Chapter 6

Conclusions

Although the extensive simulations carried out by Steve Smith, as detailed in Section 3.4, go a long way to demonstrating the effectiveness of the IPFB system, the results from the FONT experiment prove without doubt that such a system is feasible. Indeed, it must be said that the FONT system tests were a surprising success: not only was it possible to apply a rapid correction to a measured beam offset, but also to do so within the obvious limitations of the experiment (see below) was proof of the successful design of the system.

A new BPM had to be designed and fabricated to allow the measurement of the NLCTA X-band bunch train. Not only did this BPM design comfortably achieve the required design goals in making a continuous position measurement of a beam bunched at 11.424 GHz — something that, for the long-pulse beam used for FONT, had been previously achieved only with cavity BPM's — but was also able to outperform the existing NLCTA striplines in its measurement of the average bunch train position, with only rudimentary electronics. The BPM was also extremely simple in concept and construction and relatively cheap to produce, requiring virtually no modification once it had been installed onto the beamline.

The FONT experiment itself was marked by a number of successes. The BPM electronics, notwithstanding the overshoot problems documented in Chapter 4 (also see below), was able to deliver a sum and difference signal of high quality for use by the feedback electronics. The feedback circuit performed admirably, behaving exactly as expected for both the bench tests and the real beam tests: while the output of the normaliser stage was certainly limited by the bandwidth of the design, the theory of operation behind the feedback circuit was proved to be sound. The fast kicker amplifier, designed and constructed by Colin Perry, was by far the most technically challenging aspect of FONT, since a power output far in excess of the IPFB, with similar speeds, was required: yet once installed the amplifier performed without problems and enabled the FONT system to demonstrate the soundness of the IPFB concept. The contribution of all of these factors enabled FONT to demonstrate the effective behaviour of the IPFB concept in providing a rapid correction of the position of a real bunch train. The FONT experiment was able to prove the IPFB principle for both a single correction and the additional delay loop corrections.

However, there are a number of factors that contribute to the poor quality of some of the recorded data, and hence the reason why the FONT experiment should not be considered as an unqualified success. Three problems above all else hindered the operation of the FONT system:

1. The relatively poor quality of the NLCTA long pulse beam.
2. The BPM overshoot effect detailed in Section 4.8.3.
3. The long latency time associated with the FONT setup.

The problems with the NLCTA beam were twofold. Firstly, the large intra-train position variation inherent in the long pulse beam meant that the FONT system had to struggle with problems that it was not designed to fix. As mentioned at the end of Section 3.4, the IPFB system does not deal well with random bunch-to-bunch noise. It was necessary to tune the beam extensively, over a number of hours, before the position variation had been reduced to a level at which it was possible to take data. This problem was exacerbated by the relatively small voltage output of the kicker amplifier: since the full range over which FONT could operate was set by the output voltage swing of the amplifier, reducing the dipole range and increasing the input gain would enhance the relative size of any intra-pulse position variation. However, the cause of this problem is still inherent in the beam, and given the exemplary performance of the kicker amplifier the largest area for improvement should lie with the quality of the beam.

Secondly, the usable lifetime of the beam was no more than a matter of minutes. The beam quality, train length and charge distribution proved to be highly sensitive to the time of day, primarily due to the large external temperature variations which resulted in the RF delivery system requiring frequent retuning. The rapid degradation in beam quality after a few minutes severely limited the repeatability of some aspects of the tests: for the results detailed in Sections 5.5.4 and 5.5.5, the beam required retuning between each of the five dipole measurements. It could therefore be argued that each stage of these tests was not carried out with the "same" beam, since there was a continuous variation of factors external to FONT's control. As such, it is essential that for any further tests of the FONT system, the beam quality and lifetime is improved significantly to allow continuous measurements to be made over the full operating range of FONT. This was compounded by the long time taken for the data acquisition process which, primarily due to the amount of time required to read from the scopes and write to the AWG via GPIB, took almost 5 seconds to record a single beam pulse.

The problems associated with the method used to process the raw BPM output are likely to have contributed to the large intra-pulse position variation. The overshoot seen on the difference signal, as detailed extensively in Sections 4.7.2 and 4.8.3, were a result of the timing design of the BPM processor and would have been present, to a greater or lesser extent, on any measured beam position with a rapid charge variation. This essentially places any position measurement with a rapid charge variation in jeopardy. For example, the slight dip in the position signal that appears at the start of all the normalised position plots shown in Sections 5.5.4 and 5.5.5 is almost certainly a result of the overshoot phenomenon. It is therefore essential that, for any future test of the FONT system, the BPM processor is replaced with an improved version that does not experience such problems. Whether such a processor is feasible, as well as being capable of matching the performance of the current BPM processor in other areas, is a matter for future research.

The long latency time of the FONT system — some 70 ns for the full round trip time — was very close to preventing a complete test of the IPFB design. A previous design of the FONT system, with a latency of almost 90 ns, was incapable of taking advantage of the delay loop, since the entire bunch train would pass by the time that the delay loop correction could be applied, after 180 ns. Even with the improved design, detailed in Section 5.3, it was still only possible to observe part of a single delay loop correction. Although this was certainly useful, a short enough latency (or a long enough bunch train) allowing multiple passes around the delay loop would have provided absolute proof of the quality of the IPFB design. It was also necessary to push the NLCTA to its limits to get a bunch train that was long enough to observe the effect of the delay loop. It is likely that this also contributed to the unsatisfactory beam quality: while it was possible to produce a much cleaner, 100 ns long bunch train, this was of absolutely no use to FONT. For future FONT tests, assuming that the maximum train length remains limited to 180 ns, there are two ways to reduce the system latency:

1. Decrease the signal processing latency by improving the speed of the electronics.
2. Shorten the round trip time by moving the magnet assembly and BPM closer together.

Each of these options have their relative merits. Clearly (1) is more desirable, since it would require no beam line modifications of the current experimental setup. It would also keep the power requirements of the kicker amplifier to a minimum, since the full lever arm of the experimental layout is kept to a maximum. However, while it may be possible to shave a few nanoseconds off the latency by using faster chips in the normaliser circuit, or higher quality (and faster) cable it is unlikely that a large time saving — say, longer than about 10 ns — could be made purely by modifying the electronics. Indeed, the design as it currently stands might already be optimised from an electronics latency point of view, making further improvement difficult to achieve.

Reducing the distance between the magnet assembly and the BPM is a fairly simple way to reduce the system latency. Not only were the original component locations chosen to match the NLC, and to provide the maximum position offset at the BPM for a set kick strength, but the fabrication and installation of beamline components was also made cheaper and simpler: the magnet assembly could use quadrupole QD1550 as a solid reference point for alignment, as could the rather more flimsy BPM with toroid 1750. In addition, it was also only necessary to fabricate two new sections of beampipe rather than four. By moving the magnet and BPM closer together — ideally symmetrically about QD1650 — this immediately reduces the beam flight time and any cable delays for the signal return path.

However, one then runs into the problem of requiring a larger kick from the kicker amplifier to provide the same measurable beam offset: this is no easy task, given that the current amplifier was operating at its absolute limit to provide a kick that was only marginally larger than the inherent intra-train beam position variation. Increasing the power of the kicker amplifier would likely bring a corresponding decrease in bandwidth and an increase in the system latency, offsetting any attempted reduction of the system latency. Also, given that the bandwidth of the feedback circuit is lower than one would prefer (see Section 5.5.2), reducing the bandwidth further would be unhelpful. The factor that has of course been overlooked is the focusing strength of QD1650: by increasing the field strength of this magnet, the beam is more strongly defocused midway between the kicker and the BPM, enhancing the kick strength. This is also currently under investigation for future FONT tests.

A number of other problems were also not addressed, primarily as a result of the severe time constraints in which FONT had to operate. While bench tests were carried out on the normaliser stage of the feedback circuit, as detailed in Section 5.4, the performance of the attempted first-order correction to the charge signal was not rigorously tested. Given the surprising train-to-train stability of the charge profile of the long pulse beam, such a test would probably have been superfluous. In addition, the DAQ routines and the accompanying online analysis code was not optimised, contributing to the large (5 s) data acquisition period. Such optimisation may have improved the consistency of the data recorded, allowing complete tests — such as those described in Section 5.5.5 — to be carried out without the beam having to be retuned between measurements.

However, given all these problems, the success of FONT is significant. While operating under less than ideal conditions, the performance of the FONT system was remarkable. Even with the large position variation of the beam, its short lifetime, the long system latency and the inherent problems with the BPM processor it was still possible to demonstrate a feedback system that applies an iterative position correction to a beam within a single bunch train, based upon position measurements made solely upon that bunch train. This is the essential operating characteristic that sets the Smith IPFB design apart from other feedbacks: proof of the effectiveness of this design has therefore been provided in the shape of FONT.

There are a number of improvements that could be made to the FONT system. The first, as described above, would be the reduction in system latency. This would allow multiple passes around the delay loop and better assessment of the operating characteristics of FONT. Secondly, the introduction of a number of 'witness' BPM's, around the feedback BPM, would allow the position and angle of the beam to be measured. An extension of this idea would be to include a second kicker magnet, to not only resteer the beam and correct a position offset at the BPM, but also to remove any angle offset and straighten the beam through the BPM. Thirdly, it should be possible to take out the fine scale intra-pulse position variation seen on the long pulse beam using the FONT system. This would require a second AWG with the preprogrammed position profile that would also drive the kicker amplifier. Finally, there are suggestions that a solid state amplifier design could replace the tube-based kicker amplifier to provide an improved signal output and a larger drive voltage [43]. In addition, it is also recommended that an improved, possibly LabView-based, DAQ system is implemented to reduce the acquisition time.