Digital Phase Control Techniques for Accelerator Cavities.

Amos Dexter, Imran Tahir, Graeme Burt and Richard Carter

Lancaster University Engineering Department

Abstract

RF cavities used for the acceleration and deflection of charged bunches of particles in high energy physics accelerators, medical linacs and synchrotron light sources need precise phasing with respect to the arrival of bunches and to each other. The paper describes phase locking and amplitude control performance achieved for 3.9 GHz superconducting cavities in the presence of microphonics using frequency division, digital phase detection with the HMC439QS16G and digital implementation of the controller.

1. Introduction

RF cavities are used to accelerate bunches of charged particles by arranging a time varying alternating electric field to point in the direction for acceleration at the instant that bunches pass. Multiple cells are used to increase acceleration per unit length. Figure 1 shows axial electric fields in an elliptical, multi-cell acceleration cavity with respect to charged bunch position.

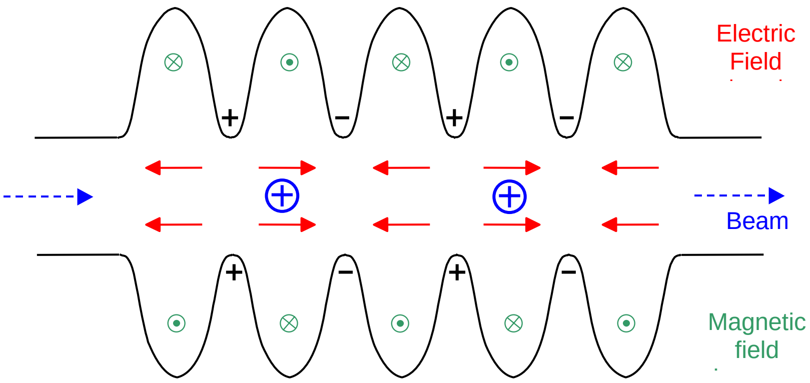

Figure 1 Fields in an RF acceleration cavity

High energy accelerators use a large number of single or multi-cell cavities. The relative phasing of RF fields in accelerator cavities determine acceleration and longitudinal bunch size. Bunches of charge particles are usually generated in trains of several hundred or several thousand. If the bunch train has a long or continuous time profile then superconducting cavities are typically used. High gradient superconducting cavities are often given an elliptical profile as shown in figures 1 and 2 to assist surface treatment and to avoid problems with multipactor (vacuum RF discharge).

Figure 2 Three cell Niobium Cavity

As well as accelerating charged bunches it is also possible to excite cavities by driving them at a higher frequency so they deflect or impart angular rotation to charged bunches. Figure 2 shows a three cell Niobium cavity with a dipole mode at 3.9 GHz and figure 3 shows the field pattern associated with dipole excitation for a five cell cavity.

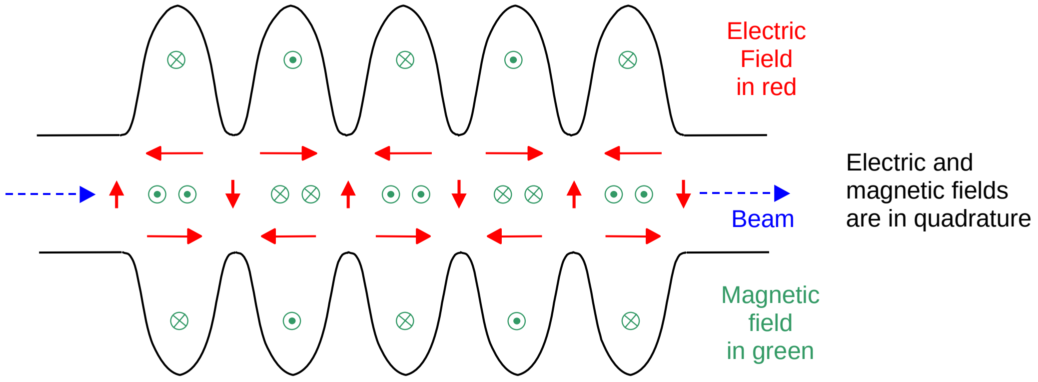

Figure 3 Dipole excitation

The intrinsic Q factor of a Niobium superconducting cavity at 3.9 GHz operated at 1.8 K can be 10^9 or more. For an acceleration cavity to be useful, power must be transferred efficiently from the source to the train of charged bunches. The bunch train can be regarded as a resistive load in the cavity and the external Q of the cavity must be matched to the beam for optimal energy transfer. For many applications, superconducting cavities are operated with external Q factors nearer to 10^6 than 10^9. For a cavity operated at 3.9 GHz and an external Q of 3×10^6 the bandwidth is 1.3 kHz.

For a cavity operated in a dipole mode so it deflects or rotates bunches, the energy transferred to the bunch is expected to be extremely small as the effort changes bunch direction with minimal change in kinetic energy. It can be seen from the field pattern for the dipole excitation that this may not be the case if the charge bunch is not quite on the central axis. On one side of the axis a charged bunch can see an accelerating field whilst on the other side a de-accelerating field. Cavities which rotate charged bunches are required for the next generation of linear colliders.

High energy physics experiments optimise energy reach by colliding particles with nearly the same energy. During 2009 the LHC synchrotron at CERN expects to collide 7 TeV protons with 7 TeV protons. Previously in 2001 CERN was colliding 100 GeV electrons with 100 GeV positrons in the LEP[1] facility. In order to collide electrons and positrons at energies much greater than 100 GeV the accelerator must have a linear format rather than a circular format. This is because synchrotron radiation rises with the fourth power of the bend radius and hence the bend radius needs to increase by a factor of 625 between 100 GeV and 500 GeV; one notes that LEP had a bend radius of 4.3 km.

The layout of the international linear collider[2] (ILC) that is being designed by a large worldwide collaboration is shown in figure 4. Opposing linear accelerators (linacs) accelerate electrons and positrons so they collider at the interaction point (IP).

---

[1] CERN website http://public.web.cern.ch/public/en/Research/LEP-en.html
[2] ILC Reference Design Report http://www.linearcollider.org/cms/?pid=1000437

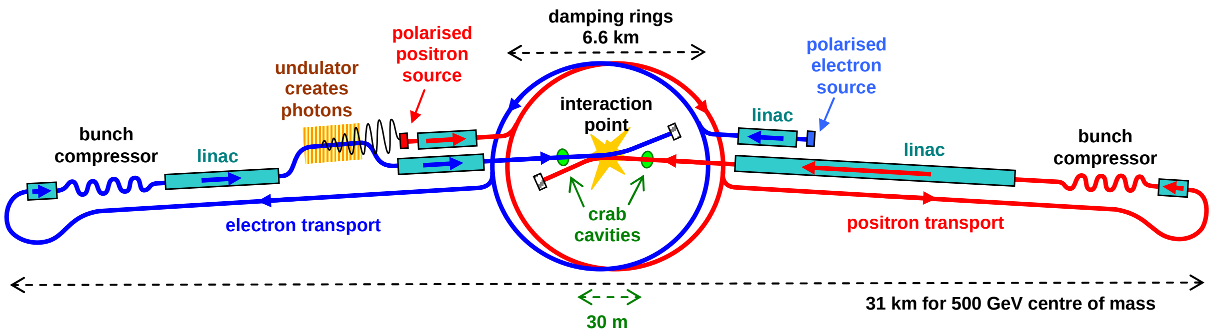

Figure 4 **Proposed layout for the ILC**

The IP is surrounded by a large detector that searches for new and exotic particles. For such a collider it is necessary for the two opposing linear accelerators (linac) to be synchronised so that collisions occur at a precise location. For the satisfactory operation of the individual linacs the acceleration fields must be precisely phased with the particle source, pulsed deflection magnets on the damping rings and the bunch compressors.

In order to maximise the number of individual particle collisions at the IP the bunches are focused to the aspect ratio of a flat tape measuring nano-meters in thickness, hundres of nano-metres in width and fractions of a millimetre in length. At the highest collision energies, design dictates that the colliding beams be inclined at a small angle to each other in the horizontal plane (about 1 degree). This is partly because it is difficult to prevent particles and their associated halos entering the opposing linac.

2. Crab Cavities

With a non zero crossing angle at the IP the maximum collision rate occurs when the bunches are rotated slightly so collisions are effectively head on. For the ILC this alignment will be performed with deflection cavities phased such that the centre of each bunch passes through the cavity at a time when it sees zero net field whilst the front and back of the bunch gets an equal and opposite kick[3]. A deflection cavity phased this way is called a crab cavity. After passage through a crab cavity a bunch should start to rotate about its centre, this is illustrated in figure 5. Any error in the phasing of the crab cavities results in bunches getting a transverse kick and hence transverse deflection at the IP. If this transverse kick is small, but identical for the crab cavities on the opposing linacs then the bunches still collide with maximum luminosity but with small transverse offset.

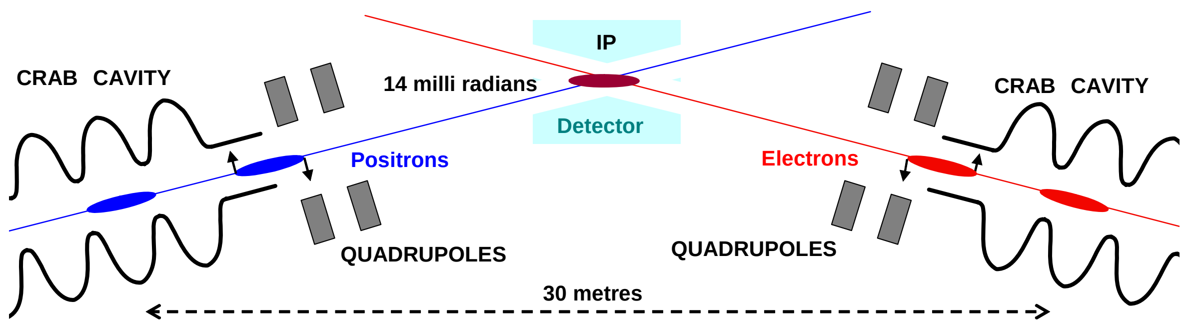

Figure 5 **Crab cavities are used to align bunches at the IP**

The correct phasing of a complex accelerator system is typically achieved by a timing distribution system whereby pulses are distributed at regular intervals along dispersionless transmission lines to

3 Burt G.C., Dexter A.C., and Goudket P., "Effect and Tolerances of RF Phase and Amplitude Errors in the ILC Crab Cavity" , EUROTeV-Report-2006-098, 2006
http://www.eurotev.org/reportspresentations/eurotevreports/index_eng.html

individual RF systems such as cavities[4]. For simple RF systems, a fixed frequency RF signal at the cavity frequency or some fixed fraction of the cavity frequency might be distributed. Phasing of individual components is then achieved by precise variation of path length with manual phase shifters.

To minimise electrical power usage, superconducting cavities were selected for the ILC main linac. The optimal bunch structure for the ILC superconducting linacs (3000 bunches 330 nano-seconds apart) makes superconducting cavities the optimum choice for the crab cavities.

The phasing accuracy requirement for the ILC linacs is about 0.1 degrees r.m.s. at 1.3 GHz, which corresponds to an r.m.s. timing accuracy of 214 fs. The phasing accuracy for the ILC crab cavities with respect to each other is 0.125 degrees r.m.s. at 3.9 GHz which corresponds to a timing accuracy of 90 fs and hence might need a dedicated control system.

The current proposal for the ILC linac timing system is to use mode locked lasers to stabilise the master clock and to use optical pulses on fibre cables to distribute the timing reference. Recent claims for the performance of optical distribution systems reach accuracies of 10 fs [5]. The synchronisation method described in this paper for the ILC crab cavities considers an RF interferometer rather than the competing and somewhat more expensive optical methods.

3. Amplitude and Phase Control of Superconducting RF Cavities

Accelerator cavities are operated with very high electric fields. Surface fields for superconducting cavities might reach 70 MVm^{-1} and for normal conducting copper cavities several times this value. The associated currents heat non-superconducting copper very rapidly and hence at these gradients copper structures must be run with very short pulses and from very high peak power sources. This constraint disappears for superconducting cavities and they are typically run with long pulses and well spaced bunches of charged particles thereby requiring lower peak power sources.

For the purpose of understanding the phase control of a cavity at its simplest level, it can be modelled as a parallel LCR circuit driven from a current source as shown in figure 6.

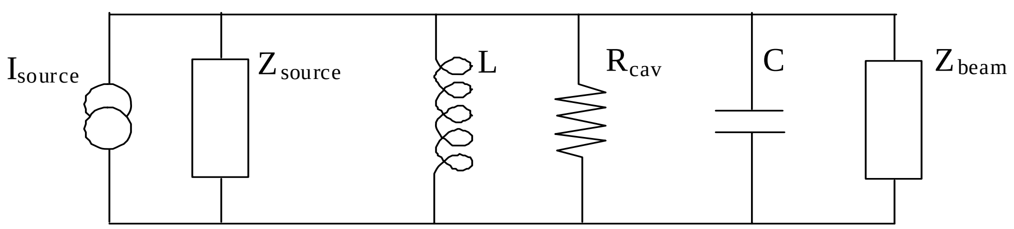

Figure 6 Equivalent circuit for a cavity mode

A beam of charged particles passing through the cavity loads the cavity and again at the simplest level of approximation can be included in the model as a parallel impedance Z_{beam} . If the beam adds power to the cavity rather than extracting power from the cavity the resistive part of Z_{beam} becomes negative. Relating Q factors to impedances we have

$$Q_{ext} = \frac{\omega U_{stored}}{P_{emitted}} = \frac{\frac{1}{2}\omega C V^2}{\frac{1}{2}\left(V^2 / Z_{source}\right)} = \omega Z_{source} C$$

4 Carwardine J., Arnold N., Lenkszus F., Saunders C., Rehlich K., Simrock S., Banerjee B., Chase B., Gottschalk E., Joireman P., Kasley P., Lackey S., McBride P., Pavlicek V., Patrick J., Votava M., Wolbers S., Furukawa K., Michizono S., Larsen R.S., Downing R., "The ILC global control system", (TUZAC01), PAC, June 2007

http://accelconf.web.cern.ch/AccelConf/p07/PAPERS/TUZAC01.PDF

5 J. W. Staples, R. Wilcox and J. M. Byrd, "Demonstration of femto-second phase stabilization in 2 km optical fibre", Proceedings of PAC07, Albuquerque, New Mexico, USA, http://accelconf.web.cern.ch/accelconf/p07/PAPERS/MOPAS028.PDF

$$Q_o = \frac{\omega U_{stored}}{P_{diss}} = \frac{\frac{1}{2}\omega C V^2}{\frac{1}{2}\left(V^2 / R_{cav}\right)} = \omega\, R_{cav}\, C$$

$$Q_{beam} = \frac{\omega U_{stored}}{P_{beam}} = \frac{\frac{1}{2}\omega C V^2}{\frac{1}{2}\left(V^2 / Z_{beam}\right)} = \omega\, Z_{beam}\, C$$

The intrinsic Q factors of superconducting cavities Q_o are typically greater that 10^9 at 1.8 K whereas beam loading is typically brought to a level where $Q_{beam} \sim 10^6$ hence the Q_o can be disregarded with respect to controlling amplitude and phase. For optimum power transfer one chooses $Q_{ext} \sim Q_{beam}$.

A big issue for controlling the phase of a superconducting cavity is that for external Q factors near to 10^6 any minute mechanical vibration of the cavity shifts its natural frequency away from the drive frequency and a phase error arises. The significance of the high Q factor is that the phase shift is proportional to the Q factor. Such phase shifts caused by mechanical vibrations in superconducting cavities are known as microphonics. The solution is to adjust the phase of the input drive as the cavity vibrates so that the phase of the cavity is correct (and staying steady). As the mechanical vibrational frequencies of typical cavities that give rise to significant displacements tend to be less than 1 kHz and more typically less than 100 Hz it is relatively straightforward to measure the cavity phase and apply a simple controller. Where mechanical movement of the cavity walls is a consequence of electromagnetic forces changing as the cavity fills, the associated de-tuning is known as Lorentz de-tuning. Typically one fills the cavity at a rate where the controller can manage Lorentz de-tuning.

4. The controller

The governing equation for the excitation of a single mode in a cavity and assuming no interaction with other modes is

$$\frac{d^2V}{dt^2} + \omega_o\left(\frac{1}{Q_o} + \frac{1}{Q_{ext}}\right)\frac{dV}{dt} + \omega_o^2\, V = \frac{2\omega_o}{Q_{ext}}\frac{d}{dt}\left\{F \exp(-j\omega t)\right\}$$

where $F = F_r + jF_i$ is the forward power in the waveguide from the source. This equation can be inferred from the circuit model in figure 6. Amplitude and phase are typically controlled by controlling the magnitude of the I and Q components $A_r(t)$ and $A_i(t)$ defined as

$$V(t) = \left\{A_r(t) + jA_i(t)\right\}\exp\left\{-j\omega t\right\}$$

which are slowly varying functions of time when the Q factors are large. Our phase control tests have to date have only used a PI controller where

$$F_r\left(t + t_{delay}\right) = c_{pr}\left(V_{sp} - A_r\right) + c_{ir}\left(\frac{\omega}{2\pi}\right)\int_{-\infty}^{t} dt\,\left(V_{sp} - A_r\right)$$

$$F_i\left(t + t_{delay}\right) = -c_{pi}\, A_i - c_{ii}\left(\frac{\omega}{2\pi}\right)\int_{-\infty}^{t} dt\; A_i$$

A differential term is not employed as it is anticipated that noise on measurements of actual cavity fields cannot be adequately suppressed for the differential term to be useful whilst retaining system response. A PI algorithm responds well to random and unpredictable system disturbances such as beam offsets. The other disturbance, namely microphonics is typically predictable on the time scale of sub milli-seconds and hence predictive, feed-forward algorithms have the potential to out perform the simple PI algorithm. Advanced control algorithms have not been considered yet as microphonic data for the final system would not be available until it has been built and characterised. Performance with a PI algorithm is a baseline performance.

The bandwidth ω / Q_e of most superconducting cavities designed to interact with a beam is invariably less than a few kHz. For a cavity control system there is typically a delay t_{delay} between measurement of phase and amplitude and corrective action. For digital control the delay time is determined primarily from the processing time. The bandwidths of amplifiers in the loop are chosen so that this delay time is not increased unduly.

For a proportional controller acting on a field value with coefficient c_p then its maximum value is limited by the point of instability determined by the in-equality $c_{pr} < \dfrac{\pi}{2 \times \text{bandwidth} \times t_{delay}}$.

This implies that the minimum fractional steady state error $\delta F/F$ of a field when a simple proportional controller is used can be roughly estimated as $\min(\delta F/F) \sim \text{bandwidth} \times T_{delay}$. This means for instance that if the cavity bandwidth is 500 Hz and we require 100 fold reduction in phase error arising from microphonics then delays in the control system can be as much as 20 μm. A permissible delay of this magnitude allows the use of a digital control system. In order to control the phase of a copper cavity with a digital control system one typically needs to include a low pass filter in the control loop to limit the bandwidth.

5. Digital Control Options

Two opportunities arise from using a digital control system. The first is to implement an intelligent controller which might have adaptive parameters, feed forward elements, smart calibration or advanced filtering of the error signal. The second opportunity is to estimate the phase and amplitude of the cavity fields by sampling a down converted waveform. The down conversion process preserves phase. Both for the LHC at CERN, Geneva and the Tesla Test Facility (TTF) at DESY, Hamburg, cavity amplitude and phase are determined by down conversion and digital sampling[6]. The TTF uses the same technology as planned for the ILC main linacs. TTF down conversion is from 1.3 GHz to 250 kHz. At TTF an adaptive feed forward controller is used. Feed forward is pulse to pulse with a 5 Hz repetition. Proportional control only is used during the pulse to correct disturbances with cannot be predicted. For the LHC down conversion is from 400 MHz to 20 MHz.

The remaining sections of this paper focus on a phase control system being developed at Lancaster University as a possible solution to phasing the ILC crab cavities. A feature of our approach is to use the Hittite digital phase detector HMC439QS16G rather than the more elaborate method of down conversion and digital sampling. Figure 7 gives a block diagram of the cavity control system being tested at this time.

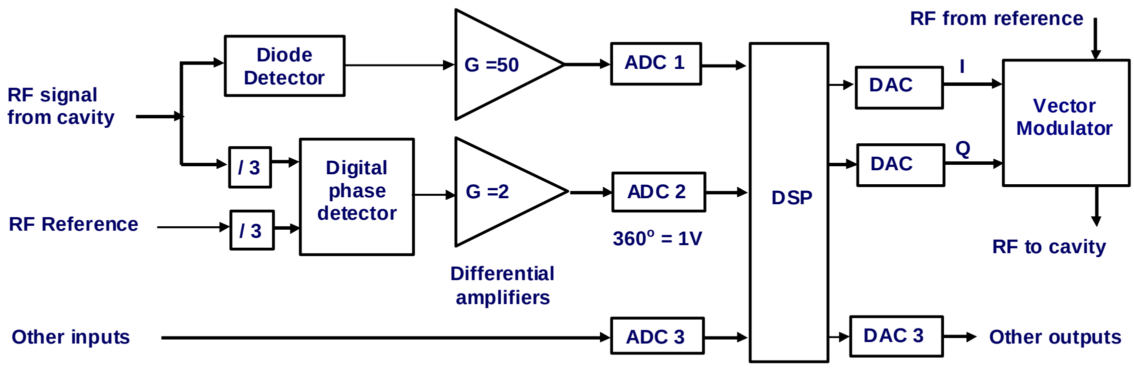

Figure 7 Digital Control Layout for Control of One Cavity

The phase noise of the HMC439QS16G at 1280 MHz is about -135 dBc/Hz and is relatively flat hence phase noise in 1 MHz bandwidth is about -80 dBc equating to r.m.s. phase jitter of 8 milli-

[6] Jalmuzna W., Napieralski A., Simrock S., "Performance of 24 Cavity Vector Sum Controller with Distributed Architecture", EPAC08 http://accelconf.web.cern.ch/accelconf/e08/papers/tupc142.pdf

degrees = 17 fs. This is quite large compared to what one might achieve with a double balanced mixer or with digital sampling of a down converted wave however it is still significantly less than the timing requirement of 90 fs. Frequencies greater than 1 MHz have virtually no effect on the cavity phase jitter performance where a superconducting cavity with a bandwidth close to or less than 1 kHz is used. Digital phase detectors only operate up to a frequency of 1.3 GHz hence they must be used with frequency dividers (*8 milli-degrees of phase jitter at 1.3 GHz implies 24 milli-degrees of phase jitter at 3.9 GHz*). The 3.9 GHz signal is frequency divided by three using HMC437MS dividers, these generate an additional 2 milli-degrees r.m.s. phase jitter at 1.3 GHz.

The configuration of figure 7 uses three 16 bit 105 MBPS ADC inputs (latency = 130 ns) and three 16 bit 40 MBPS DAC outputs (latency = 10 ns). Two of the DACs provide I and Q signals for the vector modulator AD8341 with an output noise floor of 150 dbm/Hz giving just 2 milli-degrees r.m.s. phase jitter for an input level of 0 dBm.

6. LLRF Interferometer Description

Because the crab cavity systems (one on each beam) are separated by a signal path length of at least 50 m; temperature fluctuations and vibration on the RF line between them gives significant synchronisation uncertainty and jitter. A movement of 0.1 mm for instance will give a timing shift of 333 fs which corresponds to 0.47^{o} at 3.9 GHz. In order to eliminate synchronisation errors on the RF it needs to be operated as interferometer. The complications with an RF interferometer are line losses and reflections from RF components.

If a standing wave is established on a perfect transmission line then the phase of the oscillation at every point is identical hence synchronisation is trivial. As transmission lines are imperfect i.e. waves decay, a perfect standing wave can never be established. Phase synchronisation between two points can still be achieved using phase lock loops as illustrated in figure 8.

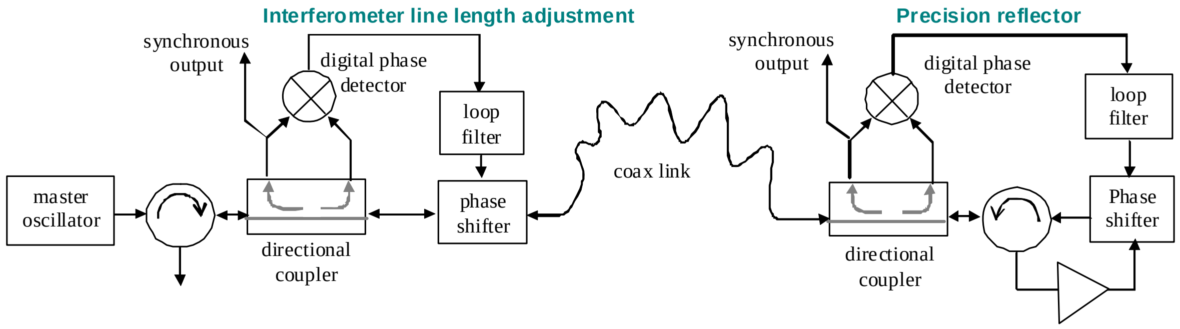

Figure 8 Interferometer giving synchronised references either end of a long coax

Assume the forward wave varies on a transmission line as $\cos(\omega t - kz)$ so that the phase at z_1 is kz_1. If the forward wave is sent back with unknown phase shift ϕ then the backward wave varies as $\cos(\omega t + kz + \phi)$ hence the phase of the backward wave at z_1 is $-kz_1 + \phi$. Using a phase lock loop the phase difference between the forward and backward waves at the point z_1 is controlled to α hence $2kz_1 + \phi = \alpha$. Using a second phase lock loop the phase between the forward and backward waves at a point z_2 is controlled to β hence $2kz_2 + \phi = \beta$. Eliminating the unknown phase shift on reflection between these two equations gives $kz_1 - kz_2 = \frac{\alpha - \beta}{2}$. Since kz_1 is the phase at z_1 and kz_2 was the phase at z_2 then the phase difference can be set between these two points with suitable choice of α and β.

With respect to figure 8 the digital phase detectors give d.c. output voltages when the RF inputs have identical frequencies. The loop filters drive both LHS and RHS phase differences to zero by varying on the LHS the coaxial line length and on the RHS the length of the loop that returns the input signal. The digital phase detectors have limited input range between -10 dBm and +10 dBm. Whilst this is sufficient for the return signal to be measured back at the LHS phase detector, one needs the return signal to be large with respect to unwanted reflections so that systematic corrections stay small. For this reason an amplifier is used on the return loop.

A difficulty with the technique described is that reflections on the line give systematic errors. Indeed our implementation uses an electronic phase shifter whose reflection coefficient varies with the phase shift it is providing. Our proposal is to correct these errors with a DSP using knowledge of line length corrections being made to maintain α and β.

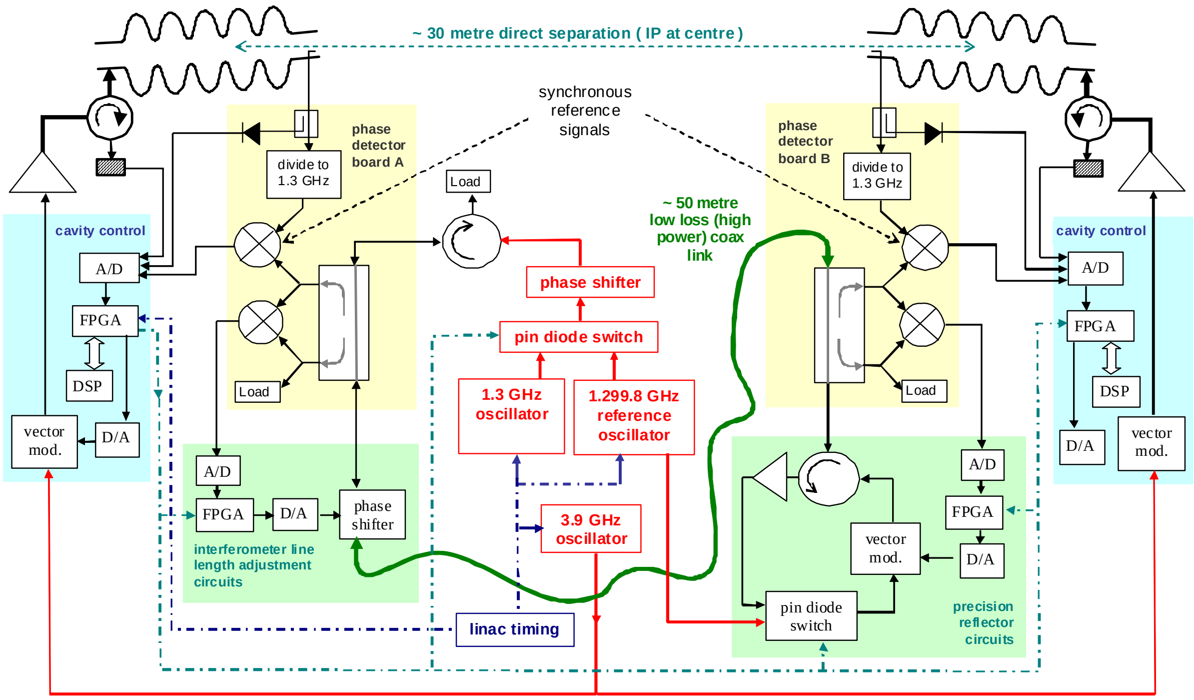

Figure 9 Full Layout for Cavity RF Synchronisation

Figure 9 shows how the full system can be implemented in a way that corrections can be applied and calibrations made. The loop filters are replaced with a digital controller composed of a ADC, an FPGA and a DAC. The voltage being applied to the phase shifter on the coaxial line determines the correction required to correct for reflection from the various RF components. A phase shifter is required on the coaxial line as waves must pass in both directions. On the precision reflector loop ideally there should only be a forward wave. A vector modulator can now replace the phase shifter in figure 9 as the FPGA can produce appropriate I and Q inputs. Use of a vector modulator allows the return signal to be amplitude modulated. This is useful in checking calibration because the part of the backward wave arriving at the phase detector adjacent to the oscillator that is amplitude modulated must come from the far end rather than from reflections at the phase shifter or the directional coupler. Using the 1.2998 GHz oscillator as a reference to the 1.3 GHz allows phases to be swung by 360^{o} giving ADC calibration.

The two reference points for phase synchronisation are the centres of the phase detectors. The distance between the phase detectors and the centre of the cavity is an uncontrolled length. To avoid additional phase jitter or error in absolute calibration these lengths should be minimised. All the components in each yellow box of figure 9 are to be mounted on a single board in a temperature controlled and vibration free environment. A prototype PCB with directional couplers, Wilkinson dividers, phase detectors, low noise differential amplifiers and the divider to reduce the 3.9 GHz signal to 1.3 GHz has been developed and tested.

7. Results to date after preliminary experiments

The locking performance of two superconducting cavities in a vertical cryostat was studied at STFC, Daresbury in August 2008. Figure 10 shows the output spectrum from cavity 2 as measured on an

Agilent Spectrum Analyser Model E4443A for its locked and unlocked states with respect to the source. When the system is unlocked spectral peaks associated with microphonics can be seen at 16 Hz, 29 Hz and 41 Hz. The latency of the control system was about 3 μs.

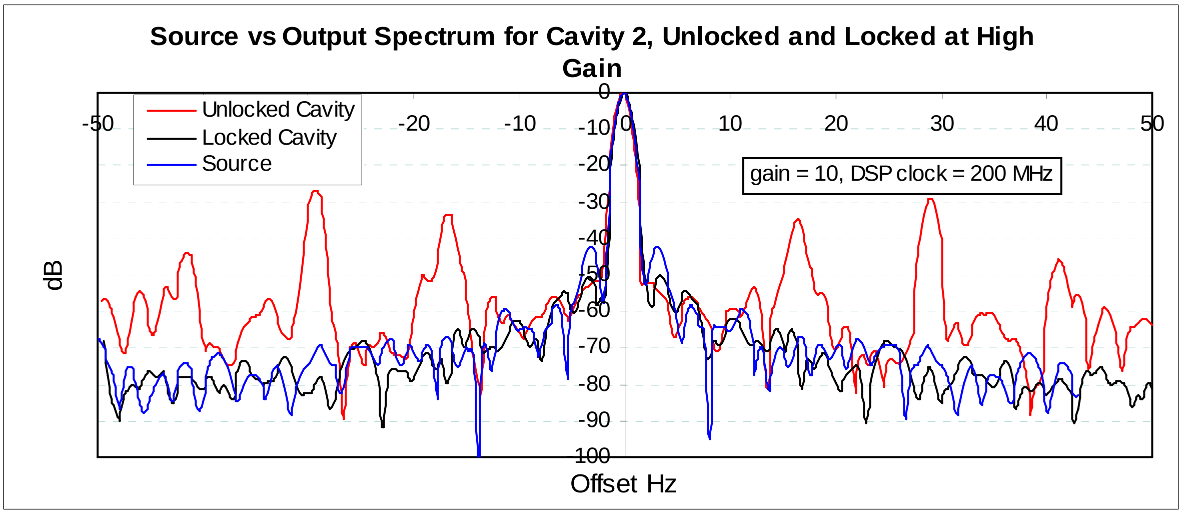

Figure 10 Cavity and source spectral output

Figure 11 compares the input drive to the cavity when it is locked to the output when it is unlocked. As one expects the frequency spectrum of the corrective action of the drive is precisely correlated with the frequency spectrum of the unlocked cavity.

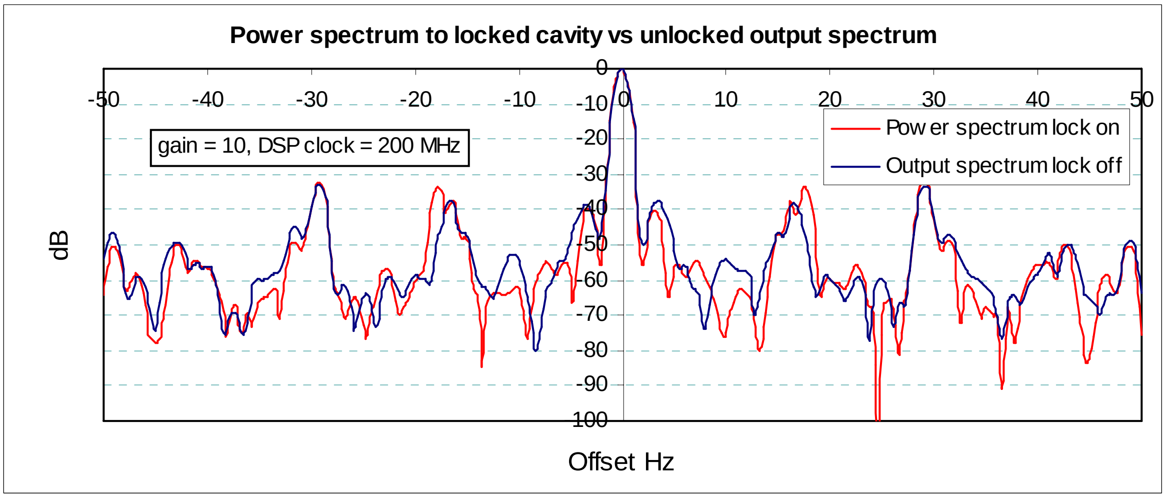

Figure 11 Locked input and unlocked out spectra compared

Figure 12 shows the locking performance as a function of gain for an increased bandwidth of 5 kHz. Interestingly figure 12 shows that at low gains the LLRF system can create additional unwanted phase noise above 500 Hz with respect to the unlocked cavity. When the gain is increased to 10 the cavity phase noise is again comparable with the source.

As well as inferring phase jitter from spectral output it was measured directly with respect to the source using frequency dividers and a digital phase detector. The phase detector was calibrated by splitting a signal from the signal generator, shifting one leg with a calibrated manual phase shifter and comparing the phases of the two legs in the phase detector after division of each. Calibration was dependent on the gain of the differential amplifiers used with the digital phase detectors. In this instance the calibration gave 7.5 mV per degree.

Measuring the phase jitter between the cavity and the source for the unlocked cavity and with a bandwidth of 500 kHz gave a peak to peak signal output of 225 mV corresponding to 30^{o} peak to peak which is about 10^{o} r.m.s.. When the cavity was locked the peak to peak noise was better than 3 mV on the same bandwidth. This implies a peak to peak jitter of 400 milli-degrees peak to peak or about 140 milli-degrees r.m.s.. This jitter includes source noise, ADC noise and some oscilloscope noise. Integration of source noise as measured on the spectrum analyser indicated that it was typically 140 milli-volts r.m.s. hence we can say is that the locking performance was substantially better than 140 milli-degrees r.m.s..

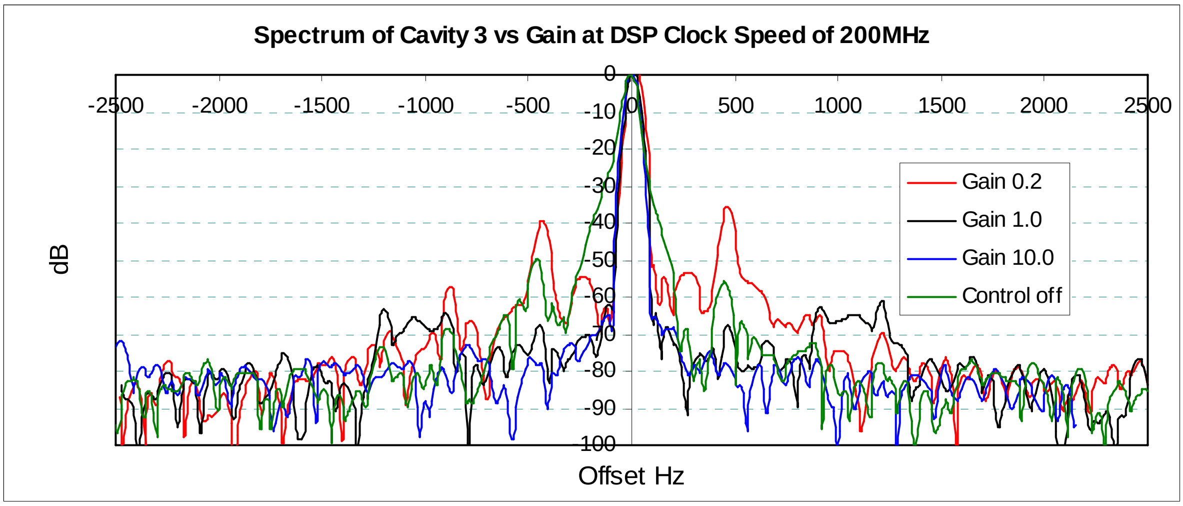

Figure 12 Locking performance verses gain

Having demonstrated lock with each of the cavities separately it was necessary to bring them to the same natural frequency to lock them together. Only one cavity had been given a tuner and problems with the design of this tuner resulted in the natural frequency of the cavity drifting by many bandwidths over a period of minutes. As both cavities were driven from the same source they had same output frequency and when one cavity's natural frequency was a few bandwidths from the drive frequency, its output was too low for the phase detectors to function correctly.

Whilst simultaneous lock was achieved on a number of occasions our skill at watching the power meter and adjusting the tuner knob was inadequate to hold the lock for a sufficient period to perform careful calibration of the double balanced mixer. In spite of this approximate results were obtained. The target phase control performance at 3.9 GHz was 120 milli-degrees r.m.s.. With the cavities locked the voltage jitter from a double balanced mixer between the cavities was nominally 50 mV peak to peak as measured on an oscilloscope over a period of 10 ms. Prior to the locking the double balance mixer was typically measuring 2 volts peak to peak. The system was reducing the jitter by a factor of 40. The microphonic level was about 4 degrees r.m.s. This meant that 50 mV peak to peak on the oscilloscope corresponded to a cavity to cavity r.m.s. phase jitter of 100 milli-degrees. The target of 120 milli-degrees had been met. For warm cavities we achieved 75 milli-degrees. Given that the effective jitter (translated to 3.9 GHz) from the digital phase detectors was probably at least 25 milli-degrees and four were in use then these alone contribute 50 milli-degrees. Adding contributions from dividers and ADCs then 75 milli-degrees is about the limit on what one might expect to achieve with this system.

8. Acknowledgements

Testing of the control system was made possible by the facilities and staff effort provided by ASTeC at the STFC Daresbury Laboratory under the management of RF Group Leader P. McIntosh. The work has been funded by STFC.