

14. PHOTODETECTOR AND ELECTRONICS.....	2
14.1 INTRODUCTION.....	2
14.2 TECHNICAL DESIGN CRITERIA	2
14.3 AVALANCHE PHOTODIODES (APDs).....	2
14.4 CARRIER BOARDS	6
14.5 APD MODULE.....	7
14.5.1 APD Module Testing.....	7
14.6 APD COOLING	8
14.6.1 Thermoelectric Coolers	8
14.6.2 Water Cooling System.....	9
14.7 LOW NOISE ASIC AMPLIFIER.....	11
14.8 SPECIAL ELECTRONICS VERSION FOR THE NEAR DETECTOR.....	13
14.9 FRONT END BOARDS	14
14.9.1 TEC Controller	15
14.9.2 Signal Extraction	15
14.9.3 Operating modes.....	16
14.9.4 Digital Signal Processing	17
14.10 POWER DISTRIBUTION.....	18
14.11 CHANGES IN THE PHOTODETECTOR AND ELECTRONICS DESIGN SINCE THE CDR	24

14. Photodetector and Electronics

14.1 Introduction

The NOvA photodetector is an avalanche photodiode (APD) operated at a gain of 100, cooled to -15°C by a thermoelectric cooler (TEC), and readout via a Front-End Board (FEB). The FEB has a low noise, $<200e^{-}$ custom ASIC amplifier matched to the APD. Over 14,000 APD / TEC / FEB assemblies are required for the NOvA Far Detector.

14.2 Technical Design Criteria

The readout of the NOvA Far Detector has two distinct tasks: (1) read out events caused by neutrinos from the NuMI beamline at Fermilab and (2) operate between spills to collect cosmic ray events for calibration and monitoring. The readout will operate in a triggerless mode to accomplish both tasks seamlessly.

The NOvA Near Detector uses similar electronics but must satisfy different constraints. The mean signal from the far end is approximately 4 times greater for the shorter modules, and there will be multiple neutrino induced events in the detector during the NuMI spill. A modified version of the basic Far Detector design that samples each channel more frequently is required.

A time-stamp generated from the kicker fire, signal S74, from the NuMI beam line will be used to determine which events occur during the $11.1\ \mu\text{s}$ single turn extraction of the Main Injector protons onto the NuMI target. The actual length of time the protons will hit the target will be $6/7$ of this, or $9.5\ \mu\text{s}$.

14.3 Avalanche Photodiodes (APDs)

The photodetector for NOvA is an avalanche photodiode (APD), shown in

Fig. 14.1. The APDs are packaged in arrays of 32 pixels and mounted on a carrier board substrate using flip-chip mounting. This device has been custom made for NOvA to optimize the fit of two fiber ends on a single pixel. The 32 pixels map directly onto the 32 cells of a single PVC extrusion module. Table 14.1 summarizes the key parameters for the NOvA APDs.

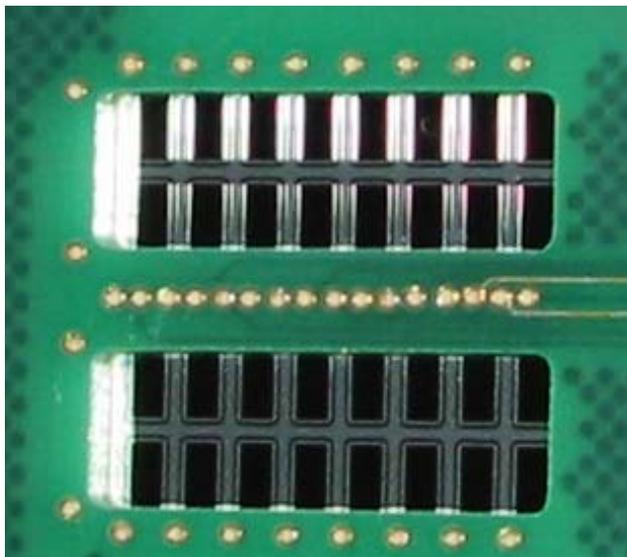


Fig. 14.1: Prototype NOvA APD mounted on carrier board.

Manufacturer	Hamamatsu
Pixel Active Area	1.95 mm × 1.0 mm
Pixel Pitch	2.65 mm
Array Size	32 pixels
Die Size	15.34mm × 13.64mm
Quantum Efficiency (>525 nm)	85%
Pixel Capacitance	10 pF
Bulk Dark Current (I_B) at 25 C	12.5 pA
Bulk Dark Current (I_B) at -15 C	0.25 pA
Peak Sensitivity	600 nm
Operating Voltage	375 ± 50 volts
Gain at Operating Voltage	100
Operating Temperature (with Thermo-Electric Cooler)	-15°C
Expected Signal-to-Noise Ratio (Muon at Far End of Cell)	10:1
APD channels per plane	384
APD arrays per plane	12
Total number of planes	1178
Total Number of APD arrays	14,136
APD pixels total	452,352

Table 14.1: Avalanche Photodiode parameters.

The general structure of an APD is shown in Figure 14.2. Light is absorbed in the collection region, electron-hole pairs are generated and, under the influence of the applied electric field, electrons propagate to the p-n junction. At the junction, the electric field is sufficiently high that avalanche multiplication of the electrons occurs. The multiplication of the current is determined by the electric field at the junction and by the mean-free-path of electrons between ionizing collisions, which depends on both the accelerating field and on the temperature. This temperature dependence occurs because the probability of electron-phonon scattering, which competes with the avalanche multiplication process, increases with temperature.

One of the operational characteristics of APDs, and, in fact, all silicon devices, is the thermal generation of electron-hole pairs which mimic the signal. Since the current from the positive carriers is amplified about fifty times less than the negative carrier current at the junction, only the current from electrons generated in the photo-conversion region (I_B), or the bulk current, needs to be considered in the noise current estimation. As it is a thermally generated current, it can be reduced by lowering the operating temperature of the APD. We will operate the APDs in the NOvA detector at -15°C to keep the noise contribution from I_B small in comparison to the front-end noise. This choice is based on measurements obtained with prototype readouts.

The amplification mechanism in the APD is itself subject to noise, characterized by the excess noise factor F , with such factors as device non-uniformities and the ratio of the positive to negative impact ionization coefficients contributing. This factor is well modeled and has been included in our signal to noise calculations.

APDs have two substantial advantages over other photodetectors: high quantum efficiency, and uniform spectral quantum efficiency. The high APD quantum efficiency enables the use of very long scintillator modules, thus significantly reducing the electronics channel count. In the wavelength region relevant to the output of the wavelength shifting (WLS) fibers, 500 to 550 nm, the APD quantum efficiency is 85%. See Figure 14.3.

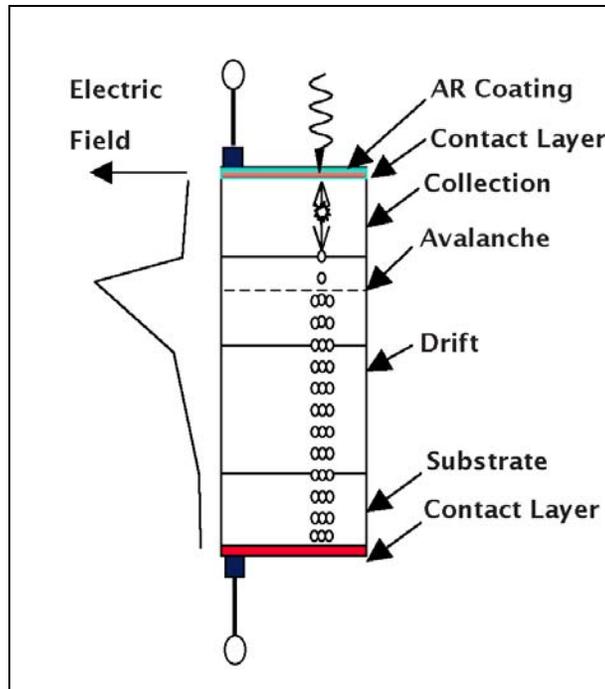


Fig. 14.2: The basic structure of a blue/green sensitive APD. Light crosses the anti-reflection coating at the surface and is absorbed in the collection region. Photoelectrons drift in the electric field to the junction where they undergo avalanche multiplication.

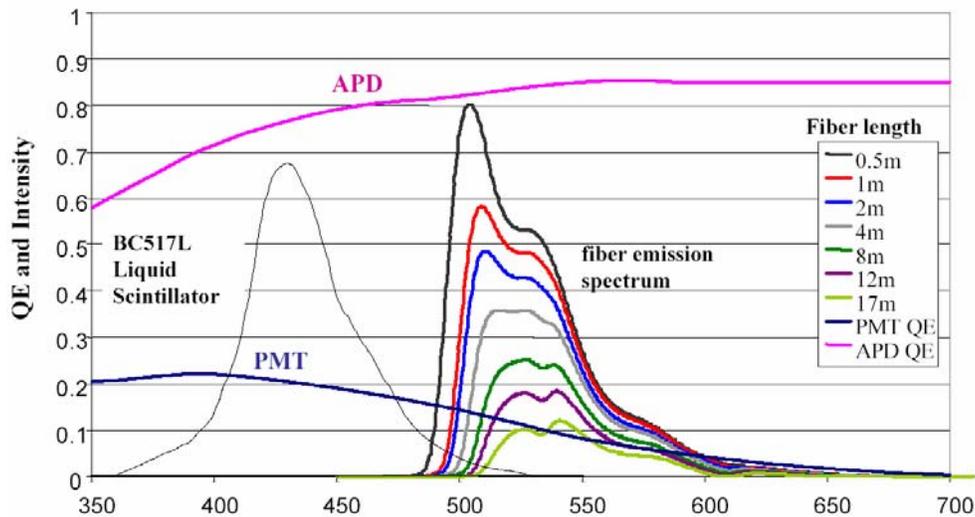


Fig. 14.3: WLS fiber emission spectra measured at lengths of 0.5, 1, 2, 4, 8, 16 m, respectively illustrating the shift of the average detected wavelength as fiber length increases. Also shown are the quantum efficiencies of APDs and PMTs (bialkali photocathode).

Hamamatsu is the only known vendor for APD pixel arrays. Hamamatsu markets a 32-pixel packaged APD with a pixel size of 1.6 mm by 1.6 mm. To maximize light output, NOvA utilizes a looped or U-shaped WLS fiber. Both ends of the looped fiber terminate on the same APD pixel. To comfortably accommodate both fiber ends, Hamamatsu has modified the pixel size and shape to match our requirements. Prototypes of the modified APDs have been provided bump bonded to an APD carrier circuit board. These devices are being tested as part of our qualification

process for the final design. Results of qualification tests for gain and bulk dark current (dark current divided by gain) are shown in Fig. 14.4 and Fig. 14.5. These channels performed well.

For initial evaluation we also purchased a number of Hamamatsu S8550, packaged APD arrays. The dark currents were consistent with expectations, and the gains were uniform from pixel to pixel on the same chip and within individual pixels. The measured pixel separation for one of the sample arrays is shown in Fig. 14.6. The fall-off on the pixel edges reflects the finite spot size used to illuminate the APD pixels.

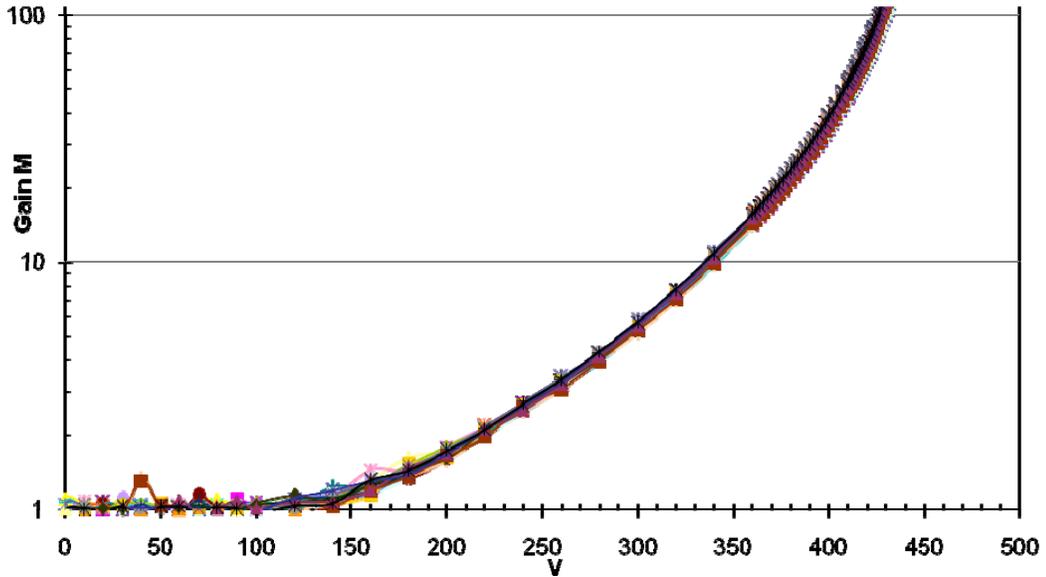


Fig. 14.4: Gain vs. applied voltage at Room Temp for 32 channels of a Hamamatsu NOvA prototype APD array.

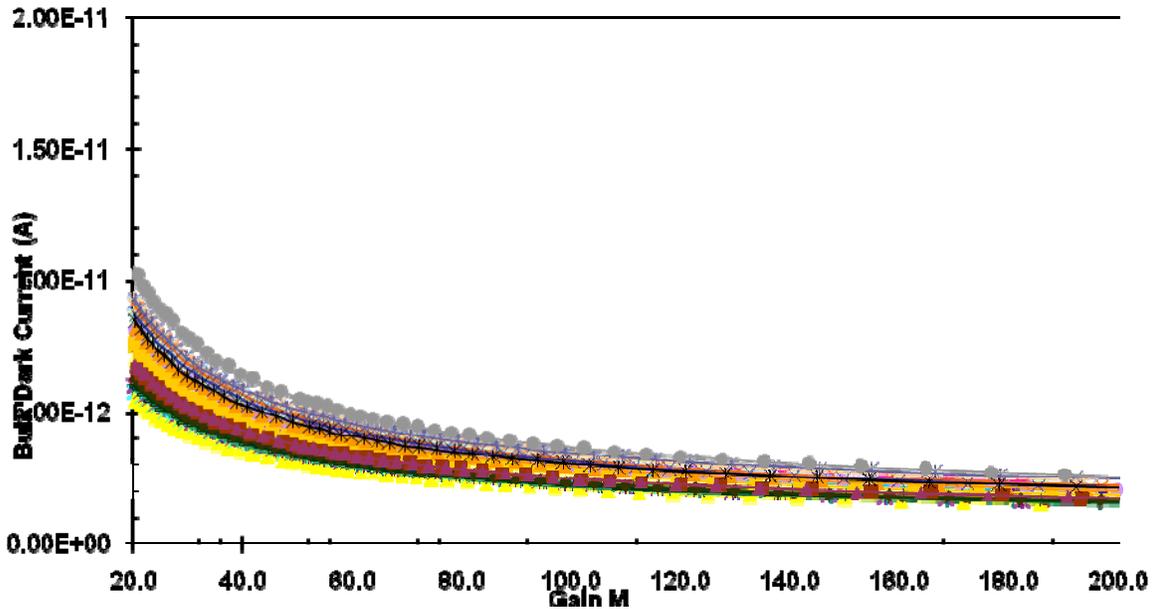


Fig. 14.5: Bulk Dark Current vs. gain for 32 channels of a prototype APD array measured at room temperature. Cooling to -15C will decrease the current by a factor of at least 50.

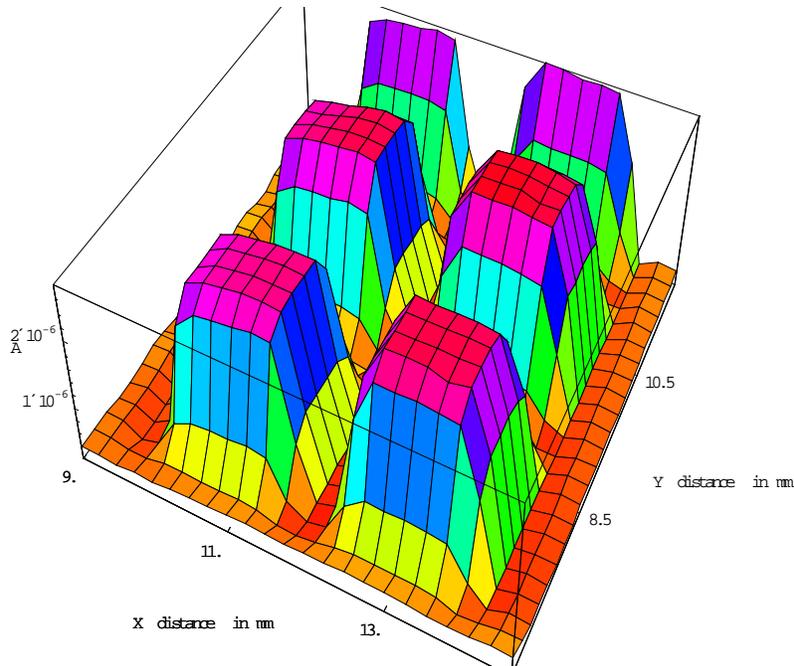


Fig. 14.6: Fine point scan across part of a Hamamatsu S8550 APD array. The fall-off on the pixel edges reflects the finite spot size used to illuminate the APD pixels.

One of the attractive features of APDs is that once they have been calibrated, the gain can be easily determined from the applied bias voltage and the operating temperature. In the NOvA detector, we will maintain the operating bias to a precision of 0.2 Volts and control the temperature to 0.5°C and thus hold the gain stability to about 3%, consistent with the pixel-to-pixel variation. The absolute calibration of temperature and voltage are not critical for the experiment, only the gain. The NOvA APDs will typically operate between 350V - 450V at a standing current of approximately 1 nA per 32-channel APD array. The high voltage system is described in Section 14.10.

14.4 Carrier Boards

The APD arrays will be mounted on a separate APD carrier board that is environmentally isolated from the other electronic components to minimize the thermal load. The mounting will be done with flip-chip technology, so the active area of the APD will face rectangular slots cut out of the PC board where the fibers will terminate. The back of the APD is protected and stiffened by a ceramic backer as shown in Fig. 14.7. The flip-chip method provides an accurate means of aligning the APD to the PC board to which the fiber connector will also be aligned. The fiber connector must accurately align the fibers both longitudinally and transversely to the APD pixels.

One of the operating requirements for the APDs is that they be kept dry. Dew-point concerns associated with the low operating temperature of the TE coolers have led us to a design where the APD and TE cooler are enclosed in a sealed environment, kept dry with silica gel as in computer disk drives.

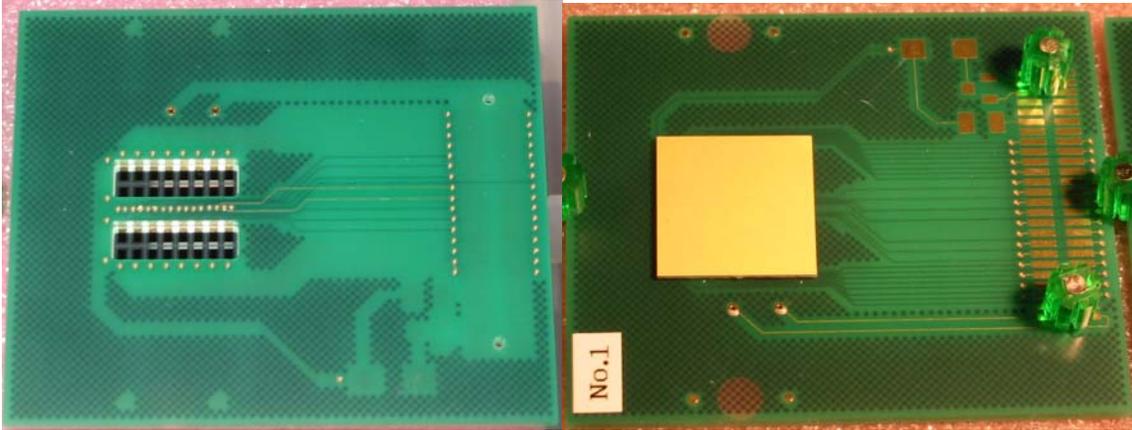


Fig. 14.7: Front and back views of an APD array mounted to a carrier board.

14.5 APD Module

The APD arrays mounted on the carrier will be integrated into a housing, and the device will be called the APD Module. The requirements on the mechanical housing for the APD arrays and associated systems required to operate the APD's are enumerated. This includes the requirements on the (1) structural, (2) thermal, (3) environmental, (4) optical and (5) electrical couplings/interfaces that must be provided by the APD module system.

The APD module provides an optical interface between a single 32-cell scintillator detector module and a 32-channel Avalanche Photodiode array, which is used to convert the optical signals from the detector module to electronic signals. The APD module thus also provides an interface to the Front End Board (FEB).

In addition to the interfaces mentioned above, the APD module must provide the means to operate the APD array in a robust and problem-free manner. A stable bias voltage must be supplied. Control and readback of APD array temperature must be provided. Exposure to light and moisture must be avoided during operation of the APD module. Ease of assembly and disassembly must be considered.

Consequently, the APD module currently consists of the following components: (a) a 32-channel APD array, (b) a "carrier" printed circuit board on which the APD array is mounted, (c) a thermoelectric cooler (TEC), (d) a heat sink for removal of heat from the TEC, and (e) an enclosure. These components are shown in Figure 14.8.

14.5.1 APD Module Testing

The APD module and APD array will be tested at the same time. An automated testing apparatus is being developed to perform this function efficiently, 6 modules at a time. The design can be scaled up to do 12 at a time if necessary. The individual pixels of the APD arrays will be tested for dark current that is out of spec, high or low, and gain as a function of voltage. These measurements will be performed at room temperature for comparison to the manufacturer tests and at operating temperature of -15C. This will test the operation of the TEC devices as well as the APD. The test will determine the operating voltage required to operate each array at a gain of 100 at the operating temperature of -15C. The results of the first test will determine the required setting of the voltage adjustment potentiometer on each carrier board. The board will be set to the correct value, and retested to confirm the proper gain is achieved.

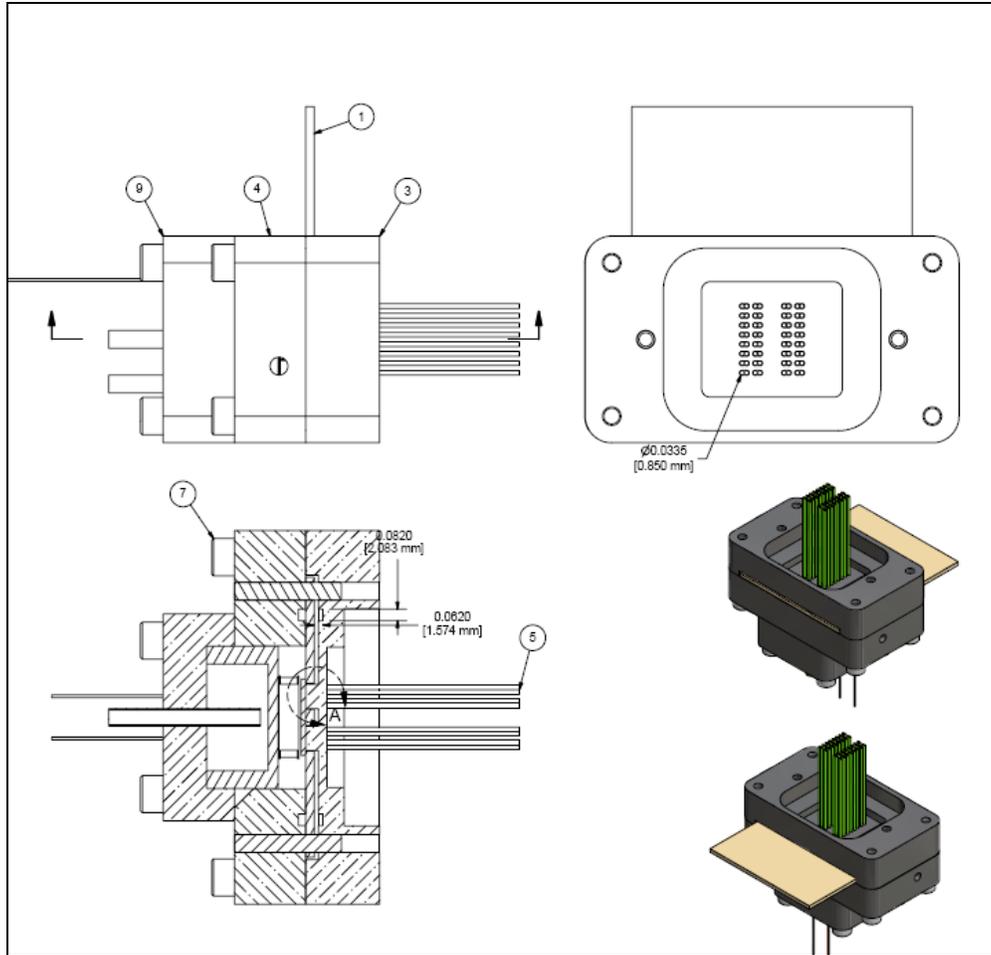


Fig. 14.8: APD Module design which isolates the APD from the environment and provides cooling via the TEC

14.6 APD Cooling

14.6.1 Thermoelectric Coolers

A single-stage TE cooler will cool each APD array. The APD module has been tested with a TEC from TE Technology, TE-31-1.0-1.3, with specifications given in Table 14.2. This has a maximum cooling capacity of 8.4W, and size to match the APD array. Measurements with a prototype APD module have shown a TE cooler electrical power requirement of less than 2W, and heat removal of less than 2W to maintain an APD temperature of -15°C . The TE cooler and power supplies have additional headroom of 3.2W and 5W respectively. The system should quickly and easily reach operating temperature. The TE cooler must not apply significant mechanical stress to the APD array, so we deploy a deformable, thermally conducting crush pad between the TE cooler and the APD array. The thermal power generated in the APD array itself is $\sim 2 \mu\text{W}$, so the thermal load will come from other components through the mechanical and electrical interconnects. The TE cooler will generate approximately 3W of heat for each 32 channel APD array.

Given the industry standard reliability of a TEC of 200,000 hours MTBF, or 22.3 years, we expect a 4%/year failure rate of the TE coolers. The APD module is designed to allow replacement of the TEC without complete removal from the detector module.

Part Number	Qmax (W)	I _{max} (A)	V _{max} (V)	dT _{Max} (C)	Width (mm)	Length (mm)	Height (mm)
TE-31-1.0-1.3	8.4	3.6	3.8	69	14.8	14.8	3.6

Table 14.2 Specification for TEC tested with the APD module.

14.6.2 Water Cooling System

Liquid cooling is a reliable and effective means of removing the heat from the TE coolers. The high heat capacity of water facilitates the removal of the required heat without requiring large temperature differences. We will use a chilled water system that provides approximately 2mL/s of 15°C water to remove the heat from each TE cooler. The water will be supplied by cooling loops installed in pairs covering each half of a pair of blocks. Each loop will service 372 TE coolers. Commercially produced compact process water chillers with capabilities similar to Neslab Merlin150 or HX150 will be installed at the head end of each loop on the highest catwalk level. The chillers will reject heat directly to air. Each chiller is capable of maintaining its water output temperature within +/- 0.1 deg C of a setpoint,

A block diagram of a cooling loop is shown in

Fig. 14.9. Water is distributed to the cooling elements via a supply and return manifold system. The system is arranged with reverse return flow to help naturally balance the flow through each element. Each loop contains a monitoring station that houses control and monitoring equipment for each loop. There are 38 of these cooling loops for the entire detector. The chillers and monitoring stations are located on both sides of the detector at 19 locations approximately every 14 feet along the length of the detector.

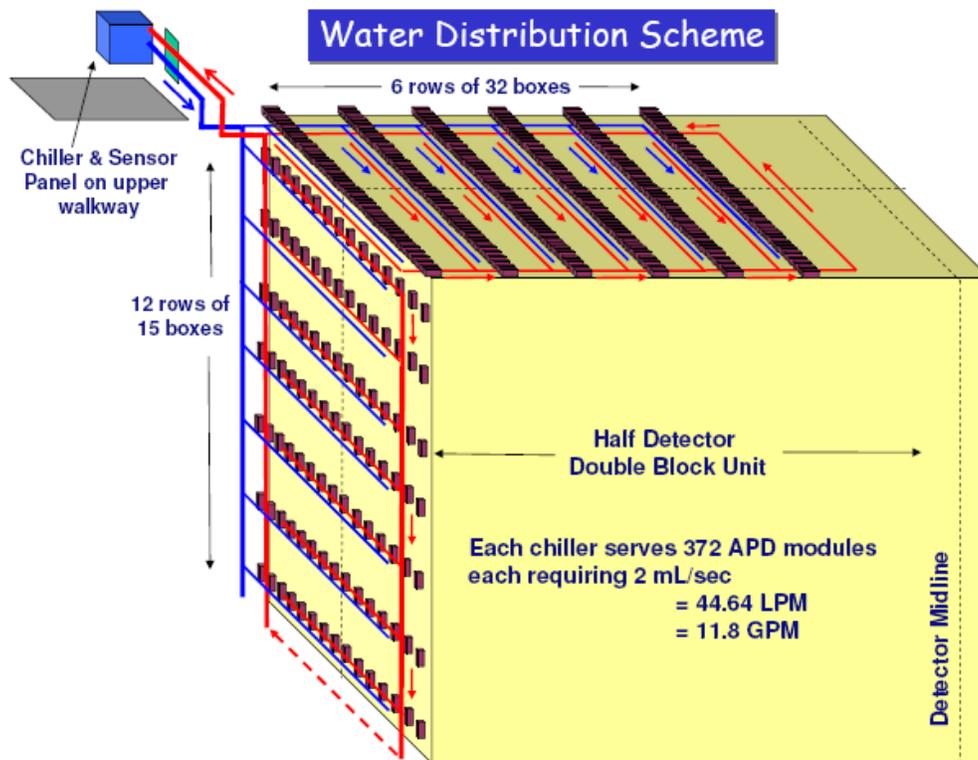


Fig. 14.9: Schematic of single water cooling loop.

There is a main distribution manifold that supplies water to the 18 different service manifold sections, 12 on the side and 6 on the top of the detector. There are two types of service manifolds due to the lower density of readout modules on the sides of the detector, but the overall features are the same. The supply and return manifolds in each location will be identical. The general features of a manifold are shown in Fig. 14.10. There is a main supply with 16 or 32 hose connections with valved quick disconnect connectors at their ends. The service manifolds are installed as a single section that corresponds to a pair of 31 plane blocks. This matches the readout electronics commissioning as well.

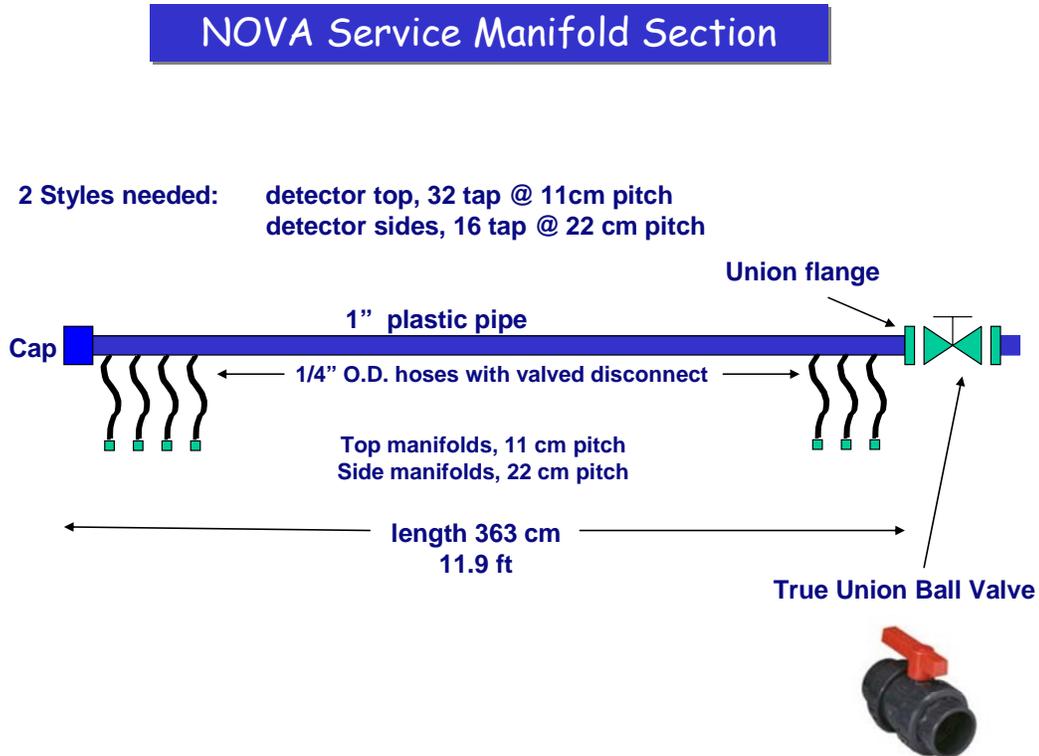


Fig. 14.10: Diagram of manifold section used for supply and return.

Supervisory control and monitoring of the cooling system will be integrated as part of the Slow Controls system. It will log historical data, manage alarms, and be the primary human interface to the cooling system. While monitoring supply and return temperatures, pressures, flow, and level in each water cooling loop, it will be able to remotely turn the chiller units on and off, and adjust the chiller temperature setpoints to optimize cooling performance. A block diagram of one of the 38 loop remote control systems is shown in Fig. 14.11. In addition, portable readout units will be available to directly monitor the sensor data and control signals. These simple electronic readouts will plug in to the sensor junction boxes at each chiller station and will not depend on the ethernet or Detector Controls computer systems.

Screens presenting realtime cooling system data will be accessible to technicians in the detector hall via a wireless network to facilitate maintenance and troubleshooting. In addition, each chiller includes a local control panel allowing manual override of remote control if necessary.

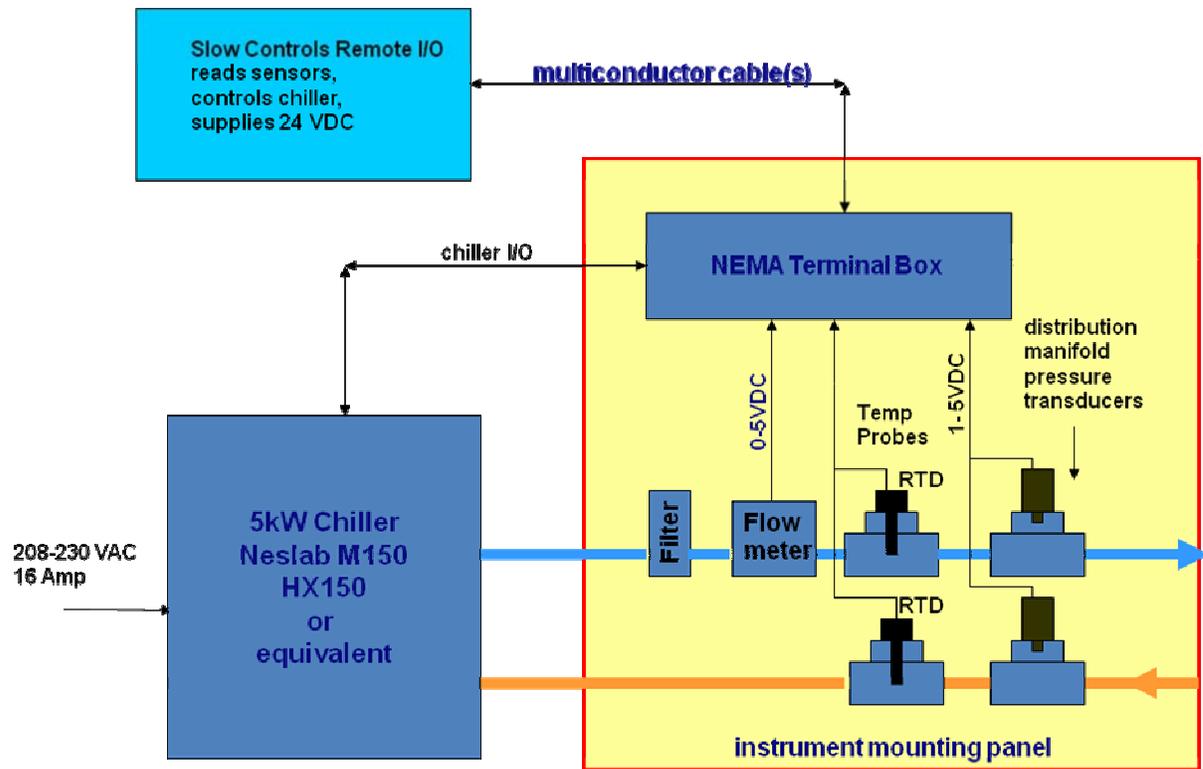


Fig. 14.11: Pump control and loop instrumentation

One important function of the control system, leak detection, will be accomplished by coolant level monitoring in the chiller reservoirs. The control system will alert operators and can automatically shut down systems to minimize damage and conserve coolant. In the worst case, a major leak resulting in rapid loss of the entire volume of a cooling loop would involve less than 50 gallons of water.

To minimize the peak electrical and building air cooling loads at startup, the control system will automatically sequence the application of power to groups of cooling loops, allowing one group to stabilize before starting another. This would keep both the peak and stabilized net heat rejected to air at about 100kW, assuming 2.5w per APD module, steady state.

Each chiller requires 1.67kW electric power. With all 38 chillers running, this would total 63.5kW. This energy, converted to heat, makes up a large portion of the heat-to-air load mentioned above. Since this chiller load is virtually all from induction motors, the power factor will be low, around 0.5, and may require corrective measures. If sequenced on in groups of 3 or 6 when possible, the 3 phase power load would remain balanced.

14.7 Low Noise ASIC Amplifier

A custom ASIC has been developed for NOvA to maximize the sensitivity of the detector to small signals from extremely long fibers in the far detector. The schematic is shown in

Fig. 14.12 with a micrograph of the device that was produced in

Fig. 14.13. Signals from individual APD pixels are processed through individual amplifier and pulse-shaping stages before being multiplexed to an ADC in sets of 8 channels. The ADC is external to the ASIC and is a commercial component. A schematic of the layout is shown in Fig. 14.14.

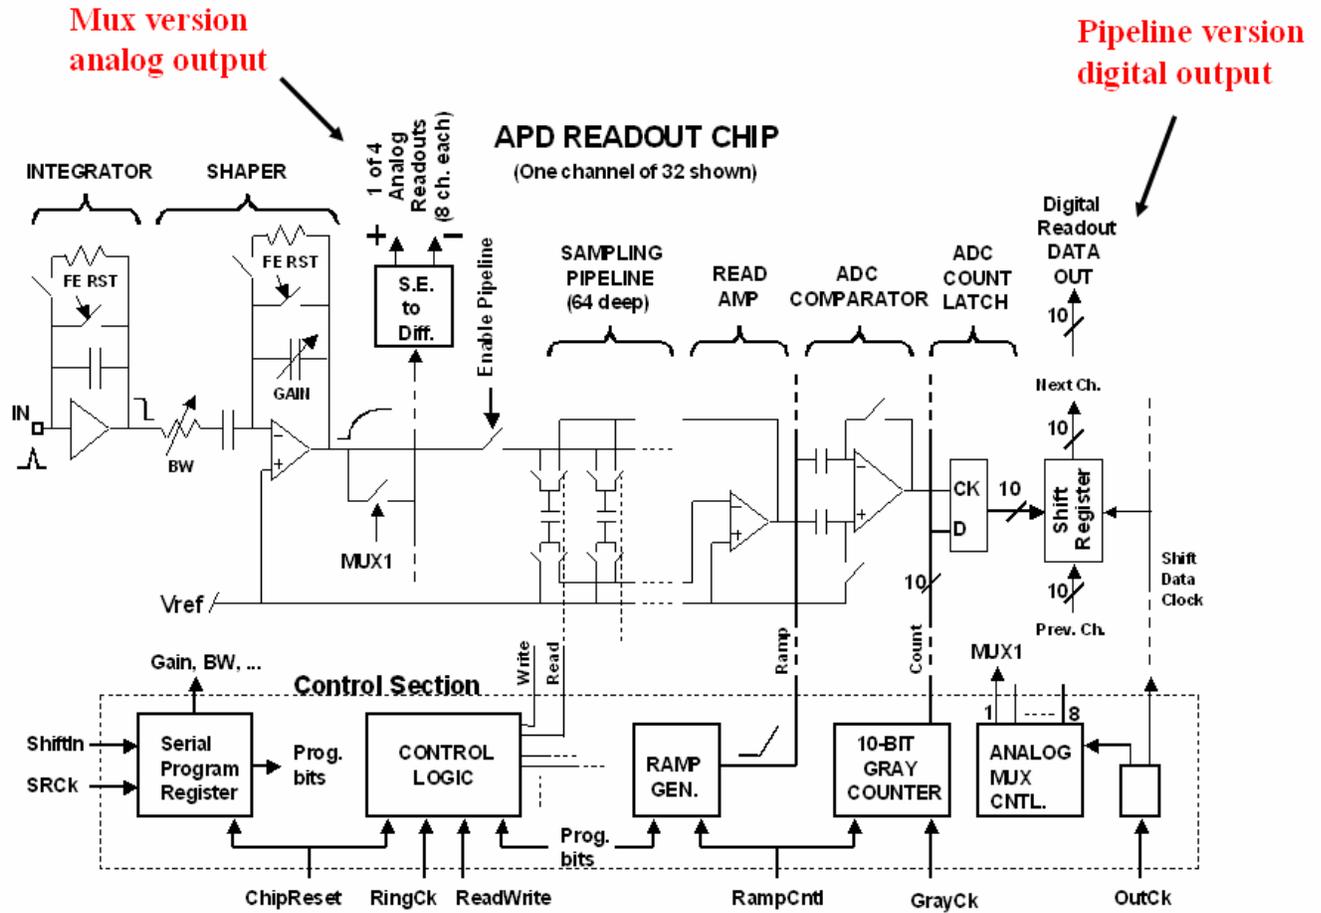


Fig. 14.12: NOvA custom ASIC schematic

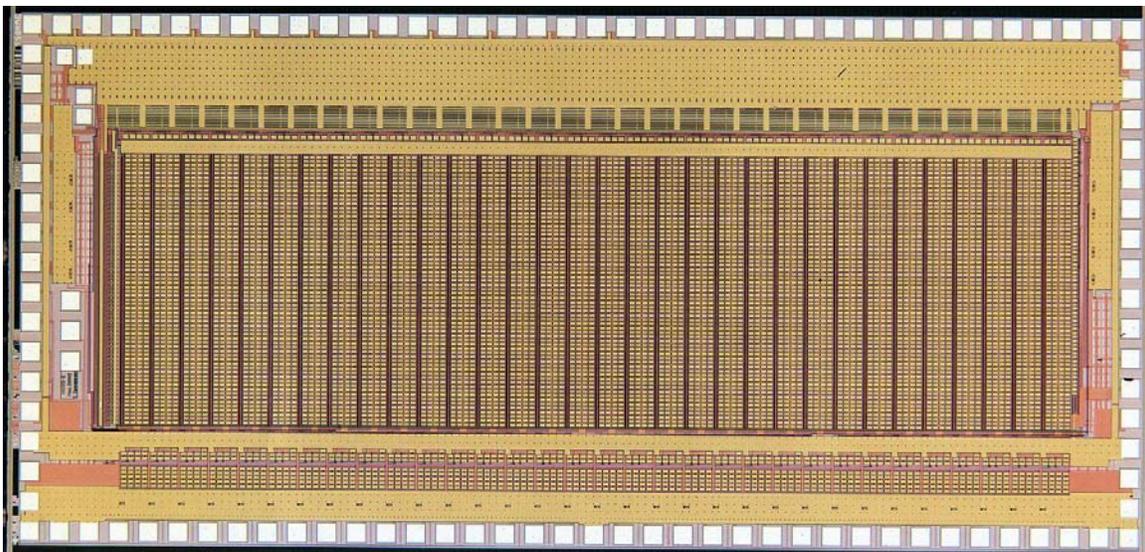


Fig. 14.13: Micrograph of prototype readout ASIC

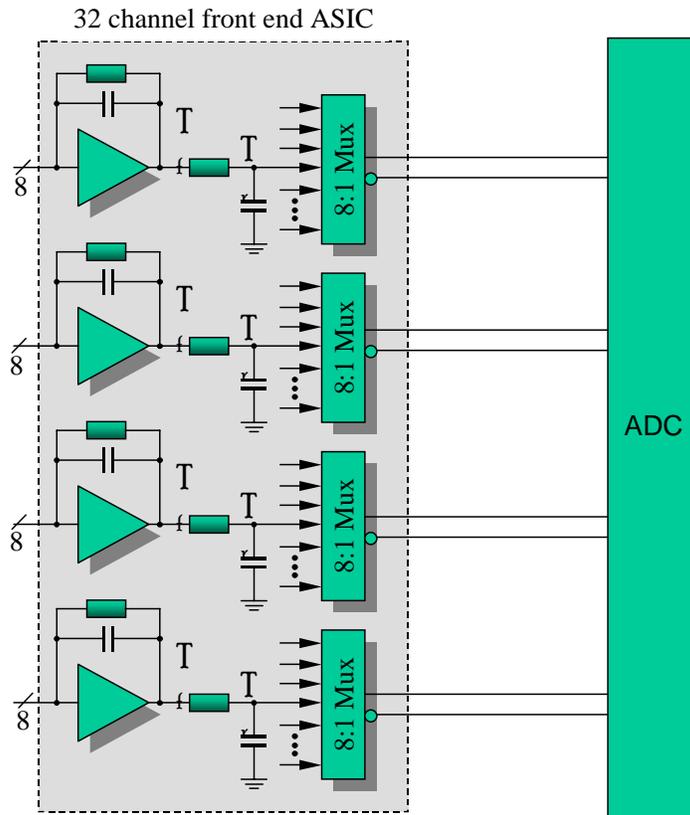


Fig. 14.14: Schematic of the Front End ASIC and external ADC for NOvA.

The computed noise level for the chip that we have designed specifically to operate with an APD with a gain of 100 and cooled to -15°C is 150 electrons. The noise measured on a prototype is 100 electrons for the optimal input transistor. Our calculations assume that this will increase to 150 electrons when packaged and mounted on the FEB. With an APD gain of 100, this 150 electrons RMS noise reduces to 1.5 photoelectrons equivalent. At -15°C we also expect 1.5 thermally generated electrons every $1\ \mu\text{s}$. The convolution of the amplifier noise with the APD noise results in a mean of 2.3 photoelectrons of noise. This is to be compared to an average photoelectron yield at the far end of an extrusion module of >20 , spread over a very short time interval. This design will have good separation of signal and noise.

14.8 Special Electronics Version for the Near Detector

The NOvA Near Detector will receive a high rate of neutrino interactions, and will also receive a large number of particles from the tunnel walls (see Chapter 6, Section 6.6.3). Detector simulations indicate that the high occupancy during a spill requires better double pulse separation than is possible using the standard electronics from the far detector. For simplicity and cost effectiveness we intend to use as many of the same components as possible in each detector. The detector will therefore use the same APD and TEC. The front-end of the low noise ASIC amplifier is also identical to that of the far detector. The ASIC will be modified from the prototype to have additional output amplifiers and multiplexing logic to operate with only 2:1 multiplexing instead of 8:1 as is used in the far detector, increasing the sampling frequency by a factor of 4. The additional amplifiers will be switched off for use at the far detector. This design will use four times as many ADCs as the far detector due to the decreased multiplexing. This will

increase the sampling rate by a factor of four. The prototype for this ASIC is currently being designed. There are not expected to be any technical problems in implementing a selectable multiplexing scheme since it has been done before to implement a selectable 16:1 or 4:1 multiplexing on the MASDA ASIC we have used for detector development. In this case it is expected that the same ASIC can be used for the far detector as well with some additional logic to disable the additional amplifiers and increase the number of channels each amplifier is multiplexing. The additional sampling and higher light yields in the smaller detector will aid in detecting and reconstructing the overlapping events in the detector.

14.9 Front End Boards

The Front-end Electronics board (FEB) is connected to the APD carrier board through a short ribbon cable. In addition to the front end ASIC and ADC, the front-end board contains a connector for interfacing to the DAQ system, TE cooler controller circuitry, DACs and ADCs for control and monitoring, and an FPGA for doing digital signal processing, I/O functions, and general board monitoring. The front end boards require low voltage to power the electronics and the TE cooler controller. The front end electronics boards will each require 3 W of clean power at 3 V. The TE coolers will require 2-5 W to be provided at 24 V. Low voltage power is discussed in Section 14.10. A sketch of the front-end board and the APD carrier board is shown in Figure 14.15. Layout of the critical components is shown in Figure 14.16.

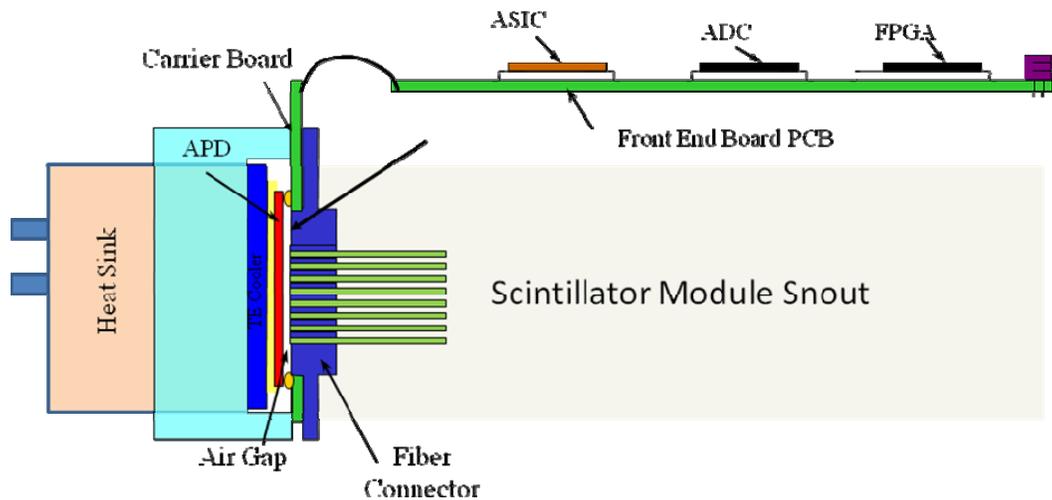


Fig. 14.15: Schematic of the APD module and the front-end electronics board showing the major components.



Fig. 14.16: Layout of the FEB. Major components are the carrier board connector location at the left, which brings the APD signals to the NOvA ASIC, which performs integration, shaping, and multiplexing. The chip immediately to the right is the ADC to digitize the signals, and FPGA for control, signal processing, and communication.

14.9.1 TEC Controller

The TEC controller circuit is a step-down switching converter that uses 24V as input and outputs a variable voltage up to 2.2A at 2.3V at an efficiency of 78%. The output is adjusted so that the thermistor on the APD package maintains a constant temperature of -15C. The circuit has been designed to minimize the output ripple so that it does not add significant additional noise to APD output. A photograph of the circuit as built and tested is shown in Figure 14.17.

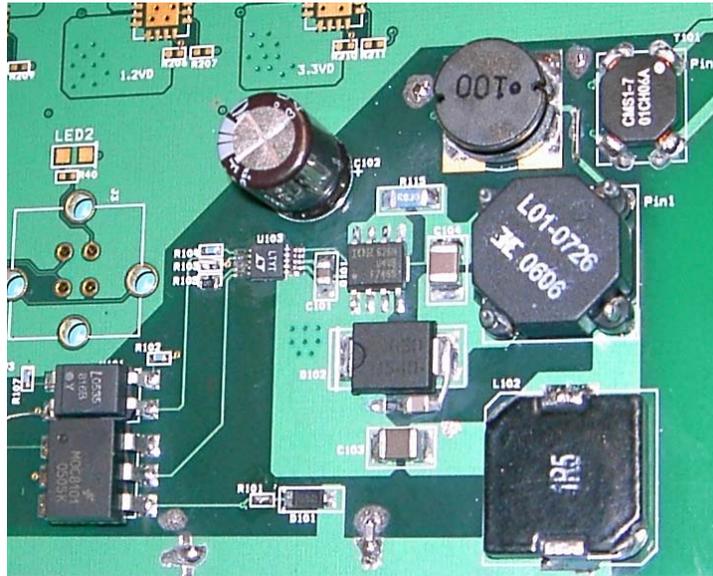


Fig. 14.17: Photograph of power circuit for TEC regulation

14.9.2 Signal Extraction

The front end electronics has the responsibility of amplifying and integrating the signals from the APD arrays, determining the amplitude of the signals and their arrival time and presenting that information to the data acquisition system (DAQ). The front-end electronics will operate in continuous digitization mode and does not require any external trigger or NUMI timing gate. Data will be time stamped and compared to a NUMI timing signal in the DAQ system to determine if the event was in or out of spill.

Data from the ADC is sent to an FPGA where multiple correlated sampling is used to remove low frequency noise. This type of Digital Signal Processing (DSP) also reduces the noise level and increases the time resolution.

The interface to the DAQ can be configured to send up to 16Mbits/s, or up to 6.2kHz/channel. The typical rate expected is due to cosmic ray muons, and will produce an average hit rate of 120Hz/cell, 2% of the link capacity. As many as 64 FEBs feed a single data concentrator, with a 1Gbit uplink speed. This limits each link to an upper limit of 15.6Mbits/s, effectively matching the maximum FEB uplink speed, with only 2% of the capacity used on average.

To prevent unnecessarily loading the DAQ system upstream links and buffer farm it is important to keep the data rate from noise to an acceptable level. The noise level of the front end electronics and the imposed threshold determines the noise data rates. Figure 14.18 shows the simulated noise rate per 32 channel APD box as a function of threshold for 2.3 effective photoelectrons of noise per channel, 10 bytes of information per hit above threshold and 10^6 time slices/second. The noise includes the APD excess noise in amplifying the bulk current. In order to keep the noise data rate on the order of the rate from cosmic ray muons (with a modest

overburden) it is necessary to impose a 12-15 photoelectron threshold. In the near detector this threshold can be increased to reduce the noise component even further since the signals are expected to be approximately 4 times larger due to the short length of the modules.

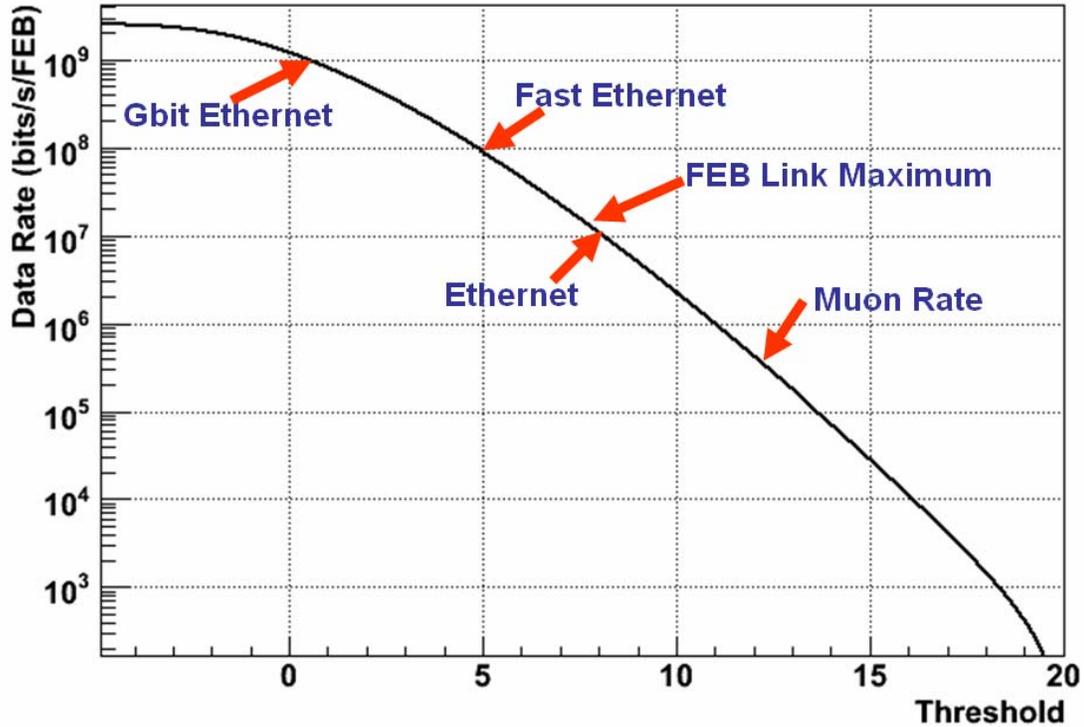


Fig. 14.18: The simulated data rate at the NOvA far detector, in bits per second, as a function of the readout threshold. Noise is due to the amplifier noise, assumed to be 1.5ENC, bulk dark current of 1.5e/ μ s, amplified by the APD with its characteristic excess noise factor of 2.5. The data rate due to cosmic ray muons, with a modest overburden, is also indicated as well as the networking technologies necessary to accommodate various rates.

14.9.3 Operating modes

The front end electronics must operate in several different modes that can be selected through the DAQ interface. These are:

- Run Mode – The Front end electronics will continuously acquire and transmit data to the DAQ system.
- Calibration mode – Data is accumulated for calibration and for determining thresholds and pedestals, noise measurements, etc.
- Test Mode – The Front end electronics must be able to simulate data and send specific test patterns to the DAQ system to check for proper operation.
- Programming mode – The front end electronics must be able to download and upload data that define its operating parameters. The front end electronics must have on-board firmware that can be reprogrammed in place via the DAQ system.

14.9.4 Digital Signal Processing

The FPGA on the FEB uses a digital signal processing algorithm to extract the time and amplitude of signals from the APD. Any signal above a pre-programmed threshold will be time-stamped and sent to the DAQ for processing. The threshold is settable at the channel level to allow different thresholds to be set depending on the particular characteristics of a given channel.

The DSP algorithm uses many samples to extract the time and amplitude of the signals. This is possible since there is very little activity in any given channel. Typically 20 samples are used in the current algorithm. The extra samples help to average out the current noise component of the signal to obtain better timing resolution than would normally be expected with 500ns digitization. It also has the effect of reducing the effective noise component by approximately 20%, giving a 20% increase in the signal to noise ratio.

A sample trace is shown for an ideal signal and one with additional noise in Fig. 14.15. A matched filtering technique is used to transform the signal to a symmetric output pulse shown in Fig. 14. The discrete points in the filter output can then be interpolated to obtain the peak with approximately four times greater resolution. The FEB will use this interpolated filter output shown in Fig. 14. to create a timestamp with a 62.5ns bin size, eight times smaller than the sampling time of 500ns.

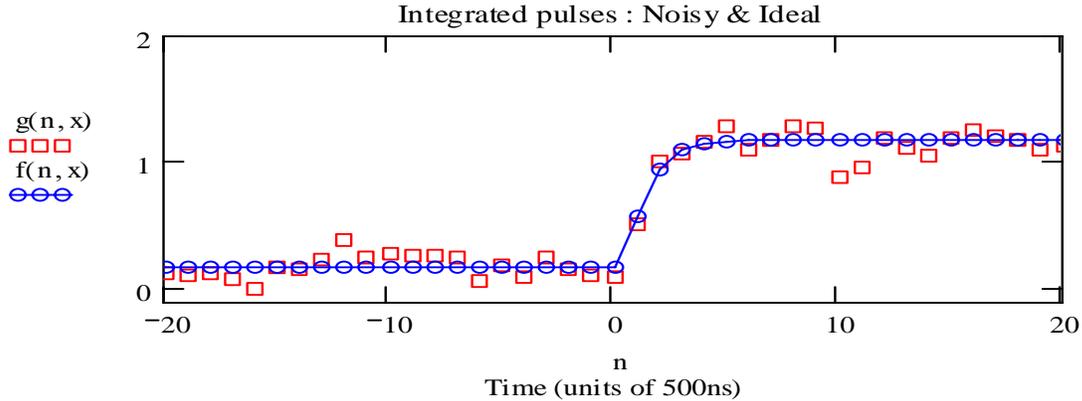


Fig. 14.15: Ideal (blue) and realistic (red) signals digitized at 500ns/sample.

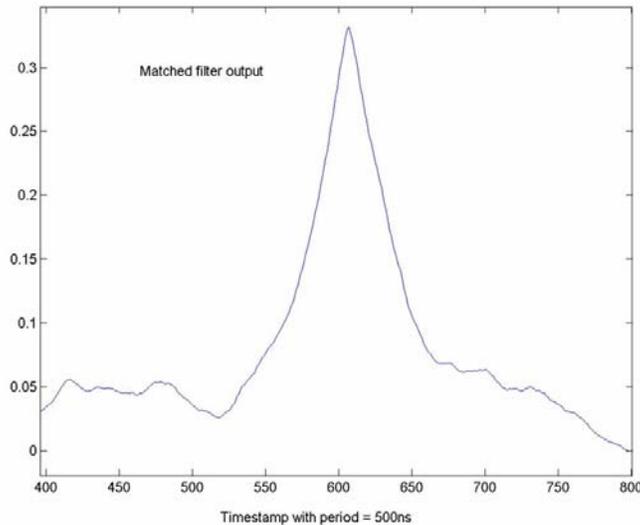


Fig. 14.20: Matched filter output used to extract time and amplitude

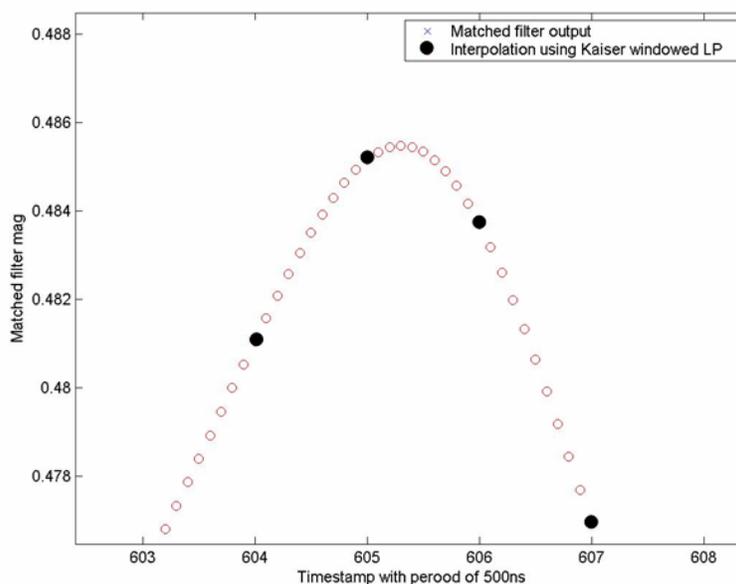


Fig. 14.21: Interpolation of filter output (filled points) showing greater precision with which the peak can be determined

14.10 Power Distribution

The NOvA power distribution system (PDS) supplies power to five different electronic components: front end boards (FEBs), avalanche photodiodes (APDs), thermoelectric coolers (TECs), data concentrator modules (DCMs), and timing distribution units (TDUs). The power distribution system consists of: 1) the power supplies and their racks, 2) the power distribution boxes that fan out the power to the FEBs, APDs, TECs, DCMs, and TDUs, and 3) the power cables and their cable trays. Two similar systems are needed: one for the near detector and one for the far detector. The large number of channels and the extent of the detectors demand that the systems must be remotely controllable. The systems must also be robust and safe. The voltage, current, and power requirements for the four components are given in Table 14.3.

Module	Voltage (V)	Current (A)	Power Distribution Box		
			Channels	Total Current	Power (W)
FEB	3.3	1.0	60-64	60-64 A	211
TEC	24.	0.30	60-64	18-19 A	456
APD	350-450	0.00004	60-64	2.4.-2.6 mA	2
DCM	24.	1.25	1	1.25 A	30
TDM	24.	1.0	1/24	1.0 A	24

Table 4.3 Power requirements for the front-end electronics.

The layout of the power distribution system is shown schematically in Figure 14.22. All power to the front-end electronics is distributed by power distribution boxes (PDBs) which are situated on the detector. A more detailed diagram of the custom power distribution box is shown in Figure 14.23. A single 6-conductor, 18AWG cable from each PDB output carries the 450V

needed by the APDs, the 24V needed by the TE coolers, as well as the 3.3V needed by the FEBs, as well as their return currents. All the electronic components and the power supplies float: ground reference is at the PDBs. A maximum of 64 FEBs and 1 DCM are served by each PDB.

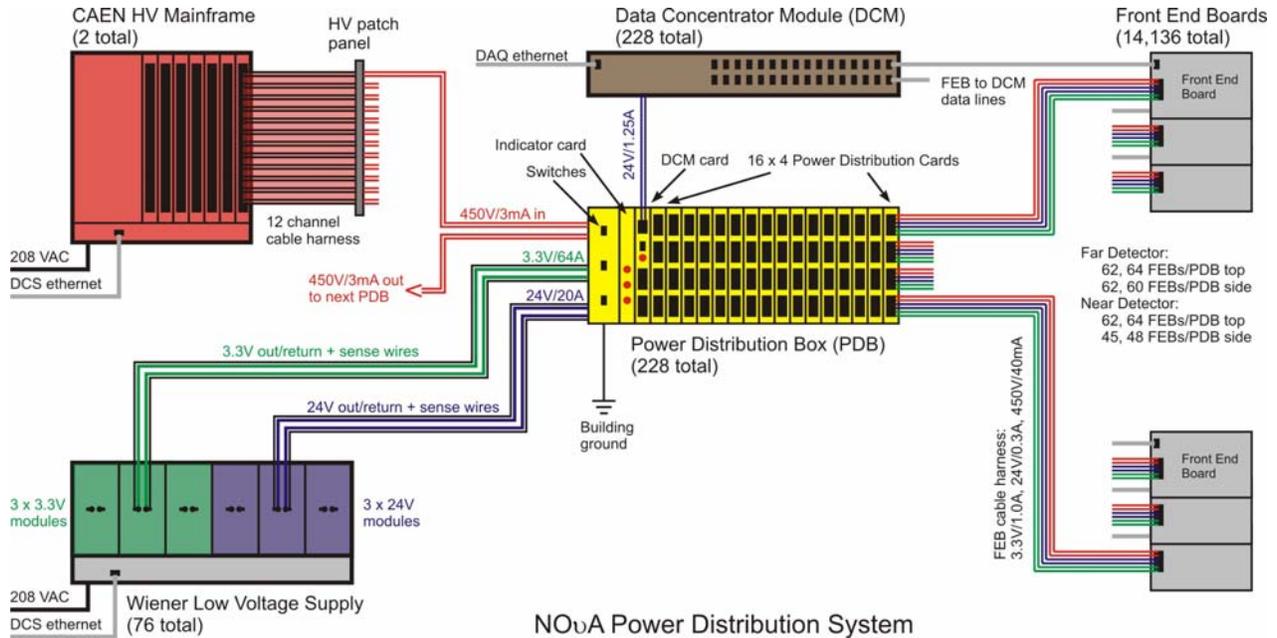


Fig. 14.16: Schematic of the NOvA power distribution system. Each power distribution box feeds 3.3V, 24V, and 450V via a single cable to a maximum of 64 front-end boards, and a single 24V line to the nearby data concentrator module. The APD voltage to the power distribution boxes is provided by CAEN A1520P cards situated in a CAEN SY1527 mainframe. The FEB, TE cooler, and DCM low voltages are provided by Wiener PL508 supplies.

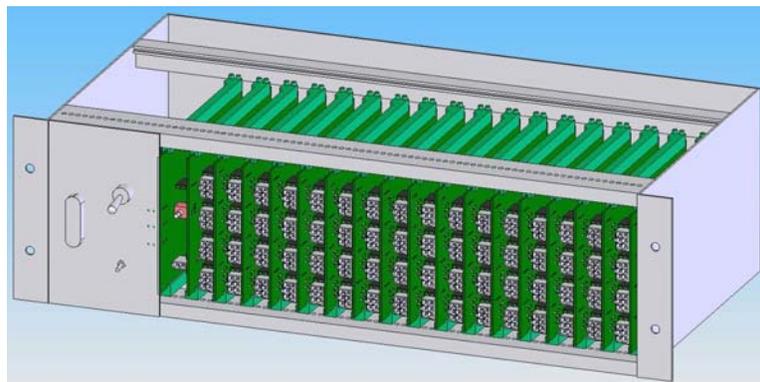


Fig. 14.23: Isometric view of a power distribution box.

The 450V to the PDBs is provided by the CAEN SY1527 Universal Multichannel Power Supply System using A1520P floating 12-channel cards. Each of the 12 channels of the A1520P is individually programmable from 0-500V, and provides a maximum of 15mA of current. A cable harness takes the 12 HV channels from each A1520P card to a nearby patch panel, and RG59 cables with SHV connectors go from the patch panel to the PDBs. Each CAEN HV channel feeds two PDBs, whose high voltage inputs are daisy-chained together.

The low voltages (3.3V and 24V) needed by the FEBs, DCMs, and TECs, are provided by the Wiener PL508 Power Supply System. Each Wiener PL508 chassis has six floating individually programmable power supply modules: three 2-7V channels, each rated to 115A and 550W, and three 12-30V channels, each rated to 23A and 550W. Hence each PL508 crate feeds power to three adjacent PDBs: the 3.3V via 2 AWG cables and the 24V via 10 AWG cables. The cables can carry the maximum rated current of the Wiener supplies without the need for fuses. The Wiener low voltage supplies have a remote sensing feature which allows the voltages *at the PDBs* to be set to their desired values. This feature is needed because the different cable lengths from the Wieners to the PDBs produce non-negligible differences in the voltage drops. The largest drops are 0.7V and 0.9 V, respectively for the 3.3V and 24V lines. Voltage drops on the 18 AWG cables running from the PDBs to the FEBs are much smaller resulting in a voltage range at the FEBs between 3.17-3.42V, which is well within the allowed range of the FEB regulators. The voltage regulators on the FEBs protect them from inadvertent application of the highest voltage the Wiener supplies can provide.

Both the Wiener and the CAEN power supplies are remotely controllable with ethernet interfaces. They have programmable trip levels and their noise and ripple specifications meet the requirements of the FEBs.

A total of 228 far detector and 10 near detector power distribution boxes need to be fabricated. To fit within the tight space constraints of both detectors the power distribution boxes employ a standard DIN 3U subrack with a custom backplane that feeds 3.3V, 24V, and 450V to individual cards. Each card provides power to 4 FEBs, allowing a maximum of 64 FEBs to be powered through one PDB. Each 3.3V and 24V output is individually fused and has LED indicators. Crate level switches power the 3.3V, 24V, and 450V lines. Transient voltage suppressors on the 3.3V and 24V lines protect the FEBs from overvoltage due to an out of regulation power supply. A special card feeds 24V power to the associated DCM. This power can be switched on and off irrespective of the status of the 24V power to the FEBs.

The fundamental unit for the power distribution system for the far detector is the 62-plane di-block. The layout of the far detector super-blocks, di-blocks, blocks and planes, is shown in Figure 14.24. There are 19 di-blocks; all but 4 having 31 horizontal and 31 vertical planes, the exceptions having 32 vertical planes and 30 horizontal planes. This results in normal di-blocks having PDBs that service 62 PDBs, and the 4 abnormal di-blocks having top PDBs feeding 64 FEBs and side PDBs feeding 60 FEBs. The layout of a far detector di-block is shown in Figure 14.25. In a normal di-block there are $12 \times 31 = 372$ vertical modules, requiring 6 PDBs, each cabled with 62 outputs, and the same number of horizontal modules, requiring 6 PDBs (3 on each side of the detector), each cabled with 62 outputs. (The side PDBs feed FEBs that are spread out over a larger region than the top PDBs.) For each PDB there is one DCM. The Wiener power supplies are mounted on racks on the upper catwalks: one for every three PDBs for a total of 4 for each di-block. Two CAEN mainframes are needed for the entire detector: one on each side at the detector midpoint. Since each CAEN output channel supplies two PDBs, 6 HV channels are needed to supply the 12 PDBs in a di-block. For the 19 di-blocks of the complete far detector, 228 PDBs are needed, 2 CAEN mainframes with a total of 10 A1520P boards, and 76 Wiener PL508 crates (Table 14.4).

The top PDBs are set on the flat part of the module manifolds, as can be seen in Figure 14.26, which shows a front view of the corner of a far detector bi-block. The DCMs are adjacent to them. The PDBs serving the horizontal modules will be mounted sideways using a commercial framing system such as Unistrut.

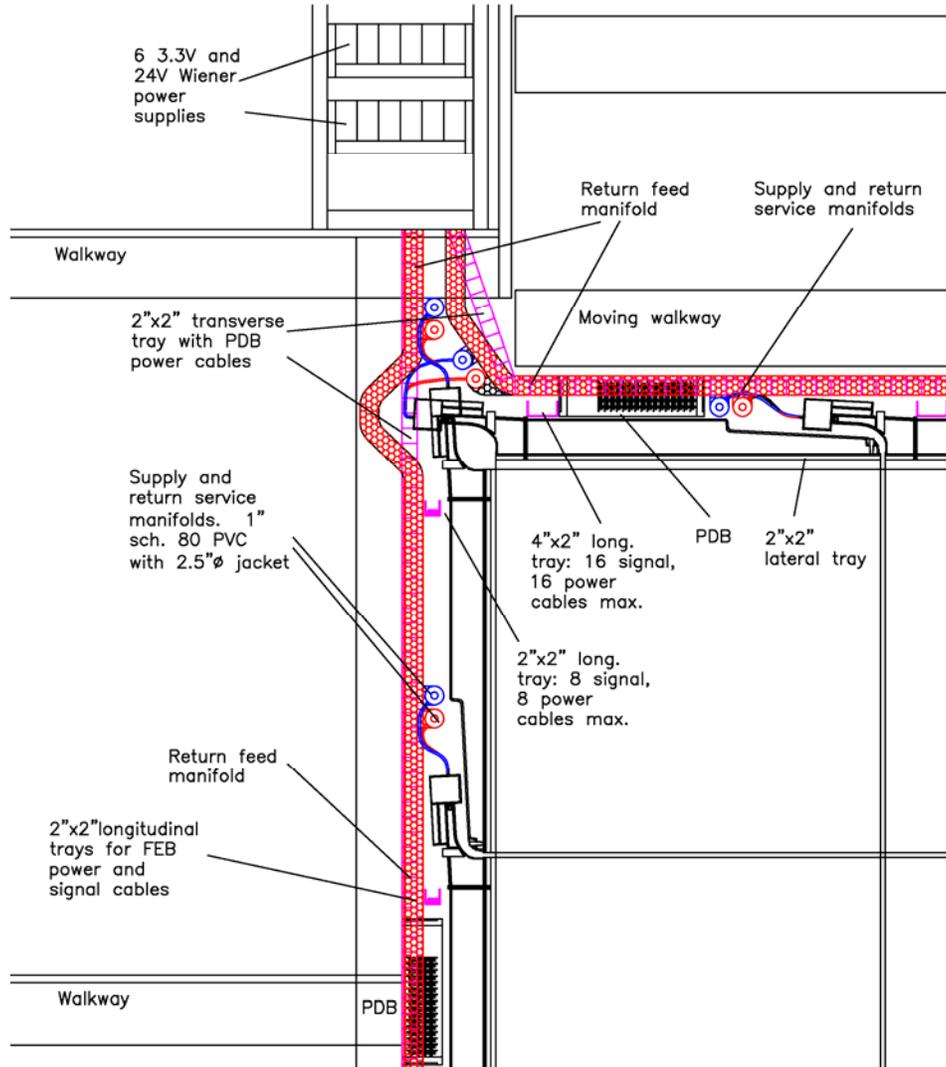


Fig. 14.26: Front view of the top corner of a far detector di-block. The PDBs fit between the module manifolds and the moving walkway on the top, and between the manifolds and the catwalks on the side. Cable trays and cooling water manifolds are shown. Not shown are the DCMs, which are behind the PDBs.

The layout for the power distribution system for the near detector is shown in Figure 14.27. The FEBs are on the top and on one side of the detector. Each side PDB feeds power to 45 (48) block type A (B) horizontal modules and each top PDB feeds 64 (62) block type A (B) vertical modules. An additional PDB feeds the 33 muon catcher FEBs. A total of 10 PDBs, 4 Wiener supplies, and one CAEN mainframe with one A1520P card are needed for the entire near detector. Unlike the far detector, where each HV channel serves two PDBs, each PDB crate in the near detector has its own HV channel.

Layout of Power Distribution System NOvA Near Detector

- 6 normal blocks
- 1 muon ranger block
- 10 Data Concentrator Modules
- 10 Power Distribution Boxes

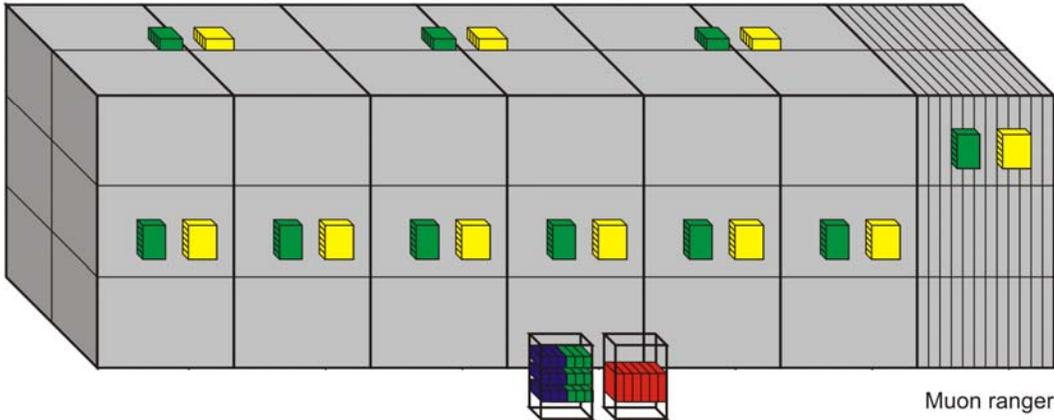
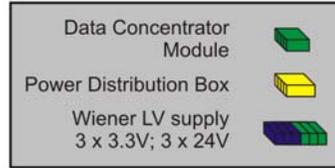


Fig. 14.27: Layout of the power distribution system for the entire NOvA near detector.

The power cables will be placed in wire basket cable trays. The layout of the cable trays for a top far detector di-block is shown in Figure 14.28. The layout of the side cable trays is similar. Cable tray and cable lengths are given in Table 14.4. The cable trays, cooling manifolds, PDBs, and DCMs have been positioned to allow access to the FEBs. A total of 14,136 6-conductor, 18 AWG, cables are needed to carry the power from the PDBs to the FEBs for the far detector.

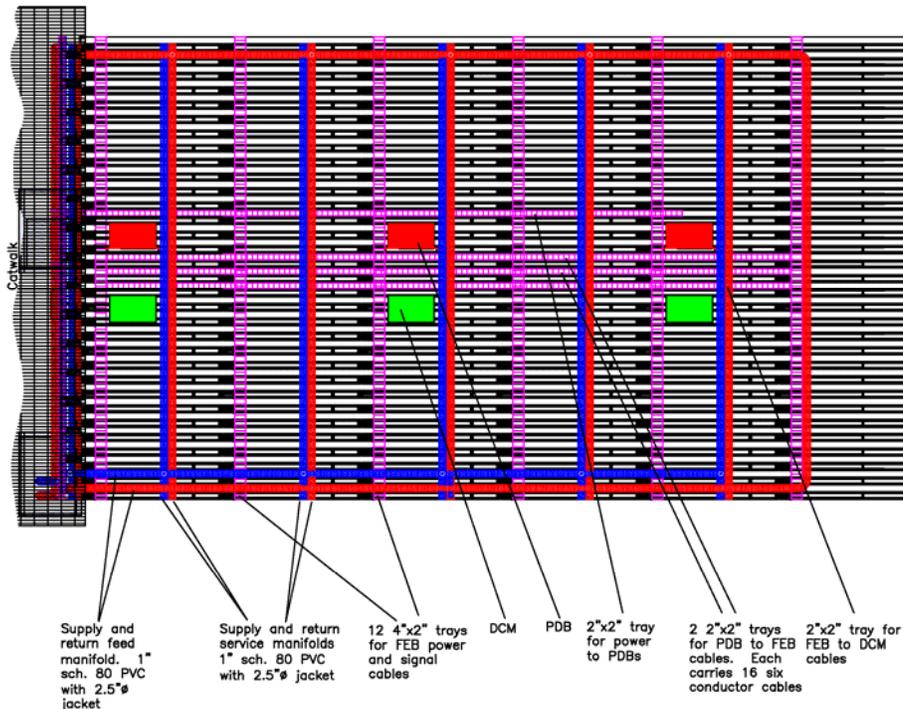


Fig. 14.25: Top view of half a far detector di-block showing the location of the cable trays and FEB cooling manifolds.

	FD	ND
Front End Boards	14,136	497
Power Distribution Boxes	228	10
CAEN A1520 boards	10	1
CAEN SY1527 mainframes	2	1
Wiener PL508 crates	76	4
3.3V, 2-conductor, 2 AWG cable: Wiener to PDB	1,513 m	66 m
24V, 2-conductor, 10 AWG cable: Wiener to PDB	1,513 m	66 m
450V, RG58 cable: CAEN patch panel to PDB	1,513 m	66 m
2-conductor, 22 AWG sense cable: PDB to Wiener	3,026 m	133 m
6-conductor, 18 AWG cable: PDB to FEB	35,634 m	1,563 m
2"x2" cable tray	3,544m	155 m
4"x2" cable tray	985 m	43 m
6"x2" cable tray	540 m	24 m

Table 14.4 Power distribution system components.

The PDBs will be fabricated and tested at the University of Virginia. Cables will be cut and terminated at Virginia. The CAEN and Wiener power supplies will be tested at the University of Virginia before being shipped out to the experiment. The University of Virginia has sufficient space to store the entire system.

14.11 Changes in the Photodetector and Electronics Design Since the CDR

Since writing the Conceptual Design Report (CDR) we have received 20 prototypes of APDs mounted on carrier boards. While this development took longer than expected to converge, test results show that the devices have performed well.

Recent development in modeling the near detector showed that there would be substantial overlapping of events in space and time since the beam spill is one single turn of the Main Injector. In order to resolve these overlaps faster sampling, and possibly different algorithms for timing extraction are needed in the near detector. The simple fix for this was to increase the number of outputs from the ASIC to reduce the multiplexing. This does not change the input stage, or the available settling time of individual outputs, so the noise performance should not change.

The water cooling scheme for removing heat from the TEC modules has evolved from a large chiller distributing process water to the entire detector to individual chillers, 1 per 31 plane block. This eliminates the need to open the water system once operation has begun, and limits the impact of any failures.