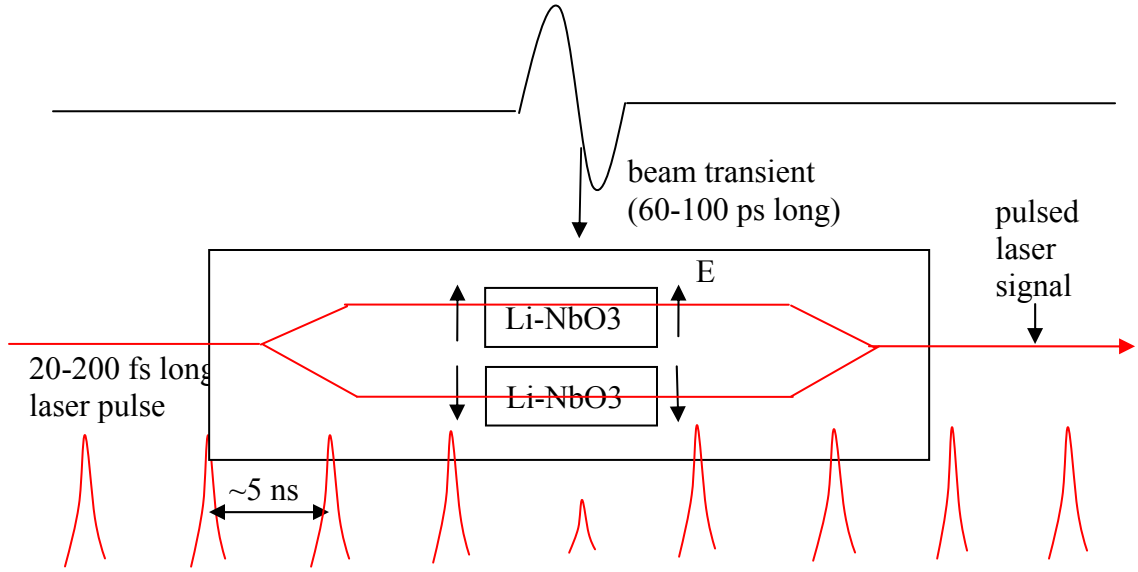

Figure 7.3.2 Mach Zehnder Electro-Optical Modulator (EOM) used to sample the zero-crossing of a beam transient pulse.

Depending on the arrival-time of the beam transient pulse with respect to the arrival-time of the laser pulse, the amplitude of the modulated laser pulse, shown in the middle of the laser pulse train in Fig. 7.3.2, will increase or decrease. To convert this amplitude change into a measurement of the pickup signal's arrival-time, one must use a laser with a very stable repetition rate that can serve as an arrival-time reference signal and the amplitude of the laser pulses must be accurately measured with a photodetector and an ADC that is clocked with a signal that is generated by the 216 MHz repetition-rate laser pulses themselves (Fig. 7.3.3).

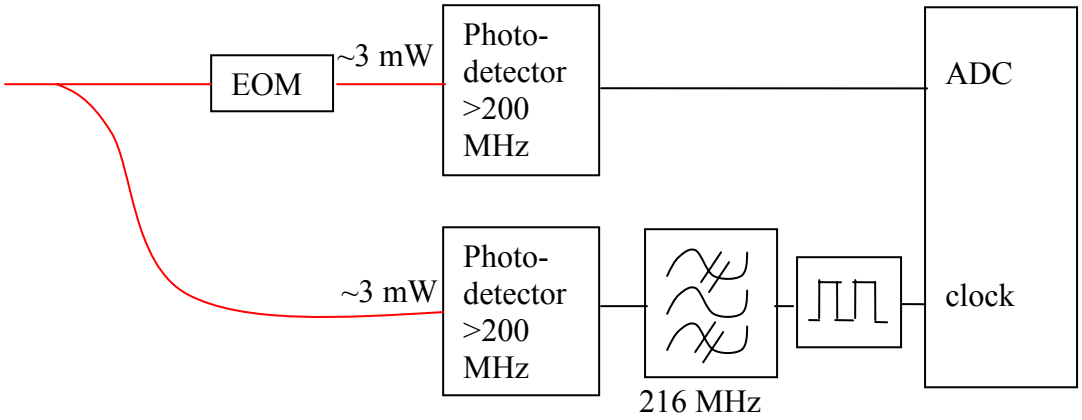

Figure 7.3.3 Measuring the amplitude of the laser pulses with an ADC that is clocked with a signal that is generated by the laser pulses themselves. A band-pass filter is used to extract a clock signal from a higher-bandwidth photodetector.

In Fig. 7.3.3, a bandpass filter with a center frequency of 216 MHz filters the signal from the photodetector. The 216 MHz signal is conditioned to make a square wave that is appropriate for the clock of the ADC. In initial experiments, an AD9510 clock-divider evaluation board and an RF bias adjustment were used to make the clock signal. In later versions, a more compact printed circuit board design was used. Within the in-house ADC-FPGA board, digital clock dividers and shifters were employed to optimize the ADC sampling time.

Once the amplitude of the laser pulse is measured, one must scan the arrival-time of the laser signal about the slope of the steeply falling edge of the beam transient pulse. This provides a calibration factor for the pulse arrival-time measurements (Fig. 7.3.4).

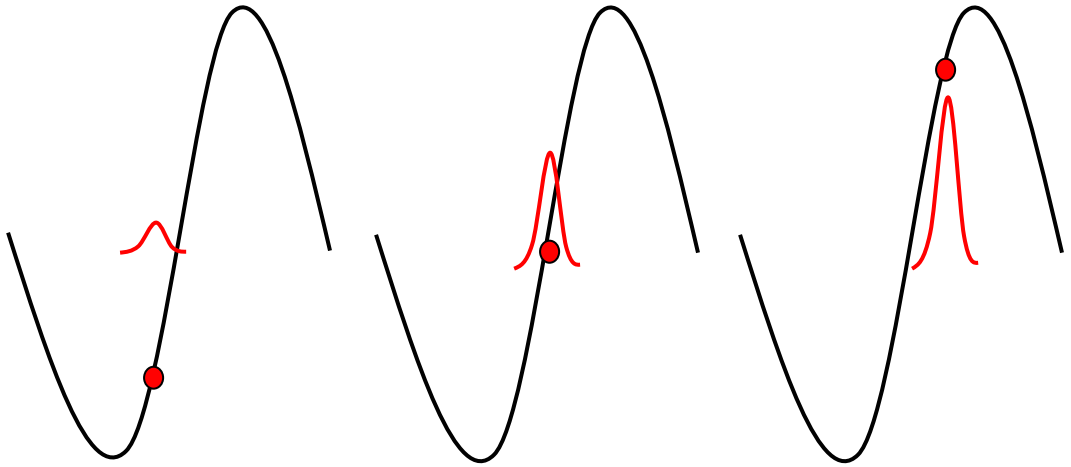

Figure 7.3.4 Calibrating the arrival-time measurement requires scanning the arrival-time of the laser pulse about the zero crossing of the beam-transient pulse. Such a scan is shown in the drawings from left to right. The black signal is the beam transient pulse from the pickup. The red dot is the sampling location. The red signal is the laser pulse. The optimal sampling location for a measurement is at the zero-crossing of the signal (middle drawing).

Whenever the beam arrival-time changes significantly, the laser pulse arrival-time must be adjusted so that it samples the signal at the zero-crossing and the measurement must be re-calibrated. Whenever beam properties change by enough to influence the amplitude of the signal, the measurement must be re-calibrated. This sampling concept was the idea of Holger Schlarb and was first implemented in [20].

7.4 Optical front-end execution

In the following figure (Fig. 7.4.1), a fiber splicing plan for the optical front-end of the chicane BPM is shown. The goal of this splicing plan is to steal a small portion of the light from a length-stabilized fiber link which has been delivered to a nearby BAM chassis, transport it over an 8 meter long fiber patch cable and adjust its polarization, amplify it, and then adjust its arrival-time at two pairs of EOMs. The tap-off point within the BAM chassis is shown in the top portion of the drawing while the splicing plan for the devices residing within the chicane BPM chassis is shown in the lower portion of the drawing. The lengths of the fibers within the BPM chassis are written above the splicing plan and the optical power levels at and insertion losses of the various components are written below the splicing plan. FC-APC connectors are shown as thick black marks, optical delay lines are labeled as ODLs, the amplifier is represented as a triangle with 4μm inside, representing that 4 μm gain fiber was used. The ~3 picosecond length of the laser pulses when they arrive in the amplifier is written above the amplifier and the transition from single mode fiber (SM) to polarization maintaining fiber (PM) is written underneath the first ODL.

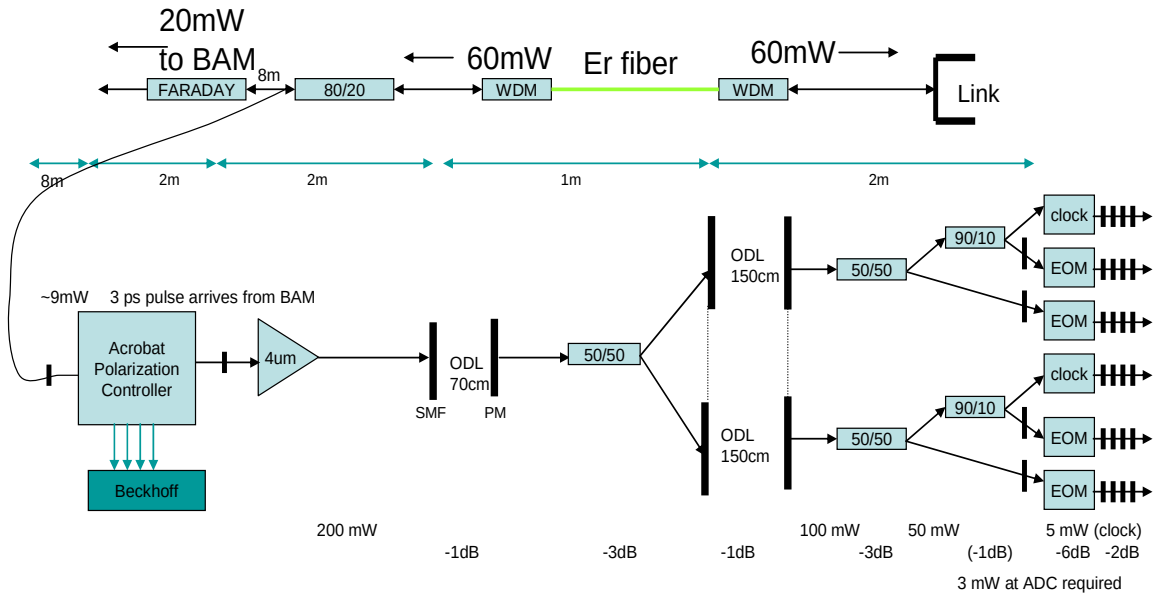

Figure 7.4.1 Chicane BPM optical front-end schematic. Laser pulses from a near-by length stabilized link are tapped off and delivered over an 8 meter long optical patch cord to the chicane BPM chassis. Within the chassis, the polarization of the incoming link light is controlled, the light is amplified and the arrival-time of the pulses at the 4 EOMs is adjusted with the Optical Delay Lines (ODLs).

The length stabilized fiber link entering the schematic in the upper right corner is a long stretch of SMF optical fiber connecting the Master Laser Oscillator (MLO) to a timing sensitive device, in this case, the chicane BPM. The pulses coming from the MLO are reflected by a Faraday rotator in the timing sensitive device and sent back to the room containing the MLO. Once there, the arrival-times of the returning pulses are measured relative to the arrival-times of pulses that have just been generated by the MLO Fig. 7.4.2. This measurement is used to adjust the path-length of the light with an Optical Delay Line (ODL) and with a piezo fiber stretcher.

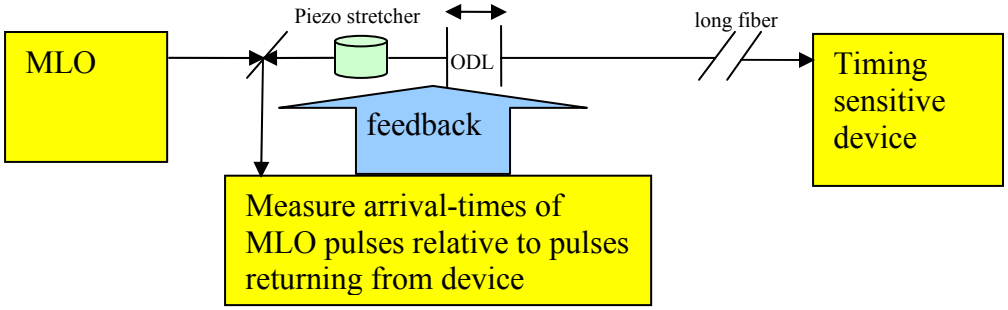

Figure 7.4.2 Length stabilized fiber link concept. The arrival-times of pulses from the MLO are compared to the arrival-times of pulses that are reflected within the timing sensitive device. This information is used to adjust the length of the fiber with a piezo fiber stretcher and an optical delay line (ODL).

For beam arrival-time measurements, it is important that the arrival-times of the pulses from the MLO are constant. For beam position measurements this is not important. Stabilizing the arrival-times of MLO pulses requires the stabilization of the optical length of the fiber over which the pulses are transported to the beam arrival time measurement (BAM) front-end. While the chicane BPM can be used as a special sort of BAM that is sensitive to the width (energy spread) of the beam, the BAM application of the optical front-end will be reserved for the next chapter.

At the end of the link (upper right corner of Fig. 7.4.1), within a nearby BAM chassis containing the link-end, the 5-10 mW pulses from the fiber link are amplified in a 60 cm long stretch of 8 μm Erbium-doped fiber which is forward pumped. This amplifies both the pulses coming from the link and the pulses which are reflected by the Faraday rotator and sent back to the fiber-link stabilization unit. The ~60 mW traveling in the direction of the Faraday rotator is split, such that 80% goes towards the BAM measurement, and 20 % goes to the chassis containing the chicane BPM front-end.

The fiber-links must not only deliver a pulse that is stable in arrival-time, but the pulses must also have an appropriate length when they reach the EOMs and when they return to the optical cross-correlator. To this end, a length of LLWBDK dispersion compensating fiber from OFS is incorporated into the fiber-link stabilization chassis. The length of dispersion compensating fiber in the present case must make the pulses short when they reach the EOMs in both the BAM chassis and chicane BPM chassis. Because the connection between the BAM chassis and the EOM chassis requires about 8 meters of fiber, an additional 8 meters of fiber is wound up inside the BAM chassis after the splitter but before the Faraday rotator. This makes the distance between the EOMs and the laser the same for both the BAM EOMs and the chicane BPM EOMs.

This additional 8 meters of fiber will not have an impact on the length stability of the BAM fiber link because all fiber from the MLO up to the Faraday rotator is length stabilized. If, however, the 8 meters of fiber connecting the BAM chassis to the chicane BPM chassis drifts relative to the 8 meters of fiber which is wound up inside of the BAM, the arrival-time of the MLO pulses at the chicane BPM front-end will not be stable and the arrival-time measurement will no longer be usable. The optical length of one meter of a standard optical fiber will drift by 60 fs/deg, 8 meters will drift by 480 fs/deg. When the chicane BPM front-end is used purely as a beam position monitor, this drift would have no impact, but if the BPM front-end is also used as a beam arrival-time monitor, this drift will cause systematic errors in the measurement. The drift can be avoided by using a special, and more expensive type of fiber called PSOF (Phase Stabilized Optical Fiber). At the moment, it is only available from one firm, Furukawa. Compared to the standard SMF-28 fiber, the optical length of PSOF does not change significantly with temperature. While using kilometers of this sort of fiber would be prohibited by the ~60-150 EUR/meter cost, 8 meters of PSOF is considerably less expensive than building an additional, ~25,000 EUR, actively length-stabilized optical cross-correlator fiber link.

Upon entering the chicane BPM chassis depicted below (Fig. 7.2.3), the polarization is adjusted with a polarization controller from BATI (labeled: “PC” in Fig 7.2.3 and “acrobat polarization controller” in Fig. 7.2.1). While the polarization of the light from the fiber-link can be adjusted in the fiber-link stabilization chassis, the polarization that is appropriate for the BAM EOMs will not necessarily be appropriate for the chicane BPM EOMs. This is the reason for the in-chassis polarization controller. It

can be placed anywhere prior to the transition from single-mode fiber to polarization-maintaining fiber. Because there is an ample power-level coming into the chassis, the polarization controller was placed before the amplifier instead of after it.

The amplifier is a 60 cm long stretch of 4 μm diameter Erbium-doped gain fiber which is pumped from both directions. It amplifies the 3 picosecond long, ~6 mW pulse up to 200 mW. This sort of large amplification would not have been advisable if the pulse length had been shorter than a picosecond. When a short pulse is amplified above 80 mW, unstable distortions in the pulse shape will be created that could have an impact on the accuracy of the sampling scheme. In order to achieve the best gain in the fiber it is advisable to wind it over the space with a large radius of curvature; wrapping it around a cylinder with a small radius can degrade the amplifier's performance. Another consideration in amplifier design is noise. To get the highest gain with the least amount of noise one should try to minimize the length of the gain fiber while maximizing the output. The rules governing this optimization vary depending on the type of gain fiber that is used, so if time allows, it is good to start with a longer stretch of gain fiber, pump it from one end, and measure the gain of the amplifier. After cutting off a few centimeters at a time, the gain curve will begin to become clear. For the best noise performance, the best place to stop cutting off lengths of gain fiber is when the gain curve starts to become linear. Of course, the length for the best noise performance may not have the gain that is needed, so a compromise has to be reached.

After the amplifier, the laser pulses enter an Optical Delay Line (ODL). This ODL will move whenever the arrival-time of the beam in the chicane changes. The fiber entering the ODL is a single mode fiber (SMF) and the fiber exiting the ODL through the collimator attached to the mobile portion of the stage is polarization maintaining (PM). It is better to put the PM portion on the moving part of the stage because if the SMF is moved, the polarization controller will have to be adjusted. While the polarization controller actuator is fast enough to react to fast changes caused by vibrating and moving fibers, this makes the task of the polarization controller harder than it needs to be. For fast polarization changes, a micro-controller or DSP can be used to adjust the polarization controller based on a measurement with a photo-detector of the un-modulated EOM output or based on an actual measurement of the polarization of the pulse using an in-line polarimeter (General Photonics). A necessity for such fast adjustments has not yet been seen for any of the systems built so-far, but when pushing toward sub-femtosecond resolution, such techniques may be employed.

After the first delay line, the fiber is split into two arms, so that half of the light goes to one end of an ODL and the other half goes to the other end of the ODL. When the ODL stage moves, the path length of one arm gets longer while the other gets shorter. This ODL stage would need to move whenever the position of the beam is changed by more than a millimeter or so. At the exits of the two arms on this ODL, the fibers are split again. Two of the four fibers will be used to sample electrical signals coming from the right-side of the pickup and the two of the fibers will be used to sample the signals from the left-side. Of the sets of two, one fiber will be used to perform the fine measurement of the arrival-time of the pickup signal and the other fiber will be split with 90% going towards the coarse measure of the pickup signal arrival time and 10% being used to generate the clock for the ADC which will be used to sample the amplitudes of the signals emerging from the EOMs.

Because fibers and EOMs are sensitive to temperature changes, the temperature of the plate upon which most of the fibers rest is actively controlled with Peltier elements. Peltier elements were chosen instead of a heating mat because they can cool as well as heat and they react more quickly to changes in the control voltage. The 4 Peltier elements are mounted underneath the aluminum plate shown in the side view of the chassis (Fig. 7.2.4). They are incorporated into a fast control loop with a Wavelength PTC temperature controller that uses a single temperature sensor mounted on the top of the metal plate in order to deliver control voltages to the Peltiers. If the Peltier is cooling the plate, heat will travel down a conducting channel towards the outside of the box. If it is heating the plate, the opposite will occur. The outside of the box is insulated from the inside of the box with a centimeter thick layer of neoprene. While the conductivity of neoprene is higher than that of air, it is a better insulator because it prevents convection from transferring as much heat from the outer box to the inner. A fan is used to cool the outer box.

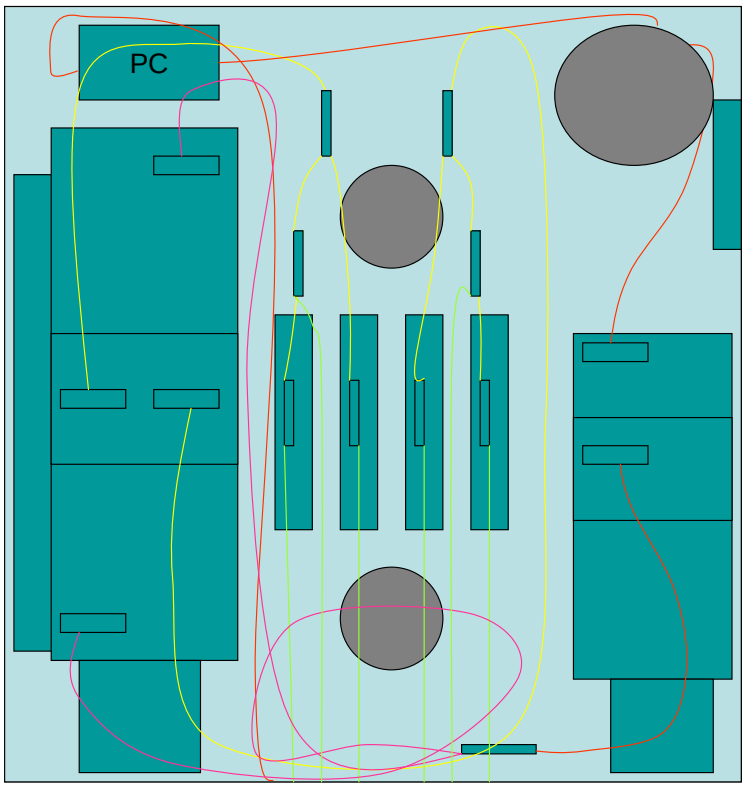

Figure 7.2.3 The layout of the fibers in the top layer of the optical front-end chassis for the chicane BPM. The stage on the left will move when the beam position changes. The stage on the right will move when the arrival-time of the beam changes. The four EOMs are in the middle of the drawing on an actively temperature stabilized plate.

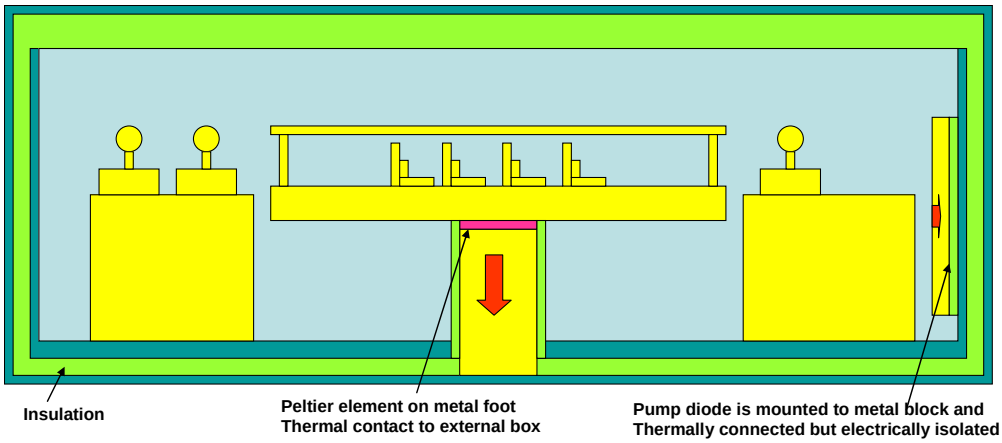

Figure 7.2.4 The side view of the optical front-end chassis. The two-layer design was made in order to facilitate the control of the temperature of the plate on which the fibers were resting.

The optical length of one meter of a standard optical fiber will drift by 50-60 fs/deg. There are approximately 2 meters of fiber along one arm of the chicane BPM optical front-end, so if the temperature of one arm differs from the other arm by less than 0.1 degrees Celsius, the beam position measurement will be affected by less than 0.5 μm. If the temperature of the chassis changes by 0.1 degrees Celsius, the beam arrival-time measurement will change by 6 fs.

The active temperature stabilization system shown in Figure 7.2.4 can keep the temperature of the box stable to within 0.003 degree C rms for slowly changing external temperatures in an air-conditioned laboratory setting. When the box is in the accelerator tunnel, the temperature changes of the plate on which the fibers rest were ~6% of the 2 degree C temperature changes measured on the outside of the box, even though the inner-box air temperature tracked the 2 degrees of tunnel temperature change. The measurement of the in-tunnel temperature stability is limited by the Beckhoff ADC to 0.03 degree Celsius resolution and it takes about a 12 hours for the temperature in the box to become truly stable after the tunnel has been opened for a maintenance day. It is anticipated that this stabilization time could be reduced if more active cooling were applied to the outside of the chassis, improving the efficiency of the removal of heat from the Peltiers. One other problem that limits the speed with which the temperature regulation loop can function is the latency between the time that a setpoint change command has been given and when the resulting temperature change is measured by the sensor. Because the controller reacts too quickly to the setpoint change, by the time the thermistor measures the resulting temperature change, the system has already over-shot the target setpoint by a large margin. In the case of a small 0.1 degree C setpoint increase, the system will heat up by more than half of a degree before it converges back to the new setpoint. This problem can be addressed by adjusting the feedback parameters of the temperature controller and by locating the feedback thermistor as close as possible to the peltier.

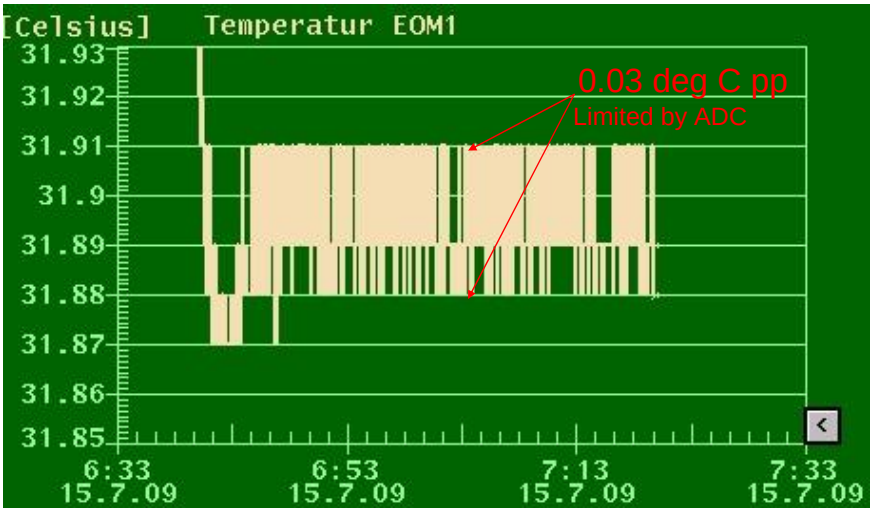

Figure 7.2.5 Effectiveness of active temperature control in the tunnel.

While the thermal stability of the system limits the long-term stability of the measurement, the resolution is limited by the bandwidth of the signal that is transmitted to the EOM. Maximum bandwidth can frequently not be used because the machine is so unstable that the beam is constantly jumping out of the measurement's dynamic range. While the measurement can achieve sub-micron resolution for large signal slopes, the measurement was only stable for signals that were attenuated so that they produced 2-4 μm resolution. If the machine becomes so stable that dynamic range is no longer a problem, the first device that needs to be optimized in order to increase the bandwidth is the combiner which reduces the measurement's dependence on vertical beam tilts. The current combiner is an in-house built device with 6 dB insertion loss. The second device that could be improved, in conjunction with a switch from SMA to type-K cables connecting the pickup to the EOM, is the Agilent N9355C limiter which can have a bandwidth of up to 50 GHz if the type-K version is used instead of the SMA version. The third device which could be improved is the EOM. The current EOM has a 3dB bandwidth of 10 GHz, but there are more expensive EOMs that accommodate above 40 GHz. The pickup is the most time-intensive component to improve, but if the pickup was re-engineered to be hollow and suspended between the two vacuum feedthroughs, the elimination of the ceramic support rings would produce a 30% improvement in the slope of the signal (Fig. 5.5.3).

Given perfect thermal stability, the device that limits the accuracy of the optical front-end more than any other is the RF limiter. The limiter, containing a combination of Schottky diodes, is used to prevent damage to the EOMs from large amounts of power. Without the limiter, the EOMs survive short pulses of even a few hundred Volts without immediate damage, however, if the beam is steered directly onto the pickup itself or if the pickup is sprayed with a shower of electrons, in the absence of an RF limiter, the EOM crystals become opaque within a matter of minutes. This was directly observed during a machine studies day when the beam orbit was dramatically off-center. The pickup in this instance was a ring-type pickup and had, therefore, a large surface area open to beam spray. Using pickups with a smaller surface area exposed to the beam could reduce the possibility for beam-spray damage of the EOM and, therefore, remove the requirement of

using the RF limiter. This might be desirable because an RF limiter produces an appreciable, yet difficult to measure, amount of AM-PM conversion. This means that the limiter can convert changes in amplitude into changes in phase. Since the phase is the quantity that we want to measure, if the AM-PM conversion is large enough, it can limit the accuracy of the measurement. Reducing the amplitude of the signal entering the limiter can also reduce this effect, but unless the bandwidth of the signal is preserved when the amplitude is reduced, this will reduce the resolution of the measurement. Measuring this AM-PM effect requires high amplitude (~60 Volts) RF signals with a frequency of several GHz.

7.5 Front-end Costs

The RF front-end and the optical front-end can deliver comparable ~5 μm resolution in their present configurations, but the optical measurement has the potential to reach <1 μm resolution while the RF front-end can only achieve 3 um resolution under best-case circumstances. The optical front-end is also immune to the effect of reflections within the pickup. It can be calibrated with an external reference and, due to reflections in the pickup, the RF measurement must be calibrated with either a scan of the accelerating gradient or of the chicane dipole current. This means that the optical measurement calibration is parasitic and the RF measurement calibration is invasive and would disturb machine operation. The optical front-end can deliver <6 fs resolution beam arrival-time measurements with length stabilized fiber-links and under the best possible circumstances, RF front-end beam arrival-time measurements could not be made better than 10-20 fs. The cost of the RF front-end is, however, about half of that of the optical front end (Table 7.5.1).

RF front-end		***Optical front-end***	
Item	*Price (EUR)*	*Item*	*Price (EUR)*
RF Vector Modulators	4,000	Optical delay stages	7,500
RF BP filters	400 each	EOMs	1,200 each
RF amplifiers	100 each	amplifiers	1,000 each
RF mixers	100 each	splitters	200 each
RF trombone	600	collimators	300 each
Temp. control	600	Beckhoff	3,000
Enclosure	100	Engineered enclosure	4,000
Struck ADC	4,000	In-house ADCs	5,000
100 m cable	600	Cabling + etc.	5,000
Total	15,000	*Total*	35,000

Table 7.5.1 Rough cost estimate of RF front-end and optical front-end for the chicane BPM. The cost of a length stabilized RF cable or optical link is not included.

8 Beam Arrival-Time Monitors

An optical beam arrival-time monitor with 6 fs resolution [20] has made a high-resolution time-of-flight energy measurement in the bunch compressor a possibility and is the only monitor system that has the accuracy to cross-check the measurements of the chicane BPM constructed in this thesis. Like the optical front-end of the chicane BPM, this monitor requires the delivery of laser pulses from the optical synchronization system. But, before describing the advantages of the newer-technology optical system, the advantages and limitations of the older-technology RF system should be clearly defined.

8.1 RF Front-end

RF synchronization of accelerator facilities has been used with great success since their inception, achieving picosecond synchronization between locations separated by several kilometers without the aid of active feedbacks to compensate for cable length changes. With the addition of an active feedback on the cable length, in-loop measurements of synchronization that is better than 100 fs have been made [46]. It is important to note that this method and any other active, point-to-point synchronization schemes can only synchronize one location with another; any locations in between these two points are not necessarily synchronized. This is because the waves may travel at different speeds in different segments of the cable.

There are several problems that may be evident in an out-of-loop measurement of the performance of an RF cable reflectrometry set-up. With an active RF cable feedback method, a signal of a few GHz is sent from a Master Oscillator (MO) over a coaxial cable to an end location at which part of the signal is reflected and sent back to the source, where the phase of the signal being generated and the phase of the signal returning can be

compared. The length of the cable can then be adjusted until the returning phase matches the sent phase, thereby removing the effects of cable drifts. Although the phase detection of the returning pulse can be made with 10 fs accuracy, 20 or more fs of noise is picked up over the long cable and the reflection of the pulse at the end of the cable is problematic due to the temperature dependence of the mismatch that produces the reflection. Any drift of the mismatch will be "corrected" by the feedback loop, but this does not accurately represent the drift of the cable and it, therefore, adds an error to the cable length stabilization. The directivity of the coupler is, perhaps, the fatal flaw in the method, because any changes in the directivity will directly impact the comparison of outgoing and returning signal phases. This reflectrometry method has not been tried extensively and the absolute limitations are not completely clear. Nevertheless, it does not look very promising. It also becomes more difficult for cables that are longer than 100 meters due to the attenuation of the signal on the cable.

If one had access to an RF reference signal with a stable arrival-time, one could measure the arrival-time of the electron bunch relative to the reference signal by taking a signal from a pickup in the beam pipe, filtering out a frequency and measuring its phase my mixing the filtered signal with the MO signal. Beam arrival-time measurements using such a technique were presented in the previous chapter. With the RF front-end described in the previous chapter, the measurement of the phase of the MO relative the arrival of the pickup signal delivered over a 30 meter long cable could be done with a resolution of 20 fs (rms) without taking into account any drift or noise on the MO signal. The arrival-time resolution can be halved by using a short cable from the pickup to the front-end chassis, halving the noise picked up over the 30 meter cable. As good as these numbers may sound, neither address the problem of delivering a stable MO signal over a large distance. They also do not address the question of costs.

RF components are frequently cheaper than optical components, but low-drift RF cables can be much more expensive than optical fibers and will have much higher attenuation (Table. 8.1.1).

	cost/meter	*drift/meter degree C*	*attenuation/km @1.3GHz*
Phase Stable Optical Fiber (PSOF)^	25 EUR	<1 fs	0.02 dB
Single Mode Fiber (SMF)^	0.20 EUR	50-60 fs	0.02 dB
Polarization Maintaining Fiber (PM)^	2 EUR	50-60 fs	0.02 dB
Cell-flex LCF12 RF cable*	3 EUR	100 fs	81 dB
Cell-flex LCF78 RF cable*	6 EUR	17 fs	46 dB
Huber+Suhner RF cable	150 EUR	3 fs	147 dB

Table 8.1.1 Costs and performance of RF cables and optical fibers. *Data from Henning Weddig. ^When contained in a bundle of 6 fibers.

- While there are RF cables with better drift properties than SMF fiber, amplifiers will frequently be required to compensate for their attenuation, and amplifiers

tend to make significant contributions to drifts as well. Couplers are another source of drifts, and can become very expensive when tolerances are tight.

- While optical signals are not significantly attenuated by transmission over long fibers, RF signals sent over long cables are.
- While optical signals are immune to electro-magnetic interference from noisy devices, RF signals are not.

The prevailing opinion at the present time is that pushing RF components below the 10 fs level would create a cost explosion as expensive cables and couplers are required. The alternative is to develop the optical technology in order to easily achieve 10 fs performance and have the possibility to later push to the sub-fs level. If the optical system is made widely available to other FEL facilities, the costs could decrease due to economies of scale. This argument does not apply, however, if the target synchronization accuracy is above 20 fs and the distances involved are much less than a kilometer. This is a regime for which RF arrival-time monitors could still be considered. Table 8.1.2 shows a cost estimate of an RF arrival-time monitor with a phase-stabilized cable.

Item	*Price (EUR)*
RF phase shifters	2,000 each
RF BP filters	400 each
RF amplifiers	100 each
RF mixers	100 each
AD8302	100 each
RF coupler	300
Engineered enclosure	600
Temp. Stabilization	3,000
Struck ADC	4,000
100 m cable	600
Total	15,000

Table 8.1.2 Rough cost estimate for a 20 fs resolution RF arrival-time monitor with an RF reflectrometry setup.

8.2 Optical Front-end

The optical synchronization system was developed in order to lock the beam to an optical reference signal which can be delivered to remote locations with sub-10 fs accuracy [2]. To this end, the beam arrival-time relative to the optical reference must be measured and the optical reference must be delivered to remote locations without loosing its phase stability. This is done with a Master Laser Oscillator (MLO) that sends pulses along fiber links to end-stations. At each end-station, a portion of each incoming pulse is reflected and sent back to the source. At the source, the arrival-times of the reflected pulses can be compared to the arrival-times of the new pulses coming from the MLO. The length of the fiber is then adjusted until the returning pulse timing matches the sent pulse timing, thereby removing the effects of temperature induced timing drifts. At this point, the principle of the optical synchronization system sounds identical to that of a purely RF synchronization system with only the word MO replaced with MLO and the word cable

replaced with fiber. The advantage of the optical system becomes apparent when the accuracy with which the detection and transport of an optical signal can be carried out is compared to that of an RF signal (Table 8.2.1). In terms of phase shifting, phase detection, EMI and vibration, the optical system has performance that is about an order of magnitude better than that of an RF system.

	Optical	*RF*
source phase noise (1kHz-10MHz)	~4 fs	~4 fs
phase shifting	~1 fs	~4 fs
phase detection	<1 fs	~10 fs
EMI + vibration	<1 fs	>10 fs

Table 8.2.1 Comparison of optical and RF systems phase noise, detection, and etc.

To summarize the advantages of the optical system:

- Attenuation is not an issue when transporting an optical signal over fibers.
- Optical signals are immune to electro-magnetic interference.
- Optical signals can contain a large bandwidth, enabling high precision arrival-time measurements of both optical pulses and electrical pulses.

A beam arrival-time measurement that uses the optical system to measure the arrival-time of an electrical pulse is depicted below in Fig. 8.2.1. It is essentially the same as the system that was described in the previous chapter for the optical front-end of the chicane BPM except that in the case of the BPM, the emphasis was on measuring the difference in the arrival-times of two beam transient pulses, whereas the beam arrival-time measurement is concerned with the arrival-time of a single beam transient pulse relative to an optical reference. When two different beam arrival-time monitor systems separated by 60 meters were measured against one another over few minute time scales, they had a resolution of 6 fs. Over longer time scales they differed by more than 30 fs. This was due to the sensitivity of the measurements to small changes in the beam shape [20].

Within the BAM front-end, a portion of the incoming laser light is reflected backwards along the fiber which delivered the light. The arrival-time of the returning pulse is measured relative to the arrival-time of a pulse from the MLO with an optical cross-correlator. Optical fiber links can also be stabilized with much less expense with an RF technique that utilizes balanced detection of photodetector signal phases [47]. While it has been shown that the RF technique and the more expensive cross-correlation method can stabilize a several hundred meter long link to within 10 femtoseconds (out-of-loop) the resolution of the RF technique has already been pushed to its theoretical limit [47], while the cross-correlator has the potential to achieve sub-femtosecond accuracy [50]. Drifts of the optical cross-correlator have not been fully understood yet and that is why the theoretical limit has not yet been reached.

A fatal error was made in two recent engineered prototypes of length stabilized optical links: one version was built and designed in-house at DESY and the other was built by a commercial vendor. This fatal error was that the correction of the length of the fiber was made for only the outgoing pulse and not for the reflected pulse. While the returning and outgoing pulses were both synchronized in the cross-correlator, the link-end was not stabilized.

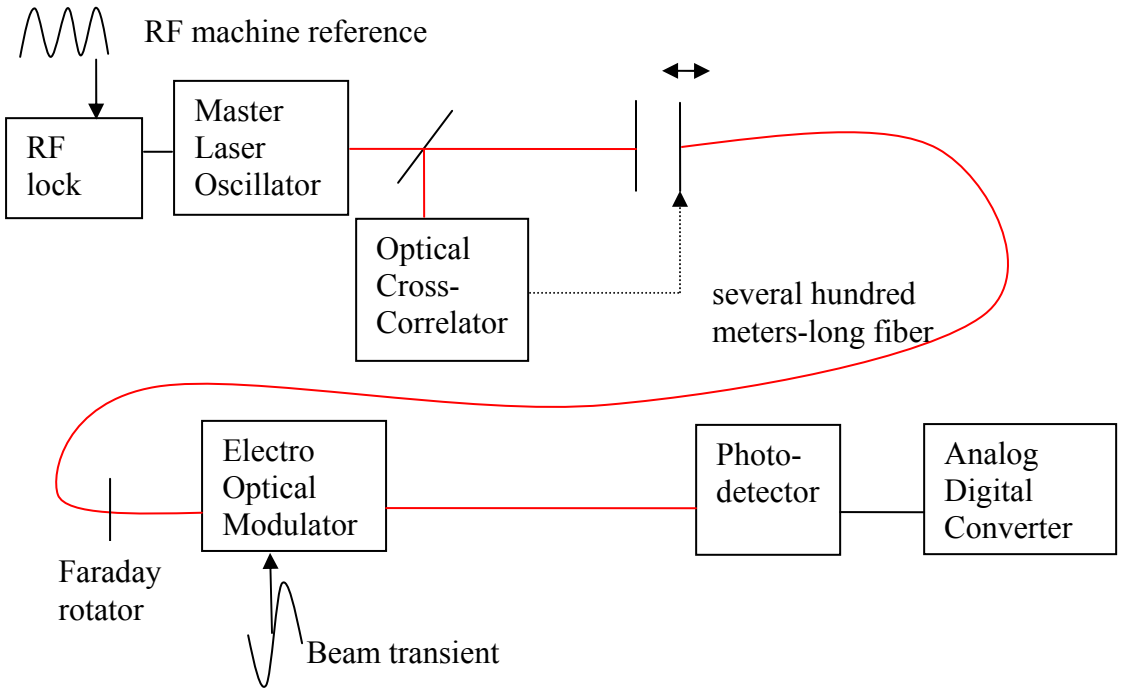

Figure 8.2.1 Beam arrival-time measurement with length stabilized fiber. The RF pulse from the pickup and the optical pulse from the MLO meet up in an EOM. The amplitude of the optical pulse exiting the EOM will change when the beam arrival-time changes.

In Fig. 8.2.2, the arrival-time of a pulse emerging from the MLO is measured relative to the arrival-time of a pulse returning from a timing sensitive device. This measurement is done in an optical cross-correlator. In an optical cross-correlator, two laser pulses with polarizations that are perpendicular to one another are sent through a dichroic mirror that reflects second harmonic light and transmits the fundamental mode. The input laser pulses are transmitted through the mirror and are sent through a PPKTP crystal which. The right side of the crystal is coated with a high reflectivity coating and an anti-reflective coating in order to reflect and transmit light pulses from the fundamental and second harmonic light. The left side of the crystal is coated with anti-reflective coatings so that it transmits both the fundamental and second harmonic light. When the pulses overlap in the crystal, second harmonic light is generated and emitted in both forward and backward directions. With the aid of dichroic mirrors, these pulses each travel to a photodetector, as the incoming light returns from whence it came.

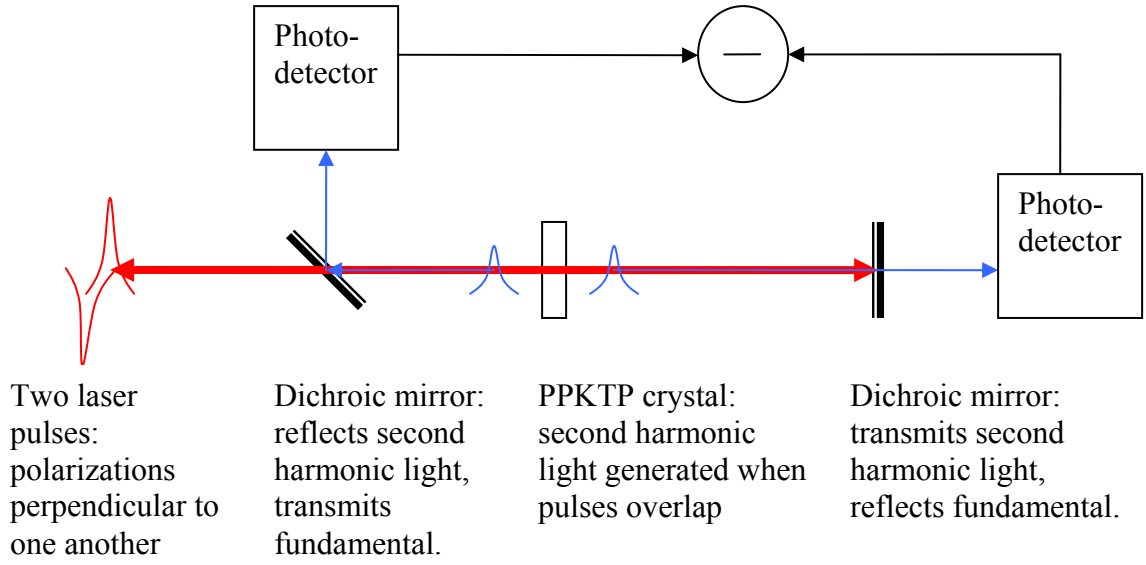

Figure 8.2.2 Balanced optical cross-correlator used to measure the difference between the arrival-times of pulses coming from and returning to the MLO [50].

In this balanced detection arrangement, when the signals at the photodetectors are balanced at about half of the maximum signal, the optical length of the fiber is stabilized with an accuracy that can be well below 10 fs [20]. Fiber length stabilization can also be accomplished with a scheme that detects the arrival of the MLO pulse and the returning pulse with RF phase measurements of the signals resulting from the laser signals impinged upon photodetectors. This is significantly less expensive than the optical cross-correlator, but the best resolution that it can achieve is about 10 fs [47].

In Table 8.2.2, an estimate of the cost of a 6 fs resolution optical front-end and a length-stabilized fiber is presented. The total cost of 55,000 EUR can be compared to a total cost of 15,000 EUR for an RF front-end and phase-stabilized cable system.

item	*Price (EUR)*
BAM front-end	30,000
Link w/OCC	25,000
Total	55,000

Table 8.2.2 Cost estimate for 6 fs resolution optical front-end and a length-stabilized fiber. (Cost for MLO distribution not included).

8.3 MLO RF-lock

All of the results quoted for the beam arrival-time measurements done with the optical front-end refer to the measurement of the arrival-time of the electron beam relative to the laser pulses from the MLO, but a beam arrival-time measurement relative to the RF reference of the machine can never be better than the lock of the MLO to the RF master oscillator (MO) that sets the reference signal for the accelerating RF. If one quotes the arrival-time of the electron beam relative to the MO instead of the MLO, one might

frequently measure differences of several picoseconds in a drifting system. This is not, however the right way to look at the problem. If everything of importance: lasers and arrival-time feedback diagnostics, are locked to the MLO, then nobody cares about their phase relationship with the MO. The MO signal is only important in so far as it keeps the accelerator stable and running, while the MLO reference is responsible for delivering fine corrections to the stability of the accelerator and maintaining synchronization with the lasers of the FEL users.

It turns out, however, that the MLO would be a useless reference without a lock to the MO. The short-term phase stability of the best MOs and MLOs can be less than 4 femtoseconds (1kHz-10MHz), but the long-term phase stability of an MLO is really quite bad. When a source with a good short-term stability is locked with a narrow bandwidth (~1kHz) to a device with a good long-term stability, the source with the good short-term stability acquires the good long-term stability of the device to which it is locked. The current plan to accomplish this is to lock the higher frequency MO to a temperature stabilized crystal oscillator with a lower frequency. This crystal oscillator would be locked to a GPS standard frequency [56]. When the MLO is locked to the MO, it will gain the long-term phase stability of the GPS standard. This is the reason that it is desirable to have an MO-MLO lock that does not drift. When it doesn't drift, then one can be certain that the long-term phase stability of the optical reference is as good as the long-term stable RF references. The setup to accomplish this lock is shown in Fig. 8.3.1.

The schematic in Fig. 8.3.1 depicts one chassis containing two identical circuits (the top-half and the bottom-half are identical). The concept was developed by several people from the optical synchronization team, but the construction and characterization of the lock was done by the author. Each circuit provides two distinct measurements of the relative phases of the MO and the MLO. One is a fine measurement mixing 1.3 GHz from the MO with 1.3 GHz generated from an MLO-based photodetector signal. The other is a coarse measurement using 216 MHz that keeps track of which bucket the fine measurement is measuring. These measurements of the relative phases of the MO and MLO are used in an ADC-DSP-DAC feedback loop in order to set the voltage of a piezo-fiber stretcher that adjusts the round-trip time of a pulse in the laser. This adjusts the phase of the MLO relative to the MO. Additional signals are provided by the chassis for the monitoring of laser power and amplification of the filtered signals. This chassis has been used to provide the RF phase measurement for the MLO-RF lock and, in parallel, to measure the out-of-loop performance of this lock.

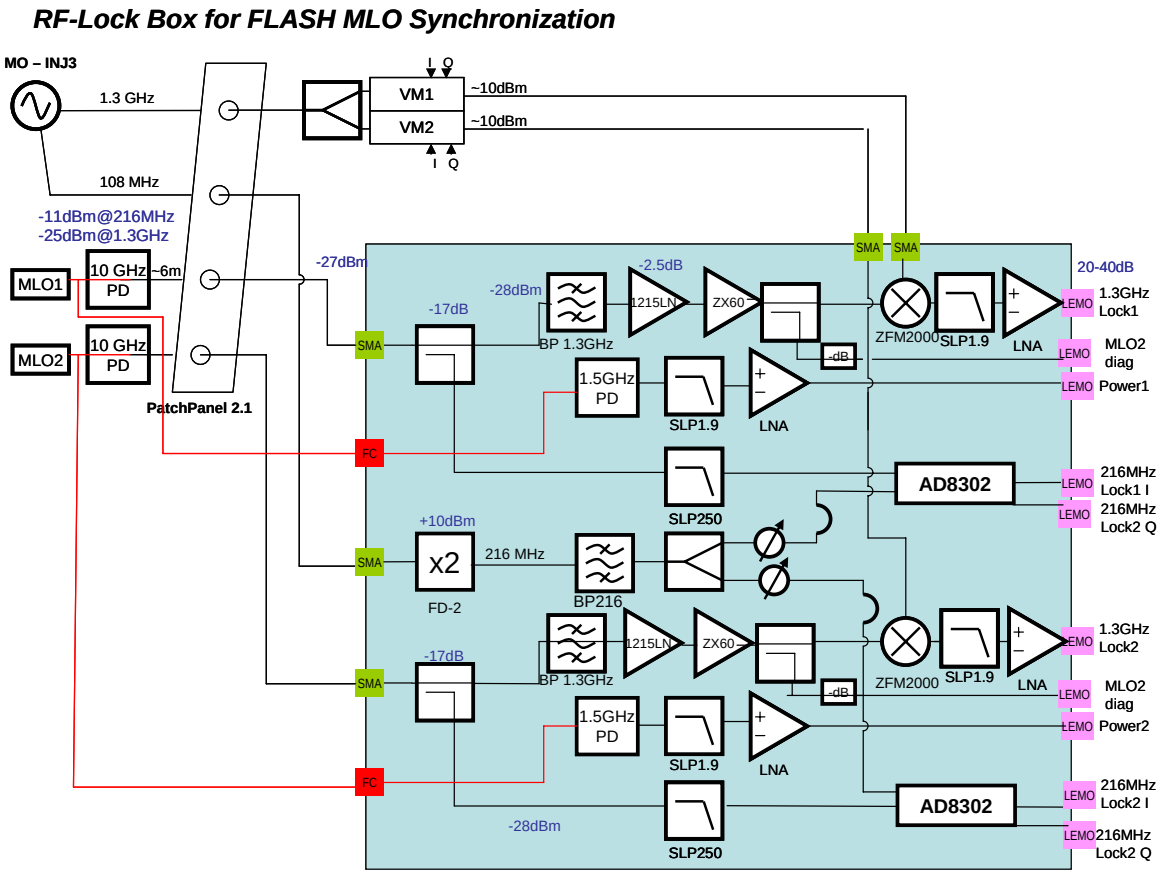

Figure 8.3.1 Schematic of MLO-MO laser-RF lock. Inputs include 1.3 GHz and 108 MHz from the MO and photodetector signals from two different MLOs. Components include vector modulators (VM1, VM2) for shifting MO phases in calibration routines, bandpass filters and amplifiers for removing and amplifying the 1.3 GHz frequency component from the photodetector outputs, mixers for generating a baseband signal from the MO and MLO signals, phase and amplitude detectors (AD8302) for a coarse determination of the MO and MLO phases with a 216 GHz measurement, and couplers for online troubleshooting of the circuit.

The strategy of the circuit shown in Fig.8.3.1 is to filter out one frequency from the frequency comb produced by impinging the pulsed MLO laser signal on a photodetector and then measuring the phase of the filtered signal relative to the MO reference. These signals are shown below in Fig. 8.3.2.

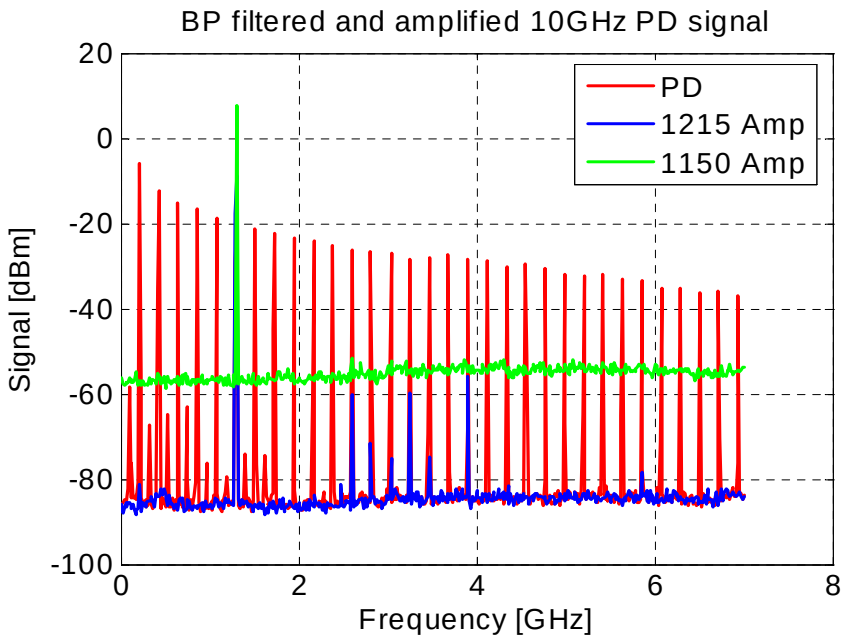

Figure 8.3.2 One frequency is filtered out of the frequency comb of pulsed laser signal on photodetector. The noise in the low frequency part of the red trace is an artifact of the spectrum analyzer.

Before worrying about the drift and noise of the photodetector, a picture of the noise density and drift of the phase measurement circuit will be made. Using a signal generator, as shown in Fig. 8.3.3, the noise contribution of the amplifiers, voltage regulators, and power supplies could be studied. By mixing the signal generator's output with itself, we can see the performance of the circuit without the influence of the noise from the photodetector or from the MO signal. The formula for measuring the K_ϕ of the circuit is also shown in the figure. The K_ϕ increases when the signal is amplified. For a low gains in the last amplifier stage (LNA), the signal-to-noise ratio of the circuit's output is poor, but above a certain gain, increasing the gain further will not produce any improvement in the signal-to-noise ratio. The gain was selected so that the signal would not saturate the ADC.

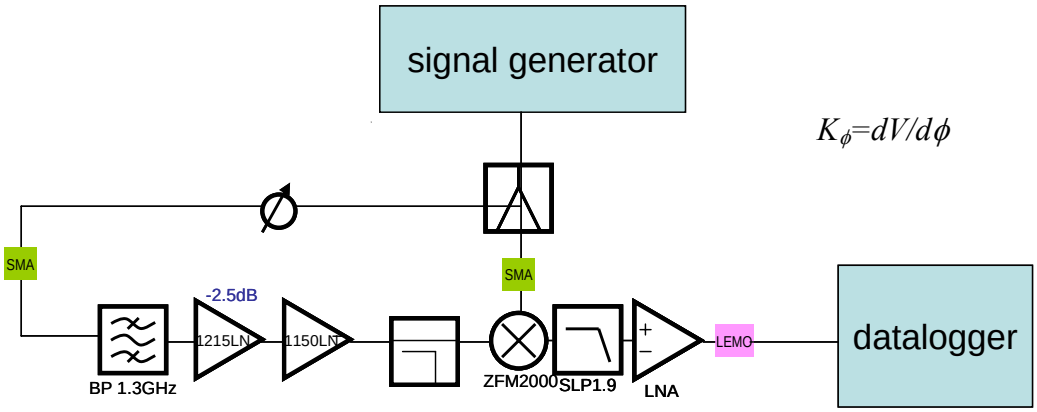

Figure 8.3.3 Setup for measurement of the mixer's K_ϕ and characterization of the spectral noise density and drift contributed by each RF component.

In Fig. 8.3.4, the spectral noise density of the signal at the exit of the low-noise amplifier (LNA) shown in Fig. 8.3.1 is plotted with a calculation of the amount of timing jitter contributed by various bandwidths. The net timing jitter of <6.5 fs is dominated by the offset frequencies above 100 kHz. Since the MLO-RF-lock bandwidth is typically only a few kHz, the noise introduced by the phase measurement will be suppressed. It will not, however, be possible at this frequency to distinguish the real timing jitter of the laser from phase detection errors. This means that the fast noise of the phase detection circuit will not limit the performance of the lock. Slow drifts must be measured separately with attention to the thermal stability of the circuit.

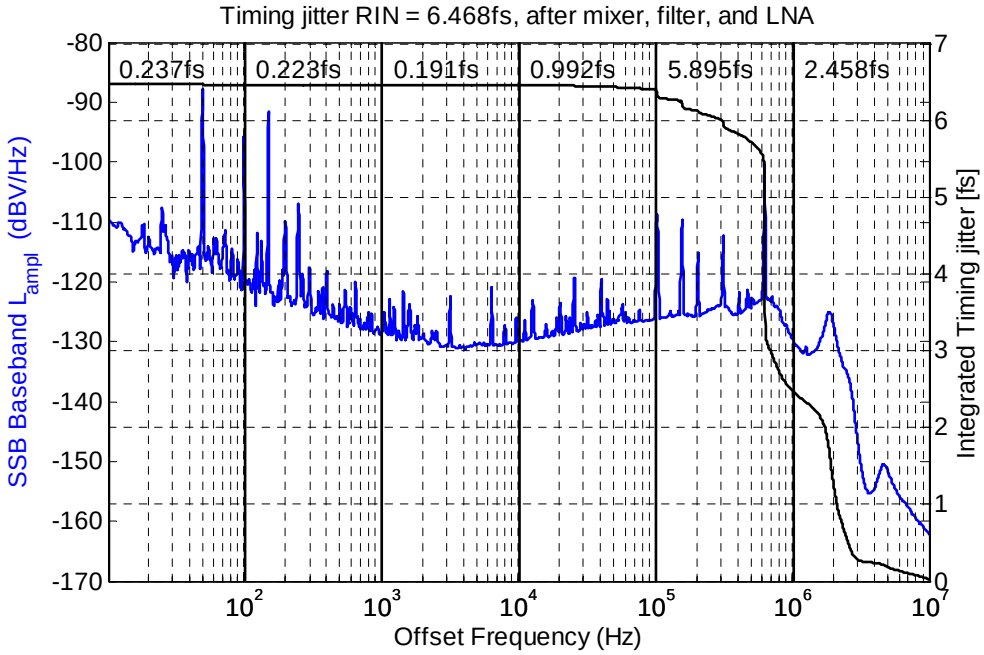

Figure 8.3.4 Spectral noise density of signal at the exit of the LNA shown in Fig. 8.3.1.

With temperature control of the RF circuit shown in Fig. 8.3.1 and amplitude control of the laser, it is conceivable that the RF lock could be stable over the long-term to below 10 fs (pkpk) without disturbances (people in the room) and 30 fs (pkpk) with disturbances. This would be the case if the lock were limited by the RF phase measurement alone (Fig. 8.3.5).

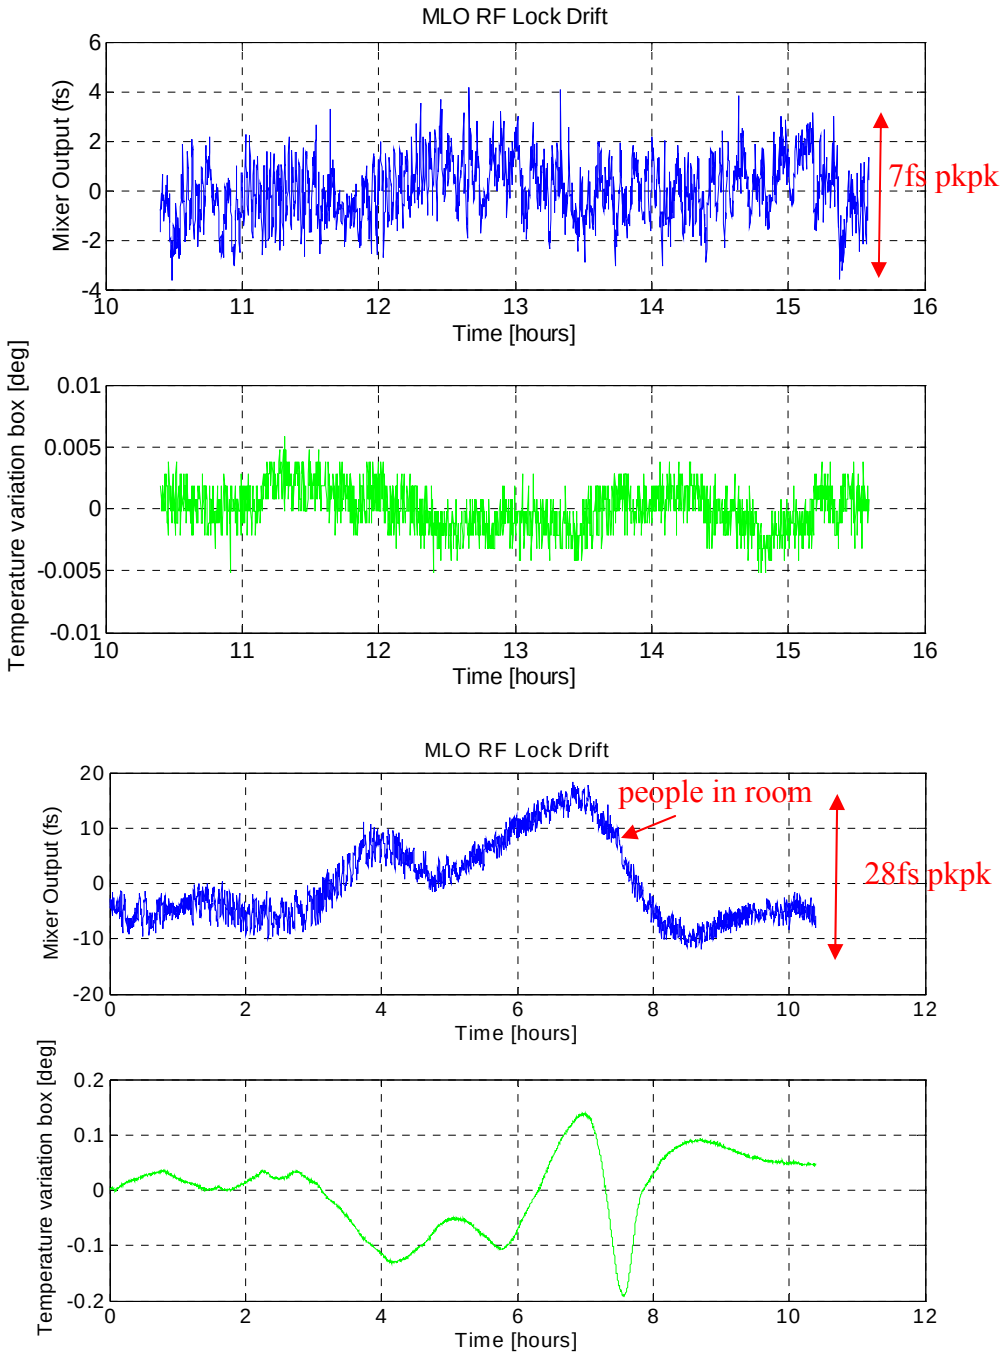

Figure 8.3.5 RF phase measurement drift with temperature control, with and without disturbances (people in room). RF signal source came from a signal generator and not from the laser.

Without temperature control, the phase measurement drift jumps to 20 fs (pkpk) (Fig. 8.3.6) without disturbances and several hundred femtoseconds with disturbances (figure not shown).

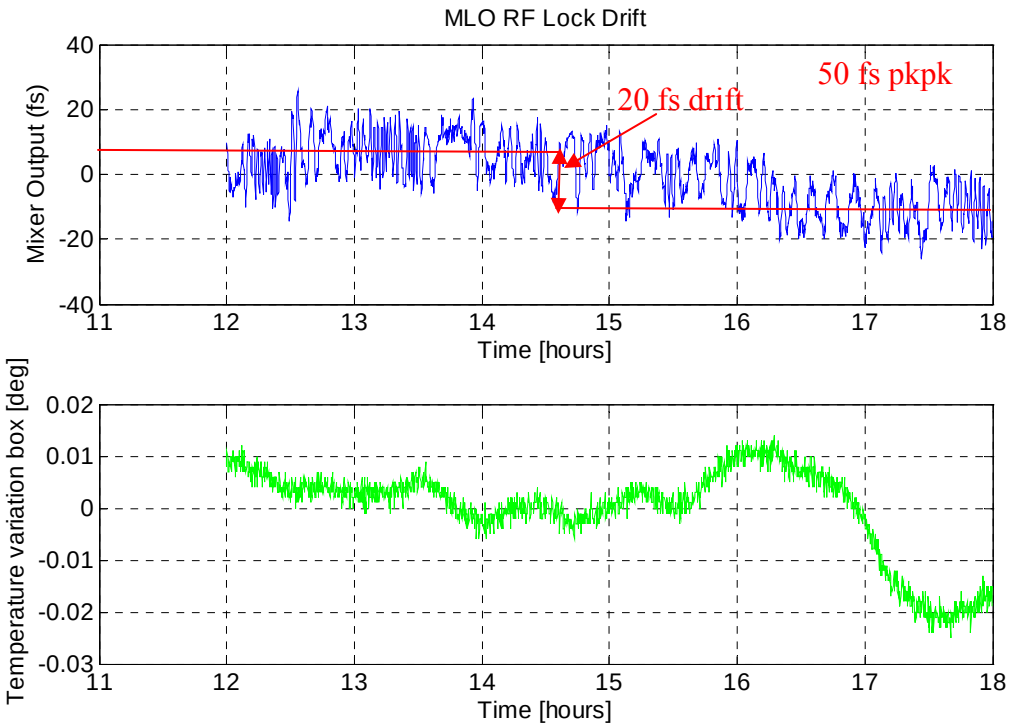

Figure 8.3.6 RF phase measurement drift without temperature control, without disturbances (people in room). Averaging over the jitter, the signal drifts by ~20 fs. Due to cables that were 2 meters longer than they needed to be, the jitter of the measurement is also larger than in the previous cases, more than doubling the jitter through cable vibrations alone.

With the RF-lock loop closed, the drift of the lock, as measured with an out-of loop measurement is 77 fs (pkpk) without disturbances (Fig 8.3.7) and several picoseconds with disturbances (figure not shown). The several picoseconds of drift that was seen when people were working in the room was due to the fact that the active temperature control system was not functioning at that time and the temperature of the room changed by a few degrees. In the measurement shown in Fig. 8.3.7, no one was in the room and the temperature of the chassis changed by about 0.1 degrees over the cours of several hours. This amount of temperature change typically causes 10 fs of phase measurement drift, (based on measurements shown in Fig. 8.3.5). At least 60 fs out of the 77 fs drift seen in the out-of-loop measurement in Fig. 8.3.7 can be accounted for by drifts of the laser amplitude. The laser amplitude drift was 0.3% (pkpk) over the course of the measurement and this would cause 60 fs of phase change in the signal emerging from a photodetector according to measurements performed in [47]. Some of the drifts observed in the out-of-loop measurement can also be accounted for by 0.03 degree change of the photodetector temperature. If the temperature of the photodetector changes by 1 degree, the phase emerging from the photodetector will change by 340 fs [48]. Based on the result from [48], 0.03 degrees photodetector temperature change would cause 10 fs of phase drift.

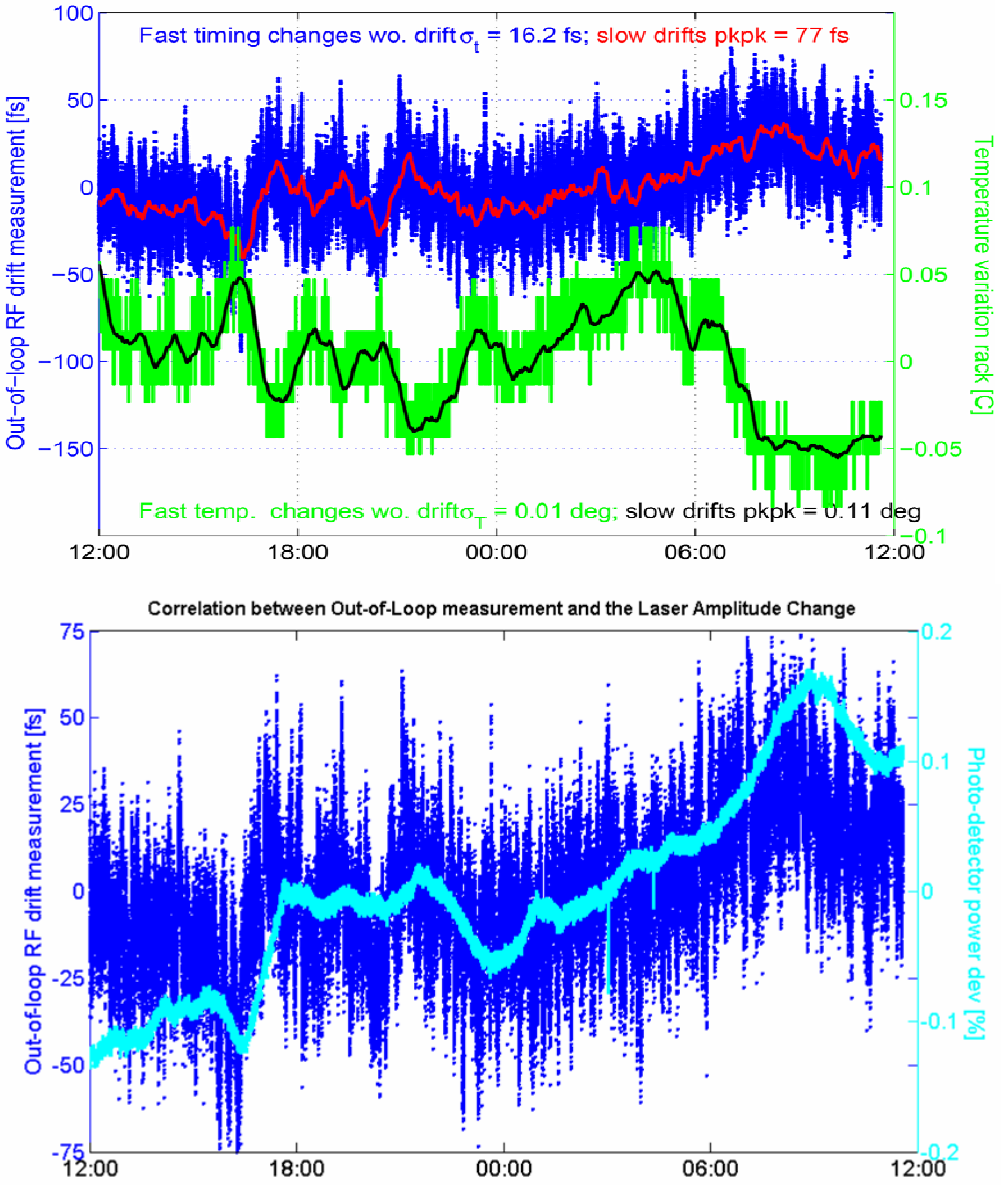

Figure 8.3.7 Out-of-loop measurement drift without temperature control and without disturbances (people in room). Laser amplitude drift during same time period.

The conclusion from this out-of-loop measurement is that with active control of the laser amplitude and of the temperature within the chassis, the RF-lock can, in principle, be made stable within 10 fs (pkpk). A DSP feedback on the laser amplitude is in the planning stages and a new, commercial MLO (OneFive) has been purchased and in preliminary tests, it had a good amplitude stability even without active feedback on the amplitude. This stability results from the sealing of the chassis containing the laser, making it insensitive to humidity changes, air-currents, and small temperature changes. All of the components in the laser-based synchronization system could benefit from such packaging.

For a frame of reference concerning the costs of these systems, the cost of an MO and an MLO with optical table and control hardware is presented in table 8.3.1. The largest cost for both MLO and MO systems is for the infrastructure, including: racks, laser table and climate control. It is possible to invest millions in climate control. PSI, for example, has made a considerable investment to stabilize the temperature of the entire 500 meter long linac within 0.1 degrees C. This reduces the drift problems associated with RF components and, compared to the optical synchronization system, represents a very different strategy to achieve good reference signal stability.

MLO-MO system	
OneFive laser	45,000
Free-Space Optics	25,000
Laser hutch	100,000
Control hardware	7,000
RF lock	5,000
MO	250,000
Total (EUR)	432,000

Table 8.3.1 The cost of an MO in a rack and an MLO with optical table and control hardware. MLO numbers come from Holger Schlarb and MO numbers come from Henning Weddig.

There is a cheaper way to do optical synchronization using CW (continuous wave) lasers and one such system was built and fully commissioned at GSI by Michael Bousonville; a similar concept was developed as a prototype at DESY by Matthias Felber. A more complicated CW optical synchronization system has been implemented at LCLS. Such systems have been shown to achieve ~50 fs synchronization, but not sub-10 fs synchronization. They are ideal for situations that involve long distances that make RF-cable costs prohibitive and where sub-10 fs synchronization is not required.

9 Synchrotron Light Monitors

Synchrotron light detectors can enable a high resolution beam position measurement after the second bend of the bunch compressor. Existing synchrotron light monitor systems at FLASH are described with respect to their limitations and capabilities.

9.1 Profile monitors

A screen and camera positioned after the third bend of the chicane can detect the synchrotron radiation produced at the third bend of the chicane (Fig. 9.1.1).

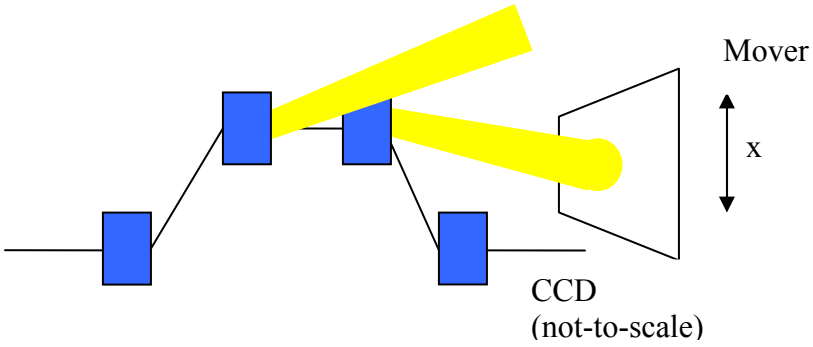

Figure 9.1.1 A synchrotron light monitor system with CCD.

The synchrotron light is detected with a CCD camera and when the pixels from the region of interest around the beam are summed together along the vertical axis, changes in the beam position can be measured by fitting a line to the sharp rising edge of the profile and checking for changes in the horizontal position of the zero crossing of the line

(Fig. 9.1.2). This method proved to be more accurate than any method taking the peak or centroid of the distribution.

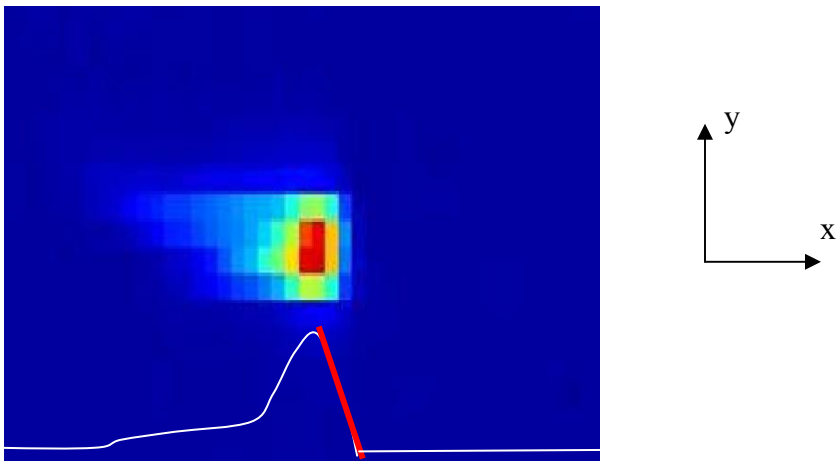

Figure 9.1.2 A picture of the beam as imaged with the synchrotron light camera. A projection of the image is shown below along with a line fit to the steep rising edge. Changes in the zero-crossing of this line give a 10 μm resolution measure of the beam position.

This method produces beam position measurements with 10 μm resolution [52]. The dynamic measurement range of the device is several millimeters and if the beam moves outside of this range, a motor must be adjusted and the beam must be manually centered on the CCD. The other problem with the monitor is that it does not have single-bunch resolution of the entire bunch train. An MCP (Micro-Channel Plate) can enable the selection of one bunch from the bunch train for analysis, but the information from the rest of the bunch-train is lost.

9.2 Photomultiplier tube monitors

Two photomultiplier tubes placed 1.7 m away from the second bend of the fist bunch compressor can provide a beam position measurement with a resolution that is slightly worse than that of the profile monitor. Because the photomultiplier signals are short in duration and are sampled with an 81 MHz ADC, they can be used to produce single bunch resolution across the whole bunch train, while the profile monitor cannot.

The Hamamatsu R5900U-00-M4 photomultipliers used in the monitor come from a HERA-B experiment and they consist of 4 photomultiplier tubes packed together in one square package with dimensions of 16x16 mm^2. Each individual tube is 8x8 mm^2 and there is 10-20% cross-talk between the tubes, meaning that photons that are measured on one side will also be partially registered by the other side. The signal from each tube is amplified and filtered with a Gaussian filter in order to give the signal a more rounded peak and increase its width. This reduces the impact of the clock jitter from the ADC used to sample the signal [53]. The synchrotron light is centered on a pair of the four tubes, such that each gets half of the light. If the beam moves to one side or the other, the signal from one side will increase and the signal from the other will decrease. If the beam

is moved away from the center of the two detectors, a motorized stage will move the detectors to where the beam is (Fig. 9.2.1).

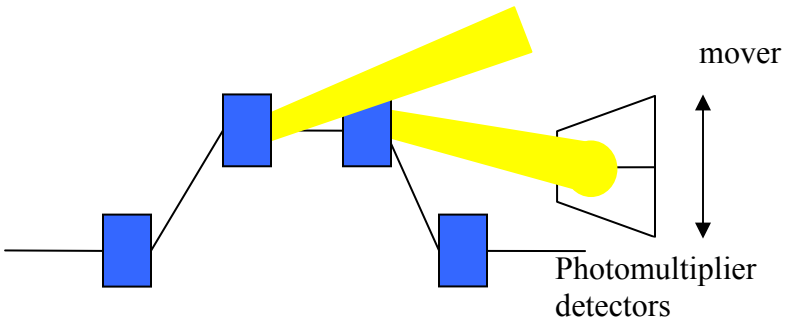

Figure 9.2.1 Two Photomultipliers used to measure the beam position in the chicane.

In this application, if the measurement resolution is limited by the number of photons that each tube detects, it is shot-noise limited. If a measurement is shot-noise limited, no matter how much a signal is amplified, the signal-to-noise ratio will not improve. To estimate the resolution of this measurement, one must first know how many photons are intercepted by the monitor.

Synchrotron radiation is emitted over a wide range of frequencies. Different frequencies are emitted with different angular distributions, but most of the power is present in an opening angle of $1/\gamma$. To accurately calculate the radiation produced in the frequency range to which the photo detector is sensitive and within the angular spread determined by the aperture of the optics leading up to the photodetector, one must integrate over the number of photons radiated per unit frequency per unit solid angle. The unit solid angle can be written in terms of its components: the emission angle and the bending angle. One can use the following formula to find the number of photons, N_λ emitted over an emission angle, θ, and a bending angle, ψ [36]

$$\frac{dN_\lambda(\omega)}{d\theta d\psi} = C_\Omega Q E^2 \frac{\Delta\omega}{\omega}\frac{\omega^2}{\omega_c^2} K_{2/3}^2(\xi)(1+\gamma^2\theta^2)\left[1+\frac{\gamma^2\theta^2}{1+\gamma^2\theta^2}\frac{K_{1/3}^2(\xi)}{K_{2/3}^2(\xi)}\right] \quad (9.2.1)$$

Where C_Ω = 1.3273e16 photons / (sec mrad^2 GeV^2 A), K is a Bessel function, Q is the beam charge, E is the beam energy, γ is the Lorentz factor, ω is the angular frequency of the radiation, and the critical frequency, ω_c, is

$$\omega_c = C_\omega \frac{E^3}{\rho} = 411.6\text{nm} \quad (9.2.2)$$

where C_ω = 3.37e18 m/(sec GeV^3) and $\rho = l_{eff}/\sin\psi$ is the bending radius. Above the critical frequency, the intensity of the radiation falls off exponentially.

If one integrates over the full emission angle, $\theta = 1/\gamma$, and bending angle, $\psi = 18°$, one is left with the number of photons emitted by a bunch as a function of frequency, ω,

$$N_\lambda(\omega) = \psi C_\psi QE \frac{\Delta\omega}{\omega_c} \frac{9\sqrt{3}}{8\pi} \int_{\omega/\omega_c}^{\infty} K_{5/3}(x)dx \tag{9.2.3}$$

where C_ψ = 3.967e16 photons / (sec mrad A GeV) [36].

The Hamamatsu R5900U-00-M4 photodetectors have a quantum efficiency of greater than 10 percent in the frequency range between 300 and 500 nm. Between 250 and 550 nm, the efficiency is greater than 5 percent. This is centered about the critical frequency of the synchrotron radiation of 411 nm. Using Eq. 9.2.3 with charge Q = 1e-9 C and energy E = 0.130 GeV, the total number of photons in the frequency range to which the photodetector is sensitive is then

$$N_\lambda(h\upsilon \in [250\text{nm} \longrightarrow 550\text{nm}]) = 9 \cdot 10^6 \tag{9.2.4}$$

This does not take into account the number of photons which are cut away by the aperture of the optics leading up to the photodetector. To find that number, one must integrate over a smaller solid angle. For an aperture of 13.6 mm at a distance of 1.6 mm from the source, the number of photons that reaches the detector is 7e6 between 250 and 550 nm and 5e6 between 300 and 500 nm. One can then estimate that the maximum number of photons that could be detected with a quantum efficiency of 10% is

$$N_D(h\upsilon \in [250\text{nm} \longrightarrow 550\text{nm}]) = 6 \cdot 10^6 \tag{9.2.5}$$

A cross-check of this result was provided by code from [55]. After the aperture, the beam diverges with an opening angle proportional to λ^2 over distance. This distance is small, but some loss of photons is still a possibility.

The shot-noise of the photo emission process is given by

$$\sigma_{N_D} = \sqrt{\mathrm{N}_D} \tag{9.2.6}$$

and the photomultiplier signal is shot-noise limited if it fluctuates by less than this amount. To determine the shot-noise limitation of the measurement, we must calculate the resolution of the monitor in terms of the shot-noise an outline of this derivation was provided by [57]. To begin, we write the normalized beam position sensitivity,

$$\mathrm{F} = \frac{\mathrm{S}_+ - \mathrm{S}_-}{\mathrm{S}_+ + \mathrm{S}_-} = \frac{\mathrm{P}\int_{x_0}^{\infty} \rho(x)dx - \mathrm{P}\int_{0}^{x_0} \rho(x)dx}{\mathrm{P}} \tag{9.2.7}$$

in terms of the probability of a photon being detected: $S_{\pm}$=N_D**quantum efficiency*. The probability of a photon being detected is also equal to the probability of a single electron emitting a photon that is detected, P, times an integral over a portion of the bunch's charge distribution, ρ. The limits of the integrals are written such that half of the beam is detected by one detector and the other half of the beam is detected by the other. The beam

position for which the beam is centered relative to the two photodetectors is written as x_0. Let us assume that there is no change in the beam profile but there is a small change of the position of the beam, Δx. We want to know how this affects the sensitivity of the monitor, so we take the derivative of the sensitivity with respect to a small change of beam position

$$\frac{dF}{d\Delta x} = -2\int_{-\infty}^{x_0} \rho(x' + \Delta x)dx' = -2\int_{-\infty}^{x_0} \rho' dx' = -2\rho(x_0 + \Delta x) \tag{9.2.8}$$

If ρ is a Gaussian distribution,

$$\rho(x_0 + \Delta x) = \frac{1}{\sqrt{2\pi}\sigma_x} \cdot e^{-\frac{1}{2}\left(\frac{x - x_0 - \Delta x}{\sigma_x}\right)}, \tag{9.2.9}$$

and then

$$\frac{dF}{d\Delta x} \le \frac{2}{\sqrt{2\pi}\sigma_x}. \tag{9.2.10}$$

Since the horizontal width of the beam is dominated by the energy spread, this means that the sensitivity function decreases for increased energy spread regardless of the detector size.

The question that we want to answer is: how big is the uncertainty of the sensitivity function when the beam position is constant; is it limited by the shot noise? Because S_+ is statistically independent of S_-,

$$\sigma_F = \sqrt{\left(\frac{\partial F}{\partial S_+}\right)^2 \sigma_{S_+}{}^2 + \left(\frac{\partial F}{\partial S_-}\right)^2 \sigma_{S_-}{}^2} \tag{9.2.11}$$

Is the uncertainty of the sensitivity function and at the point of highest resolution, $S=S_+=S_-$, so

$$\frac{\partial F}{\partial S_+} = \left[\frac{1}{S_+ + S_-} - \frac{S_+ - S_-}{(S_+ + S_-)^2} + ..\right]_{S_+ = S_-} \approx \frac{1}{2S}. \tag{9.2.12}$$

Substituting Eq. 9.2.12 into Eq. 9.2.11, we get

$$\sigma_F = \sqrt{2 \cdot \frac{1}{4}\frac{\sigma_S^2}{S^2}} = \frac{\sqrt{2}}{2} \cdot \frac{\sigma_S}{S}. \tag{9.2.13}$$

Finally, we can use this, together with Eq. 9.2.10, to write the uncertainty of the position in terms of the uncertainty of the sensitivity for a Gaussian beam,

$$\sigma_{\Delta x} = \sigma_F \cdot \frac{1}{\frac{dF}{d\Delta x}} = \frac{\sqrt{2}}{2} \cdot \frac{\sigma_S}{S} \cdot \frac{\sqrt{2\pi} \cdot \sigma_x}{2} = \sqrt{\pi}\sigma_x \cdot \frac{\sigma_S}{2S}. \qquad (9.2.14)$$

Now, since $\sigma_S = \sqrt{S}$ and $S = N$, the number of photons detected by one detector, i.e. half of the total number of photons N_D times the quantum efficiency of the detector.

$$\sigma_{\Delta x} = \sqrt{\pi}\sigma_x \frac{\sqrt{N}}{2N} = \frac{\sqrt{\pi}\sigma_x}{2} \frac{1}{\sqrt{N}} \qquad (9.2.17)$$

where

$$N = \frac{1}{2}(N_D \cdot quantum\_efficiency) = 3 \cdot 10^5$$

Substituting N into Eq. 9.2.17, the uncertainty of the position is given by the beam width times 0.0015. Taking into account the ~10% cross-talk of a pair of photomultiplier tubes, this gives an estimate for the position uncertainty ranging from 7 μm for a 4 mm wide, on-crest beam to 34 μm for a 2 cm wide, off-crest beam. Since the resolution of the monitor, as measured by comparing the relative jitter of neighboring bunches and by comparing the measurement to that of the synchrotron light camera, is 4e-5 (15 μm) for an on-crest beam and 9e-5 (30 μm) for an off-crest beam [57], only slightly worse than the best-case resolution predicted by photon statistics alone, it is fair to judge that this resolution is limited primarily by shot-noise.

The monitor has been calibrated with dipole scans, accelerator gradient scans, and motor scans. Ideally all three calibrations would match, with the dipole scan being the gold standard. While the three calibrations were sometimes within the error bars of one another, differences of 10-20% between the three different calibration routines have frequently been observed. It is suspected that this is due to a high level of energy jitter and drift that impact the repeatability of the calibrations [57].

10 Energy Measurement Benchmarking

The chicane BPM (Ch. 7), the arrival-time measurements around the chicane (Ch. 8), the photomultiplier monitor (Ch. 9), and the out-of-loop vector sum (Ch. 3) all provide measurements of the beam energy. A table comparing these measurements is shown below. The monitors constructed during this thesis are highlighted in blue.

Measurement System	**Position resolution**	**Energy resolution**	**Dynamic range**
In-loop Vector Sum (drifts)	(25 um to) 70 μm	7e-5 to 2e-4 $\pm$ 1e-2	10 cm
Out-of-loop Vector Sum (drift-free)	(25 um to) 70 μm	7e-5 to 2e-4	10 cm
BC2 BPM 1.3 GHz front-end	25 μm	7e-5	80 mm
Photomultiplier Tube Monitor	15 μm to 30 μm	4e-5 to 9e-5	2 mm
BC2 BPM 10.4 GHz front-end	(6 μm to) 10 μm	2e-5	2 mm
BC2 BPM optical front-end	2 μm	6e-6	1 mm
time-of-flight with 2 BAMs	(9 fs)	(1.5e-5) anticipated	1 mm

Table 10.0.1 Comparison of beam energy measurements in the first bunch compressor.

The highest energy resolution (ΔE/E=6e-6) is provided by the optical front-end of the chicane BPM, but it also has the smallest dynamic range. In the following sections, the measurements from each of the monitors listed above will be compared with respect to their reliability and agreement.

10.1 RF BPM Measurements

Calibration scans of the 1.3 GHz (coarse) and 10.4 GHz (fine) front-ends were presented in Chapter 7 and long-term measurements are presented below in Figs. 10.1.1 and 10.1.2. They took place over the course of several hours during relatively quiet SASE operation shifts. In Fig. 10.1.1, the setpoint of the first accelerating module is shown in black, while the coarse measurement is shown in green and the fine measurement is shown in blue. The y units are in percent energy change and the x units are in hours. The measurements of the BPM are converted into percent energy change by multiplying the beam position times the R_{16} of the chicane, in accordance with Eq. 1.10. The measurement was done over the coarse of 17 hours with both trombone and vector modulator feedbacks on, meaning that both feedbacks attempted to keep the system measuring at the zero-crossings of the two signals. The position changes measured by the trombone are subject to 100 μm errors from mechanical hysteresis and backlash (Fig. 7.2.7), nevertheless, for this 10.4 GHz measurement, the trombone changes appear to have been appropriate, since the changes measured by 10.4 GHz measurement match those of the gradient setpoint and those of the 1.3 GHz measurement. During the eight hour period in the right-hand figure, over which no gradient setpoint changes were made, it appears that the gradient regulation was drifting by 0.1 %. In the left-hand figure, it drifts by as much as 0.2% before operators react. Not all machine energy changes depend on the gradient setting of the first accelerator section, however. A change in the phase of the gun RF impacted the left-hand measurement around hour 32.

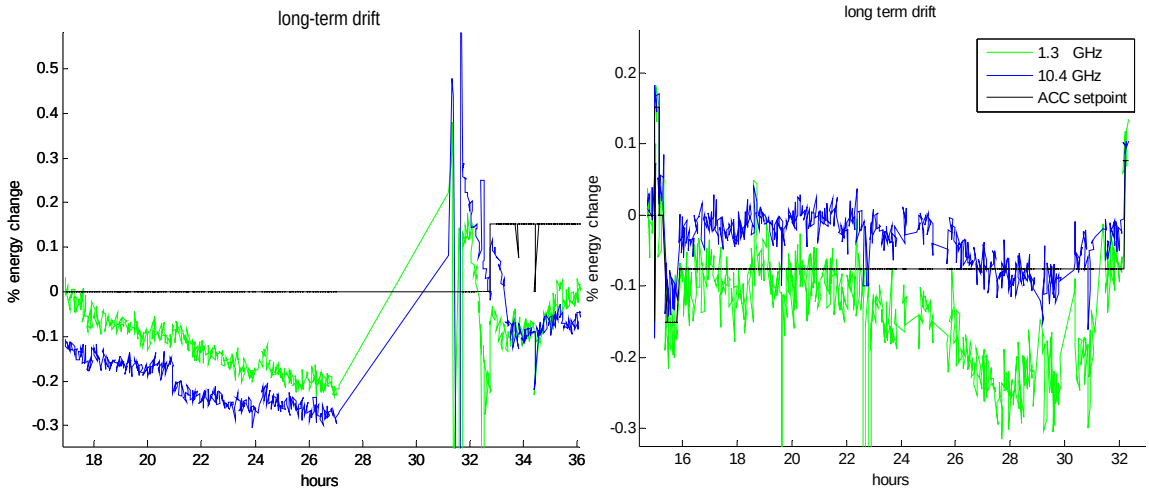

Figure 10.1.1 Measurements of energy stability in the chicane taken by the coarse and fine HF front-ends of the chicane BPM plotted with energy setpoint values from the upstream accelerating module.

10.2 Photomultiplier Tube Monitor

The photomultiplier tube monitor (PMT) provides measurements of the beam position that have a resolution that two- to six-times that of the 10.4 GHz front-end of the chicane BPM. In Fig. 10.2.1, the chicane BPM is labeled as EBPM (Energy BPM). The arrival-time is plotted in blue and the beam position is plotted in red.

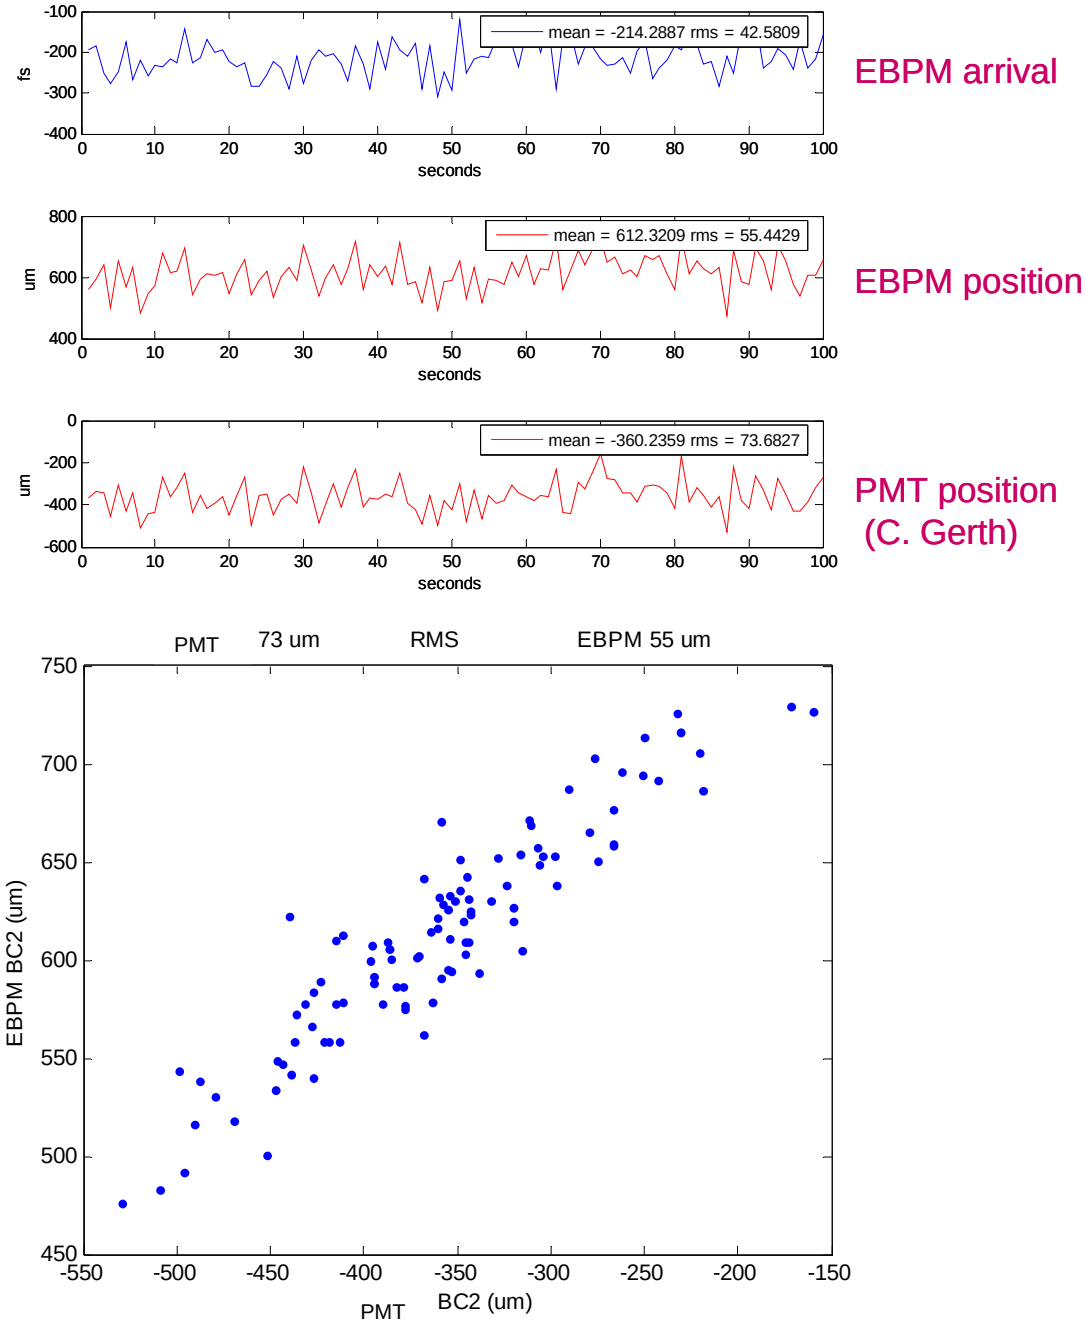

Figure 10.2.1 Correlation between the measurements of the beam position in the chicane taken by the chicane BPM (labeled EBPM) and the photomultiplier tube monitor (PMT).

Good agreement between the RF chicane BPM and the PMT BPM on a shot-to-shot basis was observed until multiple new ADCs were installed in the crate in which the chicane BPM ADC is installed. These new ADCs sent large volumes of data over the crate BUS and caused all of the devices in the crate to suffer from buffer number problems. This made it impossible to correlate any data from this crate with other devices in the machine. This was not yet a problem when the data from Fig. 10.2.1 was taken, but for all subsequent data, it was a problem.

When the PMT BPM and RF chicane BPM measurements don't agree, the reason is frequently that both measurements have a limited, few-millimeter dynamic range and when it is exceeded, a motor must be moved and the monitor must be re-calibrated. Whenever the calibration constants for these monitors are not correct, the measurements do not agree. The results from several-day measurements of the 1.3 GHz (coarse) and 10.4 GHz (fine) BPM front-ends are presented below in Fig.s 10.2.2 and 10.2.3. In the first plot, good agreement is observed between the BPMs and with the gradient setpoint. In the second plot, poor agreement is observed and was not identified by the control software.

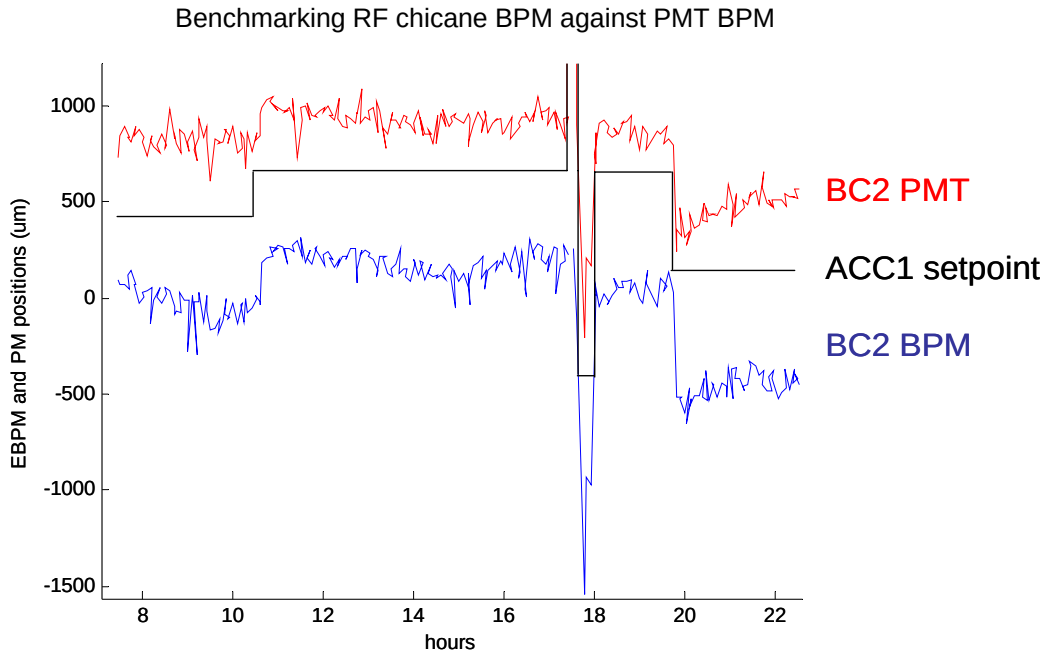

Figure 10.2.2 10.4 GHz chicane BPM front-end measurement and photomultiplier tube (PMT) BPM measurement in good agreement.

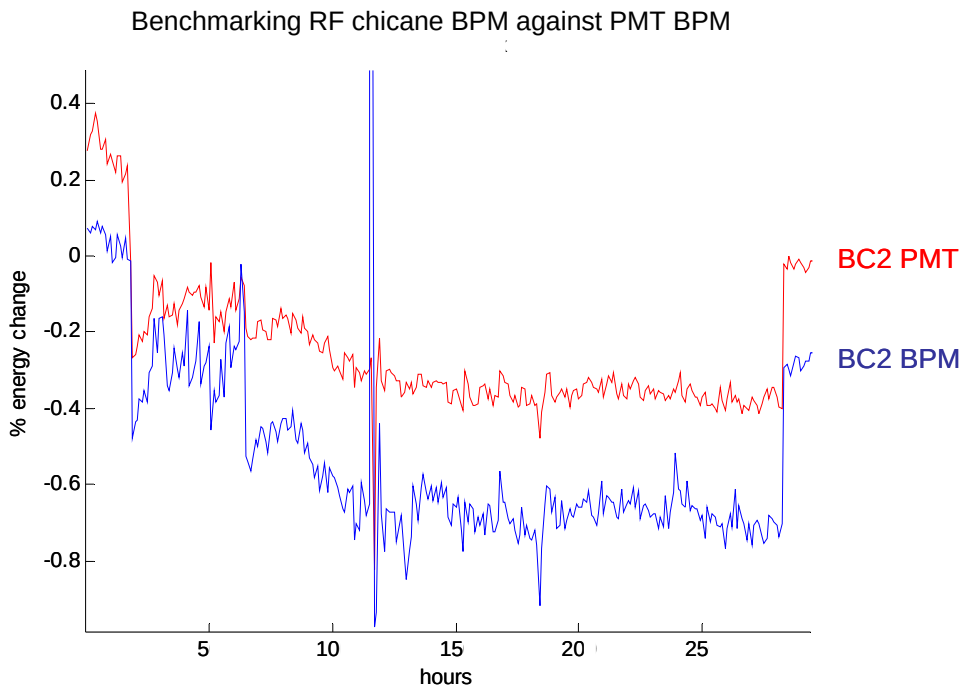

Figure 10.2.3 10.4 GHz chicane BPM front-end measurement and photomultiplier tube (PMT) BPM measurement in poor agreement. One or both calibrations are wrong.

In Fig. 10.2.3, either the calibration constant of the PMT BPM is too small or the calibration constant of the RF chicane BPM is too large. In general, the PMT BPM is calibrated by scanning the ACC1 gradient setpoint and the RF chicane BPM is calibrated by scanning an RF reference with a vector modulator. As was described in Section 7.2, this sort of calibration is subject to errors due to reflections in the pickup, unless the effect of the reflections is measured and removed. The PMT BPM has also been calibrated by scanning a the position of a stage, but when comparisons of calibrations done with ACC1 gradient setpoints, dipole current, and stage position, differences of 10-20% were observed due to a high level of energy drift that impacts the repeatability of the calibrations [56]. The result is that neither monitor provides a high level of confidence in the calibration. The author's preference would be for using the calibration of the PMT BPM done with a motor scan as a benchmark. This was not available during the studies done here.

Progress was made in terms of writing software to keep the monitors calibrated all of the time. But, because of the complexity of the algorithm required, the 10.4 GHz measurement will not be immediately incorporated into a beam energy server. The 1.3 GHz measurement, however, does not require constant adjustment and recalibration. Because of its relative reliability, it has been incorporated into an energy server that will be easy for the operators to use.

10.3 Out-of-loop Vector Sum

The out-of-loop vector sum should provide a measurement of the drifts of the in-loop vector sum (the gradient setpoint) which result from drifts of the downconverters [21]. It does not provide a measurement of drifts which occur on the long cables. Cable drifts are a concern for phase stability but less of a concern for amplitude stability, the measurement in question here. The out-of-loop vector sum should also provide energy change resolution that is comparable to that of both the PMT monitor and the RF fine front-end of the chicane BPM. It appears, however that it is subject to errors to which the PMT monitor and chicane BPM are not subject. In Fig. 10.3.1, there is a jump in the out-of-loop vector sum that is not seen on either the PMT monitor or the chicane BPM. There is also better agreement between the PMT monitor and BPM for the first setpoint change which took place just prior to the 20th hour. In the 3rd setpoint change at the 40th hour, it appears that the chicane BPM has exceeded its dynamic range and there is better agreement between the out-of-loop vector sum and the PMT monitor.

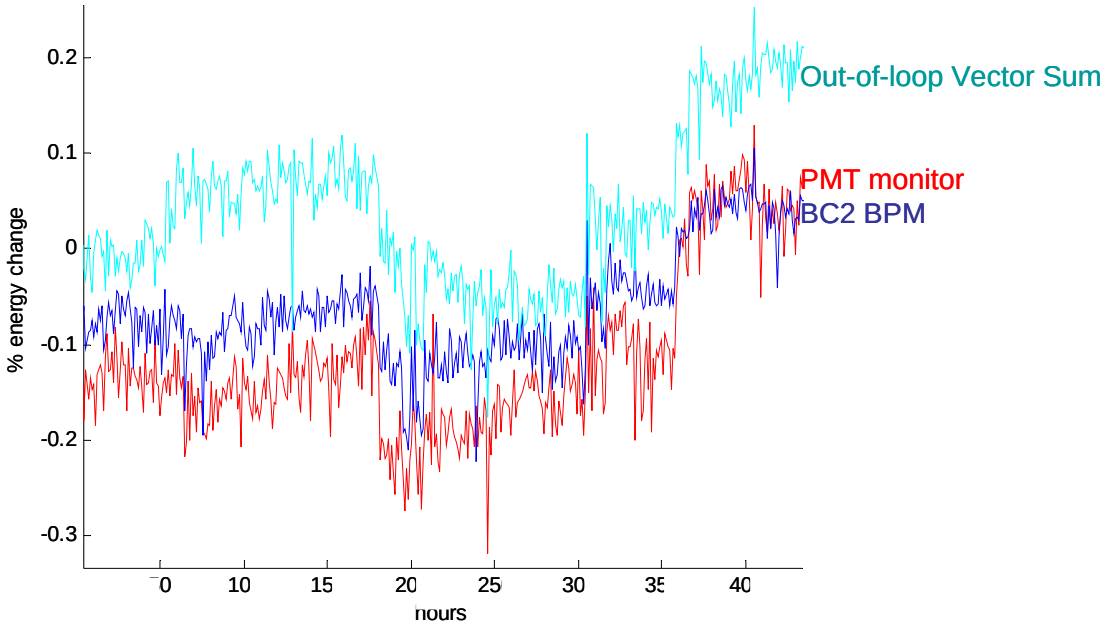

Figure 10.3.1 Fine HF front-ends position measurement and photomultiplier tube position measurement.

10.4 Optical BPM Measurements

The optical front-end was only in operation for a few days over the course of the last machine run and it was possible to calibrate it and check the resolution of the measurement. There were two separate optical front-ends constructed. One had an older version of the optical delay lines that have a limited lifetime. It did not have an active temperature control system and it was installed outside of the tunnel after a 30 meter long

cable connecting the pickup to the EOMs in the front-end. It measured 12 fs arrival-time resolution and 4 μm position resolution. The other front-end had more robust delay lines and an active temperature control system. It was installed in a lead shielded box with 2 meter cables connecting the pickup to the EOMs in the front-end. For the same beam conditions as in the out-of-tunnel case, this in-tunnel front-end measured 7 fs arrival-time resolution and 2 μm position resolution. The resolution is calculated by multiplying the accuracy with which the laser pulse amplitudes can be detected by a measurement of the slope of the pickup signal. This measurement of the slope is done by scanning the arrival-times of the laser pulses over the pickup signal zero-crossing and measuring how much the amplitude of the laser pulse changes.

The arrival-time of the pickup signal coming out of the pickup output on the inside of the chicane should change in proportion to the energy deviation times $R_{56}/2-R_{16}$ and on the outside of the chicane it should change in proportion to $R_{56}/2+R_{16}$. Facing in the direction in which the beam travels, the left output of the pickup is on the outside of chicane and the right output of the pickup is on the inside of chicane. The ratio of the change of the arrival time of the left side to the change of the arrival time of the right side should be equal to the ratio of $(R_{56}/2+R_{16})/(R_{56}/2-R_{16})$. A measurement of this ratio constitutes a check of the calibration of the arrival-time measurements done on both right and left outputs of the pickup. Such a measurement is shown in Fig. 10.4.1 and good agreement with the expected ration is observed despite the large jitter and drift of the beam.

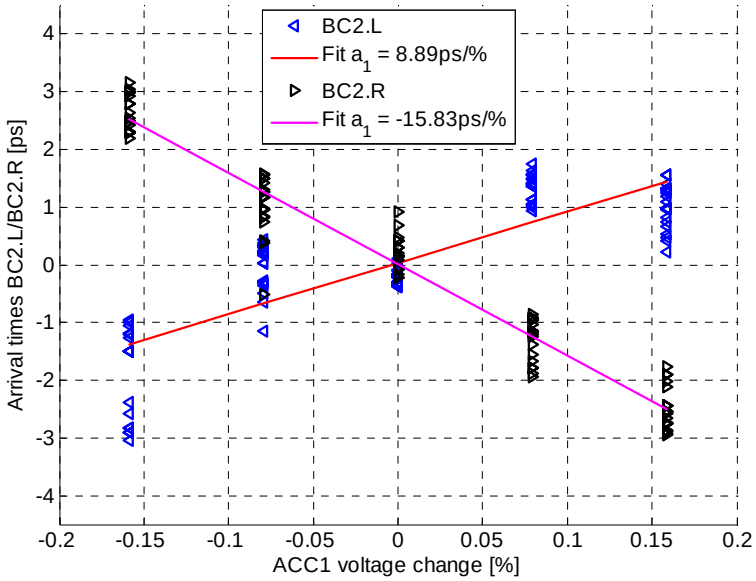

Figure 10.4.1 Verification of the calibration of optical front-end beam position measurement. The ratio of the slopes of the energy dependent arrival-time change of the signal from the left side of the pickup to that coming from the right side of the pickup should be equal to the ratio of $(R_{56}/2+R_{16})/(R_{56}/2-R_{16})$.

Both calibrations could be still be wrong in the same proportion, say, by a factor-of-two-error, or so, but checking the measurements relative to an external reference can rule that

out. This was done in two ways: by changing the setpoint of the accelerating gradient and checking that the beam position changed by an expected amount (Fig. 10.4.2, 10.4.3) and by deriving a measurement of the beam arrival-time upstream of the chicane from the beam position and arrival-time and comparing that to a measurement of the beam arrival-time measured upstream of the chicane with a button-type BAM monitor (Fig. 10.4.4).

In Fig.s 10.4.2 and 10.4.3, the accuracy of the chicane BPM optical front-end was checked by changing the energy of the beam with the accelerating gradient setpoint and measuring how much the beam position changed. In Fig. 10.4.2, the beam energy was changed by 0.3 % and the energy measured with the chicane BPM changed by a comparable amount. In Fig. 10.4.1 the beam energy measurements done with the optical front-end of the chicane BPM are compared to those done with the PMT monitor, and for a time-of-flight energy measurement done with two BAMs: one before the first bunch compressor and one after the last bunch compressor.

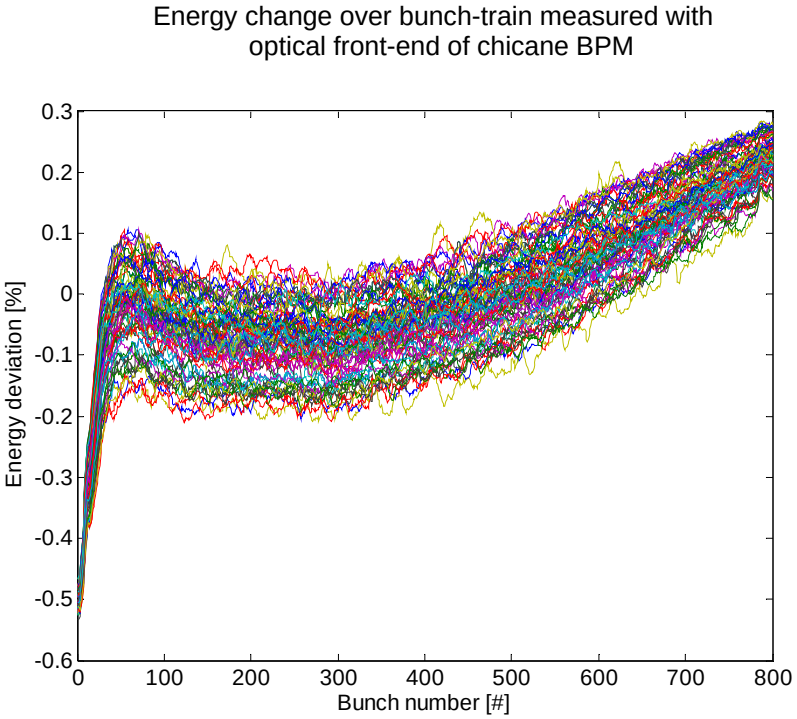

Figure 10.4.2 The beam energy was changed by 0.3 % with the accelerator gradient setpoint and the beam energy measured by the chicane BPM changed by a comparable amount. The beam was outside of the measurement range of the BPM at the end of the bunch train.

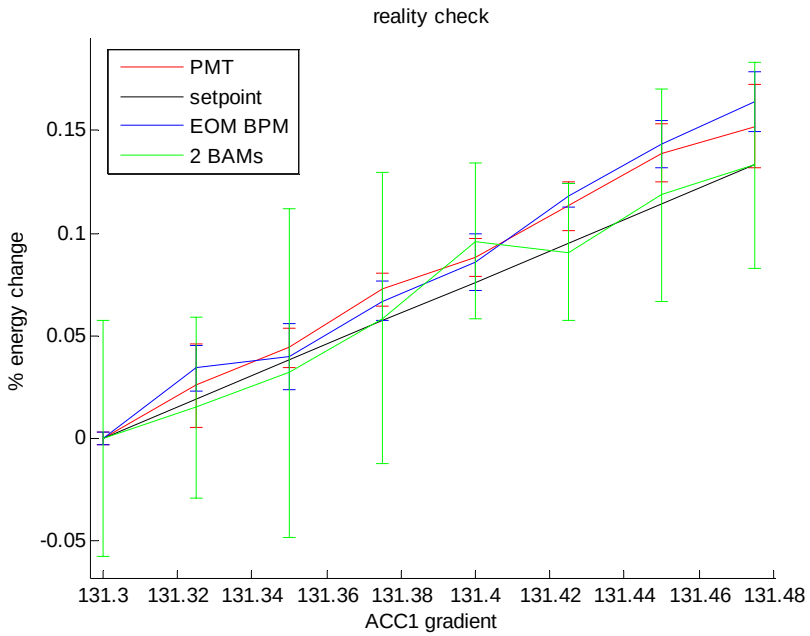

Figure 10.4.3 Optical (EOM) front-end chicane BPM measurement and photomultiplier tube (PMT) BPM measurement along with a time-of-flight measurement involving two BAMs and a line showing how the setpoint of the accelerating gradient predicted a beam energy change of 0.1%. The other measurements showed an energy change of 0.15%.

This time-of-flight measurement with two BAMs was unsuccessful due to the off-crest operation in the second accelerator section. The large jitter seen in the two-BAM measurement is due to arrival-time jitter about the slope of the RF in the second accelerator section. This jitter is not measured by monitors in the first bunch compressor. Because the machine program was dedicated to another experiment, only two sample points were taken for each step in the scan of the gradient in Fig. 10.4.3.

Another cross-check of the measurements done by the optical front-end of the chicane BPM is provided by the BAM located upstream of the chicane. The arrival-time upstream of the chicane, $t_{upstream}$, in terms of the quantities measured by the chicane BPM (x, $t_{chicane}$), is

$$t_{upstream} = \frac{R_{56}}{R_{16}} x - t_{chicane} .$$

This quantity measured with the chicane BPM is compared to the beam arrival time measured with a BAM located upstream of the chicane in Fig. 10.4.4.

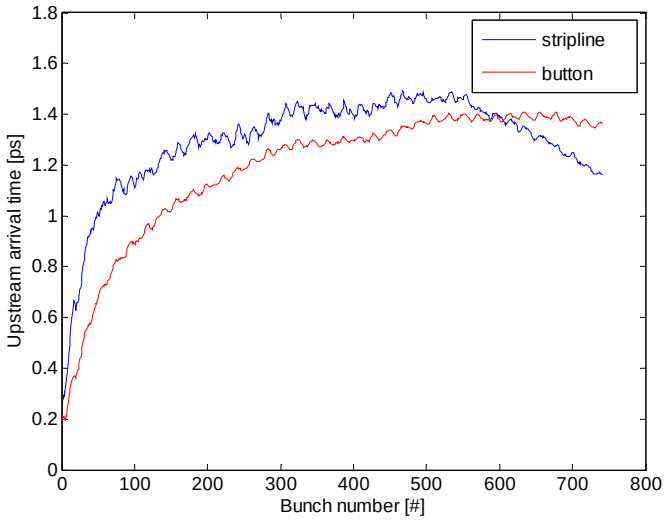

Figure 10.4.4 The beam-arrival time upstream of the chicane measured with both the transversely mounted stripline BPM installed in the chicane and with a button-type pickup BAM installed upstream of the chicane.

The ripples observed on both the BAM and BPM arrival time measurements are due to a 50 kHz oscillation on the gun phase. Both measurements were averaged over 20 shots and the more spiky texture of the BPM arrival time measurement results from problems with buffer numbers; the buffer number from one side of the pickup was not equal to the buffer number from the left side of the pickup in about 20 % or more of the cases. It appears that the buffer number problem has since been solved.

The question remains: which monitoring system should one believe? Several of the monitors described in this and past sections are plotted together below in Fig. 10.4.2. Buffer number problems aside, for all of the different monitoring systems to demonstrate agreement, they must all be properly calibrated and within the dynamic range of the measurement. Because the dynamic range of the higher resolution measurements is only a millimeter or two, these requirements are not always fulfilled and the software that is needed to diagnose whether or not the measurement is accurate becomes more complicated. As the stability of the machine is improved, the higher resolution measurements will become more interesting and valuable.

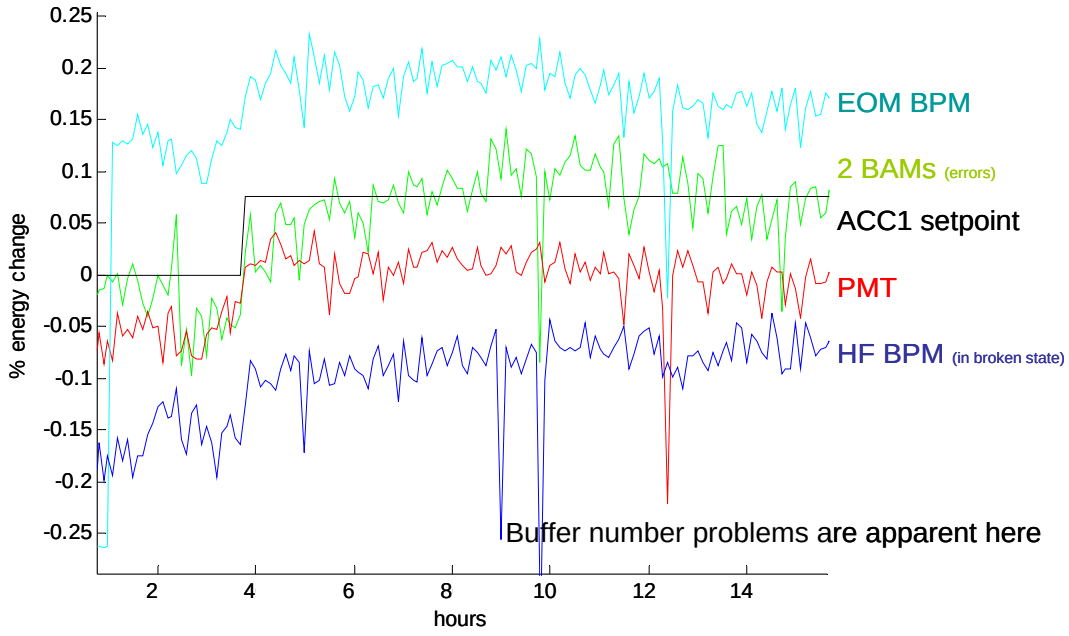

Figure 10.4.5 Optical (EOM) front-end position measurement, 10.4 GHz front-end measurement, photomultiplier tube position measurement, time-of-flight measurement involving 2 BAMs and the setpoint of the gradient are plotted together over several hours.

10 Conclusion and Outlook

Six distinct methods of measuring the beam energy in and around the first bunch compressor chicane of FLASH have been presented and compared. The various pickups and techniques have been described in detail. The chicane BPM with the optical front-end has demonstrated the highest resolution out of all of the different methods. An alternative chicane BPM front-end which used cheaper 10.4 GHz RF techniques also demonstrated acceptable, resolution. The problem with these and other high resolution systems, like the PMT monitor and the BAM, is that they have a limited dynamic range and require frequent mechanical delay-line adjustments and calibrations in order to deliver accurate measurements.

To address these dynamic range limitations, lower resolution and larger dynamic range measurements were developed for both the optical and RF front-ends. The 25 μm resolution, 1.3 GHz RF front-end is an ideal solution for situations that require a quick installation, commissioning time and no down-time due to mechanical adjustments. That is why a copy of the first prototype will be soon commissioned in the second bunch compressor of FLASH. Developing the infrastructure for another optical front-end is much more time-consuming and expensive.

The accuracy of the chicane BPM measurements was studied with respect to likely beam shapes and thermal stability of front-end systems. The monitor will be sensitive to changes in the longitudinal tilt of the beam and, unless the signals from the top and bottom pickups are combined, the monitor will also be sensitive to transverse tilts of the beam. If the chicane BPM measurement is used to measure the beam energy, corrections for incoming orbit changes must be implemented based on BPM measurements from before and after the chicane. Higher resolution BPMs were installed for this purpose. The thermal stability of the chicane BPM front-end systems has been addressed with active temperature control systems involving Peltier elements and the

resolution limitations due to the bandwidth limitations of the pickup and splitter have been described. The impact of the AM-to-PM conversion effect of the RF limiter on the accuracy of the measurement was not measured, but is of interest.

Theory and systems were described to show how the beam arrival-time and position measurements from the chicane BPM could be used to simultaneously stabilize the beam arrival-time and energy and to measure the energy spread of the beam. The merits of the alternative to this scheme, using two arrival-time monitors up and downstream of a chicane to calculate corrections to the beam energy, were also described with a view to a future in which cross-checks of all of these independent, high-resolution measurements will anchor the machine to a stable reference.

Appendix A

The derivation of Eqs. 3.2.1 and 3.2.1 will be presented here. The equations describe the arrival-time jitter of the beam after a bunch compressor. A first-order derivation of the arrival-time jitter is carried out first and an investigation of the complications posed by higher-order terms follows [56].

Starting with the equations for the energy of a particle with position z within the bunch, $E_1(z)$, and the first order energy chirp after the first acceleration section, $E_1'(z)$, we have

$$E_1(z) = E_0(z) + V_1 \cos(k_{rf} z + \varphi_1) \tag{1}$$

and

$$E_1'(z) = E_0'(z) - V_1 k_{rf} \sin(k_{rf} z + \varphi_1), \tag{2}$$

where the energy and energy chirp from before the first acceleration section are given by $E_0(z)$ and $E_0'(z)$. These terms take into account the initial correlated energy distribution entering the first acceleration section, as well as any energy variations generated by collective effects experienced by the bunch before it reaches the entrance to the chicane.

The energy dependent path-length through a chicane was derived in chapter 1 and can be expressed in terms of a nominal path length, L, a nominal energy, E_{nom}, and the R_{56},

$$T(E) = L + R_{56} \frac{E - E_{nom}}{E_{nom}} + \ldots \quad . \tag{3}$$

The first derivative of the path-length will be useful later and is given by,

$$T'(E) = R_{56} \frac{1}{E_{nom}} + \ldots . \tag{4}$$

The path-length through the chicane can be used to calculate the change in the longitudinal position of a particle of the beam relative to a nominal position. This can be described by a transformation from a coordinate system from before the chicane, z_i, into a coordinate system after the chicane, z_f (Fig. 1). The nominal longitudinal position before and after the chicane in each coordinate system is given by z=0. This position corresponds to the nominal beam energy, E_{nom}.

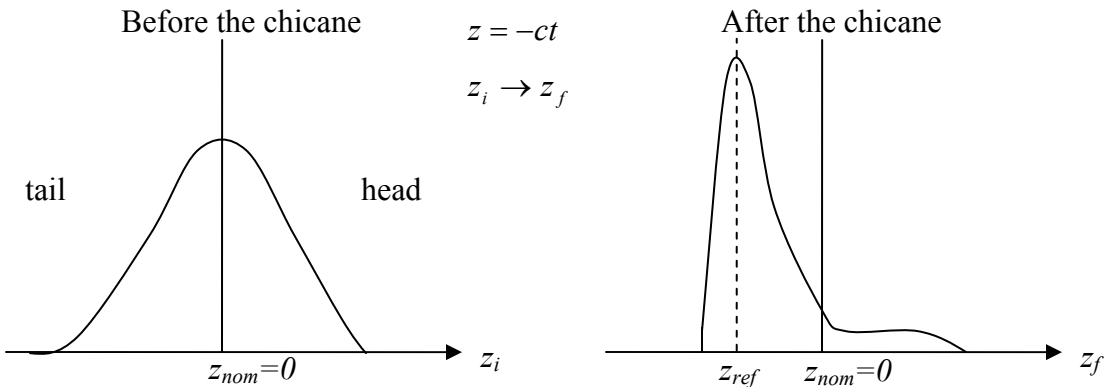

Figure 1 Coordinate system from before, z_i, and after, z_f, the chicane. The nominal longitudinal position before and after the chicane in each coordinate system is given by z=0. This position corresponds to the nominal beam energy, E_{nom}.

When the position of the center of the beam before the chicane is set equal to zero, the position of the center of the beam after the chicane might not be equal to zero. This is especially the case when higher order dispersion terms are taken into account. When this occurs, it means that the center of the beam does not have the nominal energy. In this derivation, the center of the beam will be described by a reference particle with an energy $E_{ref}=E_1(z=0)$.

A particle with position z_i before the chicane will be located at position z_f after the chicane according to,

$$z_f = z_i + T(E_1(z_i)) - T(E_{nom})\,. \tag{5}$$

Adding and subtracting the energy of the reference particle gives,

$$z_f = z_i + \left[T(E_1(z_i)) - T(E_{ref})\right] + \left[T(E_{ref}) - T(E_{nom})\right], \tag{6}$$

where the second term in brackets is a constant which vanishes if the reference energy of the bunch is equal to the nominal energy of the chicane. A Taylor expansion about z_i =0 of the first term in brackets then gives,

$$z_f \approx \left[T(E_{ref}) - T(E_{nom})\right] + (1 + T'E_1')z_i + \ldots \quad . \tag{7}$$

This result can be used to write the compression factor of a bunch which is linearly compressed:

$$C^{-1} \equiv \frac{\sigma_f}{\sigma_i} \approx \frac{\partial z_f}{\partial z_i} = 1 + T'E_1'\,. \tag{8}$$

Solving Eq. 8 for E_1' provides a relationship between the bunch energy chirp and the compression factor.

We now have the longitudinal position of a particle in terms of five free parameters,

$$z_f = z_f(z_i, V_1, \varphi_1, E_0, E_0'). \tag{9}$$

These results can be used to determine the sensitivity of the longitudinal position to each of the free parameters.

Arrival-time jitter due to voltage changes

Using Eq. 5 together with $t_f = z_f/c$, the sensitivity of the arrival-time to the accelerator voltage is,

$$\frac{\partial t_f}{\partial V_1} = \frac{1}{c_0} \cdot \frac{\partial z_f}{\partial V_1} = \frac{1}{c_0} \cdot T'\left(E_1(z_i)\right) \cdot \cos(k_{rf} z_i + \varphi_1) \tag{10}$$

When one solves Eq. 1 for $\cos(k_{rf} z_i + \varphi_1) = \frac{E_1 - E_0}{V_1}$ and uses Eq. 4 to write $T' = \frac{R_{56}}{E_{nom}}$, Eq. 10 can then be written,

$$\frac{\partial t_f}{\partial V_1} = \frac{1}{c_0} \cdot \frac{R_{56}}{E_{nom}} \cdot \frac{E_1 - E_0}{V_1}. \tag{11}$$

Making the approximations that the initial energy is much smaller than the final energy ($E_0 << E_1$), and that the final energy is close to the nominal energy ($E_1 \approx E_{nom}$), one can multiply the result by a small change in accelerator voltage, ΔV_1, to write

$$\frac{\partial t_f}{\partial V_1} \cdot \Delta V_1 \approx \frac{R_{56}}{c_0} \cdot \frac{\Delta V_1}{V_1} \tag{12}$$

Arrival-time jitter due to phase changes

Using Eq. 5 together with $t_f = z_f/c$, the sensitivity of the arrival-time to the accelerator phase is,

$$\frac{\partial t_f}{\partial \varphi_1} = \frac{1}{c_0} \cdot \frac{\partial z_f}{\partial \varphi_{1\,1}} = \frac{1}{c_0} \cdot T'\left(E_1(z_i)\right) \cdot \sin(k_{rf} z_i + \varphi_1) \tag{13}$$

Using Eq. 1, to solve for $\sin(k_{rf} z_i + \varphi_1)$, we can write

$$\frac{\partial t_f}{\partial \varphi_1} = \frac{1}{c_0} \cdot T'\left(E_1(z_i)\right) \cdot \frac{E_1'(z_i) - E_0'}{k_{rf}} \tag{14}$$

Using the approximation that $E_0' << E_1'$ and multiplying the sensitivity by a small phase change gives

$$\frac{\partial t_f}{\partial \varphi_1} \cdot \Delta\varphi_1 = \frac{1}{c_0 k_{rf}} \cdot T'\left(E_1(z_i)\right) \cdot E_1'(z_i) \cdot \Delta\varphi_1 . \tag{15}$$

Now, solving for T'E' in terms of the compression factor (Eq. 8) gives

$$\frac{\partial t_f}{\partial \varphi_1} \cdot \Delta\varphi_1 = \frac{1}{c_0 k_{rf}} \cdot \left(\frac{C-1}{C}\right) \cdot \Delta\varphi_1 . \tag{16}$$

Arrival-time jitter due to arrival-time changes prior to the accelerator section, z_i

Using Eq. 5 and Eq. 8, we can immediately write down the sensitivity of the arrival-time after the chicane to a small change of the incoming arrival-time

$$\frac{\partial t_f}{\partial z_i} \cdot \Delta z_i = \frac{1}{c_0} \cdot \frac{\partial z_f}{\partial z_i} \cdot \Delta z_i = \frac{1}{c_0} \cdot \frac{1}{C} \cdot \Delta z_i \tag{17}$$

This shows that the incoming arrival-time jitter is compressed by the compression factor of the chicane, C.

Sum of all contributions

Since we assume that the incoming arrival-time changes, voltage changes, and phase changes are statistically independent, they can be added in quadrature to find the net arrival-time jitter after the chicane, Eq. 3.x.x,

$$\Sigma_{t,f}^2 = \left(\frac{R_{56}}{c_0}\right)^2 \cdot \frac{\sigma_{V_1}^2}{V_1^2} + \left(\frac{C-1}{C}\right)^2 \cdot \frac{\sigma_{\varphi_1}^2}{\varphi_1^2} + \left(\frac{1}{C}\right)^2 \cdot \Sigma_{t,i}^2 \tag{18}$$

This is a reasonable representation of the arrival-time jitter of the beam core, but as the many approximations made in the derivation showed, it does not describe at all what the rest of the bunch is doing for non-linear compression or for operation with the 3rd-harmonic module. It can also not be used recursively for an additional bunch compression stage.

Higher order terms and the 3rd-harmonic module

To see how the 3rd-harmonic module is affects the timing jitter, more terms need to be carried in the derivation. Even without the 3rd-harminic module, the higher order terms are necessary to understand how the reference particle acts with respect to the nominal particle. Starting with the energy of a particle located at z,

$$E_1(z) = E_0(z) + V_1\cos(k_{rf}z + \varphi_1) + V_{39}\cos(3k_{rf}z + \varphi_1)$$
$$E_1'(z) = E_0'(z) - V_1 k_{rf}\sin(k_{rf}z + \varphi_1) - V_{39}\left(3k_{rf}\right)\sin(3k_{rf}z + \varphi_{39})$$
$$E_1''(z) = E_0''(z) - V_1 k_{rk}^2\cos(k_{rf}z + \varphi_1) - V_{39}\left(3k_{rf}\right)^2\cos(3k_{rf}z + \varphi_{39})$$

and the path-length through the chicane,

$$T(E) = L + R_{56}\frac{E - E_{nom}}{E_{nom}} + T_{566}\left(\frac{E - E_{nom}}{E_{nom}}\right)^2 + \ldots$$
$$T'(E) = R_{56}\frac{1}{E_{nom}} + 2T_{566}\frac{E - E_{nom}}{E_{nom}^2} + \ldots$$
$$T''(E) = 2T_{566}\frac{1}{E_{nom}^2} + \ldots,$$

one can Taylor expand $z_f = z_i + T(E_1(z_i)) - T(E_{nom})$ about z=0 and E_{nom} to get

$$z_f \approx (1 + T'E_1')z_i + \frac{1}{2}\left(T''(E_1')^2 - T'E_1''\right)z_i^2 + \ldots,$$

and then define the compression factor,

$$C(z_f) = \frac{\partial z_i}{\partial z_f} = \left((1 + T'E_1')z_i + \frac{1}{2}\left(T''(E_1')^2 - T'E_1''\right)z_i^2 + \ldots\right)^{-1}.$$

Using these terms in the derivation of the arrival-time jitter should predict the arrival time jitter of more than just the core of the bunch. More than that, it provides a tool for optimizing the setpoints of the first accelerating section and the third harmonic module so that the timing jitter is minimized. Presently, this sort of optimization is carried out with particle tracking code, but analytic solutions offer a more global picture and more flexibility in terms of their predictions for different machine configurations.

Additional bunch compression stage

If the first bunch compressor is followed by an additional accelerator section and bunch compressor, a few modifications need to be made to Eq. 18 in order to describe this

situation. First, the equations for the energy and energy chirp after the second accelerator stage are written,

$$E_2(z) = E_1(z) + V_2 \cos(k_{rf} z + \varphi_2)$$
$$E_2'(z) = E_1'(z) - V_2 k_{rf} \sin(k_{rf} z + \varphi_2)$$

Case 1: $E_1 << E_2$ and $E_1' << E_2'$
This is the case for present LCLS operation. If the second bunch compressor is located at a higher energy than the first one and the energy chirp for the second chicane is much larger than for the first one, then the approximation from Eq. 18 can be applied iteratively to the second bunch compression stage, just as it had been applied to the first bunch compression stage.

Case 2: $\varphi_2=0$
This was the case for operation of FLASH without the 3rd harmonic cavity. For this case, one can make the approximation,

$$E_2(z) \approx E_1(z) + V_2$$
$$E_2'(z) \approx E_1'(z).$$

Neglecting the non-linearity of the compression process, the path length through the second chicane is written,

$$T_2(E) = L_2 + R_{56,2} \frac{E - E_{nom,2}}{E_{nom,2}} + \ldots$$
$$T_2'(E) = R_{56,2} \frac{1}{E_{nom,2}} + \ldots$$

and a particle with a position z_1 after the first bunch compressor will be located at z_2 after the second bunch compressor,

$$\begin{aligned} z_2 &= z_1 + T_2(E_2(z_i)) - T_2(E_{nom2}) \\ &\approx .(1 + T_2' E_2') z_1 + \ldots \\ &\approx .(1 + T_2' E_2')(1 + T_1' E_1') z_i + \ldots, \end{aligned}$$

where the reference particle for each chicane has been set equal to the nominal position. This eliminates any constant offset due to injecting the bunch at an energy which is not the nominal energy of the chicane.

This can be used to derive the sensitivity of the arrival-time to the voltage, phase, and incoming timing jitter in the same fashion that was used in the single bunch compressor case. Starting with

$$\frac{\partial t_2}{\partial V_2}\cdot\Delta V_2 = \frac{1}{c_0}T_1'(E_2(z_1)\cdot\cos(k_{rf}z_1+\varphi_1)\cdot\Delta V_2$$
$$= \frac{R_{56,2}}{c_0}\cdot\frac{E_2(z_1)-E_1(z_1)}{E_{nom,2}}\cdot\Delta V_2$$
$$\approx \frac{R_{56,2}}{c_0}\frac{E_2-E_1}{E_2}\cdot\frac{\Delta V_2}{V_2},$$

where the nominal energy was assumed to be equal to the reference energy,

$$\frac{\partial t_2}{\partial \varphi_2}\cdot\Delta\varphi_2 = \frac{1}{c_0}T_1'(E_2(z_1)\cdot\sin(k_{rf}z_1+\varphi_1)\cdot\Delta\varphi_2$$
$$= \frac{1}{c_0}T_1'(E_2(z_1)\cdot\frac{E_2'(z_1)-E_1'(z_1)}{k_{rf}}\cdot\Delta\varphi_2$$
$$\approx 0,$$

where the energy chirps prior to both bunch compressors are assumed to be approximately equal,

$$\frac{\partial t_2}{\partial V_1}\cdot\Delta V_1 = \frac{1}{c_0}T_2'(E_2(z_1))\cdot E_2(z_1)\frac{\partial z_1}{\partial V_1}\cdot\Delta V_1$$
$$= \frac{R_{56,2}E_1(z_i)}{c_0E_{nom,2}}\cdot\frac{R_{56,1}\left(E_1(z_i)-E_0(z_i)\right)}{E_{nom,1}}\cdot\frac{\Delta V_1}{V_1}$$
$$\approx \frac{1}{c_0}\cdot\left(R_{56,1}+\frac{E_1}{E_2}R_{56,2}\right)\frac{\Delta V_1}{V_1}$$

where the nominal energy was assumed to be equal to the reference energy and the initial energy, E_0, is small compared to the energy after the first accelerating section,

$$\frac{\partial t_2}{\partial \varphi_1}\cdot\Delta V_1 = \frac{1}{c_0}T_2'(E_2(z_1))\cdot E_2'(z_1)\frac{\partial z_1}{\partial \varphi_1}\cdot\Delta\varphi_1$$
$$= \frac{1}{c_0}T_2'\cdot E_2'\cdot T_1'\cdot\frac{E_1'(z_i)-E_0'(z_i)}{k_{rf}}\cdot\Delta\varphi_1$$
$$\approx\left(\frac{C-1}{C}\right)\cdot\frac{\Delta\varphi_1}{c_0k_{rf}}\ \text{ where } C=C_1C_2$$

the result, Eq. 3.2.2, for the net timing jitter after two bunch compressor stages is,

$$\Sigma_{t,2}^2 = \frac{1}{C^2}\Sigma_{t,0}^2+\frac{1}{c_0}\cdot\left(R_{56,1}+\frac{E_1}{E_2}R_{56,2}\right)\cdot\frac{\sigma_{V_1}}{V_1}+\frac{R_{56,2}}{c_0}\cdot\frac{E_2-E_1}{E_2}\cdot\frac{\sigma_{V_2}}{V_2}+\left(\frac{C-1}{C}\right)\cdot\frac{\sigma_{\varphi_1}}{c_0k_{rf}}$$

Appendix B

A description of how the one can measure and remove the effect of reflections in the pickup from the RF beam position measurement follows. There is a constant ά for each pickup output that is a measure of the effect of the reflections within the pickup on the arrival-times measured at the outputs of each pickup. This is given for the right, R, and left, L, pickup outputs in terms of the beam arrival-time, t, beam position, x, and ά as

$$\Delta t_L = \Delta t_{arrival} + \alpha_L \frac{\Delta x}{c} \qquad \text{and} \qquad \Delta t_R = \Delta t_{arrival} - \alpha_R \frac{\Delta x}{c} \ . \qquad (1)$$

In terms of the beam energy and constants of the chicane, the beam arrival-time and position are given by

$$\Delta t_{arrival} = \Delta t_0 - \frac{R_{56}}{2} \cdot \frac{\Delta E}{E} \qquad \text{and} \qquad \Delta x = R_{16} \cdot \frac{\Delta E}{E}, \qquad (2)$$

where t_0 is the arrival time before the chicane. For higher energies, the arrival-time is earlier and the beam position moved to the right in the positive direction. Writing Eqs. 1 in matrix format,

$$\begin{pmatrix} \Delta t_L \\ \Delta t_R \end{pmatrix} = \begin{pmatrix} 1 & \alpha_L \\ 1 & \alpha_R \end{pmatrix} \begin{pmatrix} \Delta t_{arrival} \\ \Delta x / c \end{pmatrix} \qquad (3)$$

one can calculate the inverse of the 2x2 matrix in order to find the beam arrival-time and position in terms of the measured arrival-times and reflection coefficients,

$$\begin{pmatrix} \Delta t_{arrival} \\ \Delta x / c \end{pmatrix} = \frac{1}{\alpha_L + \alpha_R} \begin{pmatrix} \alpha_R & -1 \\ \alpha_L & -1 \end{pmatrix} \begin{pmatrix} \Delta t_L \\ \Delta t_R \end{pmatrix}. \qquad (4)$$

References

[1] P. Castro. “Beam trajectory calculations in bunch compressors of TTF2”, DESY-TECHNICAL-NOTE-2003-01.

[2] R. Brinkmann et.al.: “TESLA XFEL: First Stage of the X-Ray Laser Laboratory (Technical Design Report, Supplement)”, DESY, TESLA FEL-2002-09.

[3] W. Ackermann et al. “Operation of a free-electron laser from the extreme ultraviolet to the water window”, *Nature Photon.,* 1:336-342, 2007.

[4] A. Zholents, M. Zolotorev, W. Wan, “Generation of Attosecond Bunches”, PAC 2001.

[5] N. Stojanovic, “Status and plans of THz (IR) beamline”, FLASH Seminar, Feb 2010.

[6] Massimo Altarelli et al. “XFEL: The European X-Ray Free-Electron Laser”, Technical design report. DESY-06-097.

[7] A. Azima et.al. "Status of sFLASH, the Seeding Experiment at FLASH", Proceedings of the IPAC 2010, Kyoto, Japan.

[8] M. Krasilnikov et al. “Experimental characterization of the electron source at the photo-injector test facility at DESY Zeuthen”, PAC 2003, Portland. May 2003, p. 2114.

[9] S. Schreiber et al. “Measurement of space charge effects and laser bunch length in the TTF RF gun using the phase scan technique”, FEL 1999, p. 69.

[10] S. Schreiber et al. “Experience with the Photoinjector Laser at FLASH”, FEL 2006, BESSY, Berlin, Germany.

[11] C. Schmidt, **“**RF System Modeling and Controller Design for the European XFEL**”,** Ph.D. Thesis, University of Hamburg, 2010.

[12] K. L. Brown, “A First- and Second-order Matrix Theory for the Design of Beam Transport and Charged Particle Spectrometers”, SLAC-Report-75, 1982.

[13] J. Aurthur et al., “Linac Coherent Light Source (LCLS) Conceptual Design Report”, Tech. Report SLAC-R-593, SLAC, 2002.

[14] T. O. Raubenheimer et.al., “Chicane and Wiggler Based Bunch Compressors for Future Linear Colliders”, SLAC-PUB-6119, May 1993.

[15] Ya. S. Derbenev, et al., “Microbunch Radiative Tail-Head Interaction”, DESY, Sep 1995.

[16] E.L. Saldin, E.A. Schneidmiller and M.V. Yurkov, Nucl. Instrum. and Methods A 528 (2005) 355.

[17] P. Emma and R. Brinkmann, “Emittance dilution through energy spread generation in bending systems”, Proceedings of the PAC 1997.

[18] P. Schmueser, M. Dohlus and J. Rossbach. *Introduction to Ultraviolet and X-Ray Free-electron Lasers*, Springer, to-be-published.

[19] T. Schilcher, **“**Vector Sum Control of Pulsed Accelerating Fields in Lorentz Force Detuned Superconducting Cavities**”,** Ph.D. Thesis, Department of Physics of the University of Hamburg 1998.

[20] F. Loehl, “Optical Synchronization of a Free-Electron Laser with Femtosecond Precision”, Ph.D. Thesis, Department of Physics of the University of Hamburg 2008.

[21] F. Ludwig, C. Gerth, K. Hacker, M. Hoffmann, W. Jalmunzna, , G. Moeller, P. Morozov, C. Schmidt, “Drift Calbration Techniques for Future FELs”, Proceedings of the IPAC 2010, Kyoto, Japan.

[22] M. Hoffmann, F. Ludwig, H. Schlarb, S. Simrock, “Multichannel Downconverter for the Next Generation RF Field Control for VUV- and X-Ray Free-electron Lasers”, Proceedings of the PAC 2007.

[23] H. Schlarb et al. “Beam Based Measurements of RF Phase and Amplitude Stability at FLASH”, Proceedings of the DIPAC 2007, Venice, Italy.

[24] J. Cawardine, “Results from full beam-loading (9 mA experiment)”, FLASH seminars November 2009.

[25] H. Schlarb, V. Ayvazyan, F. Ludwig, D. Noelle, B. Schmidt, S. Simrock, A. Winter, “Next Generation Synchronization System for the VUV-FEL at DESY”, Proceedings of the FEL 2005.

[26] H. Schlarb et al., “Precision RF Gun Phase Monitor System for FLASH”, Proceedings of the EPAC 2006, Edinburgh, Scotland.

[27] H. Schlarb, “Plans for beam based feedback at FLASH”, Internal, DESY 2010.

[28] H. Schlarb et al., “Beam based measurements of RF phase and amplitude stability at FLASH”, Proceedings of the DIPAC 2007.

[29] S. Schulz et al., “Injector laser cross correlation”, Proceedings of the IPAC 2010, Kyoto, Japan.

[30] M. Dohlus and T. Limberg, “Bunch compression stability dependence on RF parameters”, Proceedings of the FEL 2005, Stanford, California.

[31] S. Wesch, **“**Spektroskopie kohaerenter Uebergagsstrahlung zur strukturanalyse von elektronenpaketen am FLASH beschleuniger**”,** Master’s thesis, University of Hamburg, 2008.

[32] M. Roehrs, “Investigation of the Phase Space Distribution of Electron Bunches at the FLASH-Linac Using a Transverse Deflecting Structure”, Ph.D. thesis, University of Hamburg 2008.

[33] C. Gahl et al., Nature Photonics 2, 165-169 (2008)

[34] K. Floettmann. Astra user manual. http: / /www.desy.de/ mpy/ astra documentation.

[35] M. Dohlus, “Field Asymmetries and Kicks”, http://tesla.desy.de/fla/publications/talks/seminar/FLA-seminar_090904.pdf.

[36] H. Wiedemann, *Particle Accelerator Physics II,* Springer-Verlag Berlin Heidelberg 1999.

[37] I. Zagorodnov, “Wakefield Effects of new ILC cavity shapes”, European Particle Accelerator Conference, Edinburgh Scotland (2006).

[38] M. Krasilnikov, et. al, “Beam-based Procedures for RF Guns”, Particle Accelerator Conference, Knoxville, Tennessee (2005).

[39] Work done together with E. Prat and C. Gerth (2008).

[40] J. Rosenzweig and L. Serafini. “Transverse Particle Motion in Radio-frequency Linear Accelerators,” Phys. Rev. E, 49(2):1599-1602, Feb1994.

[41] A. W. Chao, *Physics of Collective Beam Instabilities in High Energy Accelerators,* John Wiley & Sons, Inc. 1993.

[42] E. Medvedko, R. Johnson, S. Smith, R. Akre, D. Anderson, J. Olsen, T. Straumann, A. Young, "Stripline Beam Position Monitors for LCLS", Beam Instrumentation Workshop (2008).

[43] R.E. Schafer, "Beam Position Monitoring", AIP conference proceedings on Accelerator Instrumentation (AIP Upton, NY 1989) pp 26-55.

[44] David M Pozar**,** *Microwave Engineering*, Wiley & Sons, January 2004.

[45] F. Marcellini, M. Serio, M. Zobov, "DAPHNE Broadband Button Electrodes"*,* Frascati (January 16, 1996) Note: CD-6

[46] T. Smith, SLAC, Personal Communication.

[47] J. Zemella, "Driftfreier Detector zur Messung des Zeitversatzes zweier verschiedener Laserpulszuege", Diploma Thesis, University of Hamburg, 2008.

[48] B. Lorbeer et al., "Noise and drift characterization of critical components for the laser based synchronization system at FLASH," Proceedings of the DIPAC 2007, Venice, Italy.

[49] F. Ludwig et al., "Noise and drift characterization of direct laser to RF conversion scheme for the laser based synchronization system for FLASH at DESY," Proceedings of the PAC 2007, Albequerque, New Mexico, USA.

[50] J. Kim et al., "Long-term femtosecond timing link stabilization using a single-crystal balanced cross correlator", Opt. Lett. (2007), no. 9, 1044-1046.

[51] K. Hacker et al., "RF-lock for Master Laser Oscillator", Proceedings of the EPAC 2009, Edinburgh, Scotland.

[52] C. Gerth, "Synchrotron Radiation Monitor for Energy Spectrum Measurements in the Bunch Compressor FLASH", Proceedings of the DIPAC 2007, Venice, Italy.

[53] A. Wilhelm, "Installation and Characterization of a Diagnostic System for Bunch Resolved Beam Energy Measurements at FLASH", Diploma Thesis, Department of Physics of the University of Hamburg, August 2009.

[54] H. Wiedemann, *Particle Accelerator Physics II,* Springer-Verlag Berlin Heidelberg 1999.

[55] C. Gerth, personal communication.

[56] H. Schlarb, personal communication.

[57] A. Wilhelm and C. Gerth, "Synchrotron Radiation Monitor for Bunch – Resolved Beam Energy Measurements at FLASH", DIPAC '09.

[58] K. Hacker et al., "Demonstration of BPM with sub-5 um resolution over a 10 cm range," Proceedings of the EPAC 2009, Edinburgh, Scotland.

Acknowledgements

Many thanks must be given to all of the people on the optical synchronization team for building up the infrastructure (lasers and optical links) that made the optical BPM measurements possible. Holger Schlarb was the mastermind of the system, Axel Winter started the upgrade of the system so that it could serve more users like myself, Florian Loehl built first prototypes of many devices and ordered countless components for the optical front-ends, Sebastian Schulz built the laser and aligned the distribution section, Marie Bock and Holger Schlarb got the link up and running on time, Patrick Gessler got all of the ADCs and BAM servers going, and Matthias Felber took care of countless issues related to piezo-drivers and whatnot. Especial thanks go to Holger Schlarb for always being open to answer my questions and for explaining just about everything that is written in this thesis to me. A big thank you goes to Peter Schmueser, Holger Schlarb, Bernhard Schmidt, and Joerg Rossbach for pre-reading sections of this document, to DESY for extending my contract whenever I had another baby, and to Bernhard Schmidt for hosting me in his group and allowing me access to all of the labs and equipment and people therein. Thanks to Eduard Prat and Christopher Gerth for investigating chicane beam tilts with me. Thank you to Jan Hauschildt for the quick technical drawings, manufacture and installation of the BPM pickup, to Silke Vilcins for the drawings, manufacture and installation of the new BAM pickups, to Dirk Noelle and others for re-arranging the BPMs around the chicane, to Bernd Beyer for the technical drawings and manufacture of the mechanics for the first and second versions of the optical BPM and thank you to Matthias Hoffmann and Albert Schleiermacher for so quickly and professionally assembling the 2-layer RF chassis, the Beckhoff boxes, and the last version of the optical BPM chassis.