

TIMING AND EVENT DISTRIBUTION FOR FERMI@ELETTRA

A. Rohlev, A. Borga, G. D'Auria (ELETTRA, Trieste, Italy), L. Doolittle, A. Ratti (LBNL, Berkeley, California), M. Stettler, J. Serrano (CERN, Geneva, CH)

Abstract

FERMI@ELETTRA is a 4th generation light source under construction at Sincrotrone Trieste. It will be operated as a seeded FEL driven by a warm S-band Linac which places very stringent specifications on control of the amplitude and phase of the RF stations. The local clock generation and distribution system at each station will not be based on the phase reference distribution but rather on a separate *frequency* reference distribution which has significantly less stringent phase stability requirements. This frequency reference will be embedded in the serial data link to each station and has the further advantage of being able to broadcast synchronous machine timing signals with sub-nanosecond temporal accuracy. The phase and amplitude of the phase reference line is measured for each pulse and used to calibrate the other measurements. This paper describes the architecture used to distribute the frequency reference along with the precision machine timing and clocking signals.

INTRODUCTION

The FERMI Linac under construction at Elettra will be used as an injector for a FEL and thus has very stringent stability requirements for the cavity RF fields. The proposed low level RF (LLRF) control system will utilize several novel techniques in order to meet the RF field specifications [1]. Among these is the distribution of a frequency reference in addition to the stabilized phase reference. The frequency reference can be distributed over the same high speed data links which will connect the individual cavity controllers to a central hub. The addition of an integrated digital phase shifter and external PLL at each controller enables all of the station clocks to be phase aligned to within several 10's of picoseconds, providing very high temporal resolution for trigger transmission of and time stamping of data.

FERMI LLRF FREQUENCIES

Fig. 1 shows the necessary clocks and LO frequency for the S-band low level control system for FERMI. The entire phase error budget is 0.1° , of which, it is reasonable to assume, $1/3$ could be allocated to frequency and clock generation.




Figure 1: FERMI S-band LLRF frequencies.

This translates to a total integrated drift/jitter of 22 fs in the LO and 156 fs in the ADC and DAC clock, if their shares is split $2/3$, $1/6$, $1/6$ respectively. It will be

exceedingly difficult to maintain these values so a scheme is proposed where the distributed phase reference (Pref) is measured in the same manner as the other S-band cavity signals, but using the LO and clocks generated from a reference frequency (Fref) which has significantly looser phase stability requirements. Variations in the measured value of Pref, which is necessarily assumed to be perfectly stable, become an indication of the drift/jitter of Fref, and are used to calibrate the other inputs measured with the same clocks. If Pref is measured on a millisecond timescale, then Fref can be allowed to have sizable drift/jitter below 1 kHz. Fref can also have significant jitter at frequencies above the system bandwidth of 5 MHz without degrading the performance of the controller.

The frequency reference can be generated locally at each station from the distributed Pref signal, but this approach has limitations and is problematic: tapping off Pref, either with a coupler or splitter, would have to be done with an isolation above 75 dB in order to introduce a phase error below 0.01° . On the other hand, distributing Fref, as described below, leaves Pref untouched and also gives the LLRF stations clocks that are phase coherent to a fraction of a cycle.

F_{REF} DISTRIBUTION ARCHITECTURE

Fig. 2 shows the architecture for distributing Fref as well as the generation the local clocks and LO. The individual LLRF controllers are connected via high speed serial Gigabit Transceiver Pairs (GTPs) to a μ TCA standard module and backplane known as 'The Matrix' currently under development at CERN [2]. Each module has a single FPGA and hosts 16 serial links operating from the same clock. The backplane supports up to 12 modules. The GTPs are the main communications link between the individual controllers and the Matrix which has an Ethernet connection to a conventional workstation. The module FPGA will be driven by an external 241.6 MHz clock which is encoded into the transmitted data. At the receiver side the clock is recovered and passed through an 8-bit digital phase shifter (DLL). It then exits the FPGA and is used as the reference phase for an external PLL. A ultra low phase noise OXCO must be used to meet the jitter requirements at the LO frequency. Furthermore, the OXCO tuning range must be limited to 200 Hz so that the round trip phase length over the fiber (150 meters each way) changes by less than 0.1° over its entire frequency tuning range. The loop filter of the PLL is implemented digitally inside the FPGA. The OXCO output provides the clock for the bulk of the FPGA as well as the GTP transmitter, which are in a different clock phase domain than the GTP receiver, clock recovery, and DLL circuitry. External multipliers are used to generate the DAC clock and LO.




Figure 2: Frequency reference (Fref) distribution scheme.

PHASE LOCKING SCHEME

In the architecture shown in Fig. 2 data as well as the frequency reference (Fref) are transmitted over the same fiber link. The addition of a digitally controlled phase shifter, as well as an external PLL, allows for the phase locking of multiple systems to the same transmitted clock. Furthermore, the GTP links have a ‘loop-back’ feature providing the system with the ability to measure the round trip latency of the link in integer clock cycles. The sum of these features allow for the implementation of a global clock at each station with an absolute time resolution on the order of 20 ps. The following details this procedure.

Distributing the Reference Frequency

Fig. 3 shows the first step in the process: the distribution of Fref as the recovered clock from a serial link. Since the fiber links are of different length, and in different locations, there is no way to ensure that the phases are aligned or that they remain fixed in time.




Figure 3: Distribution of Fref. Clock phases are not coherent and can drift.

Locking to the Phase Reference

Fig. 4 shows the second step: the phase reference is measured with the recovered clock and the error controls the DLL which compensates for drifts in the fiber. The 8-

bit phase shifter has a granularity of about 1.4° or 16 ps at 241 MHz. At this stage the clocks at each station are still incoherent (since the phase of Pref is ultra stable but unspecified), however, their drift is now limited to 16 ps.




Figure 4: Locking Fref to Pref. Clock phases are fixed.

Aligning the Clock Phases

Fig. 5 shows the method that will be used to align the clock phases at each station.




Figure 5: Eliminating the fractional cycle with the DLL.

First, the GTP transceiver of the Matrix module is placed in ‘loop back’ mode which effectively short-circuits the transmitter to the receiver. This in turn presents a fixed phase to the PLL at the station. Since the OCXO has a limited range, the loop will quickly saturate. The round-trip phase of the fiber link can be seen as an

integer number of cycles, NT , plus a fractional cycle, ΔT . Next, the phase of the DLL (δ) is incremented, and at some value, δ_0 , the fractional cycle, ΔT , will either be absorbed or augmented into a full cycle. At this point the PLL will saturate to the other limit or regain control if, by chance, δ_0 happens to land within the OCXO's tuning range (0.1°) of zero phase. In either case at the point which this occurs the round trip phase of the link plus DLL is within one tick (1.4°) of zero and the fractional cycle, ΔT , has essentially been eliminated. The module's GTP can then assume normal operation with δ_0 as the nominal operating point of the DLL. As this procedure is performed at every station the resulting clocks become phase aligned. Their long term coherence is guaranteed because they are locked to Pref, however, if a phase stable signal is not available the above procedure can be repeated, as necessary, to compensate for drifts.

Setting the Time and Sending Triggers

With the fractional cycle eliminated, the only ambiguity remaining is the integer cycle delay of each link. This, of course, is measured naturally with 'loop-back', and dividing the result by two gives the one way link delay for each station. The time can now be set at each station, ideally with a resolution of 1/256 of a cycle. This is shown in Fig. 6.




Figure 6: Setting the time and sending triggers.

Triggers can be pre-delayed at the Matrix, or adjusted at each station to compensate for the individual link delays.

NOISE CONSIDERATIONS

The above procedure for aligning the clocks, as well as the distribution of Fref, relies on the assumption that the serial link, clock recovery, and DLL will not add significant jitter to the transmitted clock. This has been tested with a GTP link, over coax, between two Xilinx ML506 test boards. Fig. 7 shows the results.

Three traces are shown. The bottom trace (blue), which serves as the baseline measurement, is the 240 MHz




Figure 7: Noise measurements on GTP link with DLL.

system clock used to encode the data on the transmitter board. It has an RMS jitter of 5 ps integrated from 1 kHz to 5 MHz. The middle trace (red) is the recovered clock from the receiver board which has an integrated jitter of 11 ps. The top trace (green) is the recovered clock after going through the DLL. Its integrated RMS jitter is 21 ps, which is less than two ticks of the DLL. The GTP link and DLL add about 15 ps of jitter indicating that the phase aligning procedure should be feasible. It is expected that the actual system, with an OCXO derived clock and PLL filtering, will have even lower jitter thus allowing the controller clocks to be aligned to 16 ps, the limit of the DLL.

CONCLUSION

Presently a GTP link has been established between two FPGA test boards over coax. A phase shifter has been instantiated in the FPGA and tested with the 240 MHz recovered clock. The external PLL has not been tested nor have the ultra low noise OCXOs been acquired. This is the next step. The technique outlined here has the potential of establishing an trigger/event/data link among the individual controllers along an accelerator with a temporal accuracy in the few 10's of picoseconds. Key to this development is the advent of a central FPGA based data hub, the 'Matrix', with its multiple, phase synchronized, serial data links. This architecture promises other interesting developments along the line of medium bandwidth (10's of kHz) accelerator-wide control.

REFERENCES

- [1] J. Byrd, et al, "Plans for Precision RF Controls for FERMI@ELETTRA," PAC 2007 (WEPMN116)..
- [2] M. Stettler et al., "Modular Trigger Processing, The GCT Muon and Quiet Bit System," TWEPP 2007, Prague, Czech Republic.