The Front End Readout System for the T2K-ND280 Detectors

Antonin Vacheret, Sarah Greenwood, Matt Noy, Mark Raymond and Alfons Weber

***Abstract*– The Tokai-to-Kamioka (T2K) long baseline neutrino experiment goal is to measure ν_μ to ν_e oscillation parameters using near (280 m) and far (295 km) detectors from the beam origin in the J-PARC accelerator in Tokai. The far detector will be Super-Kamiokande. The 280 m near detectors complex (ND280) contain a number of scintillator based sub-detectors and a beam monitor. Approximately 54,000 scintillator bars will be read out by Multi-Pixel Photon Counter devices (MPPC) via wavelength shifting fibres. It will be the first time this novel type of photo-sensor is used on such a large scale. The front end readout system for the MPPC uses Fermilab Trip-t chips mounted on Trip-t Front end Boards (TFB). The TFB design and measured performance in conjunction with MPPC photo-sensors are presented.**

I. Introduction

THE Tokai-to-Kamioka (T2K) long baseline experiment goal is to measure the ν_μ to ν_e oscillation parameters θ_{13}, Δm^2 and θ_{23} [1]. The T2K beamline and detectors are shown schematically in fig. 1. A 50 GeV, 0.75 MW proton beam, generated by the J-PARC accelerator in Tokai, strikes a graphite target producing pions which decay to ν_μ. The energy spectrum of the neutrino beam can be optimised for the experiment by situating the oscillation measuring detectors, near and far, at an off-axis angle between 2 and 2.5º to the forward, on-axis, beam direction. The off-axis beam is aimed at the Super-Kamiokande far detector [2], 295 km away, passing through the near detector at 280 m (ND280). The INGRID on-axis beam monitor (scintillator/iron layers) is also sited in the 280 m detector pit. ND280 consists of an upstream π_0 detector (scintillator/lead layers) and a downstream tracker (scintillator and time projection chambers) surrounded by an electromagnetic calorimeter (scintillator/lead layers). ND280 will use the CERN UA1 0.2 Tesla magnet. For position resolution scintillator layers are segmented into bars with embedded wavelength shifting fibres (WLS) for readout, with a total channel count of approximately 54,000. The T2K experiment will begin operating in 2009.

For practical reasons (high channel count, limited access apertures) the ND280 front end readout must be located inside the magnet, and therefore a compact, magnetic field insensitive, low-cost photo-sensor is required. The Multi-Pixel Photon Counter (MPPC) meets these requirements and has therefore been chosen. T2K will be the first high energy physics experiment to use this new type of photo-sensor on such a large scale.

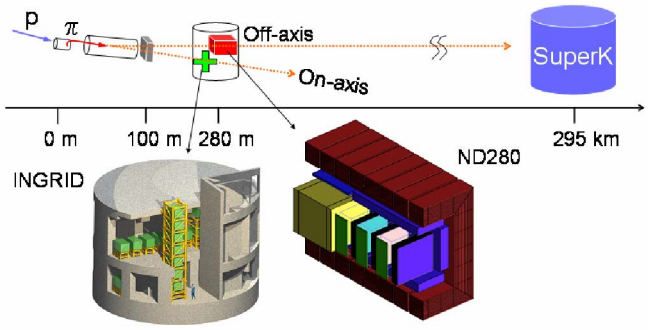

Fig. 1. Schematic view of the T2K beamline. The ND280 off-axis detector and INGRID on-axis beam monitor occupy the 280 m near detector pit. The Super-Kamiokande far detector is situated 295 km away in the off-axis direction. All distances are from the target.

II. Multi-Pixel Photon Counter

The MPPC [3] consists of a pixellated array of APDs operated in limited Geiger mode, where the reverse bias operating voltage V_{OP} is set higher than the breakdown voltage V_{BD}. The signals from all pixels are summed together on one output. Arrays of several hundred pixels, with total areas of ~ 1 mm^2, are available. A Geiger discharge is initiated in a pixel when a photon is absorbed and a hot carrier is released and accelerated by the high electric field present in the depletion region. The pixel output charge Q_{PIX} is given by

$$Q_{PIX} = C_{PIX}(V_{OP} - V_{BD}) \qquad (1)$$

where C_{PIX}, the pixel capacitance, is typically in the range 50 to 150 fF. Typical stable MPPC operating voltages [3] are ~ 1 V higher than the breakdown voltages which are, in turn, typically ~ 70 V, so from (1) it can be seen that electrons/ photon gains in the region of 10^6 can be obtained. Following the voltage drop caused by the avalanche, an individual pixel is recharged by an integrated resistor R_{PIX}, with a recovery $R_{PIX}C_{PIX}$ time constant of ~ 10 nanoseconds.

Fig. 2 shows a photograph of the active surface of an MPPC with 100 µm pixel pitch, together with an oscilloscope picture of the raw MPPC signal, demonstrating the single photon resolution and fast response. The output signal is given by the number of photo-electrons generated by the incident light (the number of pixels firing), multiplied by the avalanche gain. The combination in one sensor of photon counting with a large dynamic range, makes the MPPC suitable for low light detection as well as large intensity signals measurements.

Manuscript received November 16, 2007. This work was supported by the UK Science and Technology Funding Council (STFC).

Antonin Vacheret, Sarah Greenwood, Matt Noy and Mark Raymond are with the High Energy Physics Group, The Blackett Laboratory, Imperial College London, SW7 2BW, London, U.K. (e-mail a.vacheret@imperial.ac.uk, s.greenwood@imperial.ac.uk, m.noy@imperial.ac.uk, m.raymond@imperial.ac.uk).

Alfons Weber is with the Particle Physics Department, Denys Wilkinson Building, University of Oxford, OX1 3RH, U.K. and STFC – Rutherford Appleton Laboratory, Chilton, Didcot, OX11 0QX, U.K. (e-mail A.J.G.Weber@rl.ac.uk).

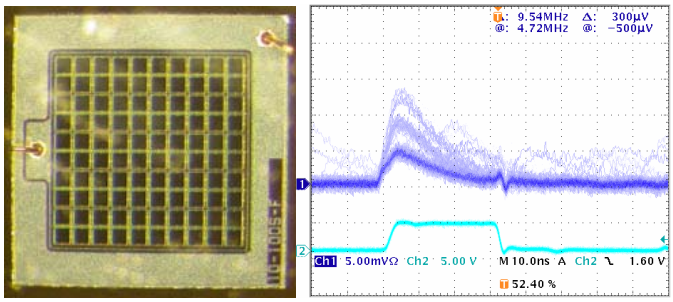

Fig. 2. A photograph (left) of the 1 mm^2 active area of a 100 pixel MPPC. Each pixel is 100um pitch. An oscilloscope picture (right) of the raw signal (blue) from a 400 pixels MPPC. MPPC avalanches are triggered by a 1 ns light pulse from a 463 nm LED source (electronic synchronization signal shown in light blue). MPPC signal amplitudes corresponding to 1, 2 and 3 photons detected are clearly visible.

It is convenient to express MPPC output signal magnitude in photo-electrons equivalent (p.e.), defined as the number of photo-electrons detected (number of pixels firing) multiplied by the avalanche gain.

MPPC linearity is limited for large signals because of the finite number of pixels [3]. In T2K we will use a 667 pixel MPPC device, with a maximum physics signal of less than 500 p.e.. It is important to preserve MPPC linearity for smaller signals, throughout the readout chain, and an electronics nonlinearity of better that 5% is sufficient, defined here as the difference between the measured and true signal amplitudes.

MPPC devices exhibit a typical dark count rate in the range 300 to 800 kHz/mm^2 at 25º C, dominated by single p.e. signals, with double p.e. signal rates an order of magnitude less.

The temperature coefficient of MPPC gain is in the region of a few % per ºC [3], and it will be important to maintain a stable operating temperature environment in the T2K experiments. If the readout noise is low enough the gain can be monitored, in situ, using the amplitude of the single p.e. dark count rate signal. To be able to resolve single p.e. signals in the amplitude spectrum, the average peak amplitude must be greater than the full width at half maximum of the distribution, corresponding to an RMS noise requirement of 0.42 p.e.. For improved accuracy we impose a more stringent criterion, a factor of 2 better, which translates to an RMS electronic readout noise specification of less than 0.21 p.e..

For typical V_{OP}-V_{BD} values of 1V giving gains in the region of 10^6 from (1) it is clear that MPPC gain is very sensitive to V_{OP}. To account for device-to-device variations it is necessary to be able to tune V_{OP} for individual devices and to adjust gains with 2% precision (for example) leads to a V_{OP} adjustment resolution requirement of 20 mV.

Track reconstruction in ND280 involves multiple subdetectors and time correlation between signals is necessary to determine track direction and to veto background signals, leading to a requirement to time-stamp MPPC signals in the readout electronics. Plastic scintillator bars in some subdetectors are several metres long, and position sensitivity can be enhanced using time information with double-ended readout. Time-stamp functionality is also required in some regions of the experiment to facilitate detector calibration with cosmic rays, to trigger readout on time coincidences between channels within a subdetector module and between modules on opposite sides of the detector. A time-stamp precision of 3 nanoseconds or less is adequate for these purposes.

TABLE I
T2K FRONT END ELECTRONICS REQUIREMENTS FOR MPPC READOUT

parameter	value
dynamic range	0 to 500 p.e.
noise	< 0.21 p.e.
nonlinearity	< 5 %
discriminator time-stamp resolution	< 3 nsec.
bias voltage trimming resolution	20 mV

Table I summarises the requirements on the MPPC front end electronics readout determined by T2K physics performance requirements and MPPC characteristics.

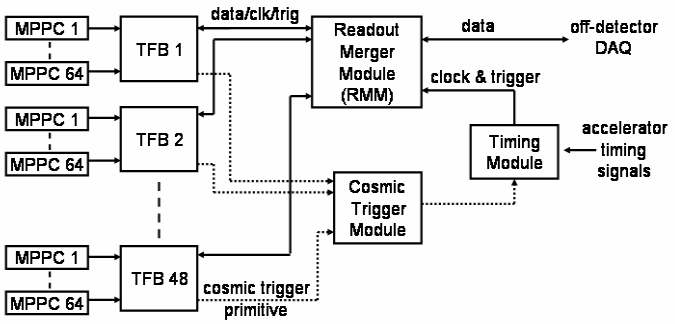

Fig. 3. The Trip-t based T2K front end electronics readout system.

III. FRONT END READOUT

The Trip-t chip, originally designed at Fermilab for Visible Light Photon Counter readout for the Tevatron D0 experiment [4], has all the features required for MPPC readout at T2K. The Trip-t based front end electronics readout system for the T2K ND280 detectors is shown in fig. 3. Each Trip-t Front end Board (TFB) instruments up to 64 MPPC sensors. The bi-directional data (control and readout) from up to 48 TFBs are processed through Readout Merger Modules (RMM) which provide the communication interface with the off-detector data acquisition system (DAQ). Cosmic trigger primitives are formed on a TFB from a time coincidence between signals from a number of MPPC channels, and transmitted to a Cosmic Trigger Module where a global trigger decision is taken. A Timing module receives the accelerator timing signals and transmits clock and synchronous trigger information to the front end via the RMM.

A. Trip-t chip

Fig. 4 shows a schematic of a single Trip-t chip front end channel. Each channel comprises an integrating preamplifier, a second amplifier stage before the analog pipeline, and a discriminator which takes the AC coupled preamplifier output following an amplifier with a fixed gain of ten. The discriminator threshold voltage is programmable and common to all channels on the chip.

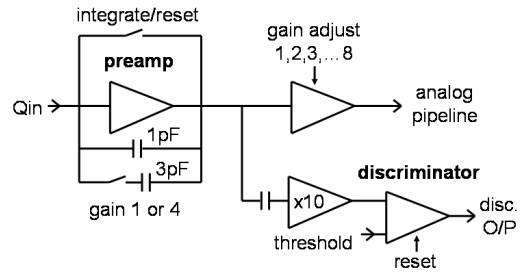

Fig. 4. Schematic of one Trip-t front end channel.

Fig. 5 shows a simplified functional block diagram of the 32 channel Trip-t chip. The analogue pipeline for each channel is up to 48 cells deep and is read out through a 32:1 analogue multiplexer. The discriminator outputs are split into two groups of 16 channels, and a digital multiplexer selects which group to be routed off chip. There is a serial digital interface to program the operational state of the internal circuitry, and a number of digital control signals are required to correctly sequence chip operation. The reader is referred to [4] for a more detailed description of the Trip-t and its operation.

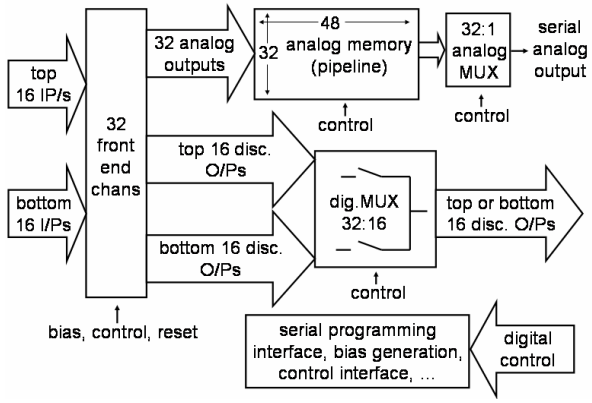

Fig. 5. Functional block diagram of the Trip-t chip. Functionality not used for T2K has been omitted.

The T2K readout system must be able to accommodate the full range of MPPC signal sizes to be encountered at T2K, while still being able to resolve and time-stamp signals at the few p.e. level (table I). Initial investigations showed that noise and threshold voltage dispersions would not allow sufficient threshold precision if the whole range were to be accommodated in a single channel, and in any case only 16 discriminator channels can be simultaneously transmitted off-chip for time-stamping. A simple solution is to divide the input signal between low and high gain channels, using 2 Trip-t channels per MPPC.

Fig. 6 shows schematically how the MPPC interfaces to the Trip-t electronics. A miniature coaxial cable connects an MPPC to the TFB. The high voltage (HV) bias is supplied on the core and the sheath carries a low voltage which can be adjusted in the range 0 to 5V, on an individual channel basis, to achieve the HV trim functionality. The MPPC signal is passively split between C_G, C_{HI} and C_{LO}, where a 10 to 1 ratio of C_{HI} to C_{LO} determines the high to low channel gain ratio. The value of C_G can be used to adjust the overall gain. For example, for a full-scale MPPC signal of 500 p.e., C_G is chosen such that the low gain channel saturates at 500 p.e.. The high gain channel therefore saturates at 50 p.e. and the discriminator threshold adjustment range is up to 5 p.e. because of the fixed x10 gain stage between preamp and discriminator (fig. 4). The capacitor values indicated in fig. 6 are matched to these signal amplitudes assuming a typical MPPC gain of $5x10^5$, corresponding to 80 fC/p.e..

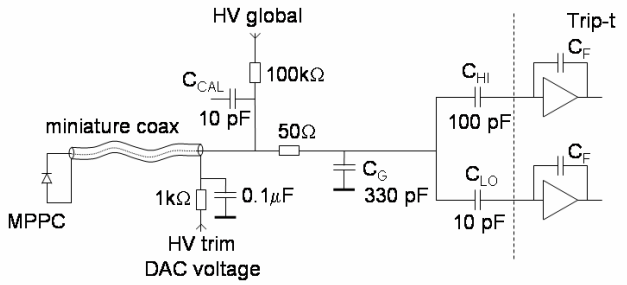

Fig. 6. Schematic of the MPPC/Trip-t interface showing the charge splitting between low and high gain channels and MPPC bias scheme.

Apart from providing good small signal resolution and large dynamic range simultaneously, the dual gain range technique also matches well with the required time-stamping functionality. Time-stamping is implemented by an external Field Programmable Gate Array (FPGA) for the 16 high gain channel discriminator outputs selected to be routed off-chip.

B. Trip-t operation at T2K

The J-PARC accelerator spill and bunch structure is not yet finalized, but a beam spill will be produced every few seconds, each spill consisting of a train of either 8 or 15 ~ 60 ns wide bunches, separated by gaps of either 540 or 240 ns (fig. 7). The physics occupancy will be low (~ few neutrino events per spill). During the spill, Trip-t operation will be synchronized to the bunch structure such that each preamplifier integration period corresponds to a single bunch, and the result of each integration period is stored in the pipeline. Any discriminators firing during the preamp integration period are time-stamped and reset during the preamp reset period.

The short inter-bunch spacing prohibits triggering and readout of pipeline samples between bunches. The proposed mode of Trip-t operation is therefore to use the pipeline to store samples for all bunches in the spill, and to continue running in the same way as long as possible after the spill to maximise the probability of catching late signals resulting from muon decay. At the end of the after-spill acquisition period the complete contents of the pipeline are read out, digitised and transferred, together with any time-stamps, to the off-detector DAQ. While the majority is pedestal data, transmitting all the analog samples allows continuous monitoring of MPPC gain using the single p.e. dark rate signals.

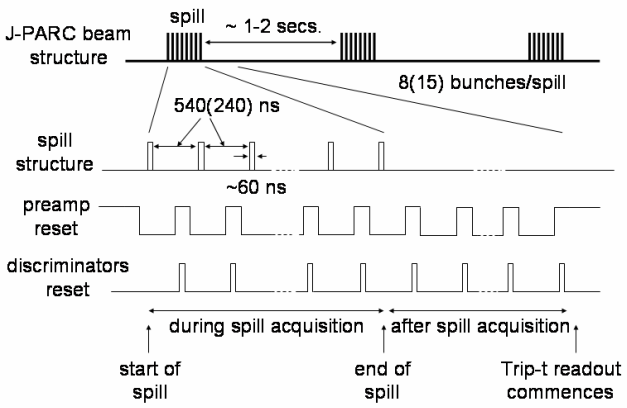

Fig. 7. T2K beam structure and corresponding Trip-t sequencing.

The maximum length of the Trip-t pipeline is 48, but the chip was not originally designed to allow the complete pipeline contents to be read out as we propose in T2K, but an acceptable workaround has been found where every other pipeline column can be used in this way, allowing up to 23 columns to be used to capture beam bunch and after spill data. Details of the Trip-t readout sequencing can be found in [4].

In the relatively long period between spills the cosmic trigger mode will be activated, where the Trip-t chips will run continuously, integrating and storing MPPC signals, operating the analog pipeline as a cyclic buffer. TFBs processing signals from subdetectors contributing to cosmic triggering will be configured to generate trigger primitives if a pre-defined combination of MPPC signals have triggered their respective discriminators within a coincidence time window. Candidate trigger primitives will be transmitted to the cosmic trigger module (fig. 3). Depending on a satisfactory combination of a number of TFB trigger primitives, the cosmic trigger module will issue a global cosmic trigger which is broadcast by the RMMs to initiate readout of all TFBs.

C. Trip-t Front end Board (TFB)

Photographs of both sides of a TFB are shown in figs. 8 and 9. The TFB is a 12 layer board (6 signal routing, 6 power/ground) with dimensions 16 cm x 9 cm. Each TFB takes 4 Trip-t chips, so can read out up to 64 MPPC sensors.

TFB operation is controlled by an FPGA (fig. 9, Xilinx Spartan-3 XC3S2000) which performs the following functions (more details in firmware section III-D):

- Trip-t configuration, calibration, analog data readout sequencing and digitization.
- Trip-t discriminator output time-stamping.
- Cosmic trigger primitive formation.
- Bi-directional serial interface to external DAQ.
- Slow control operations; MPPC bias voltage trimming, local voltage and temperature monitoring.

MPPC sensors are connected to the TFB using miniature coaxial cables which plug into the four arrays of connectors that can be seen in fig. 8. The signals pass through the board to the opposite side (fig. 9) where the gain range splitting and bias components of fig. 6 are situated. The resulting signals are fanned in to the Trip-t inputs on internal routing layers, screened above and below by ground planes to minimise interference pick-up. Eight 8-channel 8-bit DACs (AD5308) can also be seen in fig. 9, which provide the 0 to 5V programmable trim voltage, with the required 20 mV resolution (table I), to each MPPC via the coaxial cable sheath connection.

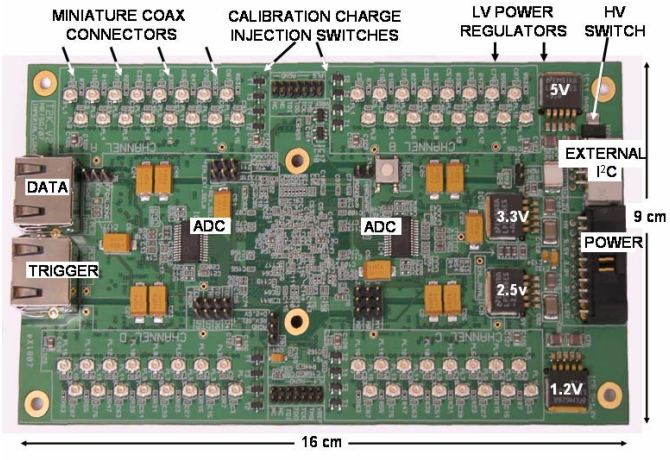

Fig. 8. Photograph of TFB top surface.

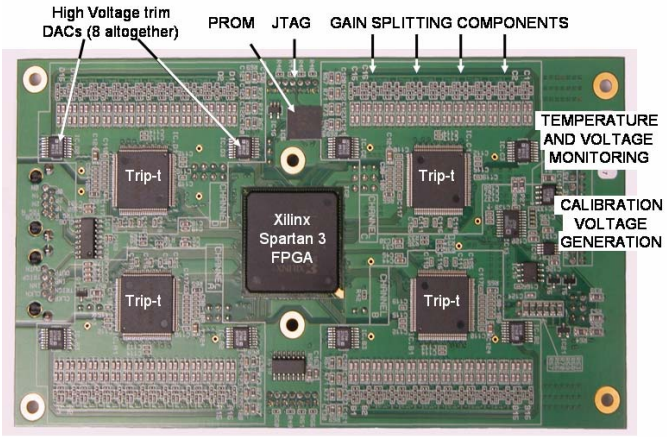

Fig. 9. Photograph of TFB bottom surface.

Four serial Trip-t analog output streams are digitised by two 10-bit dual channel ADC chips (AD9201). The Trip-t output multiplexer clock and hence ADC digitization rate has been chosen to be 5 MHz to allow ample time for Trip-t output samples to settle.

To perform electronic calibration of the readout chain a circuit is provided which injects a charge into all MPPC readout channel inputs. Four MOSFET switches are associated with each Trip-t (see fig. 8), allowing inputs to be exercised 4 channels at a time. One side of capacitor C_{CAL} in fig. 6 is charged to a voltage level up to 5 Volts, provided by a 12-bit programmable DAC, and then discharged by the MOSFET switch, thereby injecting a negative signal of up to 50 pC into the channel input.

A section of the TFB devoted to local monitoring can be seen in fig. 9, where a slow ADC (AD7998, 12-bit, 8 channel) is used to monitor the levels of low voltages on the board,

MPPC high voltage bias and current, and also the voltage level generated by the calibration DAC. This ADC interfaces to the FPGA via an I^2C bus, which is also used to communicate with an on-board temperature sensor device (LM92). The I^2C bus is also taken to an edge connector (fig. 8) to allow external temperature sensors to be read out, but can also be used to communicate with any device conforming to the I^2C standard.

The circuitry on the TFB uses four low voltage levels; 5V, 3.3V, 2.5V and 1.2V. These are provided by on-board low drop-out linear voltage regulators, and to minimise power consumption the levels provided at the external power connector are just sufficient to satisfy the headroom requirements of the regulators; 5.5V, 3.8V. 3.1V and 1.7V. The MPPC high voltage bias also enters the TFB through the same power connector to supply all the MPPC channels on the board (HVglobal in fig. 6).

A JTAG interface is provided to allow a default FPGA power-up configuration to be stored in an on-board PROM (fig. 9). The default configuration can be subsequently modified via the external communication interface.

Signals are transmitted to and from the TFB at LVDS levels over cables plugged into RJ45 type connectors to be seen on the left hand side of figure 8. Both Data and Trigger connectors carry 4 twisted pair signals. The functionality of signals on the Data connector will be described in the TFB firmware section III-D. The Trigger connector uses only one of the pairs to transmit the cosmic trigger primitive off-board.

D. TFB Firmware

The serial data interface to the TFB comprises three upstream TFB inputs (100 MHz clock, data and trigger), and one downstream data output. These lines are AC coupled and a Manchester Encoding scheme is used to provide DC balance.

A custom, full-duplex, packet-based protocol is used over the link. Packets comprise a 10-byte header, variable length payload and a 16-bit cyclic redundancy check (CRC). The header contains a unique 12-bit TFB identifier, 4-bit pipe ID, payload length, and status information. The unique identifier is stored in the User-ID field of the FPGA configuration PROM on the TFB. The pipe ID field determines destination and source for upstream and downstream data packets respectively, allowing the logical multiplexing of different types of packets over the same physical layer. When entering the TFB, an upstream packet is deserialised, the CRC calculated, the pipe ID read and the destination selected.

A 16-bit read/write data and address bus provides access to the user configurable firmware parameters. This supports register value reading and writing, status readout, and access to all the peripheral components on the board via the appropriate hardware communication protocol (HV trim DACs, temperature readout, calibration DAC, monitoring ADC and Trip-t programming interface).

In-spill operation is managed by a finite state machine (FSM) that sequences the four Trip-t chips during the filling and readout of the analogue pipeline, providing all the necessary signal transitions with programmable timings, and synchronising the readout with the on-board ADC. All four chips are sequenced by a single instance of this FSM.

During readout each Trip-t provides the data from 34 channels; the 32 active channels (16 high gain, 16 low) plus data from two extra dummy pipeline channels above and below the active pipeline. The 10-bit ADC data from the 34 channels from each preamplifier integration period for all four Trip-t chips is packed into 85 16-bit words. A beam spill thus generates 23 of these 85 word ADC data packets.

Time-stamps are generated from the Trip-t discriminator outputs using an array of D-type latches internal to the FPGA clocked on the four different quarter phase shifted edges of the 100MHz clock. This gives an effective sampling rate of 400MHz, or a time resolution of 2.5ns. Sixty-four arrays are required for the 16 discriminator outputs of each of the four Trip-t chips. A great deal of care has been taken with the placement of these D-type latches to reduce systematic variations in the bucket sizes due to signal and clock routing delays in the FPGA. The quadrant in which the discriminator fires is encoded using a 2-bit binary scheme and combined with a 32-bit binary counter value to give a 34-bit time-stamp that wraps approximately every 43 seconds. This time-stamp can be reset synchronously through the trigger interface, permitting the synchronization of all TFBs in the system.

Time-stamps are stored in a pipeline inside the FPGA that mimics the Trip-t pipeline, allowing both spill and cyclic operation. The pipeline pointers advance once per integration cycle, and a bit is set if a discriminator fires and generates a time-stamp. Once a spill has been encoded, the pipeline is examined in a sequential manner and zero-suppressed. The 34-bit time-stamp value plus a 6-bit TFB channel address, plus a 5-bit field to indicate in which preamplifier integration period the time-stamp occurred, gives an individual time-stamp data packet size of 3 16-bit words (with 3-bit zero padding).

The overall time-stamp data volume per spill will vary. A data packet (packet length plus time-stamp data payload) is produced for each channel, even if the packet length is zero, so a total of 64 time-stamp packets are always produced. A mask exists to allow noisy channels to be masked, to avoid inflating the overall data volume with bad data. A packet is produced for a masked channel, but its payload is always empty.

IV. TFB Measured Performance

Approximately 20 TFBs have been produced to allow functionality and performance to be verified, and to satisfy requirements of T2K collaborators developing the different sub-detectors using the Trip-t based MPPC readout system. Consistent board-to-board performance has been observed and typical experimental results are presented in this section.

The TFB control and readout system for the T2K experiment is shown in fig. 3. For bench tests a simplified system is used (fig. 10), where a single TFB can be controlled and read out by a PC using a custom FPGA based interface circuit to emulate RMM functionality. A typical DAQ transaction would be where the PC requests a spill-type data event, whereupon the RMM emulator generates a start-of-spill trigger to the TFB and receives, formats and returns the resulting data to the PC.

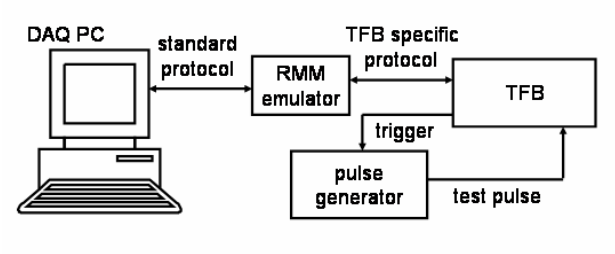

Fig. 10. TFB test bench system.

To characterize TFB performance it is often necessary to trigger an external pulse generator synchronously with the preamplifier integration periods to provide Trip-t input signals at the appropriate time. A convenient way to do this, shown in fig. 10, is to use one of the spare outputs on the trigger connector to provide an external trigger to a pulse generator.

A. Gain, Linearity and Calibration

In the T2K experiment single p.e. dark rate signals will be used to measure MPPC gain. An absolute calibration of the electronic readout chain is not required, but the single p.e. signals only cover a very limited region of the overall signal range so there is a requirement to be able to characterize, and correct for, nonlinearity over the full signal range.

The internal calibration facility allows signals to be injected into all TFB input channels to verify functionality and to measure any nonlinearities in the electronic signal processing chain. The value of charge injected is calculated using the programmed calibration DAC voltage and the 10 pF values of the charge injection capacitors (C_{CAL} in fig. 6).

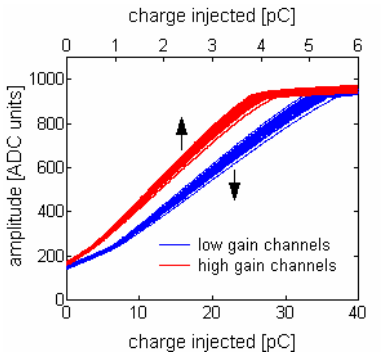

Fig. 11. Dependence of output amplitude on injected signal for all 64 low and high gain channels on a TFB, measured using the on-board calibration circuitry.

Fig. 11 shows typical high and low gain curves for all 64 channels of a TFB measured using the internal calibration circuit. The high and low gain channels can be seen to saturate at approximately 4 and 40 pC injected charge respectively. The spread in gains can be attributed to the 5% tolerance capacitor values used for the gain setting and calibration components used in fig. 6. There are obvious nonlinearities in the curves in fig. 11, particularly where channels start to saturate, but also at lower signal magnitudes. The particularly noticeable change of slope around 6 pC for low gain channels can be attributed to a change of gain occurring when high gain channels saturate. After saturation, no further charge can flow into the high gain channel Trip-t preamplifier, so the distribution of charge between C_G, C_{LO} and C_{HI} (fig. 6) changes. Any signals larger than the high gain channel saturation level can only be shared between C_G and C_{LO}.

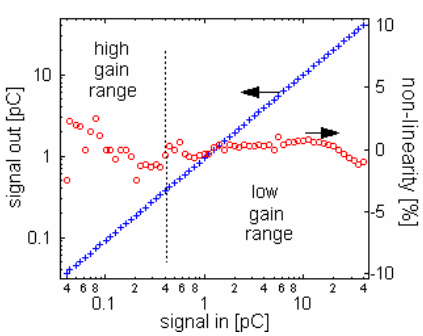

Fig. 12. Typical Trip-t single channel linearity. Blue crosses show dependence of output amplitude on externally injected signal after correction for non-linearity using a look up table generated from the internal calibration data. Nonlinearity dependence on injected signal amplitude is also plotted. If $S_{out}(Q)$ is the signal out for input signal Q, and fit(Q) is the corresponding result given by a straight line fit to all data, then the nonlinearity plotted here is defined as $[S_{out}(Q) - fit(Q)]/Q$ expressed as a percentage.

The data in fig. 11 can be used to generate a look-up-table for high and low gain channels, to allow nonlinearities to be corrected for. Fig. 12 shows the corrected response achieved for a known charge injected on a single channel. To simulate as closely as possible the situation where a signal is generated by an MPPC, a 20 cm miniature coax cable was plugged into a TFB input channel and the charge was injected at the remote end. The input charge magnitude was swept to cover both gain ranges and the composite response in fig. 12 is obtained by combining data from high and low gain ranges after nonlinearity correction. The effectiveness of the nonlinearity correction can also be seen in fig. 12, where residual nonlinearities of just a few percent remain, which we attribute to small differences in response when the simulated signal is injected at the TFB end of the coax cable while the calibration charge is injected at the remote end.

B. Noise

An upper limit of 0.21 p.e. on the acceptable readout system noise is specified in table I. The MPPC interface component values for the TFB (shown in fig. 6) were chosen to achieve a gain of 10 ADC units/p.e. for a representative MPPC gain of $5x10^5$. This leads to an RMS noise requirement of less than 2.1 ADC units for the TFB high gain channels

Fig. 13 shows histograms of the RMS noise levels measured for all 64 channels on a TFB. The high gain channels exhibit slightly higher noise because they have a higher effective input capacitance due to the gain splitting capacitor values, and a trend towards increased noise for longer integration times is also evident. It can be seen that noise levels are in all cases below 2 RMS ADC units, well within requirements.

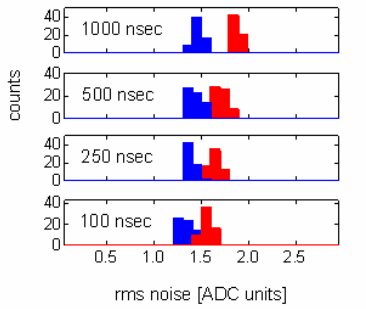

Fig. 13. Histograms of RMS noise for all low and high gain channels on a TFB for different preamplifier integration periods. High gain channels are shown in red, low gain in blue.

C. *Results with MPPC*

The test setup of fig. 10 can be used to verify TFB performance with MPPC sensors by using the TFB output trigger line to generate a short (typically ~ 10 ns) synchronous LED light pulse which can be coupled optically to an MPPC.

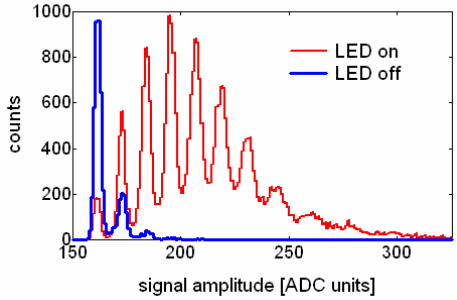

Fig 14. HPK MPPC spectra acquired with a TFB channel, obtained with and without an LED generated light pulse during the 250 ns Trip-t preamplifier integration window.

Fig. 14 shows an MPPC spectrum obtained with and without an LED light pulse incident on the MPPC during the Trip-t integration period. Discrete peaks corresponding to different numbers of photoelectrons generated in the MPPC are clearly visible, indicating an acceptable noise level. The spectrum with no LED light arises due to the MPPC dark rate, where the pedestal, single and double p.e. peaks are visible. These spectra will be generated offline during physics data taking to calibrate and track changes in MPPC gain.

Fig. 15 demonstrates the operation of the HV trimming circuitry where the voltage supplied by the HV trim DAC for the MPPC channel subtracts from the 72.5V HV global supply to the TFB. The nominal operating voltage for this particular MPPC is 70V at a gain of 7.5×10^5, which corresponds to the spectrum where the gain is 120 fC/p.e.. The 8-bit HV trim DACs currently used on the TFB allow the HV to be trimmed with a precision of 20 mV, as specified in table I.

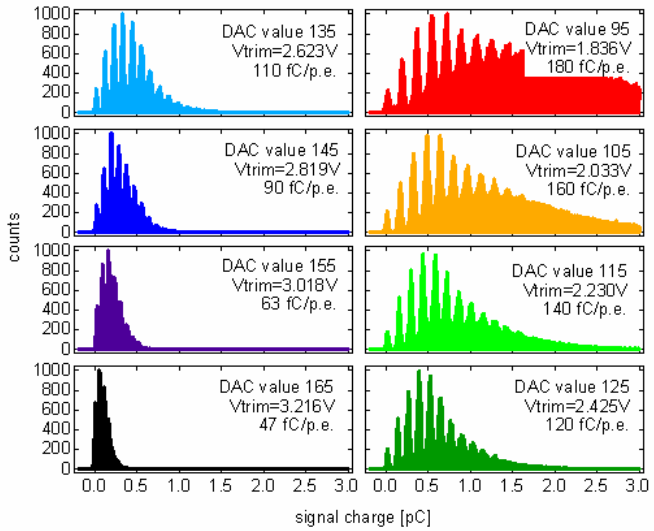

Fig 15. HPK MPPC spectra acquired with a TFB channel, showing the effect of HV trimming on the MPPC gain. The global TFB high voltage bias is set to 72.5 V and the voltage corresponding to the HV trim value DAC subtracts from that value.

D. *Discriminator resolution and Time-stamping*

The time-stamping functionality implemented in the TFB FPGA allows to register a firing discriminator with a resolution of 2.5 ns. An individual discriminator will fire if its input signal exceeds a programmed threshold. The threshold is common to all channels on the chip, so any intrinsic channel-to-channel threshold mismatch will determine the effective resolution of the discriminator performance. Apart from channel mismatch, noise will also degrade the discriminator threshold resolution. In T2K the discriminator threshold will be set to limit firing on dark rate signals.

Discriminator resolution can be characterized by a turn-on curve, measured experimentally by injecting a signal of fixed magnitude for a fixed number of times, and counting the number of times the discriminator fires while sweeping the threshold over a range from where the discriminator never fires to where it fires every time the signal is injected.

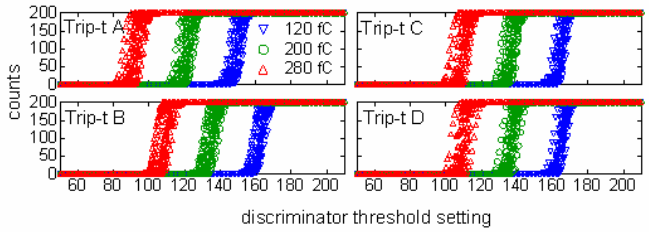

Fig 16. Discriminator threshold scans for all 16 channels on all 4 Trip-t chips on a TFB. Three separate measurements for 1.5, 2.5 and 3.5 p.e. equivalent external injected charge are shown (assuming 5×10^5 electrons/p.e.)

Fig. 16 shows discriminator turn-on curves for all 64 channels on a TFB, separated into groups of 16 by Trip-t, for three different values of charge injected using the TFB calibration feature. For an MPPC with a gain of 5×10^5 these signal sizes would correspond to 1.5, 2.5 and 3.5 p.e.. Systematic differences between chips are of no significance because different thresholds can be programmed. The channel-to-channel spread in the discriminator turn-on curves can be

seen to be ~ 0.2 p.e., and the clear separation between curve families for different signal levels demonstrates that it will be possible to discriminate against signals at the single p.e. level. In other words, a threshold can be defined where double p.e. signals will always fire a discriminator and single p.e. signals will not.

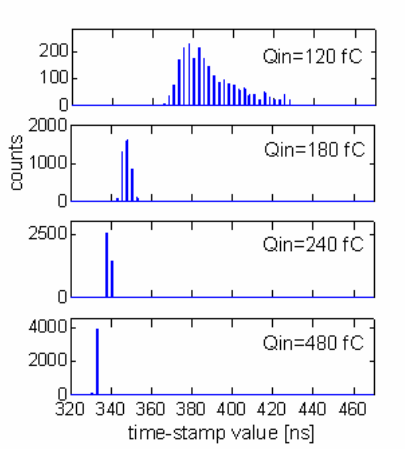

Fig 17. Histograms of time-stamp values for one Trip-t for four values of externally injected charge in the range 120 to 480 fC, with a discriminator threshold setting of 120 fC. Time-stamps generated using one calibrate line feeding four channels on one Trip-t, for 1000 triggers. Each bar represents one 2.5 ns time bin.

The time-stamping performance is demonstrated in fig. 17 which shows histograms of time-stamp values generated by signals of different sizes in the range 120 to 480 fC, where the average discriminator threshold, determined from fig. 16, has been set to 120 fC. The data in fig. 17 are generated using a single calibrate line (firing four channels), for 1000 triggers. When the signal size is the same as the threshold the distribution of time-stamp values is broad, and reflects the spread in the discriminator turn-on curves in fig. 16. For larger signal sizes the distributions get increasingly narrower, becoming confined to one or two 2.5 ns time bins only.

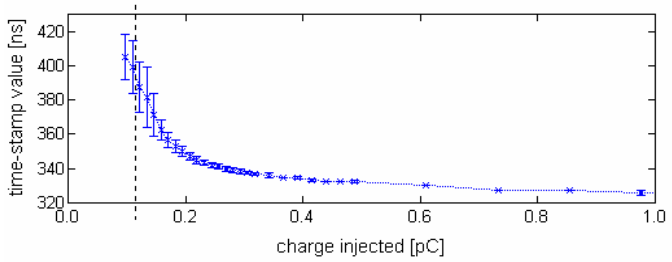

Fig 18. Timewalk curve; the average time-stamp value dependence on signal size. Points represent the mean of the time-stamp histograms (fig. 17). Error bars shown are ± one standard deviation. The vertical dashed line shows the position of the discriminator threshold at 120 fC.

Fig. 18 shows the time-walk curve corresponding to data taken in an identical fashion to that in fig. 17. Time-walk occurs because of the finite rise-time of the signal at the Trip-t preamplifier output and can be corrected for using the time-walk curve and the amplitude of the signal. Once the signal exceeds approximately twice the value of the discriminator threshold setting in fig. 18 the uncertainty in the value of the time-stamp reflects only the 2.5 ns resolution of the time-stamping functionality.

E. Power and other functionality

Table II shows the measured TFB power consumption broken down by supply rail. The total power level of 3.5 Watts translates to 55 mW / MPPC channel, leading to an on-detector overall power consumption of ~ 3 kW. The high voltage MPPC bias power consumption is negligible.

TABLE II
TFB POWER CONSUMPTION

external supply [V]	after regulation [V]	TFB circuitry supplied	current consumption [A]	power consumption [W]
1.7	1.2	FPGA core	0.26	0.44
3.1	2.5	Trip-t & FPGA I/O	0.6	1.9
3.8	3.3	ADC and FPGA I/O	0.05	0.2
5.5	5.0	ADC and HV trim DACs	0.18	1.0
				3.6

The circuitry responsible for monitoring voltages, currents and temperature has been tested and its performance is consistent with the specifications of the individual commercial components described in section III-C.

V. CONCLUSIONS

The Trip-t Front-end Board (TFB) which will be used to read out novel Multi-pixel Photon Counter sensors in the T2K experiment has been produced. The functionality and performance has been measured and is well matched to the experimental requirements.

Large scale production of TFBs will begin in 2008. Approximately 1000 will be required for the T2K experiment, which will begin operation in 2009.

ACKNOWLEDGMENT

We thank Ray Yarema and Paul Rubinov of Fermilab for providing access to the Trip-t chip, and for all their help and advice on how to make best use of it for our application.

REFERENCES

[1] Y. Itow et al, "The JHF-Kamioka neutrino project", arXiv:hep-ex/106019, 2001.
[2] S. Fukuda et al, "The Super-Kamiokande detector", A501, pp. 418-462, 2003.
[3] M. Yokoyama et al, "Development of Multi-Pixel Photon Counters", arXiv:physics/0605241, 2006.
[4] J. Estrada, C. Garcia, B. Hoenison, and P. Rubinov, "MCM II and the Trip chip," D0 note 4009, Fermilab-TM-2226, Dec 2003.