Start-up Synchronization Procedure of the CMS Tracker

K. Gill and N. Marinelli
March 2001

Abstract
A scheme is proposed that will achieve relative synchronization of the APV chips throughout the whole tracker to the level of ~1ns in several simple steps, based on the measurement of the delays in arrival at the FEDs of tick signals generated at the APVs. This procedure can be carried out during start-up, without the requirement of having particle beams. After correcting for particle time-of-flight delays, all that is then required to bring the whole tracker into absolute sychronization, with respect to particles emerging from collisions, is an adjustment of the phase of the LHC clock into the Tracker and confirmation of the latency setting.

1. Introduction and overview of the synchronization procedure

Sensor modules, with their front-end electronics, are arranged in rings from the point of view of clock, trigger, reset and other control signals[1,2]. For optimum performance of the Tracker the whole system has to be well synchronized, such that all the APVs sample signals from incident particles from the same event, with a resolution of ~1ns with respect to the peak of the signal.[1] This level of synchronization can be achieved in several simple steps, which will be described in the following sections and illustrated using results from a preliminary simulation of the effect of the various delays that signals will encounter in the system.

The basic parts of the overall system are illustrated schematically in Fig. 1, containing all the elements that affect the overall synchronization. TTC signals (clock, trigger and reset) are sent to the Tracker partition(s), and on to the FEDs and FECs using digital optical links. The FECs pass on the TTC signals, in addition to slow control commands, to the various control rings, again using digital optical links.

In the tracker volume the timing and control signals are then received at digital opto-hybrids that communicate over copper cables to and from the attached control rings. The CCU modules, also linked by copper cables within the rings, then distribute signals to (and collect signals from) the front-end hybrids that contain the APVs and the analogue optohybrids that contain the laser transmitters.

While receiving a clock signal, but in the absence of trigger signals, the APVs generate only synchronization pulses, otherwise known as 'tick marks' (or 'ticks'), every 1.75μs. The ticks are approximately 8mA in amplitude, passed as an 800mV differential signal into the analogue optical link. The ticks can then be digitised at the FED if the FED is triggered. The ticks are output after a fixed delay following the arrival of a reset signal at the APV[2]. Therefore they can also be used as a precise probe to determine the relative timing offsets between different APV chips.

The first step in the synchronization procedure at the start-up of the Tracker involves measuring the relative delay of the APV ticks arriving at the FED. This is illustrated in Fig. 2. For the APVs in a given control ring, the relative delay between signals arriving at the FED is due to any difference in the analogue optical fibre lengths, added to the difference in the delay in the signals going around the control ring to the particular APVs.

Once the relative delays in the arrival time of the ticks at the FED have been measured, the contribution from the length of the analogue optical link can be subtracted, assuming that the length of fibre is known. This leaves the delay due to the position of the APV in the given control ring. The programmable delays at the various PLLs around each ring can then be set to compensate for the measured time-offsets at the front-end chips in each ring. This involves setting the delays at each PLL to match the propagation delays of signals to the last group of APVs in the control ring. This step will bring all the chips in a given control ring into synchronization.

---

[1] This timing resolution is critical when the chips are used in 'deconvolution' mode. Each 1ns deviation from the optimum sampling point leads to ~4% loss of amplitude. The synchronization can be less stringent, typically ~5ns, when operating with the APVs in 'peak' mode.

[2] We assume that reset commands arriving at the FEC are processed in a reproducible manner that maintains the synchronicity of the system once all the required delays at the front-end have been programmed.

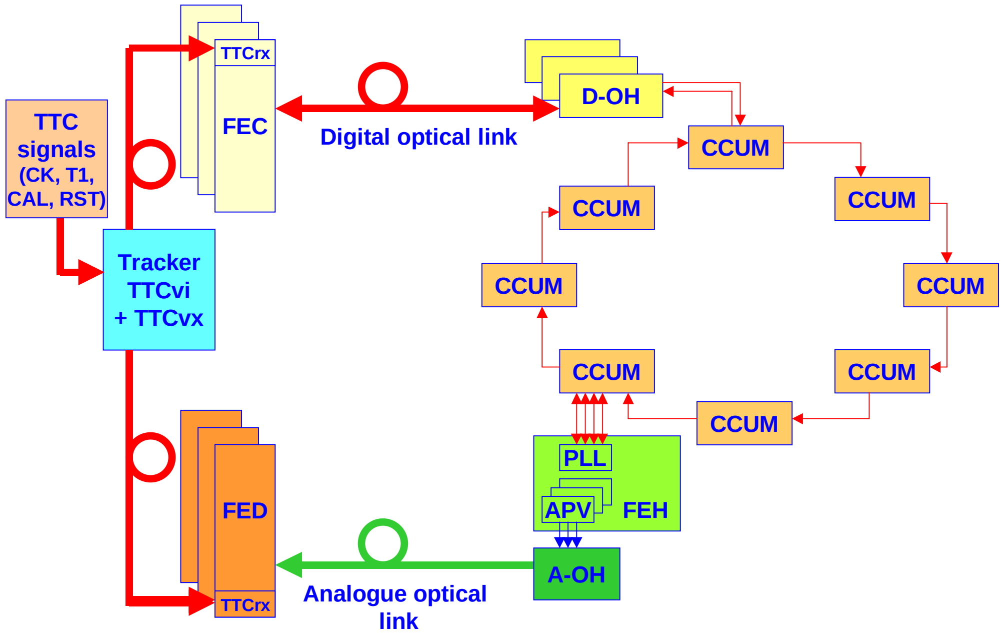

Figure 1: Schematic diagram of the tracker control and readout system showing the elements that affect the synchronization. The TTCrx and PLLs all contain programmable delays that are used to synchronize the system.

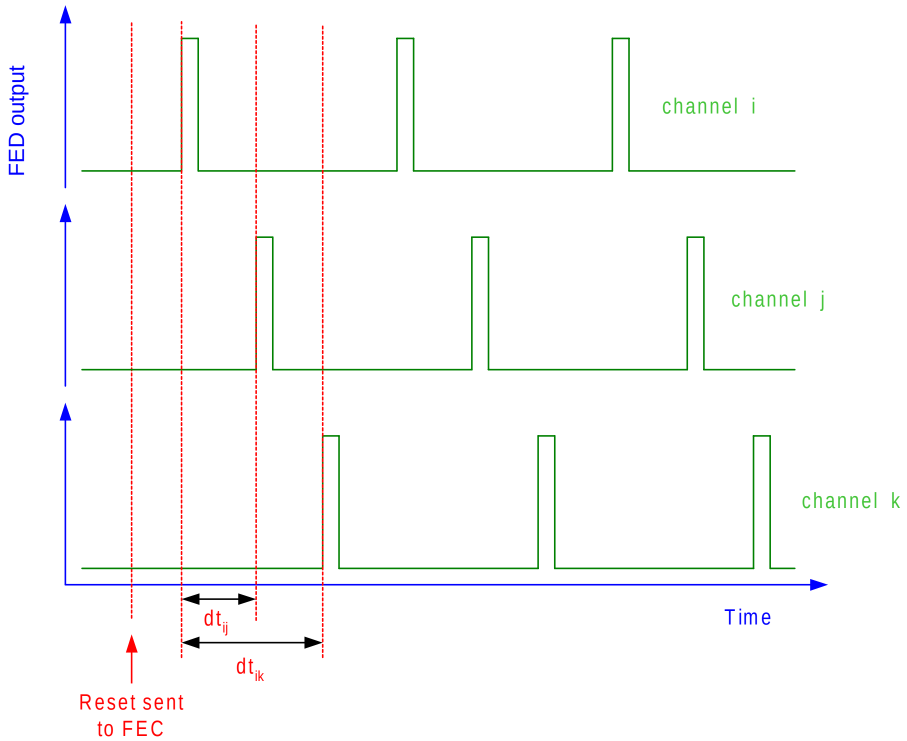

Figure 2: Schematic illustration of tick pulses arriving at a FED from APVs in a given control ring. The differences in arrival times (e.g. dt_{ij} and dt_{ik}) are due to the relative differences in the lengths of the analogue optical links combined with the delays between the APVs at different positions around the control ring.

Having carried out this procedure separately on all the control rings in the Tracker the method can then be extended in a simple manner to synchronize between the different rings throughout the whole Tracker.

The delays that remain between APV chips on different control rings are due to the differences in fibre lengths in the digital optical links transmitting TTC signals to the FECs (and FEDs), and then from the FECs to the control rings inside the Tracker. The lengths of optical fibre in the digital links must therefore also be known, in order to compensate for the delays that they introduce.

In principle, given the fibre lengths, the relative delays between channels can be calculated and then compensated at the adjustable programmable delays available at each FEC, FED and at the PLLs around the control rings. However to avoid errors, this part of the procedure should be carried out by repeating the scan of tick arrival times at the FEDs, at the level of one channel per ring (since each ring is already synchronous). In this way the delays between rings can be measured, compared with the expected values as a cross-check, and then compensated using the various programmable delay settings. Again, the compensation step involves setting the appropriate programmable delays so that all channels match the channel that has the longest delay in the system.

Finally a further scan of the ticks should be made to check that the only remaining delays in the arrival of ticks at the FEDs correspond precisely to the contributions due to the analogue optical links. After following all of the above steps, all the APVs in the Tracker should be emitting tick marks synchronously.

It should be noted that for each change of programmable delay that affects the output of a given APV, the clock skew on the corresponding FED channel must also be readjusted. This is so that sampling still occurs at the correct point on a signal level, allowing the optical link output enough time to settle.

It is assumed that the lengthes of the optical fibres in the system will be known to a precision of ~10cm. These can be measured either by the firm that cuts the fibre into prescribed pieces, or by OTDR[3] measurements on the fibre after installation in the Tracker. A 10cm precision in the measurement introduces a systematic uncertainty of ~0.5ns per fibre segment.[4]

In contrast, it should also be noted that the length of the electrical cables used to connect to the CCUMs, and between CCUMs, do not necessarily have to be precisely known. This is because the tick marks will provide an accurate measurement of the delays due to the copper cables (added to any propagation delays of the clock signals across the front-end hybrids). However, if there are any significant differences in the lengths of cable used to connect between the front-end hybrids and the analogue opto-hybrids, then the delay contribution from these cables should be known to reasonable (sub-ns) precision.

The remaining step that can be made in the absence of the beam is to apply a time-of-flight correction to compensate for the expected delay in the arrival of particles at the sensors. Once the beam becomes available and collisions occur, a scan of the APV latency setting, plus a scan of the phase of the clock into the Tracker, will be the final steps in order to achieve absolute synchronization of the Tracker.

2. Start-up synchronization method

2.1. Step 1 - Measurement of the relative 'fine' delays using the FED skew

This first step of the synchronization is actually an important part of setting up the optical links, which must be carried out as one of the first items in starting up the Tracker, following power-up and cooling. It is a requirement that the analogue levels being transmitted through the links have settled to a certain level before the FED samples the signal. The FED clock skew is therefore optimal, with respect to the settling of signals transmitted by the optical link, if it is set to coincide near, but not too close to, the end of a sample. This is shown in Fig. 3, which illustrates sampling of data around two interleaved APV tick-marks.

Setting up the FED clock skew to sample correctly the output of the optical links therefore also provides an opportunity to determine, at the same time, the 'fine' delays required at the PLLs in order to synchronize the signals being output from the APV. It should be noted that even if the synchronization of the Tracker is eventually performed in a different manner,[5] the measurement of the ticks remains essential to the setting up of the analogue optical links.

---

[3] OTDR = Optical Time Domain Reflectrometry

[4] It will be necessary to also measure carefully the various short pigtail lengths during module construction, in order that these fibre segments do not add significantly to the uncertainty in the delay in the analogue links.

[5] e.g. by using a laser beam to generate independent timing reference signals at the sensors.

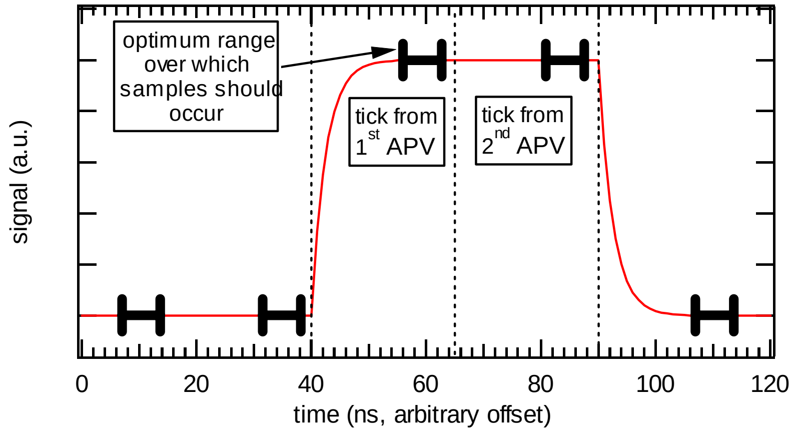

Figure 3: Tick-mark profile and illustration of optimum FED sampling point, that allows for sufficient settling of the analogue signal output from the optical link. Two ticks are shown interleaved from neighbouring APVs where each APV is 'ideal', generating ticks with the same amplitude and baseline.

This first synchronization step is done by reconstructing the shape of the tick marks in time, which can be achieved by sweeping the clock skew on all FED channels, and acquiring non-zero suppressed data at each skew setting. The FED needs to be operated in scope mode and triggered with a random (or software) trigger signal. It is assumed that to avoid having APV frames sampled with the ticks, the APVs triggers can be suppressed such that triggers are only sent to the FED. However, if it is unavoidable that the APVs are triggered, the samples that contain the header and data-frame could be discarded, by appropriate cuts applied in the software, so that only the part of the APV output containing ticks is analysed.

By histogramming the sets of samples taken at each FED skew setting, an image of a tick can be built up in several parts. This is illustrated in Fig.4, which shows simulated results of the sampling of APV ticks arriving in one FED channel, for each clock skew setting[6].

The rising edge of the tick can be traced between the histograms in Fig. 4. The FED clock skew setting that contains samples near the mid-point of the rising edge can then be used as an accurate measurement of the relative 'fine' delay from channel to channel. This fine delay will be combined with the 'coarse' delay, measured in the next step, to determine the overall delay between APVs around each ring.

The data in Fig. 4 are based on a simulation of 2.5ns FED skew steps, as in the current FED PMC. To measure the fine delay between ticks more accurately, the data could be interpolated between different FED skew steps in order to improve the measurement of the mid-point of the tick rising edge. However, if 1ns steps are available at the FED skew, interpolation would not be necesary, and this granularity would match the fine adjustment available in the programmable delay settings.

---

[6] More details of the simulation are given in the Appendix. The data in Fig. 4 are based on simulated FED sampling of ticks that have the same shape as those in Fig. 3, with nominally 100cts amplitude, 2.5cts average noise (due to the analogue optical links) and the accumulated jitter from five elements each with 0.2ns average jitter. It should be noted that in the final system, once the links have been correctly set up, the ticks will have an amplitude of ~650cts.

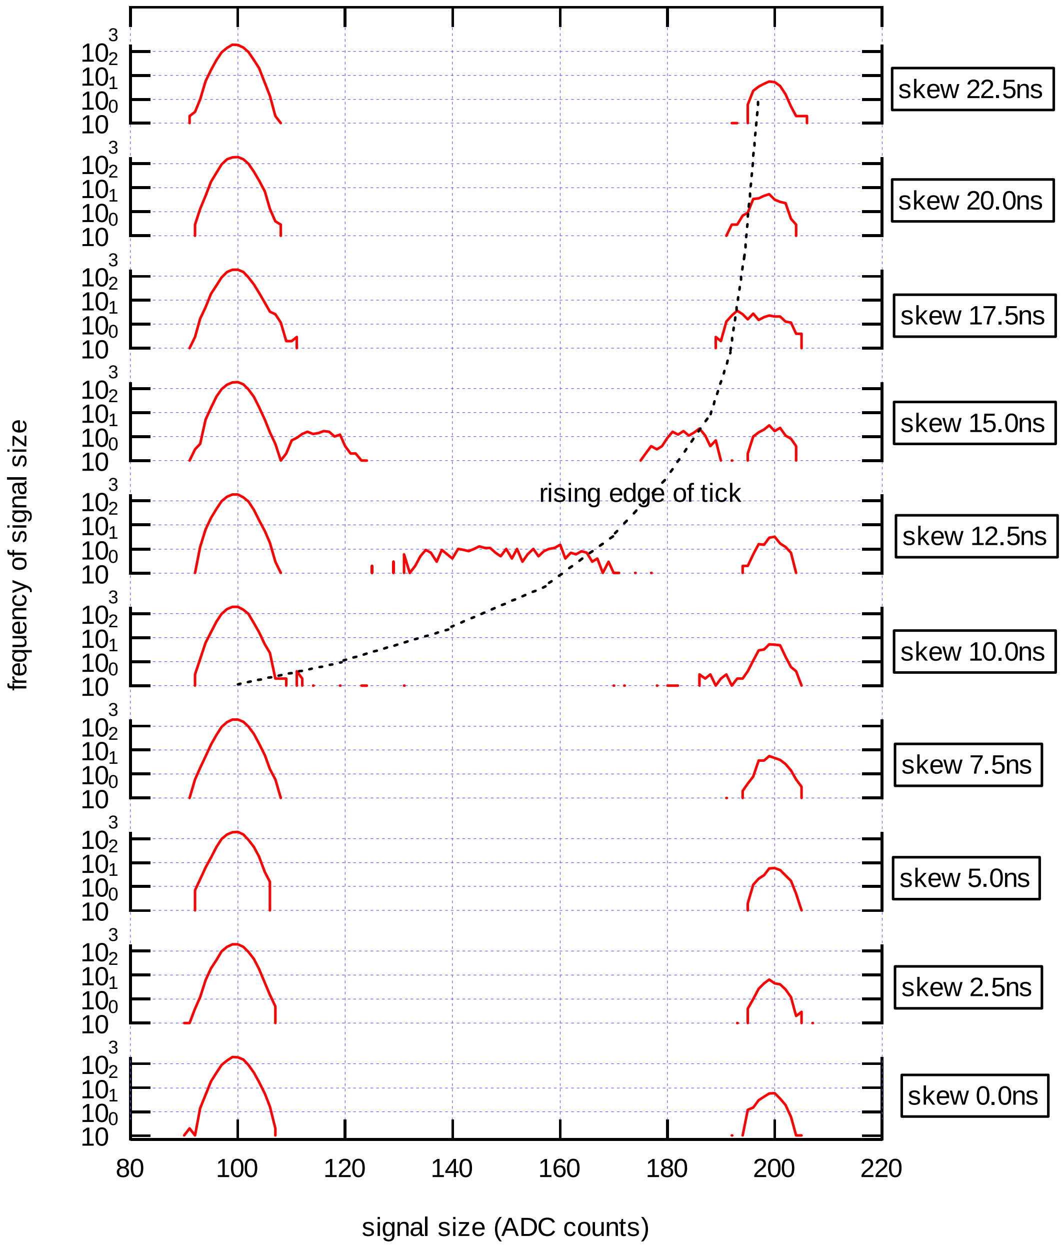

Figure 4: Histograms of APV output ticks in a simulated system where the FED skew step size is 2.5ns.

2.2 Step 2 - Measurement of the relative 'coarse' delays of ticks arriving at the FED

In the previous step the fine delay, i.e. the FED skew corresponding to the rising edge of the ticks, was measured for each channel. In this step the 'coarse' delay is measured, corresponding to the delay between the arrival of ticks at each FED in units of 25ns.

One way of measuring this coarse delay is simply to record the number of the sample for each FED channel corresponding to the peak due to the first tick measured at the FED following a trigger. This is illustrated in Fig. 5.

Again the FED needs to be operated in scope mode, with non-zero suppressed data. Ideally no triggers should be sent to the APVs. However, if this is not possible, the presence of APV frames would not affect this method since the sample number corresponding to the start of the digital header could be recorded instead of the first tick.

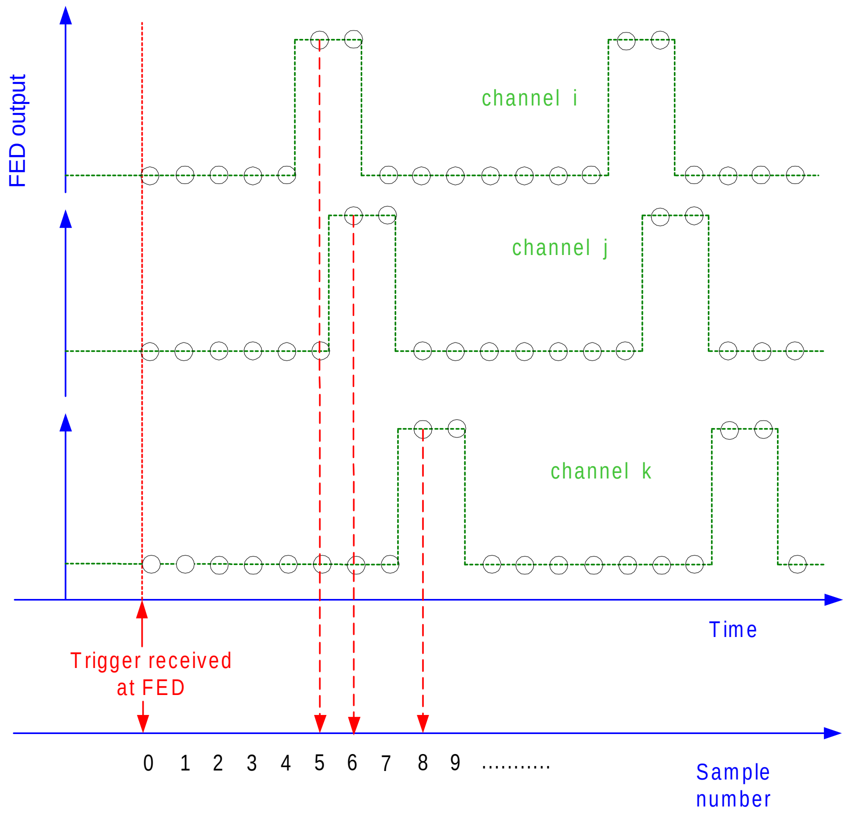

Figure 5: Schematic illustration of the 'coarse' relative delay measurement. The sample number where the first tick is found is recorded for each FED channel. In this example, channels i, j and k, give first peaks in the data frame, that correspond to ticks, in sample numbers 5, 6, and 8 respectively.

This part of the procedure was simulated for a part of the system containing one FED, one FEC, and two control rings, with details of the simulation and preset delay parameters due to fibre and electrical cable lengths, given in the Appendix. The output from the simulation of this part of the procedure is shown in Fig. 6. In this plot the left-hand axis is the sample number where the first peak in the signal, due to a tick, was found and the right hand axis is the optimum FED skew setting with respect to the optical link (defined as 7.5ns before the end of an analogue level). In this particular configuration, the effect of the simulated delays due to the fibres and the electrical cables (with increasing length around each control ring) is observable as an increase in the FED clock skew setting that eventually leads to an increment in the sample number each time the maximum FED skew setting is passed. There is a discontinuity after channel 48 when the output moves from APVs in the first control ring to the second ring.

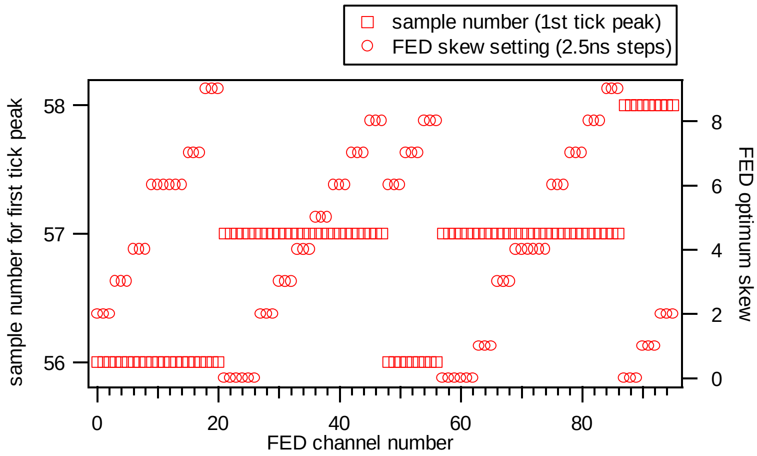

Figure 6: Sample number corresponding to the first tick peak measured in the simulation.
The configuration is such that the tick delay increases around a given control ring.
The FED skew setting (optimised with respect to the optical link settling time) is also shown.

2.3 Step 3 - Calculation of relative delays at the APVs around each ring

The relative delay between APV channels around a given ring is therefore calculable at this point. It is the difference in the number of samples between where ticks were detected on different FED channels (multiplied by 25ns) added to the difference in the FED clock skew settings (multiplied by the skew step size which is 2.5ns in the simulation).
This is illustrated by the simulation results shown in Fig. 7, based on the data shown in Fig. 6. FED channels 1-48 are shown, i.e. the relative delays of ticks arriving at the FED from APVs in the first control ring. Channel 1 of the FED is used as the reference when calculating the relative delay for APVs in the first ring, such that the delay in FED channel 1 is defined to be zero.
The relative contributions from the different lengths of fibre in the analogue optical links are also shown.[7] These contributions are based on the difference in lengths between each channel and the length of the fibre connected to FED channel 1, where each 10cm of additional fibre adds 0.5ns delay. The lengths of fibre used in the simulation are listed in the Appendix.

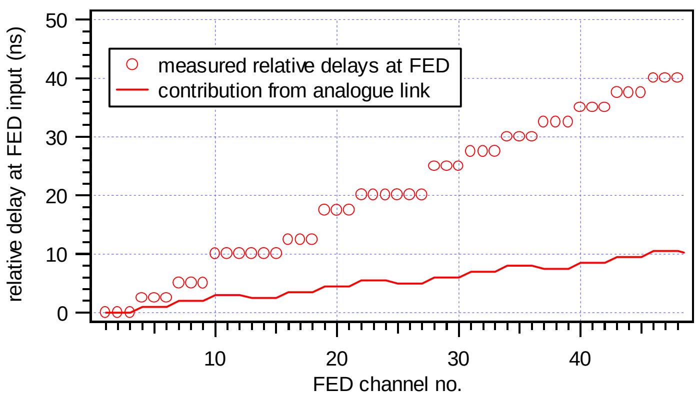

Figure 7: Delays measured from channel to channel at the FED, along with the contribution due to the differences in the lengths of fibre in the analogue optical links.

[7] No uncertainty was included yet in the length of the fibre. If it is not actually measured in the final system this may contribute up to ~0.5ns uncertainty in the delay, given that the lengths are specified typically to be within 10cm of a prescribed value.

The relative delay contribution due to the fibre can then be subtracted, leaving the relative delay between signals output at the APVs, i.e. the delays at the APVs due to their different locations around the control ring.
In Fig. 8 the delays measured in the simulation for the APVs around the first control ring are compared to the actual delays between CCUs that were input into the simulation. In all cases the measured delay is within ~1-2ns of the delay expected for that channel. The level of precision is dominated by the 2.5ns FED clock skew step size, such that the accuracy of the relative timing measurement would improve if finer FED skew steps were used.

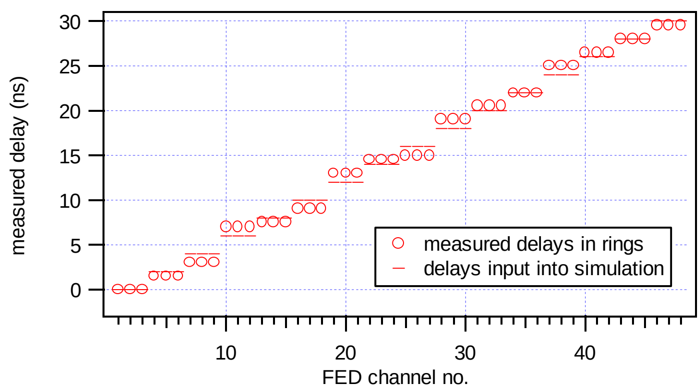

Figure 8: Relative delays measured in the simulation of the procedure for the APVs in the first control ring that are connected to FED channels 1-48 (after subtraction of the relative delay contribution due to the analogue optical link). The lines in the plot indicate the relative delay between channels that was input into the simulation.

2.4 Step 4 - Compensation of the delays and synchronization of the entire Tracker

The coarse and fine delay settings available at the PLLs on the front-end hybrids can then be adjusted in order to synchronize the APVs on a given ring, based on the previous fine and coarse measurements of the tick delays. The PLL settings should be adjusted such that all the chips are aligned in time with those connected to the CCU that is the furthest around the ring, i.e. the module that receives the TTC signals with the longest delay. It should be noted that any adjustments made to the fine delays at the PLLs have to be matched by adjustments of the FED clock skew, otherwise the sampling point will not be in the correct position. It would therefore be ideal if the FED skew step matched the PLL fine delay step, i.e. 1ns intervals.
A safe method would then include a check to make sure that the PLL adjustments have been done correctly, simply by re-measuring the delays of the ticks into the FEDs, i.e. redoing the first two steps, checking that the APVs are indeed synchronized.
The previous step of the procedure described in Section 2.2 also provides the key to synchronizing the whole system. Once the APVs on individual control rings have been synchronized, then the remaining step is to synchronize from ring-to-ring. This can be acheived by sending hardware triggers to the (whole) Tracker and measuring the relative delays of tick signals into each FED. This requires the analysis of only one channel per ring, again measuring the sample number corresponding to the first tick peak.
To synchronize the whole Tracker in this step it is necessary to know the length of all of the optical fibres transmitting the TTC signals to the FEDs and FECs, so that their contribution to the delays can be compensated. If they are all of the same prescribed length then the situation is straightforward and no correction should be needed. However if different fibres lengths are used to send TTC signals to different FEDs (or to different FECs), then the relative delay contributions due to these different lengths must also be compensated. This can be done using the programmable delays on the TTCrx in the FEDs and FECs

and, if necessary, using the PLLs around the control rings. Again, any adjustments made to the fine delays affecting the signal propagation to the APVs have to be matched by adjustments of the FED clock skew.
It should also be emphasised that the arrival of APV signals at the FEDs will not necessarily be synchronous since the fibre lengths in the analogue optical links are not identical. The essential point is that the synchronization has been made at the level of the output of the APVs.

2.5. Summary of procedure

A summary of the overall procedure described in the preceding sections is illustrated in Fig. 9.

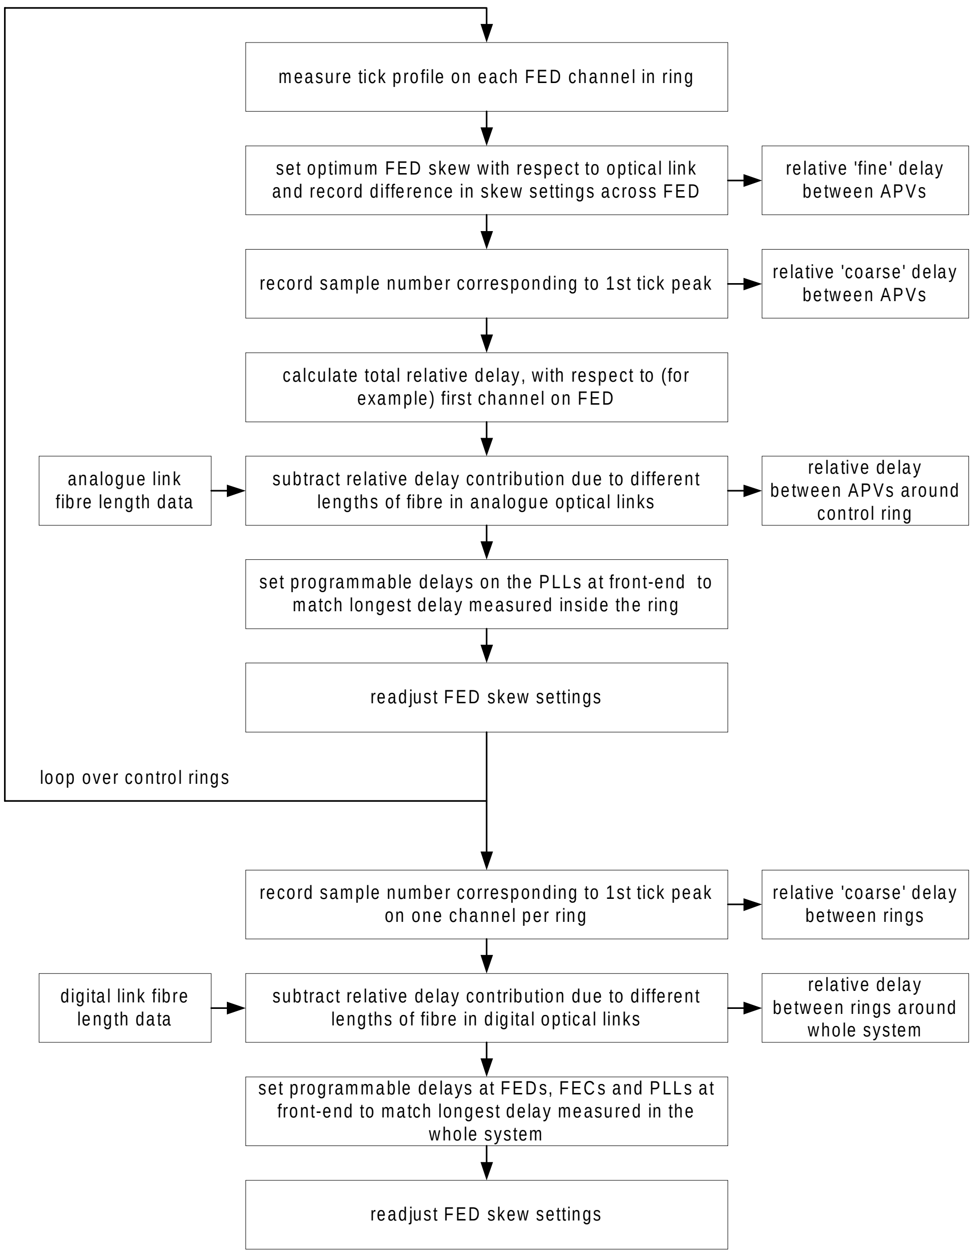

Figure 9: Overall procedure to bring APVs into relative synchronization at start-up of the Tracker.

3. Conclusion

The method described allows precise relative synchronization of the APVs, such that all chips send out ticks synchronized to within ~1-2ns. The method uses the APV ticks as a timing probe, whilst scanning the FED sampling skew. The lengths of all the optical fibre channels in the system have to be known in order to calculate the delay settings required at the PLLs on the front-end hybrids and at the TTCrx units on the FEDs and FECs. The APV chips in each ring are synchronized first of all, and then the different rings are brought into synchronization.

In principle the whole Tracker can therefore be synchronized in the absence of a colliding beam. A few steps are still required in order to achieve absolute synchronization with the LHC collisions. A simple correction for time-of-flight delays can be made based on the location of the sensor. A scan of the phase of the clock into the (whole) Tracker has to be made, once the beam is available, in order to precisely align the system in time with the collisions. In addition, a scan of latency settings should be made in order to be sure that events are read from APV pipeline cells containing data from the correct bunch-crossing.

The precision of the synchronization procedure outlined is limited by the size of FED skew steps and/or uncertainty in the fibre lengths. If 1ns skew steps are available at the FED, matching the fine delay steps in the PLLs, then the resolution of the measured delays of the ticks would be <1ns. The synchronization precision would then be determined by the systematic uncertainty in fibre lengths, which is likely to be ~1ns.

It should be emphasised that this scheme adds little overhead to the calibration activities that are already necessary steps in starting up the system. In particular the precise measurement of the ticks and the setting of the appropriate FED sampling skew are important steps in the setting up the analogue optical links, which has to be done very early in the commisioning of the Tracker.

References

[1] CMS Tracker TDR

[2] 'A system for timing distribution and control of front end electronics for the CMS tracker', A. Marchioro et al., 3rd LEB Workshop, 1997.

Appendix: Preliminary simulation of the measurement of delays in the system

The simulation of the effect of delays in the system includes the following elements and parameters, with references to acronyms given in Fig. 1:

1 FED

96 input fibre channels
10bit resolution
1024 samples at 40MHz stored per trigger
2.5ns clock skew steps
20m fibre for input of TTC signals

1 FEC

2 outputs connected via 60m digital optical link to digital optohybrids (D-OH)
20m of fibre for input of TTC signals

2 D-OHs

each connected to a control ring with 16 CCUMs

2x16 CCUMs

each connected to front-end hybrid (FEH) containing
1 PLL
6 APVs
40cm electrical cable between each CCUM around the ring (and between D-OH and the first CCUM).

192 APVs

with identical tick characterictics:
800mV tick amplitude
equal output digital baseline levels
with output from neighbouring pairs of APVs multiplexed onto 1 analogue optical link channel

96 analogue optical link channels

gain = 0.15 (compared to gain=1 in final system)[8]
noise contribution = 2.5ADC counts (mean) at the FED
settling time of 15ns (to 1%)
transmitters arranged on optohybrids with 3 lasers transmitting signals from 6 APVs
fibre lengths: these are based on the total lengths of the 96-way cable segment from the compact patch panel (in the HCAL crack) to the FED, added to the 12-way ribbon cable from the distributed patch panels around the tracker to the compact patch panel, plus the pigtails between the distributed patch panels and the lasers.
The individual fibre segments have lengths of:
96-way cable = 60m
12-way ribbon = 6.5-10m (increasing in 0.5m intervals over the 8 ribbons)
1-way pigtails = groups of 12 fibres are divided into 4 sets of 3 lengths, increasing from 0.2m to 0.8m in 0.2m steps.

Timing uncertainty and jitter

No uncertainty in the fibre length is included thus far.
The jitter in the system is simulated at a preliminary level as the random combination of 5 separate elements each contibuting 0.2ns jitter.

---

[8] This low gain of the optical link was simulated in order to be able to observe the noise contribution clearly in the results, e.g. in Fig. 4.