

Transmitter Protection System Manual

This is an abbreviated manual for the TPS (Transmitter Protection System) for the Alcator C-Mod Lower Hybrid transmitter. This system was developed under DOE SBIR funding. For a complete background description and system overview, please refer to the SBIR final report (appended). For complete details of the COTS model calculations, please see that separate document (appended).

Although this system was designed specifically to meet the requirements of the Alcator C-Mod LH system, it was implemented in a fashion to enable reconfiguration for a variety of transmitter protection architectures. This manual distinguishes the general features as well as the specific implementations for Alcator C-MOD LH.

It is important to keep in mind that the Alcator C-Mod LH system presently includes an array of 12 klystrons, soon to be expanded to 16. These are physically mounted on “carts” of four klystrons each, and the carts share many resources. The TPS follows this architecture, and each TPS unit services four klystrons.

- **analog outputs** – Each of the 24 analog input signals are scaled and buffered in the input section, then piped separately (i.e. with independent buffer amplifiers) to both single-pin LEMO monitor testpoints as well as to digitizer outputs on a single 32-channel 68scsi connector. In addition, ten other analog signals generated internally are similarly presented to the dual outputs. Since only 32 channels are available on the LEMO testpoints and the digitizer connector, jumper selects choose alternate outputs on two of the 32 output channels.
 - ❖ *As configured for Alcator C-Mod, two of the spare analog channels are not even stuffed for use, and the jumpers are set to ignore these. The alternate signals -- two of the four collector temperature channels as read by the ASQUIRT muxed links are used for these outputs.*

- **fault comparator levels** – The 16 standard output channels each have three fault comparators. One of these three is hardwired to a negative reference voltage, with the intention of catching faults if current transformer signals are wired with the wrong polarity. The second of the three is a conventional fault comparator, with a 0-10V positive reference adjustable via an onboard trimpot... intended to be set as a backstop fault which should never be exceeded. The third comparator – intended for general fault catching use – is wired to an independent DAC reference, where the DAC output level can be set over the Ethernet interface. All of the fault comparator outputs are wired to CPLD (complex programmable logic device) chips which can be reconfigured in firmware as needed.
 - ❖ *As configured for Alcator C-Mod, each set of three comparators are wired to a single latch bit in the CPLDs. These latches are SET by any of the fault comparators tripping, and cannot be cleared until a RESET signal is received AND the fault condition is clear during the RESET.*

- **digital signal inputs** – There are a large number of fast fault signals which are received on dedicated F/O. All of these F/O receivers are on the RTM (rear transition module) assembly of 3 boards, which connect through the cPCI backplane to the main board with the fault CPLDs. These signals wire directly to the CPLDs for maximum speed of response.
 - ❖ *As configured for Alcator C-Mod, each signal feeds a dedicated latch, which similarly cannot reset until the RESET signal is received while the fault condition no longer persists. In addition, the F/O dedicated fault lines are defined as light=good, dark=fault. Four of these are used for "elbow arc detectors", four for "circulator arc detectors", one for a fault output from the CPS (coupler protection system).*

- **control signals** – In addition to the above fault signals, there are a number of other signals – both input and output – which wire directly to the CPLDs for TBD logic.
 - ❖ *As configured for Alcator C-Mod, these control signals embed a sequencer which controls the application of high voltage and RF input to the LH klystrons. The particular logic of this system is arbitrary -- i.e. it is defined for the C-Mod system, but is entirely defined in firmware coding of the fast CPLDs. This control logic incorporates both fast (dedicated fiber) and slow (PLC via multiplexed DSQUIRT) control signals, thus the TPS serves as an interface for the PLC to define the state of the fast fault response engine.*

- **muxed F/O I/O** – The TPS system is heavily dependent on distributed I/O. This I/O is remotely managed in digital and analog off-board circuits which multiplex a number of signals into an information packet, and transmit it over a single fiber at moderate repetition rates. A receiving mate takes the information on this fiber and re-creates the same analog or digital signals on the far end, thus serving as a non-galvanic link with copper at both ends. There are several varieties of these multiplexed I/O units¹:
 - **ASQUIRT01** – a two channel full-duplex 10 kHz analog link, with each unit transmitting and receiving two signals, with analog range of 0 to +10V and 12-bit precision².
 - **ASQUIRT02** – a 16 channel full-duplex 1 kHz analog link, with each unit transmitting and receiving 16 signals, with analog range of 0 to +10V and 12-bit precision.
 - **DSQUIRT02IO** – a 16 channel full-duplex 5 kHz digital link, with each unit transmitting and receiving 16 signals³.
 - **DSQUIRT02II** – a 32 channel half-duplex 5 kHz digital link, with each unit converting 32 channels copper to fiber.
 - **DSQUIRT02OO** – a 32 channel half-duplex 5 kHz digital link, with each unit converting 32 channels fiber to copper.
- ❖ *As configured for Alcatel C-Mod, one DSQUIRT02IO and one DSQUIRTII are thus virtually incorporated for communications link to the PLC -- thus we have slow control digital signals with 16 inputs (PLC to TPS controls) and 48 outputs (TPS to PLC status reports).*

- **Ethernet comm** – an Ethernet port is provided on the TPS front panel. This port is wired through a 3rd party interface which converts data full duplex data streams to a UART, which pipes to one of the on-board CPUs. In general, this is completely software configurable.
 - ❖ *As configured for Alcatel C-Mod, this port talks to an MDSplus script which sets fault thresholds and other configuration parameters before each shot, and downloads the results of the data logger post-shot.*

- **power systems** – The TPS uses the +5V cPCI backplane power only. There are two separate 12W +5V to +15V switching boost regulator supplies, and two separate 6W +5V to -15V supplies on the TPS which provide local analog power rails, in addition to some intermediate reference voltages.

¹ The TPS wires three F/O connectors (two TX and one RX) from the RTM panel to a CPU instead of a CPLD, expressly for the purpose of creating “virtual DSQUIRTs” inside of this CPU. These signals do not appear on copper within the TPS, but only appear virtually for use in control algorithms. Similarly, two F/O RX are on the front panel of board 2, acting as receivers for one ASQUIRT01 and one ASQUIRT02, with the virtual ASQUIRT receivers in a different onboard CPU.

² The ASQUIRT options include both 0-20mA and 0-10V input scaling options.

³ The DSQUIRT signal levels are compatible with Allen Bradley PLC I/O, opto-isolated 10-30V, active-hi sourcing outputs and active-hi sinking inputs.

- **inhibit timing** – Four dedicated signals are piped around within the TPS CPUs and CPLDs, intended to be used for inhibit timers. These are driven by one of the cpus, and are wired to pins on that cpu which can directly drive via hardware clocks without incurring the latency of software response.
 - ❖ *As wired for Alcator C-Mod, two of the CPLD control lines – those which effectively represent "HV ON" and "RF ON" are brought to this same CPU, and wired to pins which can be set for generating interrupts on pin transitions. We use these to start and stop the timer output pins to generate programmable delays on rising or falling edges (or both) of either HV ON or RF ON. These are used to mask the occurrence of any of the fault comparators that we wish. For C-Mod, we mask (for example) beam current CTs for a short time (a few microsec) after the rising edge of HV_ON. We only mask the MDS settable fault comparator, not the hardwired (fixed reference) comparator. We thus can set the fault trip point closer to the operating point, and be immune to short over-current surges that occur while charging the capacitance in the cable and klystron cathode feedthrough. Similarly, we mask body currents during rising and/or falling edges of HV_ON or RF_ON (i.e. "rabbit ears"). We always mask only the MDS settable comparator, never the fixed comparator, thus we never leave our fault protection "naked" even for short intervals.*

- **COTS logic** – The distribution of analog signals between two circuit boards was intentionally partitioned to place all beam current channels and all RF power channels on the same board. This board has additional hardware: dedicated ADCs and DACs talking to a CPU which is not used for any other purpose.
 - ❖ *Although this is a very flexible additional kit, it is built for a specific purpose: modeling the klystron collector temperatures for the COTS (collector over-temperature system). The need for the COTS is due to the inability of the klystron to run full beam power without any RF output for full-length pulses....**the klystron cooling design assumes that the beam is spent, and the max pulse duration must be derated if RF output is less.** We cannot safely provide for collector temperature fault protection by simply measuring the temperature -- the surfaces are too deeply embedded in the klystron, and the downstream water temperature is too delayed for safety. Instead, we do a realtime prediction of the collector temperature based on calculating the rate of temperature change: $(P_{in}-P_{out})/C$ (power in minus power out over heat capacity). P_{in} is quite simple: $I_{beam} \times V_{beam}$. P_{out} is RF power plus power removed by water cooling. We obtain the beam voltage via the ASQUIRT01 input (already digitized), and we locally digitize each of the four beam current channels onboard (in parallel with external MDSplus data system digitizing). The water cooling information (flow rate, inlet temperature, pressure) are brought in via the ASQUIRT02 inputs, and we similarly digitize the RF forward power signal from the same diode signals used for reflected power fault ratios. The COTS calculation is performed in realtime, and is not limited to "shot" times: The four COTS channels are continuously updated at 1 kHz intervals. The model COTS temperature is generated on local DACs, and is presented as additional output channels. Similarly, the ASQUIRT02 dataset includes downstream collector water readings, which are copied to DAC channels for output. Fault comparators for the four COTS signals are generated by software, then fed to the CPLD control logic like any other comparator signals.*

More Details

Now that we have surveyed the key subsystems within the TPS, we will revisit them one at a time and provide greater depth, specifically as implemented for the Alcator C-Mod LH system.

- **ANALOG INPUTS & FAULT COMPARATORS** – The 24 analog inputs are split between board “FRONT1” and board “FRONT2”.
 - ❖ Four channels on FRONT1 service the slow body currents (one for each klystron). These are signals produced by Bergoz CT-100mA current transformers, which give 10 mA/V into high impedance, oriented for positive signals. The gains are configured for x1, thus the internal analog signal levels, as well as the corresponding analog outputs, have a calibration of 10 mA/V, or 100mA full scale. The trip points are nominally set to 80mA, with 90mA on the unmasked fixed comparators. The filter capacitors at the inputs give about 100 μ s filtering. These signals are wired to 2-pin LEMO connectors J1-J4. The programmable comparators are masked by HV_ON leading and trailing edges programmed by delay channel #3.
 - ❖ Four channels on FRONT1 service the fast body currents (one for each klystron). These are signals produced by Stanganese CT model 2-1.0, which gives 1 A/V into high impedance, oriented for positive signals. The gains are configured for x50, thus the internal analog signal levels, as well as the corresponding analog outputs, have a calibration of 20 mA/V, or 200 mA full scale. The trip points are nominally set to 80mA, with 180 mA on the unmasked fixed comparators. The filter capacitors at the inputs give about 1 μ s filtering. These signals are wired to 2-pin LEMO connectors J5-J8. The programmable comparators are masked by HV_ON leading and trailing edges programmed by delay channel #2, and by RF_ON leading and trailing edges programmed by delay channel #4.
 - ❖ All six body current comparators for each of the four klystrons are programmed to operate a single fault latch in the CPLD logic matrix firmware. This latch thus represents a “body current fault for tube A”, for example, with no further information on which of the six comparators caused it.
 - ❖ A ninth channel on FRONT1 services the total ground current measurement for the cart, provided by a Pearson 101 Current Transformer, with a scaling of 100 A/V into high impedance, oriented for a positive signal. The gain is set to 5x, giving a 20 A/V scaling, or 200A full scale. The comparator trip point is nominally set to 60A, while the unmasked fixed comparator is set to 100A. The filter capacitors at the input gives about 1 μ s filtering. This signal is wired to 2-pin LEMO connector J9. The programmable comparator is masked by HV_ON leading edge only, programmed by delay channel #1. All three comparators feed a single fault latch in the CPLD.
 - ❖ Connectors J10-J12 on board FRONT1 service spare analog channels, with identical topology to the other CT signals. Only one of these, J10 (Spare1) is fully populated on the board, and thus can be implemented with only firmware revisions. The comparator outputs are wired to the CPLDs, but are ignored in the present revision of firmware.
 - ❖ Four channels on FRONT2 service main beam currents (one for each klystron). These are signals produced by Bell RS-100 current transformers, which gives 40 mV/A into high impedance, oriented for positive signals. The gains are configured for x5, thus the internal analog signal levels, as well as the corresponding analog outputs, have a calibration of 5 A/V, or 50 A full scale. The trip points are nominally set to 12A, with 22 A on the unmasked fixed comparators. The filter capacitors at the inputs give about 1 μ s filtering. These signals are

wired to 2-pin LEMO connectors J1-J4. The programmable comparators are masked by HV_ON leading edges only, programmed by delay channel #1.

- ❖ In addition to being wired to the two analog outputs (test points and digitizer connector) with an internal scaling of 5 A/V, the beam currents are scaled and presented to an onboard 12-bit digitizer with a 0-4095 count range of 0-25A. These inputs are sampled at 1 kHz intervals each for use in the COTS calculations.
 - ❖ Eight channels on FRONT2 service the RF diode inputs. These are arranged in pairs for forward and reflected power, one pair for each klystron. These are generated from -50dB MDSL loop directional couplers, oriented for negative signal, typically 0 to -2V. A gain of -2.5x is applied, giving a typical 0 to +5V range on the internal signal path for comparators and outputs. The forward power signals (only) are also scaled to present a range of 0 to -3V *on the raw signal* to the 12 bit digitizers, for 1 kHz sampling to the COTS calculation. These forward/reflected signal pairs are presented to 2-pin LEMO pairs J5/J6, J7/J8, J9/J10, and J11/J12.
 - ❖ There is only one comparator supplied for each pair of diode inputs. This comparator output is high if the adjusted forward diode signal is greater than the reflected diode signal. The forward signal is adjusted with an AD633 analog multiplier, which applies a gain between 0 and 1, and an offset of 0.5V. The gain is remotely programmable by a DAC output, programmed by an entry in the MDSplus tables. We thus generate a fault signal if the reflected power exceeds a programmed fraction of the forward power, and the offset prevents nuisance trips at low or zero power. Each comparator feeds a single latch within the CPLD fast fault matrix, which are presently programmed to generate SOFT faults, which only shutoff RF, but not shutdown HV.
- **ANALOG OUTPUTS** – Each of the 24 analog input signals are scaled and buffered in the input sections as described above, then routed with independent buffer amplifiers to both single-pin LEMO testpoints, and to the 32-channel digitizer connector (SCSI68). We fill out the additional 32 output channels (LEMO testpoints and digitizer channels) as follows:
 - ❖ Four channels monitor the internal model for the klystron collector temperature (COTS), with a scaling of 0-10V mapping to 0-250C,
 - ❖ Four channels monitor the measured klystron collector temperature (downstream, delayed), with the same scaling (this is an echo of four of the ASQUIRT02 channels),
 - ❖ One channel monitors the beam voltage, as reported by the ASQUIRT01 receiver, with a scaling of 0-10V mapping to 0-60kV,
 - ❖ One channel monitors the total system current, as reported by the ASQUIRT01 receiver, with a scaling of 0-10V mapping to 0-500A.
 - ❖ The astute reader will realize that we have accounted for 34 of the 32 channels. The FRONT3 board is fed all 34 of these signals, but has only 32 LEMO1 testpoints, and 32 channels to the digitizers. There are two jumpers which select which channels we ignore. One jumper selects between the Spare analog input #3 and the measured klystron temperature for channel C, while the other jumper selects between the Spare analog input #2 and the measured klystron temperature for channel D.

- **CIRCULATOR ARC DETECTORS** – Four of the F/O inputs presented at the RTM are for waveguide arc detectors mounted in the circulators. These signals are defined as light=ok, dark=fault, and are generated by an interface circuit with an auto-clear function. Each signal is routed to a latch in the CPLDs, which contribute to the SOFT FAULT condition (which cuts RF but not HV).

- **CPS (COUPLER PROTECTION SYSTEM)** – One of the F/O inputs presented at the RTM is for faults generated by the independent (and similarly complex) coupler protection system. This system also has an internal reset logic. This fault is latched and contributes to a SOFT FAULT condition, but the *fault condition which triggers the latch* is masked (inhibited) unless both HV and RF are on.

- **ELBOW ARC DETECTORS** – Four of the F/O inputs presented at the RTM are for additional waveguide arc detectors situated at the waveguide elbows. These are simple 3rd party devices which fire an SCR on a detected optical signal. The fault signal from the detector is itself latched, thus resets become problematic, as we must separately reset the device *and* the TPS internal CPLD latch. This is further complicated by the reset mechanism for the SCR – we pull the SCR output down, thus starving the SCR and allowing it to turn off – but the reset action itself results in the appearance of a fault condition. We must thus relax our usual logic of not clearing the internal TPS latches while the fault condition exists (dangerous), or provide an entirely separate signal path to clear the SCRs before we attempt to clear the internal latches. We opted for the latter, but this was a late addition to the TPS system, thus is not implemented cleanly. We presently generate a single 4ms pulse within the TPS on the rising edge of the combined OR of any of the Elbow Arc faults, which feeds to the PLC rack over a spare DSQUIRT channel. The PLC has the additional authority to reset any of the Elbow Arc SCRs during testing – so we combine this new signal with each of the four PLC signals (via diode voting).

MISC TOPICS

Additional features of the implementation configured for Alcator C-Mod LH must appear under the miscellaneous category. Generally, these are software features, and can be reconfigured quite flexibly, if used at all, for any other applications of the TPS hardware.

SEQUENCING – HV ENABLE

Built into the fast (CPLD) programmable logic matrix, we have a sequencer for enabling the HV system.

Unfolded into words, the sequence is:

- 1) The PLC requests a HV ENABLE
- 2) The TPS passes this request to Box #1 (one/cart) if there are no HARD faults, and if FAULT_BYPASS is low.
- 3) Box #1 checks door interlocks, activates Ross relays, and when relay switches are ok, responds with a signal.
- 4) The TPS passes it to Box #7A (there is a HEARTBEAT hardwired interlock on this signal)
- 5) Box #7A (one/system) waits until all carts have passed this signal, then echos it back to all carts & enables HV.
- 6) The TPS now has HV ENABLED (as should all TPS in the system). The info is passed to the PLC.

Note that Box #7A serves to synchronize the separate TPS units on different carts – if any of them have a hard fault or otherwise remove HV Enable permission, it is shut down for all, but without cycling the ross relays on the carts that didn't have any faults.

SEQUENCING – HV ON

HV Enable merely represents HV being dangerous, not on. To turn on HV, the sequence continues:

- 1) The ACS requests HV be turned on (this is called Trigger Gate)
- 2) If TPS is ARMED, there are no CTRL OK faults, RESET is not happening, HV is enabled, and ACS requests HV, the TPS passes this request on to Box #7A (there is a HEARTBEAT hardwired interlock on this signal)
- 3) Box #7A waits until all carts have passed this signal, then it turns on the HV supply.

Note that Box #7A again serves to synchronize separate TPS carts. HV ON is a system function.

SEQUENCING – RF ON

The sequence to turn on RF is similar. It goes like so:

- 1) The TPS maintains a permissive for RF (RF_IS_OK), which is the logical “AND” of ARMED, \RESET, \SOFT_RESET, \SOFTFAULT, \HARDFLT, \CTRLXFLT, \WGSWINH.
- 2) The ACS requests RF be turned on. If RF_IS_OK is true, this gets passed on to Box #7A (there is a HEARTBEAT hardwired interlock on this signal).
- 3) When Box #7A receives this signal from all carts, it returns it back to all carts. This generates the internal signal RF_IS_ON within the TPS.
- 4) Any of the four channels may have the PLC control bit for RF ENABLE set or not set. For any of those where the bit is set, the signal is now sent to switch on the PIN diodes to pass RF to the klystron. (there is a HEARTBEAT hardwired interlock on each of these four signals).

Note that again, Box #7A serves to synchronize RF. If there are any soft-faults, even if only on one klystron, the synchronization action of Box #7A will serve to shutdown RF Pin diode switches on all TPS units for all carts.

SEQUENCING – FAULT BYPASS

For calibration purposes, we sometimes wish to enable RF operation without HV (at low power levels). This was previously done by wiring in interlock jumpers, a particularly unsafe practice. We now have a new signal, called FAULT BYPASS. This is set by the PLC, and shipped to the TPS on a dedicated fiber. When this is high, all HV activity is prohibited (enable and on). Most of the RF permissives are additionally bypassed, so that the net effect is that any PIN diode can be switched on and off with the PLC RF enable bit for that channel. All fast fault protection logic is bypassed, but this is with the understanding that this is all low power calibration work – and is deemed safe since HV operation is precluded.

SEQUENCING PROTECTION - HEARTBEAT

A number of critical output signals are protected by a heartbeat circuit on the TPS. This circuit uses a 100 Hz oscillator generated on one CPU, which is then passed daisy-chain through all cpus and CPLDs on the TPS assembly. The pass-thru within the cpus is performed by polling code (not interrupt code) at the outermost loop level, and thus will halt if any hangup occurs within the code.

On exiting the last chip, this signal is averaged, and then routed to two comparators. If it drifts off the middle towards either rail, hardwired logic gates shutdown any HV ENABLE, HV ON, or RF ON activity.

The heartbeat “ok” signal (i.e. the logical signal that says that the average is off the rails) is sent out to the PLC on a dedicated fiber (i.e. not dependent on virtual DSQUIRT communications).

(Note that this is an entirely separate function from the CPU watchdogs. The watchdogs are intrinsic functions built into the ARM7 microprocessors (of which there are 3 on the TPS), which must be “fed” regularly. Any code hangup which results in a failure to “feed the dog” results in a hardware reboot of the cpu. Any reboots result in faults and shutoff of HV and RF. The watchdog timeout is set to 1 second (and cannot be much shorter, as the flash-write function takes a significant fraction of a second, during which the dog cannot be fed).

SOFT FAULTS

SOFT faults are defined as those which require instantaneous shutdown of the RF, but not the HV. These generally result from waveguide arcs, as detected by excessive reflected power, elbow or circulator arc detectors, or CPS arc detectors. The occurrence of any SOFT fault will turn off the RF_IS_OK bit, which will shutdown RF (i.e. shutoff the group signal to Box #7A as well as directly shut down all PIN diodes).

Waveguide arcs usually clear in tens of milliseconds, so there is logic to wait a short time, then attempt to clear the soft fault. There is a separate SOFT FAULT RESET signal which is generated after a programmable delay, and which will clear any soft fault latches provided that the fault condition has disappeared. Normally, this will result in RF being re-enabled after a short time.

The delay before this reset is applied is programmable, as is the number of attempts before the TPS gives up and declares a HARD fault, which shuts down all HV for the duration of the shot.

The other danger with SOFT faults is that there is more power dissipated on the klystron collector when the HV is on but RF is off – there is no RF power extracted, so the spent electron beam heats the copper at a greater rate. The specification of the klystrons says that the maximum pulse duration must be derated if the full unspent beam is incident on the collector. This is the motivation behind the COTS circuit, which maintains a predicted model temperature of the water in the klystron collector cooling system, to detect dangerous overheating.

HARD FAULTS

HARD faults are defined as those which require instantaneous shutdown of the high voltage. These generally result from tube arcs or control failures. The tube arcs show up as overcurrent comparator trips on any of the current transformers. Control failures can occur from loss of communications with one of the ASQUIRT or DSQUIRT modules, failure of one of the local on-board TPS power rails, or failure of the ACS “CONTROL OK” signals which are brought in on dedicated fibers.

Two additional situations can result in a HARD fault. One is that the SOFT-RESET-LIMIT is reached, i.e. that we have attempted to reset a SOFT fault more times than the pre-programmed limit allows. The other is that the COTS calculation shows that the collector temperature of one of the klystrons has passed the trip point. Both of these are essentially escalations whereby a SOFT fault turns into a HARD fault.

Nothing is perfectly black and white, however. Two types of hard faults do NOT actually produce a HARD fault, they merely shut down HV ON (Trigger Gate) and RF ON, but they do not shut down HV ENABLE. These are the SOFT-RESET_LIMIT and the loss of ACS CONTROL OK.

PLC I/O

There are 64 bits of digital communication between the PLC and the TPS, communicated over the two DSQUIRT (multiplexed digital interface transceivers). One of these is a 32-bit one-way link, reporting 32 status bits to the PLC from the TPS. The other is a full-duplex 16/16 bit two-way link, reporting 16 more status bits to the PLC, and receiving 16 control bits from the PLC to the TPS.

The 48 status bits transmitted are listed below in the tables for the MDSplus HF command, which echos the same information (48 status bits) on an MDS query.

The 16 control bits received from the PLC are:

- HV_EN_IN_PLC (HV enable request) – channel 0
- RF_ENA_IN_PLC (RF enable for A) – channel 1
- RF_ENB_IN_PLC (RF enable for B) – channel 2
- RF_ENC_IN_PLC (RF enable for C) – channel 3
- RF_END_IN_PLC (RF enable for D) – channel 4
- RESET_IN_PLC (RESET) – channel 5
- \B7AFLT (Box #7A fault, COTS override) – channel 6
- spares – channels 7-15

FRONT PANEL LEDs

There are two clusters of four LEDs on the front panel of the TPS.

The bank of four LEDs on the middle-level (with the two fibers, the beam currents, and the RF diodes) are for the COTS functions. They are red/green for each of the COTS channels, showing the fault status. They also do parallel duty to show synchronization with the ASQUIRT transmitters – one pair of two lights shows the synch status for the ASQUIRT01, the other for the ASQUIRT02.

The bank of four LEDs on the level-1 board (with the body currents and Ethernet connector) show the status of the operation:

LED1 = green if ARMED, blank if not ARMED (never red)

LED2 = green if HV is enabled, blank if not (never red)

LED3 = green if HV is ON (trigger gate), blank if not. red if there is a HARD fault

LED4 = green if RF is ON, blank if not. red if there is a SOFT fault

ARMED

The TPS maintains an ARMED bit, which is a state bit significant for operation permission and for use of the LOGGER. Under normal operation (i.e. not FAULT BYPASS), the HV ON (Trigger Gate) and RF ON outputs require that the TPS be ARMED. Note that the HV ENABLE does not require ARMED, so that the Ross Relays will not be cycled every shot unless fault conditions require it.

ARMED also is crucial for operating the LOGGER. When not ARMED, we can answer to all MDS commands, which may involve sufficient latency that we won't properly log events. When ARMED, we lock out most MDS communications and concentrate heavily on timing crucial events.

In the normal course of a shot, ARMED is set by MDS command prior to the start of the shot, but after all the initialization script has completed. The TPS is now awaiting the rising edge of MDS START to begin the event timer clock. At the falling edge of MDS START, ARMED is turned off, and the TPS again responds to MDS commands.

If the TPS is ARMED and the MDS START never appears, it will disarm after a timeout period. Similarly, if the MDS START hangs high and never goes low, the TPS will disarm after the same timeout. The timeout period is programmable via the MDSplus parameters.

MDS CLOCK and START

These two signals are received via F/O inputs on the RTM. They are routed to the LOGGER cpu, which needs them for the LOGGER timing...and the START is also routed to the MASTER which uses it to perform a disarm at the end of the START pulse.

As an additional service, both signals are echo'd out via front panel LEMO-1 connectors, for routing to an adjacent digitizer.

RESET

SOFT fault resets are discussed above. RESET is a signal which attempts to clear all fault latches (provided that the originating fault situation is gone). RESET is the logical "or" of RESET requests by MDS, PLC, front-panel pushbutton, or POWER ON one-shot. All are indistinguishable from each other.

REPROGRAMMING

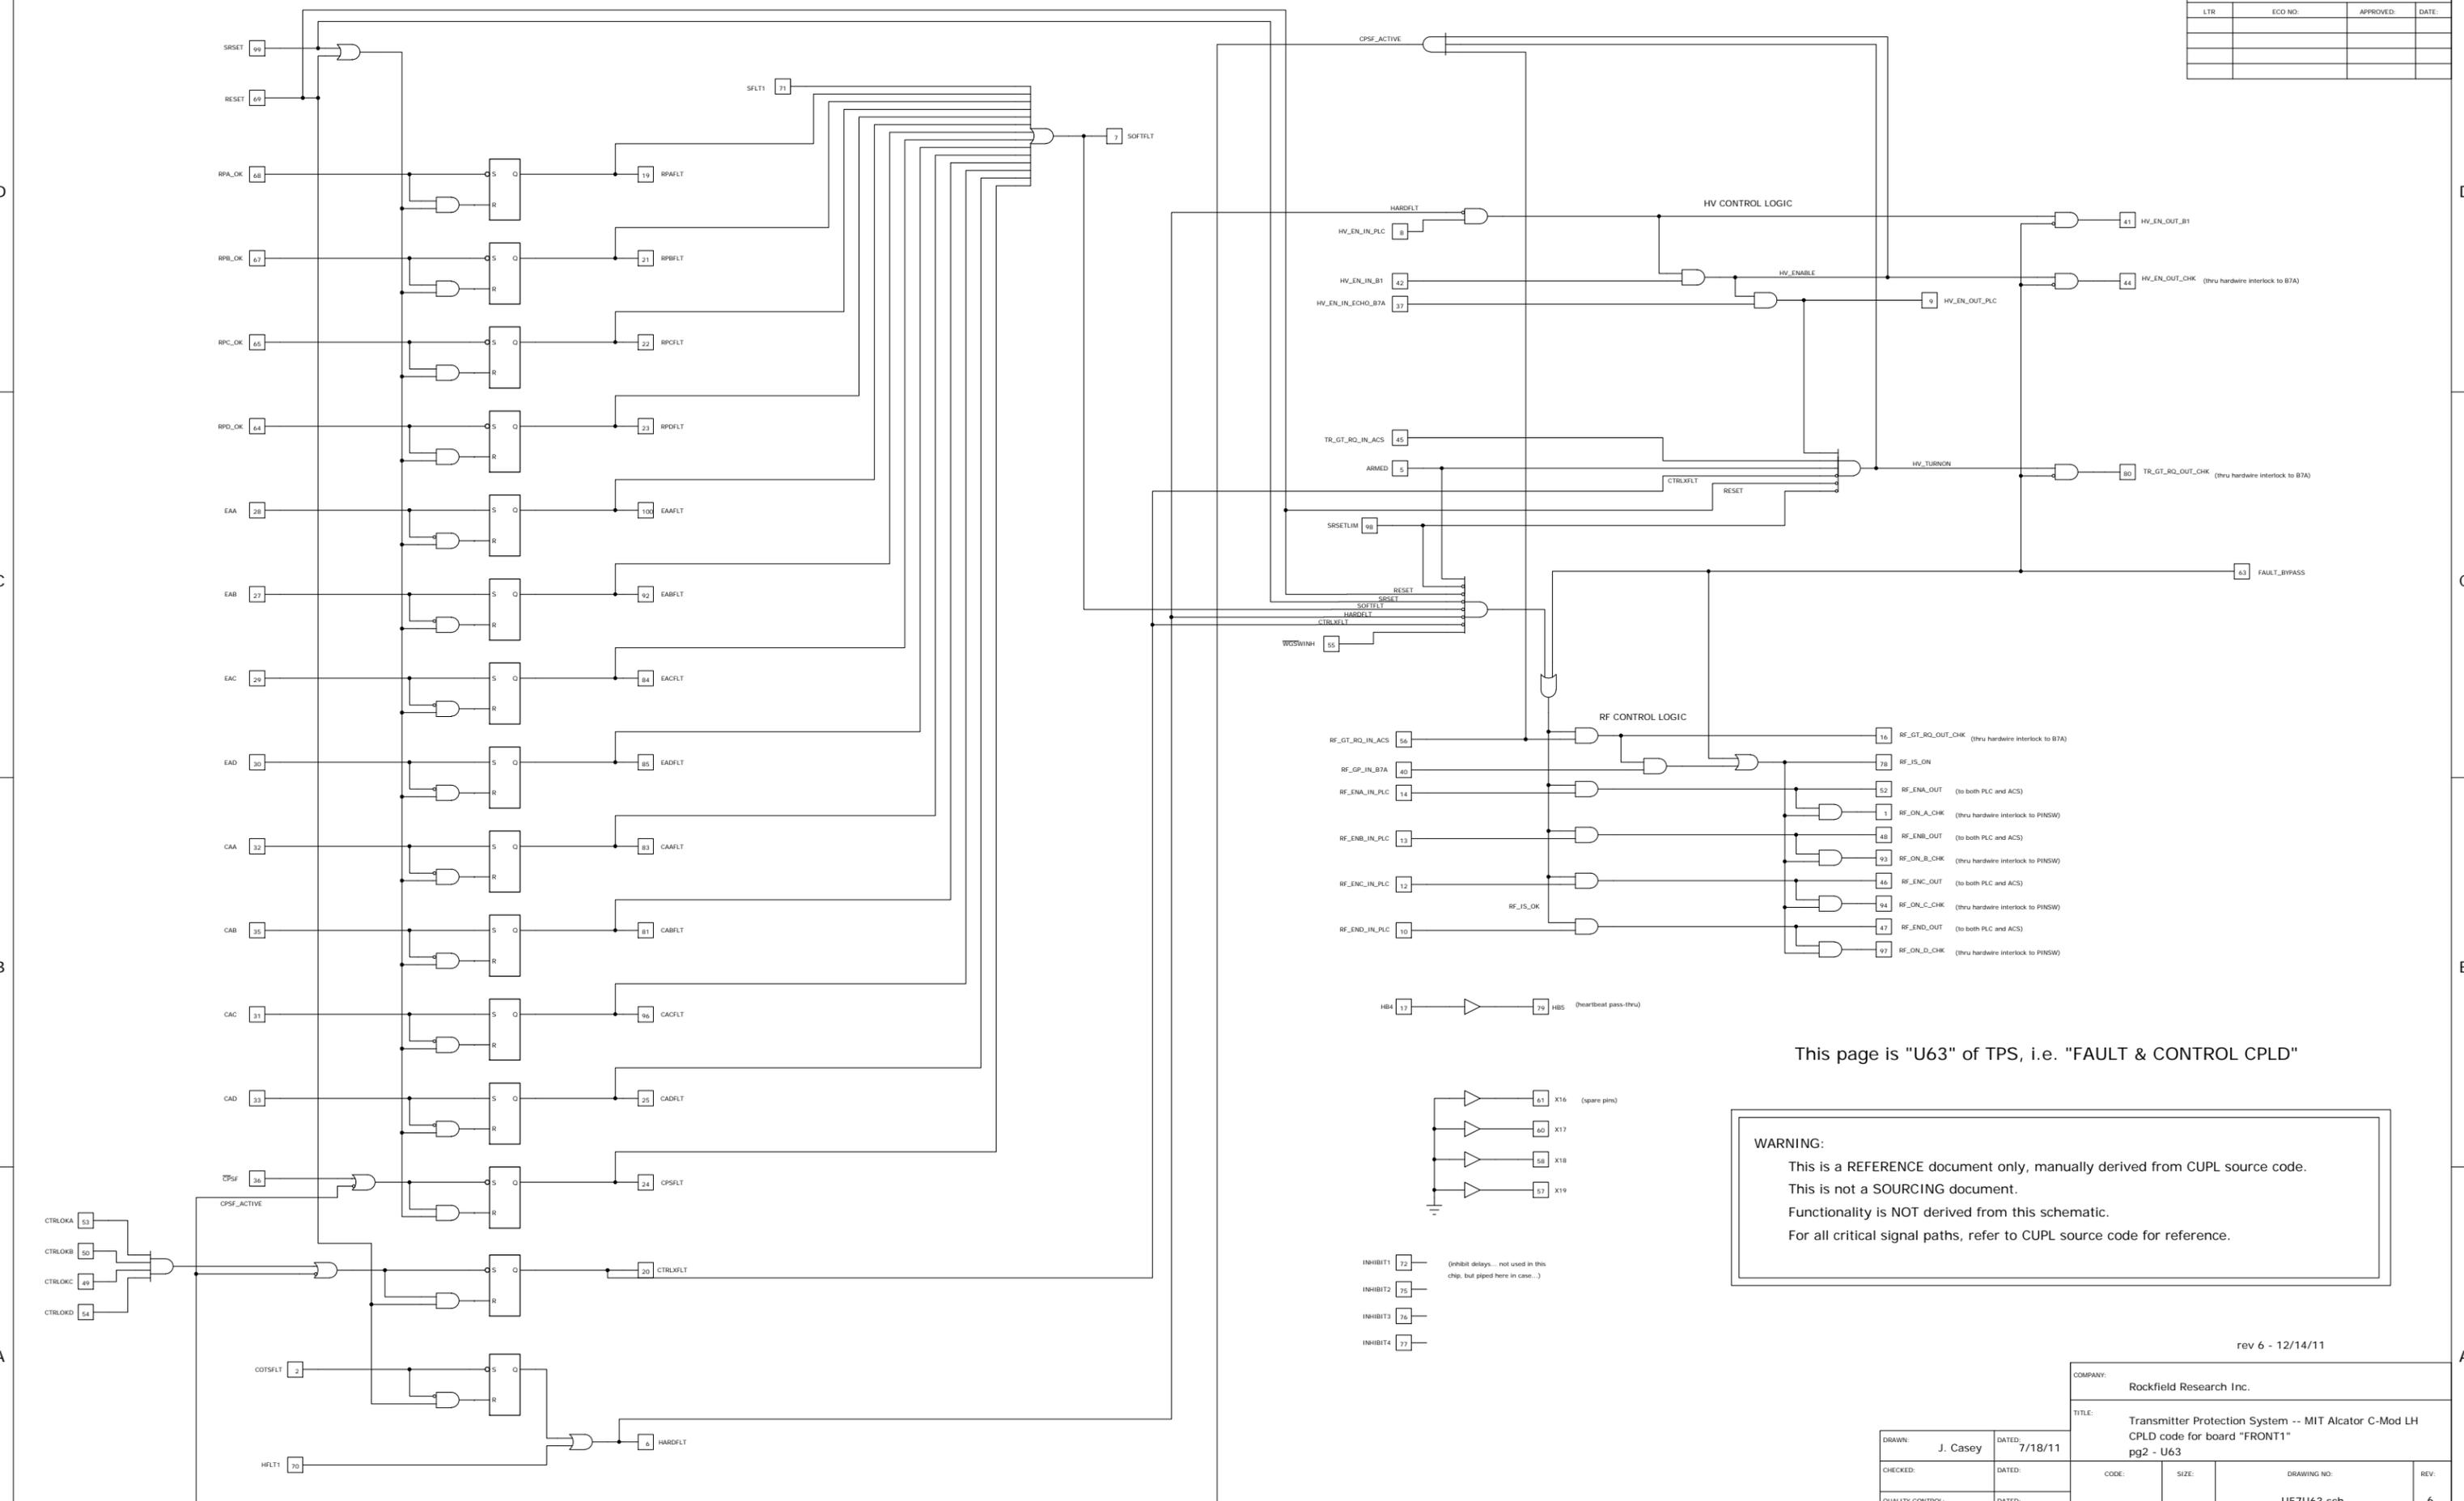
The TPS has five "intelligent" chips onboard....three NXP ARM7 microprocessors and two Atmel CPLDs. Each of the microprocessors has a dedicated SIP8 connector, to which a custom programming signal may be brought in (this interfaces to one of the chip UARTs, as well as two logic lines for programming, and talks to a serial port on a PC for downloading assembled code). Each of the CPLDs has a dedicated 2x3 header bringout the JTAG pins for programming with any universal programmer.

The ARM7 processor on the COTS board (front board 2) has one additional constraint – it needs both available UARTs for ASQUIRT communications, so one must do double duty for programming as well. There is an onboard jumper which must be moved for programming, and moved back for operation.

(Note that one should ALWAYS perform a FD and FF command to initialize factory operation coefficients after programming.)

Contact Rockfield Research for programming details.

REVISION RECORD			
LTR	ECO NO.	APPROVED:	DATE:



This page is "U63" of TPS, i.e. "FAULT & CONTROL CPLD"

WARNING:
 This is a REFERENCE document only, manually derived from CUPL source code.
 This is not a SOURCING document.
 Functionality is NOT derived from this schematic.
 For all critical signal paths, refer to CUPL source code for reference.

COMPANY:		Rockfield Research Inc.			
TITLE:		Transmitter Protection System -- MIT Alcatel C-Mod LH CPLD code for board "FRONT1" pg2 - U63			
DRAWN:	J. Casey	DATED:	7/18/11	CODE:	
CHECKED:		DATED:		SIZE:	
QUALITY CONTROL:		DATED:		DRAWING NO.:	U57U63.sch
RELEASED:		DATED:		REV.:	6
SCALE:					SHEET 2 OF 2

TABLE #1 – Analog Signal Pathways

Signal	Raw Scaling	Internal Scaling	Input	Test Point	SCSI68 conn
I_{kA} (beam cur A)	25 A/V	5 A/V	F2-J1	F3-J2L	F3-J1.3 (00)
I_{kB} (beam cur B)	25 A/V	5 A/V	F2-J2	F3-J2U	F3-J1.4 (01)
I_{kC} (beam cur C)	25 A/V	5 A/V	F2-J3	F3-J3L	F3-J1.5 (02)
I_{kD} (beam cur D)	25 A/V	5 A/V	F2-J4	F3-J3U	F3-J1.6 (03)
I_{BFA} (fast body cur A)	1 A/V	20 mA/V	F1-J5	F3-J4L	F3-J1.7 (04)
I_{BFB} (fast body cur B)	1 A/V	20 mA/V	F1-J6	F3-J4U	F3-J1.8 (05)
I_{BFC} (fast body cur C)	1 A/V	20 mA/V	F1-J7	F3-J5L	F3-J1.9 (06)
I_{BFD} (fast body cur D)	1 A/V	20 mA/V	F1-J8	F3-J5U	F3-J1.10 (07)
I_{BSA} (slow body cur A)	10 mA/V	10 mA/V	F1-J1	F3-J6L	F3-J1.11 (08)
I_{BSB} (slow body cur B)	10 mA/V	10 mA/V	F1-J2	F3-J6U	F3-J1.12 (09)
I_{BSC} (slow body cur C)	10 mA/V	10 mA/V	F1-J3	F3-J7L	F3-J1.13 (10)
I_{BSD} (slow body cur D)	10 mA/V	10 mA/V	F1-J4	F3-J7U	F3-J1.14 (11)
P_{fwd-A} (fwd pwr A)	(nonlinear)	raw x2.5	F2-J5	F3-J8L	F3-J1.15 (12)
P_{fwd-B} (fwd pwr B)	(nonlinear)	raw x2.5	F2-J7	F3-J8U	F3-J1.16 (13)
P_{fwd-C} (fwd pwr C)	(nonlinear)	raw x2.5	F2-J9	F3-J9L	F3-J1.17 (14)
P_{fwd-D} (fwd pwr D)	(nonlinear)	raw x2.5	F2-J11	F3-J9U	F3-J1.18 (15)
P_{refl-A} (refl pwr A)	(nonlinear)	raw x2.5	F2-J6	F3-J10L	F3-J1.19 (16)
P_{refl-B} (refl pwr B)	(nonlinear)	raw x2.5	F2-J8	F3-J10U	F3-J1.20 (17)
P_{refl-C} (refl pwr C)	(nonlinear)	raw x2.5	F2-J10	F3-J11L	F3-J1.21 (18)
P_{refl-D} (refl pwr D)	(nonlinear)	raw x2.5	F2-J12	F3-J11U	F3-J1.22 (19)
COTS_A (temp model A)	(internal)	25 C/V	–	F3-J12L	F3-J1.23 (20)
COTS_B (temp model B)	(internal)	25 C/V	–	F3-J12U	F3-J1.24 (21)
COTS_C (temp model C)	(internal)	25 C/V	–	F3-J13L	F3-J1.25 (22)
COTS_D (temp model D)	(internal)	25 C/V	–	F3-J13U	F3-J1.26 (23)
T_{collA} (temp meas A)	300F ⁽⁴⁾	25 C/V	ASQ02 (14)	F3-J14L	F3-J1.27 (24)
T_{collB} (temp meas B)	300F	25 C/V	ASQ02 (15)	F3-J14U	F3-J1.28 (25)
T_{collC} (temp meas C)	300F	25 C/V	ASQ02 (01)	F3-J15L	F3-J1.29 (26)
T_{collD} (temp meas D)	300F	25 C/V	ASQ02 (00)	F3-J15U	F3-J1.30 (27)
I_G (cart gnd cur)	100 A/V	20 A/V	F1-J9	F3-J16L	F3-J1.31 (28)
I_{TOT} (total cur)	50 A/V	50 A/V	ASQ01 (2)	F3-J16U	F3-J1.32 (29)
V_K (beam voltage)	6 kV/V	6 kV/V	ASQ01 (1)	F3-J17L	F3-J1.33 (30)
SP₁ (spare #1)	tbd	tbd	F1-J10	F3-J17U	F3-J1.34 (31)
SP₂ (spare #2) ⁽⁵⁾	tbd	tbd	F1-J11	F3-J15U	F3-J1.30 (27)
SP₃ (spare #3) ⁽⁶⁾	tbd	tbd	F1-J12	F3-15L	F3-J1.29 (26)

⁴ The collector temperature measurements are provided to the ASQUIRT02 transmitter as 4-20mA signals, and the ASQUIRT02 inputs are configured for 0-20mA, thus the 300F range comprises 20% to 100% of the 12-bit dynamic range.

⁵ SPARE #2 replaces T_{coll} channel D with a jumper selection.

⁶ SPARE #3 replaces T_{coll} channel C with a jumper selection.

TABLES #2 – Alternate Debugging Monitor Signals

The MDSplus settable parameter “CK” allows us to customize what data are presented on the analog outputs which are normally reserved for the echo of the klystron downstream (delayed) output temperature measurements.

For a setting of **CK=0** (default), the data is as presented above in TABLE #1 (repeated here):

Signal	Raw Scaling	Internal Scaling	Input	Test Point	SCSI68 conn
T_{collA} (temp meas A)	300F	25 C/V	ASQ02 (14)	F3-J14L	F3-J1.27 (24)
T_{collB} (temp meas B)	300F	25 C/V	ASQ02 (15)	F3-J14U	F3-J1.28 (25)
T_{collC} (temp meas C)	300F	25 C/V	ASQ02 (01)	F3-J15L	F3-J1.29 (26)
T_{collD} (temp meas D)	300F	25 C/V	ASQ02 (00)	F3-J15U	F3-J1.30 (27)

For a setting of **CK=1**, we show the relevant inputs used for COTS temperature calculation for tube A:

Signal	Raw Scaling	Internal Scaling	Input	Test Point	SCSI68 conn
T_{in} (inlet temp)	100C	100C	ASQ02 (12)	F3-J14L	F3-J1.27 (24)
T_{collA} (temp meas A)	300F	25 C/V	ASQ02 (14)	F3-J14U	F3-J1.28 (25)
P_A (outlet press A)	232 psi	232 psi	ASQ02 (03)	F3-J15L	F3-J1.29 (26)
Q_A (flow A)	100 gpm	100 gpm	ASQ02 (02)	F3-J15U	F3-J1.30 (27)

Likewise, for a setting of **CK=2**, we show relevant inputs used for COTS temperature B:

Signal	Raw Scaling	Internal Scaling	Input	Test Point	SCSI68 conn
T_{in} (inlet temp)	100C	100C	ASQ02 (12)	F3-J14L	F3-J1.27 (24)
T_{collB} (temp meas B)	300F	25 C/V	ASQ02 (15)	F3-J14U	F3-J1.28 (25)
P_B (outlet press B)	232 psi	232 psi	ASQ02 (05)	F3-J15L	F3-J1.29 (26)
Q_B (flow B)	100 gpm	100 gpm	ASQ02 (04)	F3-J15U	F3-J1.30 (27)

Likewise, for a setting of **CK=3**, we show relevant inputs used for COTS temperature C:

Signal	Raw Scaling	Internal Scaling	Input	Test Point	SCSI68 conn
T_{in} (inlet temp)	100C	100C	ASQ02 (12)	F3-J14L	F3-J1.27 (24)
T_{collC} (temp meas C)	300F	25 C/V	ASQ02 (01)	F3-J14U	F3-J1.28 (25)
P_C (outlet press C)	232 psi	232 psi	ASQ02 (07)	F3-J15L	F3-J1.29 (26)
Q_C (flow C)	100 gpm	100 gpm	ASQ02 (06)	F3-J15U	F3-J1.30 (27)

Likewise, for a setting of **CK=4**, we show relevant inputs used for COTS temperature D:

Signal	Raw Scaling	Internal Scaling	Input	Test Point	SCSI68 conn
T_{in} (inlet temp)	100C	100C	ASQ02 (12)	F3-J14L	F3-J1.27 (24)
T_{collD} (temp meas D)	300F	25 C/V	ASQ02 (00)	F3-J14U	F3-J1.28 (25)
P_D (outlet press D)	232 psi	232 psi	ASQ02 (09)	F3-J15L	F3-J1.29 (26)
Q_D (flow D)	100 gpm	100 gpm	ASQ02 (08)	F3-J15U	F3-J1.30 (27)

It is important to remember that ALL of the ASQUIRT02 input signals are 4-20mA signals, read by a 0-20mA input. The data is transferred as 12-bit (0-4095), then output to test points and digitizer as 0-10V signals. The range of the data thus comprises only the 20%-100% of the dynamic range, or 2V-10V. This is NOT true for the temperature measurements (the default signals), where we have taken the extra step to rescale the 0-10V range to the same 0-250C range used by the COTS temperature outputs.

Finally, for a setting of **CK=5**, we echo the internal calibrated forward power, which is calculated from the forward power diode measurement and the cubic polynomial fit coefficients which are fed in via the MDSplus entries. The internal power measurements are calculated and stored internally as integers in units of mW. In order to present these on a 12-bit DAC output, they are right-shifted by 2^{17} , thus the 0-10V range maps to 0-536.74 kW.

Signal	Raw Scaling	Output Scaling	Input	Test Point	SCSI68 conn
P_{fwd-A} (fwd pwr A)	(nonlinear)	53.674 kW/V	F2-J5	F3-J14L	F3-J1.27 (24)
P_{fwd-B} (fwd pwr B)	(nonlinear)	53.674 kW/V	F2-J7	F3-J14U	F3-J1.28 (25)
P_{fwd-C} (fwd pwr C)	(nonlinear)	53.674 kW/V	F2-J9	F3-J15L	F3-J1.29 (26)
P_{fwd-D} (fwd pwr D)	(nonlinear)	53.674 kW/V	F2-J11	F3-J15U	F3-J1.30 (27)

TABLE #3 – MDSplus COMMANDS

The Ethernet interface provides a means for MDSplus to run a setup script prior to each shot, setting all comparator thresholds, calibration coefficients, and operating values to settings out of the data tree. Similarly, a readout script run after the shot can download and archive the data logger results.

All commands are two upper case characters, perhaps with parameters, terminated by <CR>. Parameters are passed as ascii strings representing unsigned integers. All parsing errors return a “?”. All valid commands return an echo (in lower case) followed by the relevant parameter. A valid command followed by an unparseable parameter results in NO CHANGE to the parameter, but instead the return string will show the present setting.

It is not actually necessary to send the commands as upper case, nor does it matter if spaces and such are inserted – incoming data is shifted to uppercase and all generic whitespace characters are removed before parsing.

The MDSplus commands which are used for setup, maintenance, and overall operation are:

Description	Command	Parameters	Return	factory default
ARM TPS for a shot	A+	(none)	a+	
DISARM ⁽⁷⁾	A-	(none)	a-	
set MDS DISARM timeout	AT	1-9999 (minutes)	at####	5 (5 min)
set debugging parameter CK	CK	0-5	ck#	0
set all parameters to factory default	FD	(none)	fd	
burn current setpoints to flash ⁽⁸⁾	FF	(none)	ff1 / ff2 / ff3	
set ASQUIRT01 comm address	Q1	0-7	q1#	0
set ASQUIRT02 comm address	Q2	0-7	q2#	7
set DSQUIRTIO comm address	Q3	0-7	q3#	1
set DSQUIRTII comm address	Q4	0-7	q4#	2
RESET	RR	(none)	rr	
set soft-reset retry count limit	RN	0-500	rn###	5
set soft-reset retry delay	RT	5-150 (ms)	rt###	20 (20ms)
test watchdog, master chip U58	T1	(none)	reboot1	
test watchdog, logger chip U64	T2	(none)	reboot2	
test watchdog, cots chip front2	T3	(none)	reboot3	
enable eXXtra reset at MDS START	XX	0-1	xx#	1 (enabled)
query firmware version	VV	(none)	vv#####	

⁷ DISARM is the only accepted command once ARMED. Normally, once ARMED, the TPS will automatically DISARM after the MDS-START signal has gone high followed by going low. (MDS-START high when not ARMED results in a fault).

⁸ FF will burn the current setpoints to flash, so that they become the defaults on subsequent power-up. Note that FD does not do this. Restoring the TPS to factory defaults requires a FD followed by a FF.

The MDSplus commands which load calibration and fault threshold setpoints are:

Description	Command	Parameters	Return	factory default
set fast body curr flt thresh A	BA	0-4095 (0-200mA)	ba####	1639 (80 mA)
set fast body curr flt thresh B	BB	0-4095 (0-200mA)	bb####	1639 (80 mA)
set fast body curr flt thresh C	BC	0-4095 (0-200mA)	bc####	1639 (80 mA)
set fast body curr flt thresh D	BD	0-4095 (0-200mA)	bd####	1639 (80 mA)
set slow body curr flt thresh A	BE	0-4095 (0-100mA)	be####	3278 (80 mA)
set slow body curr flt thresh B	BF	0-4095 (0-100mA)	bf####	3278 (80 mA)
set slow body curr flt thresh C	BG	0-4095 (0-100mA)	bg####	3278 (80 mA)
set slow body curr flt thresh D	BH	0-4095 (0-100mA)	bh####	3278 (80 mA)
set COTS heat capacity A	CA	50-200 (*100 J/K)	ca###	110 (11 kJ/K)
set COTS heat capacity B	CB	50-200 (*100 J/K)	cb###	110 (11 kJ/K)
set COTS heat capacity C	CC	50-200 (*100 J/K)	cc###	110 (11 kJ/K)
set COTS heat capacity D	CD	50-200 (*100 J/K)	cd###	110 (11 kJ/K)
set diode calibration coeffs A	DA	##### four numbers: uns int 0-4095 uns int 0-30k uns int 0-60k uns int 0-9k	da#####	412 8050 14400 465
set diode calibration coeffs B	DB	(same)	db...	same
set diode calibration coeffs C	DC	(same)	dc...	same
set diode calibration coeffs D	DD	(same)	dd...	same
set ground current flt threshold	GG	0-4095 (0-200mA)	gg#####	1229 (60A)
set inhibit A delay (HV_ON ldg edge for beam cur)	IA	5-999 (us)	ia###	10
set inhibit B delay (HV_ON both edges fst bdy cur)	IB	5-999 (us)	ib###	10
set inhibit C delay (HV_ON both edges slw bdy cur)	IC	5-999 (us)	ic###	10
set inhibit D delay (RF_ON both edges fst bdy cur)	ID	5-999 (us)	id###	10
set beam current threshold A	KA	1000-4095 (5-20.48A)	ka#####	2400 (12A)
set beam current threshold B	KB	1000-4095 (5-20.48A)	kb#####	2400 (12A)
set beam current threshold C	KC	1000-4095 (5-20.48A)	kc#####	2400 (12A)
set beam current threshold D	KD	1000-4095 (5-20.48A)	kd#####	2400 (12A)
set COTS trip multiplier A ⁹	MA	20-50 (%)	ma##	28 (28%)
set COTS trip multiplier B	MB	20-50 (%)	mb##	28 (28%)
set COTS trip multiplier C	MC	20-50 (%)	mc##	28 (28%)
set COTS trip multiplier D	MD	20-50 (%)	md##	28 (28%)
set RF ratio fault threshold A	RA	0-4095 (0-100%)	ra#####	819 (20%)
set RF ratio fault threshold B	RB	0-4095 (0-100%)	rb#####	819 (20%)
set RF ratio fault threshold C	RC	0-4095 (0-100%)	rc#####	819 (20%)
set RF ratio fault threshold D	RD	0-4095 (0-100%)	rd#####	819 (20%)
set spare chan #n flt threshold	Sn	n=1-3, 0-4095 (tbd)	sn#####	410
set COTS trip safety margin A	TA	5-30 (degC)	ta##	15 (15C)
set COTS trip safety margin B	TB	5-30 (degC)	tb##	15 (15C)
set COTS trip safety margin C	TC	5-30 (degC)	tc##	15 (15C)
set COTS trip safety margin D	TD	5-30 (degC)	td##	15 (15C)

⁹ The multiplier is expressed as $M = (T_{cots} - T_{in}) / (T_{wet} - T_{in})$ in units of percent

The MDSplus commands which are used for post-shot readout are:

Description	Command	Parameters	Return	factory default
“hoo faulted”? (i.e. readout all status bits normally sent to PLC to determine what fault latches are presently set)	HF	(none)	fault code, see below	
how many events?	EN	(none)	en#### count of events	
event dump readout	ED	(none)	event dump data, see below	

COMMAND “HF” DETAILS

The command HF produces a hex readout, consisting of 12 hex characters. This represents 48 status bits, which are an exact copy of the 48 status bits which are shipped out in realtime to the PLC over the DSQUIRTIO and DSQUIRTII interfaces.

The least significant 4 chars represent the 16 bits of transmitted DSQIO output, which are (from LSB to MSB):

HV_EN_OUT_PLC (0)	CTRLOKB (8)
RF_ENA_OUT	CTRLOKC
RF_ENB_OUT	CTRLOKD
RF_ENC_OUT	\RPAFLT
RF_END_OUT	\RPBFLT
\IGFLT	\RPCFLT
WGSWINH	\RPDFLT
CTRLOKA (7)	\IKAFLT (15)

The most significant 8 chars represent the 32 bits of DSQII output, which are (from LSB to MSB):

\IKBFLT (0)	COTSB_OK (16)
\IKCFLT	COTSC_OK
\IKDFLT	COTSD_OK
\IBAFLT	\CTRLXFLT
\IBBFLT	\PWRFLT
\IBCFLT	\CPSFLT
\IBDFLT	\SPIFLT
\CAAFLT	RF_IS_ON
\CABFLT	ELBOW_ARC_RESET
\CACFLT	HV_EN_IN_ECHO_B7A
\CADFLT	FAULT_BYPASS
\EAAFLT	\SOFTFAULT
\EABFLT	\HARDFULT
\EACFLT	\SRSET
\EADFLT	RESET
\COTSA_OK (15)	ARMED (31)

COMMAND “ED” DETAILS

The command ED produces a block readout of up to 40 events. Each event is configured as a 30 hex character encoding, consisting of:

- 3 hex character id tag (a 12 bit number, id from 0-2047)
- 18 hex character status word (72 data bits...this is extendable to 96 bits if needed)
- 8 hex character timestamp (32 bits, count of MDS clock at 1 MHz since MDS-START edge)
- 1 character terminator <LF> (ascii 10)

The packet of 40 events is thus smaller than a single ethernet packet, aiding in flow control.

The first ED will start at event #0. Successive ED calls will continue the list at #40, #80, etc. The last ED call will not be padded, but will result in less than or equal to 40 events. Additional calls to ED will restart the next block with event #0 again. The script should first perform an EN command, in order to know how many events to expect.

The on-board memory can only hold 2048 events. If there are more than 2048 events, the counter will continue to increment, but the additional events will not be saved (i.e. the memory will always represent only the first 2048 events). If there are more than 9999 events, the EN counter will return ##### as a numerical error code.

The 72-bit status word is comprised of (from LSB to MSB):

IKxFLT, 4 channels	bits 0-3	PLC_RESET	bit 43
IBxFLT, 4 channels	bits 4-7	RESET_BUTTON	bit 44
RPxFLT, 4 channels	bits 8-11	ARMED	bit 45
EAxFLT, 4 channels	bits 12-15	FAULT_BYPASS	bit 46
CxFLT, 4 channels	bits 16-19	HV_EN_IN_PLC	bit 47
COTSxOK, 4 channels	bits 20-23	HV_EN_OUT_B1	bit 48
CTRLOKx, 4 channels	bits 24-27	HV_EN_IN_B1	bit 49
CTRLXFLT	bit 28	HV_EN_OUT_CHK	bit 50
IGFLT	bit 29	HV_EN_OUT_PLC	bit 51
SPxFLT, 3 channels	bits 30-32	HV_EN_IN_ECHO_B7A	bit 52
PWRFLT	bit 33	TR_GT_RQ_IN_ACS	bit 53
CPSFLT	bit 34	TR_GT_RQ_OUT_CHK	bit 54
COTSFLT	bit 35	WGSWINH	bit 55
HARDFLT	bit 36	RF_ENx_IN_PLC, 4 channels	bits 56-59
SOFTFLT	bit 37	RF_ENx_OUT, 4 channels	bits 60-63
SRSETLIM	bit 38	RF_GT_RQ_IN_ACS	bit 64
SRSET	bit 39	RF_GT_RQ_OUT_CHK	bit 65
RESET	bit 40	RF_GP_IN_B7A	bit 66
POR	bit 41	RF_IS_ON	bit 67
MDS_RESET	bit 42	RF_ON_x_CHK	bit 68-71

Questions and comments should be directed to Jeff Casey at Rockfield Research Inc., see header for contact information, or email to casey@rockfieldresearch.com.

APPENDIX A – COTS Documentation

On the following pages, you will find a separate document which details the calculations of the COTS feature of the TPS..

APPENDIX B – SBIR Final Report

Following the COTS documentation, you will find the final report from the DOE SBIR contract, which presents a more general overview of the background, use, and features of the TPS, as well as many of the peripheral circuits used in the Alcator C-Mod TPS system.

Version:

This document was previously released as a DRAFT, in August of 2011. Since then we have had many changes in the scaling of various coefficients, as well as in the algorithm for determining the fault threshold temperature. **Changes are flagged in red.** This document is up to date as of 1/12/12, firmware revision v010.

TPS COTS Protection

The new TPS system includes a Collector Over Temperature System (COTS), which does a real-time prediction of the water temperature inside the klystron collector, and allows operation closer to the danger limit than would otherwise be safe if we relied on external temperature sensors alone. This document shows the details of the calculation.

The COTS algorithm is coded in ARMASM to run in realtime on an NXP (Phillips) LPC2103 ARM7 microprocessor. No operating system or high level compiler is used, thus increasing speed and reliability (at the cost of coding effort.) This is a relatively fast 32 bit micro, but has no floating point math skills. Some pre-processing has been done in coding, and will be necessary in converting from MDS+ setpoints to shot-loadtime setpoints in order to streamline the calculations. No divide or transcendental math are performed in realtime – all such calculations are pre-fit to polynomial expansions over the range of interest, allowing exclusive use of hardware instructions for multiply, accumulate, and shift.



Digitized Inputs

The COTS calculation is performed on the 2nd (middle) of three boards in the front TPS assembly, which is a stack of three boards using three slots of a cPCI crate. There is one TPS assembly on each cart, thus each one services four klystrons. The beam current transducers and the RF diodes are processed on this same board for proximity to the COTS usage.

Beam Current

Each of the beam current (cathode current, klystron current) signals is generated by a FW Bell RS-100, with a raw signal calibration of 40 mV/A. It is oriented for positive signal under normal operations.

The analog processing section of the COTS board performs a clean-up (transient glitch protection, 1MHz RC filter), then buffers this through a 25 MHz op-amp (LT1358) with a differential gain of 5x. (The signal is treated differentially on input from a two-pin LEMO connector, although it is realized that the signal is technically single-ended). This buffered signal is fed to three comparators for fault detection (one to detect backwards-oriented installation, one with a pot-settable trip-point which is always active, and one with an MDS+ settable trip-point which is inhibited for a brief inhibit



delay after the HV is turned on). The signal is then fed to several analog points of interest: 1) buffered and passed to a dTacq digitizer connector, 2) buffered and passed to a single pin LEMO testpoint, 3) routed to a local 8-channel 12 bit multiplexed digitizer. This last is digitized with a 12 bit scaling of 0-25A. The COTS microprocessor polls each ADC channel at 1 kHz.

RF Diodes

Each of the RF diode pairs (forward, reverse), are provided by loop directional couplers with a default orientation for a negative signal. The directional couplers have a 50 dB attenuation, and the signal paths have an additional 15-20 dB attenuation. Typical signal amplitudes are about -2V.

Each detector input is again cleaned up (transient protection, 1 MHz RC), then buffered through a -2.5x gain differential amplifier, producing a positive signal with typical amplitude of about +5V. Each is similarly sent to the dTacq digitizer and to a single pin LEMO test point, and a ratiometric fault sensor is applied to each forward / reflected pair. This entails using an analog multiplier (AD633) with an MDS+ settable multiplication factor between 0.0 and 1.0 to attenuate the forward signal (and applying a 0.5V positive offset as well), and then faulting (via comparator) if the reflected power exceeds this calibrated fraction of the forward signal.

In addition, the forward signals (but not the reflected signals) are locally digitized by the same 8 channel 12 bit digitizer as the beam current. **The scaling of this is such that the 12-bit range corresponds to 3.00V of the raw diode signal.** This is also polled by the cpu at 1 kHz.

Acquisition Rates

The main loop of the COTS code polls this digitizer at 8 kHz, with interleaved sampling of the four beam currents and forward powers ultimately sampled at 1 kHz each. The COTS calculation is called at 4 kHz intervals, with interleaving between the four klystron channels – and each klystron COTS calculation is done as soon as its relevant signals are acquired; thus the COTS calculation for each klystron is updated at 1 kHz. Benchmarking shows that the entire calculation update for a single klystron is performed in 26 μ sec.

ASQUIRT Inputs

The TPS system makes extensive use of ASQUIRT and DSQUIRT modules, which are very similar to the in-house AFOL and DFOL links that Alcator C-Mod already uses. They differ in that they are fully digital, with complex messages passing via F/O both ways. These messages include dip-switch settable addressing, checksums, firmware version numbers, and other failsafes to help prevent erroneous operation of equipment. The ASQUIRT01 units transmit and receive 0-10V signals at 12 bit resolution, with two channels updated at 10 kHz. The ASQUIRT02 units transmit and receive either 0-10V or 0-20mA signals at 12 bit resolution, with 16 channels updated at 1 kHz.

One ASQUIRT01 unit placed in the LH HV control area of the Alcator C-Mod cell transmits the power supply voltage on channel 1, with a 12-bit range of 0-60 kV. The total supply current is on channel 2, with a 12-bit range of 0-500 A. The ASQUIRT01 units are made with four parallel F/O outputs (but only a single F/O input). This ASQUIRT01 unit feeds to all 3 (4 after the upgrade) LH klystron carts. There is no corresponding ASQUIRT01 mate to receive this signal – the F/O receiver is on the COTS board, and the signal is decoded by a “virtual ASQUIRT01” routine within the COTS cpu. The ASQUIRT01 timing is controlled by the sending unit, and the receiving decoder is polled, thus is asynchronous with the COTS calculations. When complete and valid message packets are received, the “outputs” (buffers in SRAM) and status flags are updated. Timeout on receive, receive errors, or other faults revert the outputs to values of zero, and error flags are set.

Four ASQUIRT02 units are used at each cart. One normal ASQUIRT02 pair monitors the RTD temperature sensors on the klystrons, with these signals received only at the PLC station, and dealt with normally there. One more ASQUIRT02 – modified to have dual parallel F/O outputs (but no F/O input to copper output section)

monitors fourteen signals associated with the klystron cooling on that cart. One fiber goes to a second ASQUIRT02 receiver at the PLC station, while the other goes to a “virtual ASQUIRT02” receiver on the COTS board, again decoded by the COTS cpu. This is treated similarly – the reception is polled and timing is controlled by the sender, thus asynchronous with the COTS calculation. Valid received packets are decoded and the messages are updated (packets are actually transmitted at 8 kHz, with two channels sent per packet). Any fault or timeout on reception zeros all received signals and sets fault bits. The channels are:

- Channels 14,15,01,00 – Klystron A,B,C,D outlet water temperature, 0-300 deg F
- Channels 02,04,06,08 – Klystron A,B,C,D water flow rate, 0-100 gpm
- Channels 03,05,07,09 – Klystron A,B,C,D outlet water pressure, 0-232 psig
- Channel 10 – Inlet water pressure, 0-362 psig
- Channel 12 – Inlet water temperature, 0-100 deg C
- Channels 11,13 – spares

All of these transducers are 4-20mA, and the ASQUIRT02 inputs are set for 0-20mA range, so the above transducer ranges correspond to 820-4095 counts

Since both ASQUIRT0x units operate asynchronously, updating the “outputs” – which are really just buffer memory in the cpu – whenever valid data is received, this data is treated as a simple lookup table for the COTS calculations, supplying the most recently received data.

ANALOG Outputs

The COTS system will produce a calculated model temperature for the water in the klystron collector cooling loop for each tube. The goal is to treat this as a realtime signal, and to latch a fault if it exceeds a given threshold. For verification purposes, we also convert this model in realtime to an analog signal and ship it out to the dTacq digitizer for acquisition by MDS+. These signals are updated at 1 kHz (as the COTS calculation is performed), and scaled with a 0-10V range corresponding to a 0-250C temperature. The analog outputs are generated by a muxed 12-bit DAC. (These analog signals are also available as single-pin LEMO testpoints on the TPS unit).

Additional analog outputs are provided for convenience in benchmarking performance of the system. We echo the two ASQUIRT01 signals (beam voltage, total current) to 0-10V signals with the same ranges (0-60kV, 0-500A respectively), updating these as the ASQUIRT01 packets are received, and sending to both dTacq digitizer pins and LEMO testpoints.

We also echo the *measured* klystron outlet water temperatures which are received from the ASQUIRT02 transmission. These are similarly sent to dTacq pins and LEMO test points. These are rescaled to have the same range as the COTS temperature signal – 0-10V range corresponds to 0-250C, and are updated as ASQUIRT02 packets are received.

For debugging, a new feature allows an MDSplus settable parameter (“CK”, for Cots Kludge) to modify what data is sent out over these last four outputs. For default/nominal settings, the outputs echo the measured klystron outlet water temperature as originally specified. For settings of 1-4, the four outputs instead show all four relevant ASQUIRT02 inputs for one of the four klystrons (the A/B/C/D temperature monitor channels instead show Tin-meas/Tout-meas/Pressure/Flow). For a setting of 5, the four outputs show the result of the calibrated RF diode signals, as digitized and converted to real units via the coefficients and algorithm discussed below. For all cases the outlet temperatures are in same units as COTS: 0-10V corresponds to 0-250C. For all other water parameters, the signals are copied from the ASQUIRT02 received data (0-4095 counts), scaled to 0-10V. For the RF diode calibrated power calculations, 0-10V corresponds to 0-536.74 kW.

COTS Calculation

The COTS calculation is quite simple. We determine $Power\ In - Power\ Out$, and divide by the heat capacity of the klystron collector cooling system to get the dT/dt , i.e. the rate of change of temperature. The power in is simply the beam voltage multiplied by the beam current. The power out is the sum of the forward RF power as determined by a calibration curve on the forward RF diode, and the power removed by the cooling water. The latter is determined by the difference between the COTS temperature and the measured water inlet temperature, and multiplied by the water flow rate. Thus, for a timer interval dt :

$$dT = \frac{I_{beam} V_{beam} - P_{rf} - Q(T_{COTS} - T_{in})}{k} dt .$$

RF Diode Calibration

The first part of the calculation requires us to calibrate the forward power diode signal. The C-Mod team¹ maintains calibration measurements, fit to a cubic with a calibration for the 50dB attenuator used. The form of these calibrations is defined as:

$$P = \frac{(a_1 V^3 + a_2 V^2 + a_3 V + a_4) * 10^{(-a_5/10)}}{10^6} ,$$

where V is the raw diode voltage in Volts, and which is **positive**, and the power P is given in kW.

Some representative values for the coefficients a_1 through a_5 for the forward diodes were recently given as:

~~79.1565 651.3133 256.1627 4.9602 50.0000~~
~~21.1117 580.2051 176.3012 5.2716 50.0000~~
~~20.2457 540.8515 161.5038 3.8028 49.9000~~

and we were instructed to accept these as typical.

With the corrected fits using positive polarity signal, and with the fit more properly constrained for a zero coefficient in a_4 , we are now instructed that typical data are:

1.83 33.07 46.28 0 50.2
1.54 32.46 38.26 0 50.3
1.74 29.08 41.10 0 50.0
1.94 34.00 43.85 0 50.0

Note that this can be reduced somewhat, as the above calibration does not really represent five independent constants, but only four. If we fold the attenuator calibration into the other constants we get a simple cubic expansion. Furthermore, our forward diode signal is digitized with a 12 bit (0-4095) range ~~of 0-2.456V~~, **of 0-3.00V, and we must remember that we have removed the sign flip in the process.** Rolling this all together, we now get something with a more tractable form:

$$P = b_1 N^3 + b_2 N^2 + b_3 N + b_4 .$$

N is now the number of raw counts of the digitized RF signal. (We will keep the unusual ordering of the coefficients, vs. the standard nomenclature where the subscript matches the exponent, to match Atma's calibration documents.)

Applying a few tricks to make this work well for 32 bit unsigned math, we can rewrite this as:

¹ Personal communication, Atma Kanojia

$$P = -c_1 \left(\frac{N}{2^6} \right)^3 + c_2 \left(\frac{N}{2^4} \right)^2 - c_3 N + c_4 \quad ,$$

with N being the raw RF diode counts over a 0-4095 range, and the power P given in mW, the calculation and all intermediate steps maintain maximum accuracy while staying within a 32 bit range. All four of the calibration coefficients are now defined as strictly positive, and the sign is maintained in the form of the calculation.

We will load the coefficients c_1 through c_4 for each klystron via MDS+ tables. These can be loaded prior to each shot, or written into flash memory within the cpu to become power-up default values – confirmatory writing prior to each shot is of course recommended.

The conversion from Atma’s calibrations to these numbers is as follows:

$$c_1 \equiv a_1 * 10^{\left(\frac{-a_5}{10}\right)} * \left(599.75 * 10^{-6} \mu\text{V} / \text{count} \right)^3 * 2^{18}, \text{ typical} = \text{default} = 226,$$

and c_1 will be constrained to be a positive integer $< 2^{12}$.

$$c_2 \equiv a_2 * 10^{\left(\frac{-a_5}{10}\right)} * \left(599.75 * 10^{-6} \mu\text{V} / \text{count} \right)^2 * 2^8, \text{ typical} = \text{default} = 5400,$$

and c_2 will be constrained to be a positive integer $< 2^{14}$.

$$c_3 \equiv a_3 * 10^{\left(\frac{-a_5}{10}\right)} * \left(599.75 * 10^{-6} \mu\text{V} / \text{count} \right), \text{ typical} = \text{default} = 11800,$$

and c_3 will be constrained to be a positive integer $< 2^{16}$.

$$\text{And last, } c_4 \equiv a_4 * 10^{\left(\frac{-a_5}{10}\right)}, \text{ typical} = \text{default} = 465000,$$

and c_4 will be constrained to be a positive integer $< 2^{20}$.

We perform this calibration at the beginning of each COTS calculation, and save the result, which is an unsigned integer in units of one mW per count.

The rather significant changes in the scales of the coefficients arrived at by v010 has been accommodated by simply removing the lower limit of the coefficient entry constraints, and raising the upper limit to that noted above which is required to prevent integer overflow. Calibration checks are ongoing as of this writing.

Temperature model

The temperature model is maintained in a similar fashion – we roll the intermediate coefficients up into single unsigned integers, with care to maintain the range of all intermediate terms as well as the final result. We repeat the COTS calculation at 1 kHz, so the value of dt is fixed. The form of the calculation is:

$$\Delta T = k_1 N_V N_I + k_2 N_Q N_{Tin} - k_3 P_{cal} - k_4 N_Q N_{COTS} \quad .$$

We have already discussed most of these terms.

- N_V is the raw counts of the ASQUIRT01 beam voltage measurement, calibrated at 60 kV = 4095 counts.
- N_I is the raw counts of the locally digitized beam current, calibrated at 25 A = 4095 counts.
- N_Q is the raw counts of the ASQUIRT02 flow rate, calibrated at $N = (Q/100\text{gpm}) * 3276 + 819$.
- N_{Tin} is the raw counts of the ASQUIRT02 inlet temperature, at $N = (T/100C) * 3276 + 819$.
- P_{cal} is the previously calibrated RF power, at 1 mW = 1 count

- N_{COTS} is the model outlet water temperature, which we will define at 0-250 C = **22 bits, or 59.6 $\mu\text{K}/\text{count}$.**
- ΔT is the incremental correction to T_{COTS} , thus is calibrated at the same **59.6 $\mu\text{K}/\text{count}$.**

The calculation is performed as a chain multiply-accumulate-shift with 64 bit accuracy. The additional accuracy is necessary in that very small temperature changes must be allowed for the small (1 ms) time intervals, particularly at the tail end of the exponential decay as the tube cools back down towards the inlet temperature. The four constants, k_1 through k_4 are actually a factor of 2^{32} oversized, so that the extra accuracy is on the correct (LSB) end of the calculation, and so that we may simply harvest the top half of the result at the end of each calculation. (64 bit, or two word, multiply and multiply-accumulate instructions are also primitives for the ARM7 processor.)

Note that there is really only one constant that plays into any of this, all the rest is fixed conversion factors. If we know the heat capacity of the klystron collector coolant system (**in units of 100J/C**) $\equiv C_0$, then

- ~~$k_1 = 50356096 / C_0$~~
- ~~$k_2 = 138448237 / C_0$~~
- ~~$k_3 = 562950 / C_0$~~
- ~~$k_4 = 34603591 / C_0$~~
- $k_1 = 64455802 / C_0$
- $k_2 = 177213743 / C_0$
- $k_3 = 720576 / C_0$
- $k_4 = 346036 / C_0$

Until we have better numbers, the initial default value for C_0 will be assumed to be **11,000 J/C, or a value of $C_0=110$** . We will be able to determine this much better by observing the decay curves for T_{COTS} and T_{OUT} at the end of a shot. The value for C_0 will be maintained by MDS+ for each klystron. As with the RF coefficients, this can be written into flash to replace the power-up default, or can be preloaded by MDS+ prior to each shot for surety. On receipt of the command to overwrite a new value for C_0 , the COTS code will calculate and store the intermediate values for k_1 through k_4 , so that they do not have to be recalculated repetitively in realtime.

The recent realization that the heat capacity units were wrong, and that a typical value was 11,000 J/C rather than 110 J/C, prompted changes in the code. The internal resolution of the temperature model was increased from 15 bits to 22 bits, but the philosophy of doing the delta-T calculation in 64 bit resolution and “harvesting” the top half of this at the end still sufficed. Coefficients were modified to scale this – the 50356096 value, for example, must be 100x smaller to accommodate the larger effective C_0 , but doing only this results in quantization error of small asymptotic temperature ramps, as we are quantized to dT/dt in units of 7 K/sec. By changing the model to 22 bits, we thus divide by 100 then multiply by 128 to accommodate the additional 7 bits, thus transforming 50356096 to 64455802. Our quantized slope is thus 60 mK/sec, which is quite satisfactory, and slightly better than we had before. (Note that k_4 only changed by the division by 100, since it represents a calculation where the internal COTS temperature is on both sides of the equation, thus it doesn’t scale with changing scale of the temperature).

This calculation is continuously iterated at 1 kHz, maintaining the COTS model temperature. Clearly, with no signal for beam current or RF power, the COTS temperature should always converge quickly to the inlet water temperature. Note that we maintain no synchronicity with the HV or RF controls which are separately maintained elsewhere within the TPS hardware – this calculation proceeds blindly and continuously as long as the TPS is powered, with no knowledge of whether or not a shot is taking place.

Fault Threshold

The previous calculation generates a model for the outlet water temperature, thus making a calculated proxy for the actual internal water temperature (which we cannot measure). We would like to set a trip-point threshold which we use to shutdown the HV for the remainder of the shot – one high enough to allow the largest possible window of run parameters (particularly “picket-fence” operation with reasonable duty cycle of RF low or off) – but also within safe bounds to avoid damage to the klystron.

Literature from CPI has given us guidance on the water temperature at the skin layer adjacent to the copper surface in the collector – typically 25% higher than the bulk water temperature (in deg C), thus we need to depress the fault threshold below 100C to account for this. Pressurization of the water, on the other hand, elevates the boiling point, which thus elevates the fault threshold. Finally, prudence dictates that we allow an additional safety margin from the actual boiling point.

~~In equation form, this yields:~~

$$T_{fault} = T_{boil}(P) * k - \Delta T_{margin} \quad . \quad XXX$$

Greg Wallace has recommended an improved fault-determining algorithm, based upon the CPI determinations of the “wet” temperature on the critical surfaces of the klystron collector. We still use a “multiplier” (k), as well as a “safety margin” (ΔT_{margin}), both of which are entered via the MDSplus tables, but both are now relative to the the inlet temperature:

$$T_{fault} = (T_{boil}(P) - T_{inlet} - \Delta T_{margin}) * k + T_{inlet} \quad .$$

A fault condition is defined as the internal model for the COTS water temperature exceeding this fault threshold.

We’ve already discussed how the pressure is measured in the suite of signals received over the ASQUIRT02. This pressure measurement can be used to first calculate the boiling point of the water at the elevated pressure. Reducing steam table curves to a cubic fit, and using our usual technique for multiply-accumulate-shift with single 32-bit word internal accuracy, we use the massaged formula:

$$T_{boil} = \frac{\left[6312 \left(\frac{N_p}{2^7} \right)^3 - 18580 \left(\frac{N_p}{2^5} \right)^2 + 45371 N_p + 55246848 \right]}{2^{15}},$$

Where N_p is the raw pressure measurement (in counts) with the 819 count floor subtracted off (i.e. 0-3276 counts = 0-232 psig), and the boiling temperature is in units of 4095 counts = 250 C. All of these coefficients are fixed. This fit gives better than 1% error over much of the range 100-200C, with the worst error (about 3C) right at 0 psig.

The factor k has a default value of 0.8, and will be maintained by MDS+, either loaded every shot or overwritten in flash memory for future power-on defaults. It is entered as an integer, to be divided by 4096, e.g. the default value is actually 3277. (The MDS+ value loaded must be a positive integer < 4096).

The factor ΔT_{margin} is also maintained by MDS+. It is recommended that this have a minimum value of 5C, due to the errors in the boiling point approximation, and prudence dictates that it actually be considerably larger. The default will be 15C. This is entered as an integer on the 250C = 4095 count scale, thus the default is 246.

This fault temperature threshold is updated in realtime on every cycle that the COTS temperature is updated, and used for determining whether or not a fault condition has occurred.

Fault Logic

The COTS cpu produces five fault (status) bits. There is one bit for each klystron (COTSAOK, COTSBOK, COTSCOK, COTSDOK) which give the status of the COTS temperature model relative to the fault threshold. If the fault condition is tripped for one of these, it will LATCH off, and will not reset until the main cpu on the TPS has issued a RESET command, *and* the over-temperature condition no longer exists.

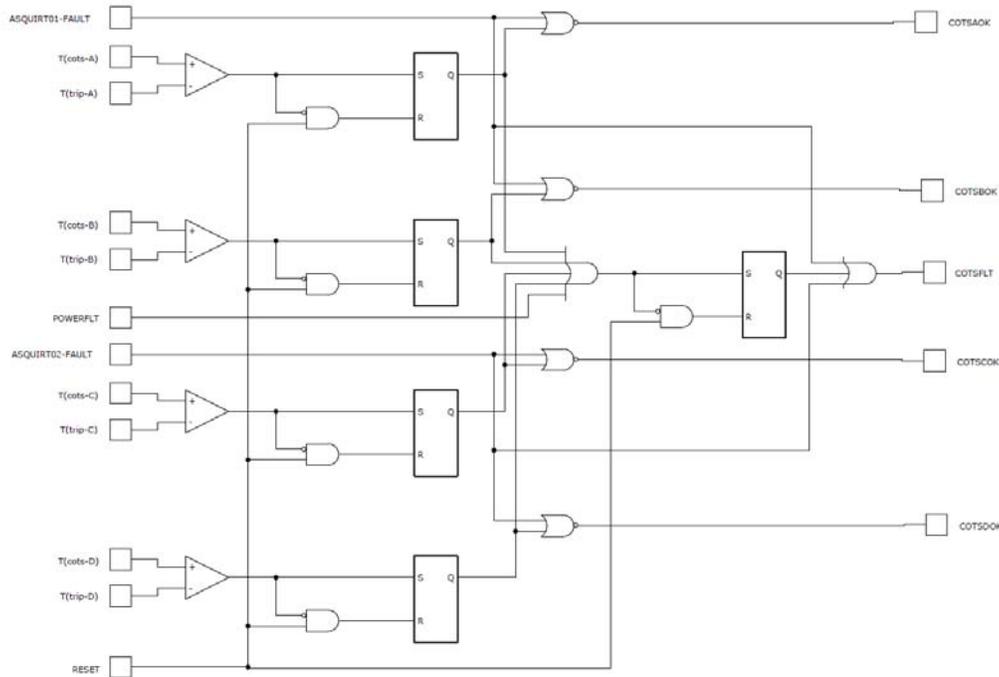
These four bits are also hardwired directly to four bicolor LEDs on the front of the COTS board, and piped to the layer1 board of the TPS, where they are monitored by both the EVENT LOGGER and by the DSQUIRT communications which echo them to the PLC.

In addition, there is a single master fault bit of opposite polarity, COTSFLT. It is latched on if any of the COTSFLT latches go on, or if the local power supply monitor (this board only) shows a fault. It also will only reset on a master RESET with all fault conditions clear. The COTSFLT bit is transmitted to the layer1 board of the TPS, where in addition to the EVENT LOGGER and DSQUIRT communications, it is used in the CPLD logic matrix to generate a “hard fault” which takes down all HV immediately, and for the duration of the shot.

We maintain a different logic for the faults on the ASQUIRT receiver routines. These will fault out if the receivers timeout without a valid packet, or if addresses or checksums don't match, or if there is a format error in the received packet. If either of these receivers fault out, we express the COTSFLT, *but without latching*. Similarly, we turn the COTSAOK and COTSBOK bits false *without latching* if the ASQUIRT01 faults out, and we turn the COTSCOK and COTSDOK bits false *without latching* if the ASQUIRT02 faults out.

This serves the following purpose:

- we generate a COTSFLT which inhibits all HV if there is any ASQUIRT failure, which is as it should be, as the COTS model temperature is no longer valid,
- we have an unlatched LED indicator at the front panel to indicate the validity of the communications link with the ASQUIRT01 and ASQUIRT02 senders – these will flicker red to green as the fibers are plugged in and the communications synchronize.



Testing

NOTE: All of the tests below were performed for the August 2011 writeup, and the coefficients and algorithms are in some cases superceded by the changes noted above.

For tests, we load in all the defaults mentioned above *except* for the fault threshold coefficients. For those, we assume a k multiplier of 1.0, and a margin of zero. We will thus begin with the fault trip point equal to the boiling point calculation.

In addition, we bypass the ASQUIRT02 readings, and superimpose fixed numbers for the following:

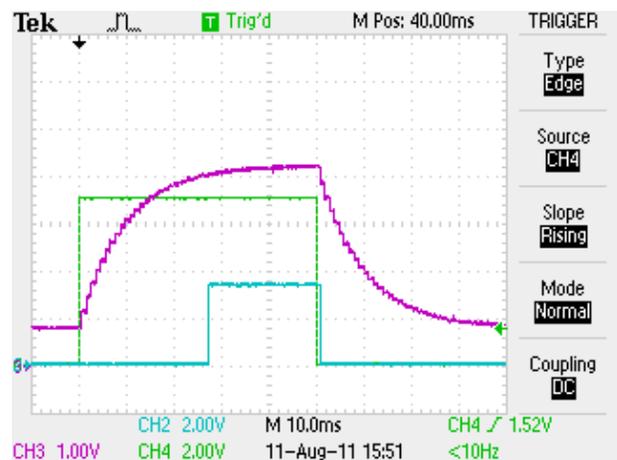
- $T_{in} = 20C$
- $Q = 38 \text{ gpm}$
- $P = 0 \text{ psig}$ (thus the fault trip should be 100C)
- $T_{out} = 40C$

Similarly, we bypass the ASQUIRT01 readings, and superimpose fixed numbers for

- $V_{beam} = 50 \text{ kV}$

We will apply no signal to the forward RF signal. Our only variable for this first test will then be the beam current. We will apply a signal from an arbitrary waveform generator to the I_{beam} input for channel A, and simultaneously monitor the analog COTS signal. We will also temporarily suspend the latching nature of the COTS fault bit, and echo the COTSA fault out to a spare pin on the cpu, which we also monitor.

In the shot at the right, we see a signal on channel 4 (green) applied to the I_{beam} input of the COTS board for



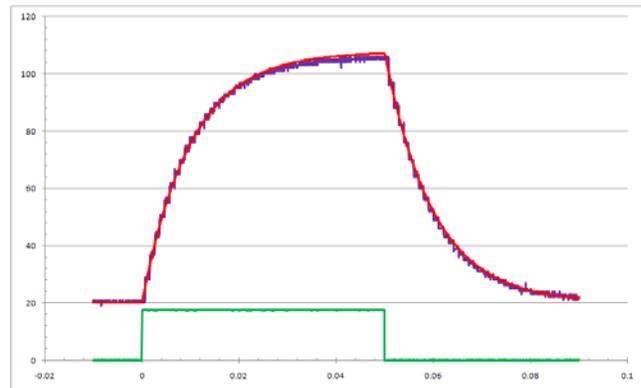
channel A. It is a pulse 700 mV in height for 50 ms. This corresponds to 17.5 A, or at 50 kV, to 875 kW of input power.

The steady state cooled temperature has no free parameters at all - only the transients depend on our assumptions for the heat capacity. Applying the steady state relation that $V*I = Q\Delta T$, we quickly calculate an expected temperature rise between inlet and outlet of 88 degC.

The purple trace (channel 3) is the COTS model temperature for channel A, which rises from 800 mV to 4.20 V. At 10V = 250C, this translates as rising from a base of 20C to a steady state of 105C, for a rise of 85C. Check (close enough).

The blue trace, channel 2, is monitoring the temporary test point which echoes the (unlatched) \COTSAOK bit. Careful inspection shows that this is acting correctly – as a comparator for the COTS temperature relative to a reference of 100C. Check.

If we capture the same data to a spreadsheet and run the transient model, we can verify the dynamics of the rising and falling exponentials. In this case, we have the default coefficients derived from a heat capacity of 100 J/degC. The green and purple traces are the same (now in Amps and degC, respectively), while we have added a red trace calculated from first principles with the above numbers. Check.



Other contributions to the calculation are checked by, one at a time, removing the debug code which overwrites the input values from the source, and applying a signal to the appropriate input (in the case of the ASQUIRT0x modules, to the inputs on the remote sending unit).

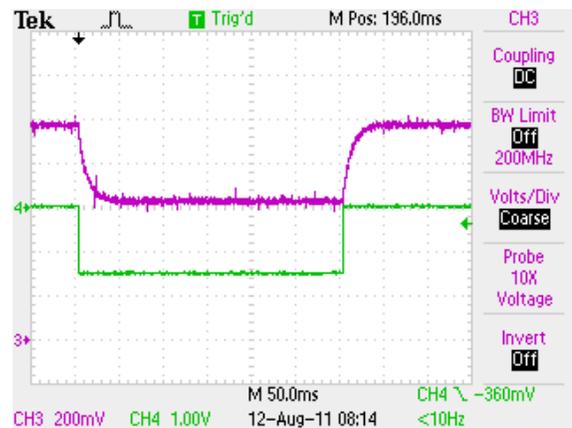
- For a 10 psig signal (overwritten), the boiling point should rise to 116C. At $k=1.0$ and zero margin, we now observe the fault threshold at 116C
- Entering $k=3277$, for a factor of 0.8, we should observe the fault threshold at 93C. Check.
- Leaving $k=0.8$, we raise the margin to our 15C (246), which should drop the fault threshold to 78C. Check.
- Leaving $k=0.8$ and margin of 15C, we eliminate the overwrite of the measured pressure and use the ASQUIRT02 value. At 5 mA, corresponding to 14.5 psi, the boiling point should be elevated to 121C. That should raise the threshold to 82C. 81C observed. Check. At 15 mA, corresponding to 160 psi, the boiling point should be elevated to 187C. That should raise the threshold to 135C. 132C observed. Check.

We now leave zero signal into the pressure input, which should leave the boiling point at 100C (this however is the point of worst case error in the cubic fit, and it instead is about 103C). For k and margin as set, the fault trip point should now remain fixed at 67C. Check.

- We now remove the debug overwrites for all five of the temperature readings over ASQUIRT02. Supplying a signal to the input temperature channel, we now correctly see the baseline for the COTS model temperature shift up and down. At 8mA, corresponding to 25.3C, we correctly see a 25.3C baseline.
- Next, we remove the debug overwrites for all four of the flow readings over ASQUIRT02, and apply a signal to the remote ASQUIRT02. This, unsurprisingly, is the most profound change. At 18mA signal, which corresponds to 87 gpm cooling, a 100 ms 750 mV pulse into the I_{beam} input brings the temperature to

only 40.5C (from a 0C baseline). 750 mV corresponds to 18.75A, or 938 kW at our fixed 50 kV. At 87 gpm, the steady state power balance requires a 40.8 degree ΔT . Check. At the other end, we crank the cooling way down, for only a 5mA signal, corresponding to 6.8 gpm, we must crank the beam current way down so as not to saturate the COTS model at 250C. Similarly, it takes much longer to reach steady state. We fix on a 300 mV pulse for 300 msec. This corresponds to 7.5A of beam current, or 375kW. The steady state power balance requires 210 degrees of ΔT . From a 0C baseline, we observe 210C. Check.

- Next, we put back the debug overwrites for flow at 38 gpm, and we remove the overwrites for beam voltage from ASQUIRT01. Applying 5V to the ASQUIRT01 input corresponds to 30kV. Pulsing the I_{beam} input with 500 mV (12.5A), or 375kW should give us a 37.4C temperature rise at 38 gpm. We measure 37C. Applying 9V to the ASQUIRT01 input corresponds to 54kV. For the same beam current, we should have 675 kW, and a 67.3C temperature rise. We observe 68C. Check.
- Finally, we apply a *negative* pulse to the forward RF diode input. This has the curious and unphysical interpretation of RF power being extracted when there is no beam, thus magically cooling the klystron. In the shot at right, we are applying a steady 8mA to the ASQUIRT02 input temperature terminal, thus giving us an elevated baseline of about 25C. We apply a 300 ms negative pulse of 1.5V (the green trace), which should calibrate to be 85 kW of RF power out. At our fixed 38 gpm, this should give us a ΔT of 8.5C. We observe 8.4C. Check.



Summary

These tests will be used to check out each of the boards delivered, and repeated over all 4 channels on each board. The COTS board is only part of the TPS assembly. The main TPS board serves the all-important function of monitoring all fault signals for conditions which require shutdown of RF, HV, or both. We process the RF and Beam Current signals on the COTS board because these signals are needed for local digitization for the COTS algorithm. In addition, these signals require fast comparators to sense fault levels, and the digital outputs are transmitted promptly back to the main board.

Additional functions that are provided by the COTS cpu, besides those discussed above, include:

- Controlling the multi-channel DACs which generate the MDS+ programmable fault-thresholds for beam current and RF reflected power ratio,
- Pass-thru of a master ‘heartbeat’ signal for watchdogging the TPS intelligencia,
- Communications (over a dedicated SPI port, which the COTS cpu is a slave to) with the master cpu on the main TPS board, and interpretation of commands passed by MDS+ (e.g. modification of setpoints, etc)
- Read/write of FLASH memory as needed, or as directed by MDS+ command.

The latter point requires some additional background. This microprocessor contains both SRAM and FLASH memory within its 32 bit address space. The assembled code is written to FLASH, and fixed calibration constants are embedded within the code. Modifiable calibration constants are kept in a separate block of FLASH memory which is copied to a mirrored SRAM block at powerup. The SRAM values are used for reference during code execution. Two external commands from MDS+ (via the main communications cpu on the layer1 board) will service this FLASH memory. One will generate a set of factory default coefficients, and

load these into the SRAM block. Another will copy the SRAM block to FLASH. Note that on first programming this cpu, these two commands will be executed at the factory. In future use, MDS+ may choose to overwrite various calibrations (which will overwrite the SRAM values) and then execute the SRAM to FLASH transfer, which will embed the new values for use on all subsequent power-up defaults. The other chips will inhibit the SRAM to FLASH write while a C-MOD shot is active, as this shuts down all cpu activity for about 250 ms, hence is unsafe during any HV or RF operation.

The algorithms and operation of the entire assembly will be discussed in another document. This document is being released in draft form at this time (August 2011) for review and discussion with the LH team.

Questions and comments should be directed to Jeff Casey at Rockfield Research Inc., see header for contact information, or email to casey@rockfieldresearch.com.

APPENDIX B – SBIR Final Report

On the following pages, you will find the final report from the DOE SBIR contract, which presents a more general overview of the background, use, and features of the TPS, as well as many of the peripheral circuits used in the Alcator C-Mod TPS system.

U.S. Department of Energy SBIR Final Report

Phase 1 SBIR (2007)

Phase 2 SBIR (2008)

Phase 2 Supplemental Program (2010)

“Improved Controls for Fusion RF Systems”

P.I. Jeffrey A. Casey, Ph.D.

8 November, 2011

Rockfield Research Inc.

11010 Onslow Court

Las Vegas, NV 89135

(702) 487-6970

Grant Award # DE-FG02-07ER84762

1.0 Executive Summary

The DOE Fusion Energy Sciences program supports a number of objectives with the common goal of expanding the science of heating and controlling thermonuclear fusion plasmas for energy production. Among these objectives are further development of several methods for plasma heating and current drive using radio frequency (RF) energy - including Lower Hybrid Current Drive (LHCD). The promise of LHCD is high, as it is a means by which tokamak plasmas can be sustained non-inductively.

For typical tokamak parameters, a LHCD system must supply several MW of power at a few GHz frequency - and for current drive, this must be presented to the plasma as an array of sources with tightly constrained phase relationships. The controls are necessarily complex, requiring sophisticated timing, phasing, protection, and monitoring. As the understanding of fusion plasmas evolves and the experiments become more complex, the demands of the controls for such a LHCD system evolve as well. In parallel, as the numbers of subsystems in a fusion plasma experiment increase, the demands of reliability for each system must increase as well.

Existing fusion experiments have kept pace with, or often led, improvements in modern electronics to meet these demands of complexity and reliability. These improvements, however, are not always applied uniformly across all subsystems. In many cases, older technology is still in use for reasons of budget, priority, or simply the availability of legacy equipment. As the world fusion program moves towards ITER, the motivation is high to modernize all such complex controls to yield the optimum in utility, reliability, safety, and ultimately cost.

In this effort, we have addressed the specific requirements for the integrated systems



Figure 1. The suite of five complete TPS core units (front assemblies and rear transition module assemblies) on the test bench after final testing, ready to ship to MIT.

controlling an array of klystrons used for LHCD. **The immediate goal for our design was to modernize the transmitter protection system (TPS) for LHCD on the Alcator C-Mod tokamak** at the MIT Plasma Science and Fusion Center (MIT-PSFC). Working with the Alcator C-Mod team, we have upgraded the design of these controls to retrofit for improvements in performance and safety, as well as to facilitate the upcoming expansion from 12 to 16 klystrons. The longer range goals to **generalize the designs in such a way that they will be of benefit to other programs** within the international fusion effort was met by designing a system which was flexible enough to address all the MIT system requirements, and modular enough to adapt to a large variety of other requirements with minimal reconfiguration.

1.1 Extended Background

The legacy LH system on C-Mod was installed earlier using conventional hardwired protection logic. There were several known disadvantages with this type of a system, offset by the simplicity of adapting known circuitry from other similar (e.g. radar) facilities. Difficulties included inaccessibility of controls for parameter adjustments (e.g. fault thresholds), long runs of cables in and out of HV areas (both increasing noise pickup and adding risk of personnel safety), and large physical footprint in the experimental cell. The goal of this program was to adapt modern electronics capabilities and experience learned in other modern phased array radar facility design to directly address these issues, creating a compact and reliable superset of the existing controls functionality. As an additional bonus, the space savings resulting from retrofitting the existing 12 klystron array with these upgraded controls considerably eases the expansion of the facility in the upcoming addition of four more klystrons for a 16 klystron LH system.

This Phase 1 effort and the initial stages of the Phase 2 effort were used to fully specify the requirements of the upgraded controls. Initially, we left an open topic for consideration as to whether to use a fully digital architecture or a hybrid analog/digital signal path – a very early decision was made to opt for the hybrid architecture, as this is inherently more robust as the fault comparators are on robust and independent signal paths, thus key equipment safety sensing circuits are implemented in a nearly failsafe manner.

Another key early decision was that the architecture would service four klystrons per unit, following the “cart” topology of the physical installation. The sheer quantity of signals that must be serviced in a four cart protection system, however, is prohibitive, particularly for a compact central control unit. That, plus the desire to reduce long cable runs to the greatest extent possible led to another early decision – we would modularize a great deal of the input/output hardware, and site a much of it at remote stations which multiplex groups of signals and transmit this over fiber-optic isolated communication lines.

During the middle of the Phase II program, we thus finalized design of a number of different I/O stations which would serve as isolated communication links, and moved into manufacturing of these even before the remainder of the system was fully specified...this gave a parallel path for both design and manufacturing, and helped constrain the final design by adhering to the chosen architecture of the I/O systems.

The core transmitter protection system (TPS) for four klystrons (one cart) was chosen to reside in a cPCI crate topology for co-location with the fast dTACQ digitizers used for data acquisition in the LH system. We did not choose to utilize the digital backplane of the cPCI

crate as a computer backbone *per se*, but instead used it as a mechanical structure and power supply. Communications to the Alcator C-Mod data system (MDSplus) was over an on-board Ethernet port directly installed in the TPS hardware, thus bypassing the backplane of the cPCI crate. We did make use of the expansion architecture of cPCI – which allows for a “rear transition module” (RTM), plugged into the backside of the crate with direct access to auxiliary contacts on the backplane connectors. This is usually used for I/O expansion, and in similar spirit, we used this for nearly all F/O connections, feeding these strictly digital signals across the cPCI backplane to the main TPS unit.

1.2 Program Schedule Modifications

During the second half of the program, we had undergone a number of delays within the Alcator C-Mod program, not the least of which was the temporary downchecking of the motor/generator flywheel, resulting in both calendar delays and unavailability of the MIT half of the team for design of some of the interfaces, originally assumed to lie on the MIT side of the collaborative effort. As the Phase II effort drew to a close, we submitted a proposal, granted by DOE, for an additional year of funding. This supplementary program served both to extend support during installation and integration, and to shift the load of design and manufacturing of a number of these auxiliary interface circuits – customized for various remote substations to provide local signals, filter and sense inputs, and standardize to the remote I/O format for transmission to and from the TPS central unit. During the first half of this additional year, these aux interfaces were designed, manufactured, tested, and shipped to MIT for early phases of the install.

In the second half of the extension year, we finally accumulated all the minor specification changes we had defined along the way and froze the design of the central TPS units. The six board assemblies were fabricated, embedded firmware written, and hardware delivered to MIT. The completion of this install, coinciding with tokamak downtime in late summer and fall of 2011, provided the opportunity to uninstall the legacy LH control systems and complete the install of the new controls. As of this writing (early November 2011), all of this equipment has been delivered to MIT and installation is ongoing, with initial *in situ* operation to begin phasing in during December 2011 and January 2012.

Rockfield Research has committed to fully support this installation and deployment, in spite of the technical termination of contract during mid August 2011. In fact, underestimate of the scope of the auxiliary circuits led to spending out the contract during June/July 2011, and Rockfield Research has continued on internal funding during subsequent efforts. The separate financial final report details the breakdown of the extent of post-contract spending.

1.3 Program Impact & Commercialization

This effort has developed a general purpose modular protection system for multi-VED (vacuum electronic device) control systems. The specific implementation delivered to MIT for the Alcator C-Mod LH controls were configured to fit that application, yet the hardware architecture design is quite open-ended, and heavily software configurable.

Typical protection systems for high voltage VED systems are custom designed and application specific. These circuits are critical – in some cases providing the only line of defense

against single-point failures for VEDs with replacement costs of several hundred thousand dollars each.

We have designed a general purpose multi-VED system which is a superset of the usual capabilities of such protection systems, and is very reconfigurable. In the fusion community, application for any such heating system is likely to greatly reduce the overhead development costs of installation, reduce time to operation, and improve both personnel and equipment safety.

It is the Rockfield Research Inc. business model to offer sale and support of a limited smorgasbord of specialized high end electronics. The manufacture and testing of such items can be largely outsourced, thus the **primary product is the support consultation and intellectual property invested in the design**. A span of a few years with zero sales will not damp the enthusiasm for continued support of such products – a sentiment unlikely to be shared by conventional electronics manufacturing and sales firms. The TPS system (and auxiliary multiplexed communications devices) have now been added to the Rockfield Research catalog, and will be supported for future customers as needed.

2.0 Architecture of the TPS

Figure 2 on the following page shows the updated block diagram of the one-cart (four klystron) TPS. The distributed nature of the TPS topology is apparent in this view – the central vertical block is the core TPS, which contains the fault logic, analog signal buffering, fault comparators, timing, and Ethernet interface. Much of the I/O, however, has been located at remote locations – sited adjacent to the source (or receivers) of the appropriate signals, with isolated multiplexed F/O links between. These remote I/O links are noted and labeled in red here.

This architecture went through many iterations in the early part of the Phase II program, and it includes the detailed specification of the interface between Rockfield Research and MIT – in regards to both the signals *and* the responsibility for interfaces. The TPS was initially agreed to comprise the core unit plus the I/O links shown and labeled “Rockfield”. Several blocks, labeled “MIT” represented custom interface circuits (usually quite small and application specific) to be built by MIT.

These links, seven in total, were deferred by MIT due to resource limitations. One of them (HV Interface box) was designed by the MIT team, as it was needed for upgrades to the legacy LHCD system with the refurbished CPI tubes. The remaining hardware links were split between the MIT team and additional effort by Rockfield under the supplemental funding extension, and are discussed in more detail in later sections.

2.1 TPS Core Unit

The TPS core unit resides in a cPCI crate adjacent to the fast digitizers, with one TPS and one digitizer per cart (4 klystrons). This unit serves as the “brains” in the transmitter protection, receiving instructions from MDS-plus over Ethernet before each shot, monitoring status lines and PLC permissives in realtime during shots, and managing fault status with sub-microsecond response time. Even with the intense effort at distributing I/O among remote stations, the I/O to this unit is a significant constraint to size and function. It is designed as a triple-width cPCI unit, with additional I/O brought in through a triple-width RTM (rear transition module) in the back of the cPCI crate. Features include:

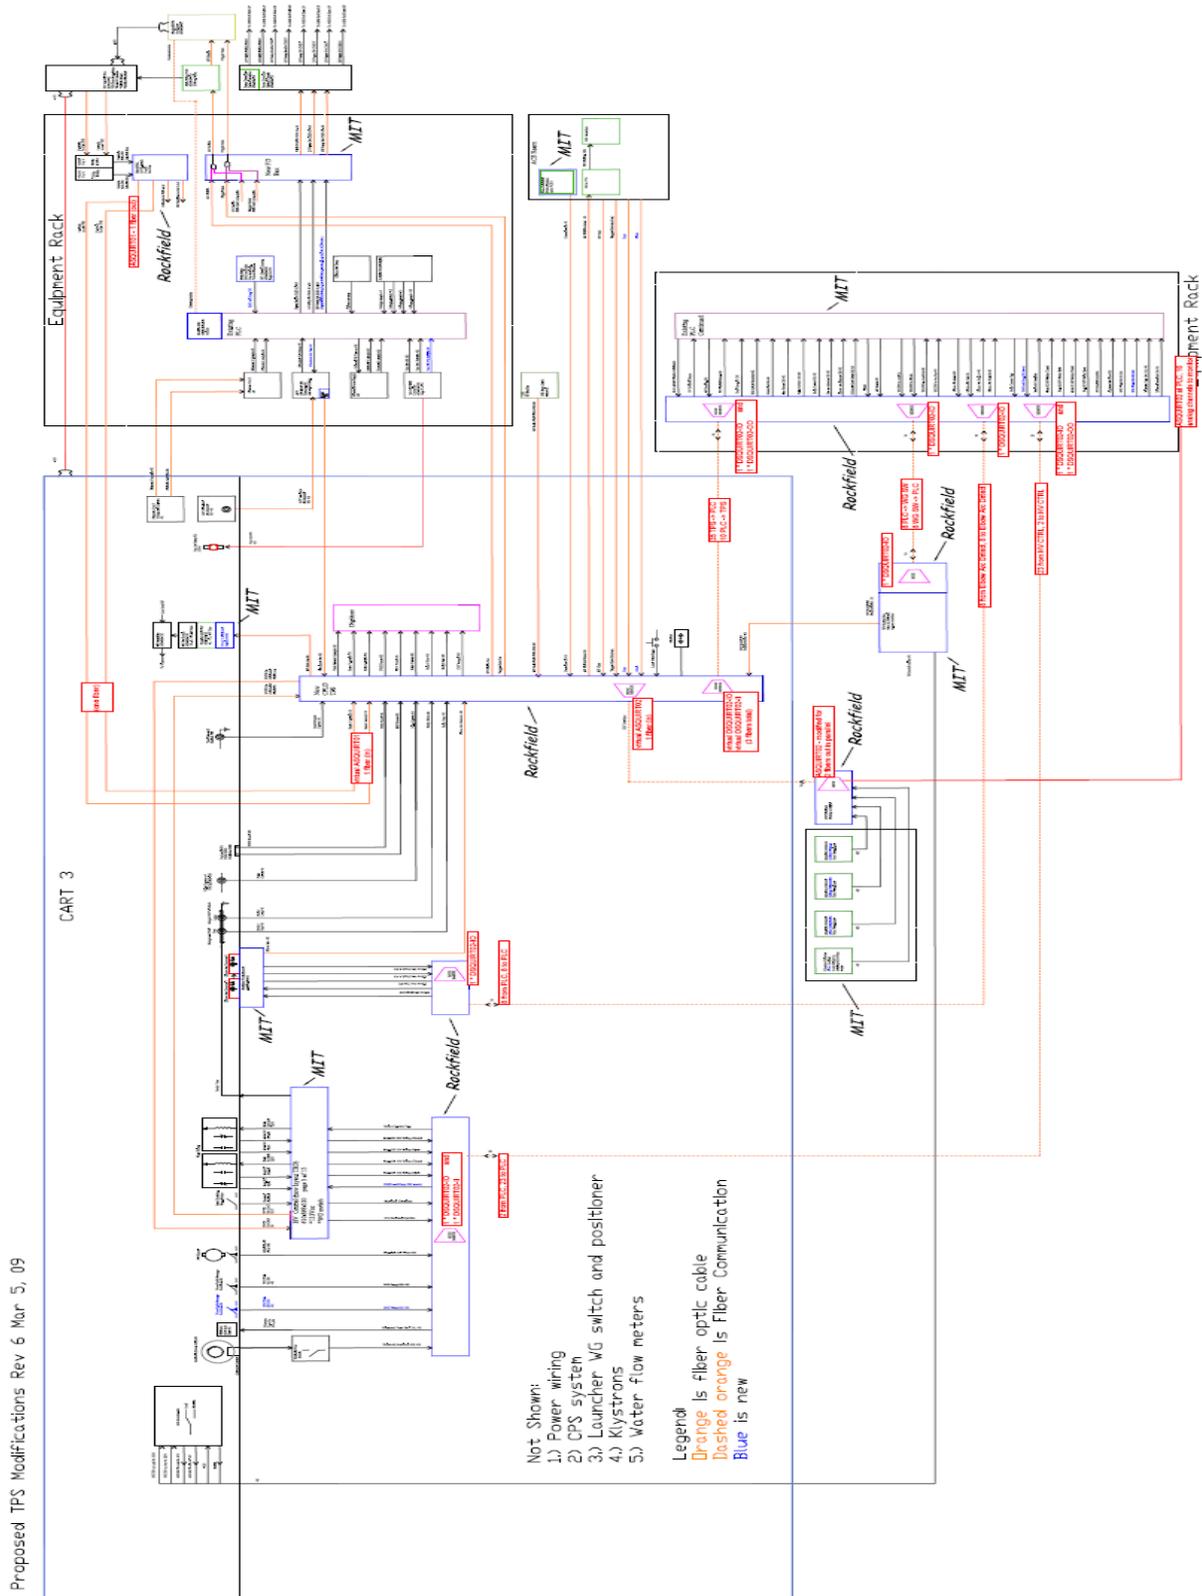


Figure 2. Block Diagram of a 4-klystron “cart” with TPS and interfaces.

- The main front board has 12 channels of analog input – these service two body currents per klystron, one ground current per cart, and three spare analog inputs. Each analog input has three comparators with programmable thresholds and inhibit delays.

- The main board (Figure 3) also has the fault FPGA matrix, the communications cpu, auxiliary memory for event logging, inhibit timing logic, Ethernet interface to MDS+, and cPCI backplane connections for power and interface to the RTM boards.

- The center front board (Figure 4) has 12 additional analog input channels – these service the one beam current per klystron, plus forward and reflected power sensors for each klystron. The beam current is processed similarly to body currents, while the forward and reflected power are compared ratiometrically for faults. Fiber optic inputs for the multiplexed analog data highways are on this board as well.

- In addition, the center board *digitizes* the beam current and forward power (albeit at relatively slow rates) for use in the COTS calculation. This is combined with the water temperature and flow measurements and the beam voltage measurement (received via analog links) for predictability of the realtime collector temperature. An analog temperature predictor is generated as a new data signal (to external digitizers), and sent to a fault comparator with programmable fault threshold.

- The third front board (Figure 5) is an analog repeater board, which echoes



Figure 3. Five of the main boards are shown above. Each of these boards services 12 analog channels (eight body currents, 1 ground current, and 3spares). In addition, there is a repeater port for MDSplus start and clock signals, an Ethernet port, and reset button and LED status display. Both the master CPU and Event Logger CPU, in addition to the programmable fault logic matrix are resident here.

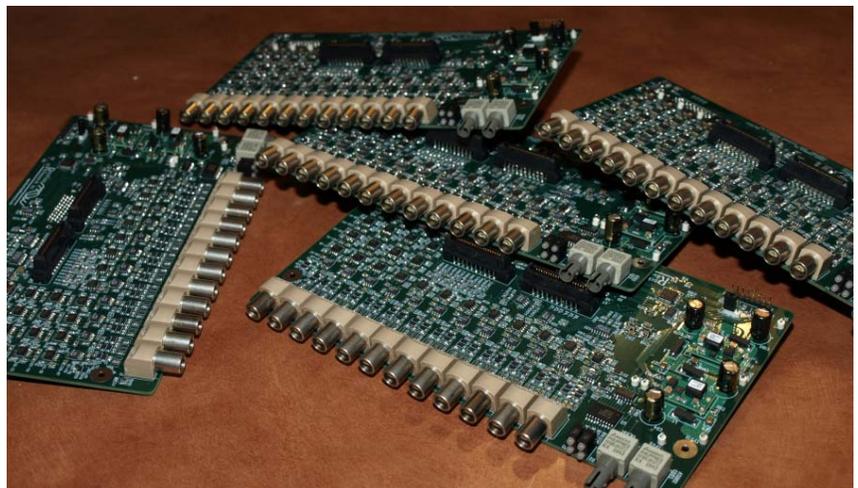


Figure 4. Five of the middle (COTS) boards are shown above. Each of these boards has 12 analog channels (four beam currents, and four pairs of RF diode inputs), as well as F/O inputs for the two ASQUIRT inputs. The COTS cpu is onboard here as well.

30 separate analog signals to front panel testpoints as well as to a 68-SCSI connector with pinout to match the input to the d-tAcq digitizer.

- The triple RTM module (Figure 6) has 42 F/O connectors – 36 of which have pre-allocated functionality, and 6 of which are spares (and configurable as either input or output).

Board-to-board mating uses hard connectors (not cables), which comprise part of the 3D assembly of boards. One of the three boards in each set (front and back) mate to the cPCI backplane for power and communications.

On the following page, Figure 7 shows the overall topology of the 6-board core TPS unit. On the left are the “front” board stack, from the bottom up:

- Main board, with Ethernet communications, fault logic, half the analog signals, data logger, and cPCI backplane connection to RTM modules,
- COTS board (middle), with the COTS cpu, and half of the analog signals,
- Monitor board (top), with buffered monitor points for debugging and SCSI168 connection to the fast digitizer.

On the right side of Figure 7, we see the three-board stack of the RTM, which contains little besides 42 fiber-optic I/O channels. All of the on-board intelligence to interface to these signals resides on the main “front” board.

In the following sections, we will discuss many of the key features of the core TPS unit, broken down by function rather than physical board location.



Figure 5. Five of the analog repeater boards (the third, or top board of the front assembly) are shown above. Analog signals, scaled and buffered are brought onto this board via board-board connectors, repeated via on-board buffers separately to single-pin LEMO test points and to a SCSI168 connector to mate with fast digitizer.

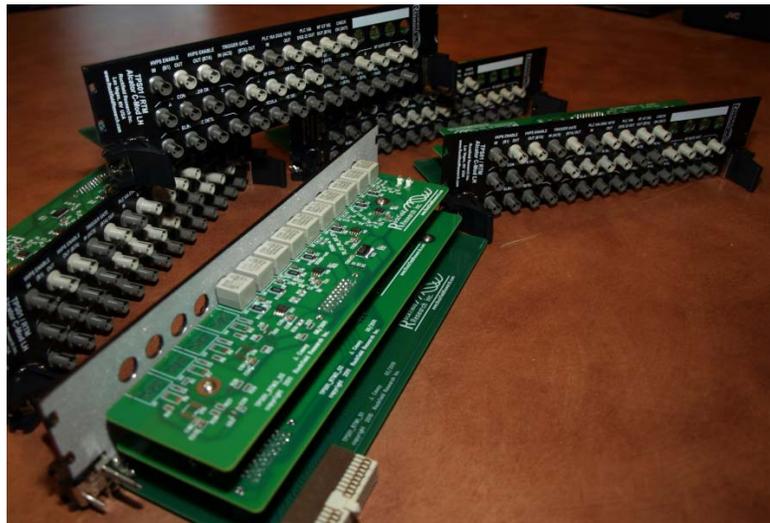


Figure 6. Five of the six three-board RTM assemblies manufactured are shown above. The lowest of the three boards mates to the cPCI backplane, as well as mounting board-to-board connectors to the other two boards. The panel hosts 42 fiber-optic I/O channels.

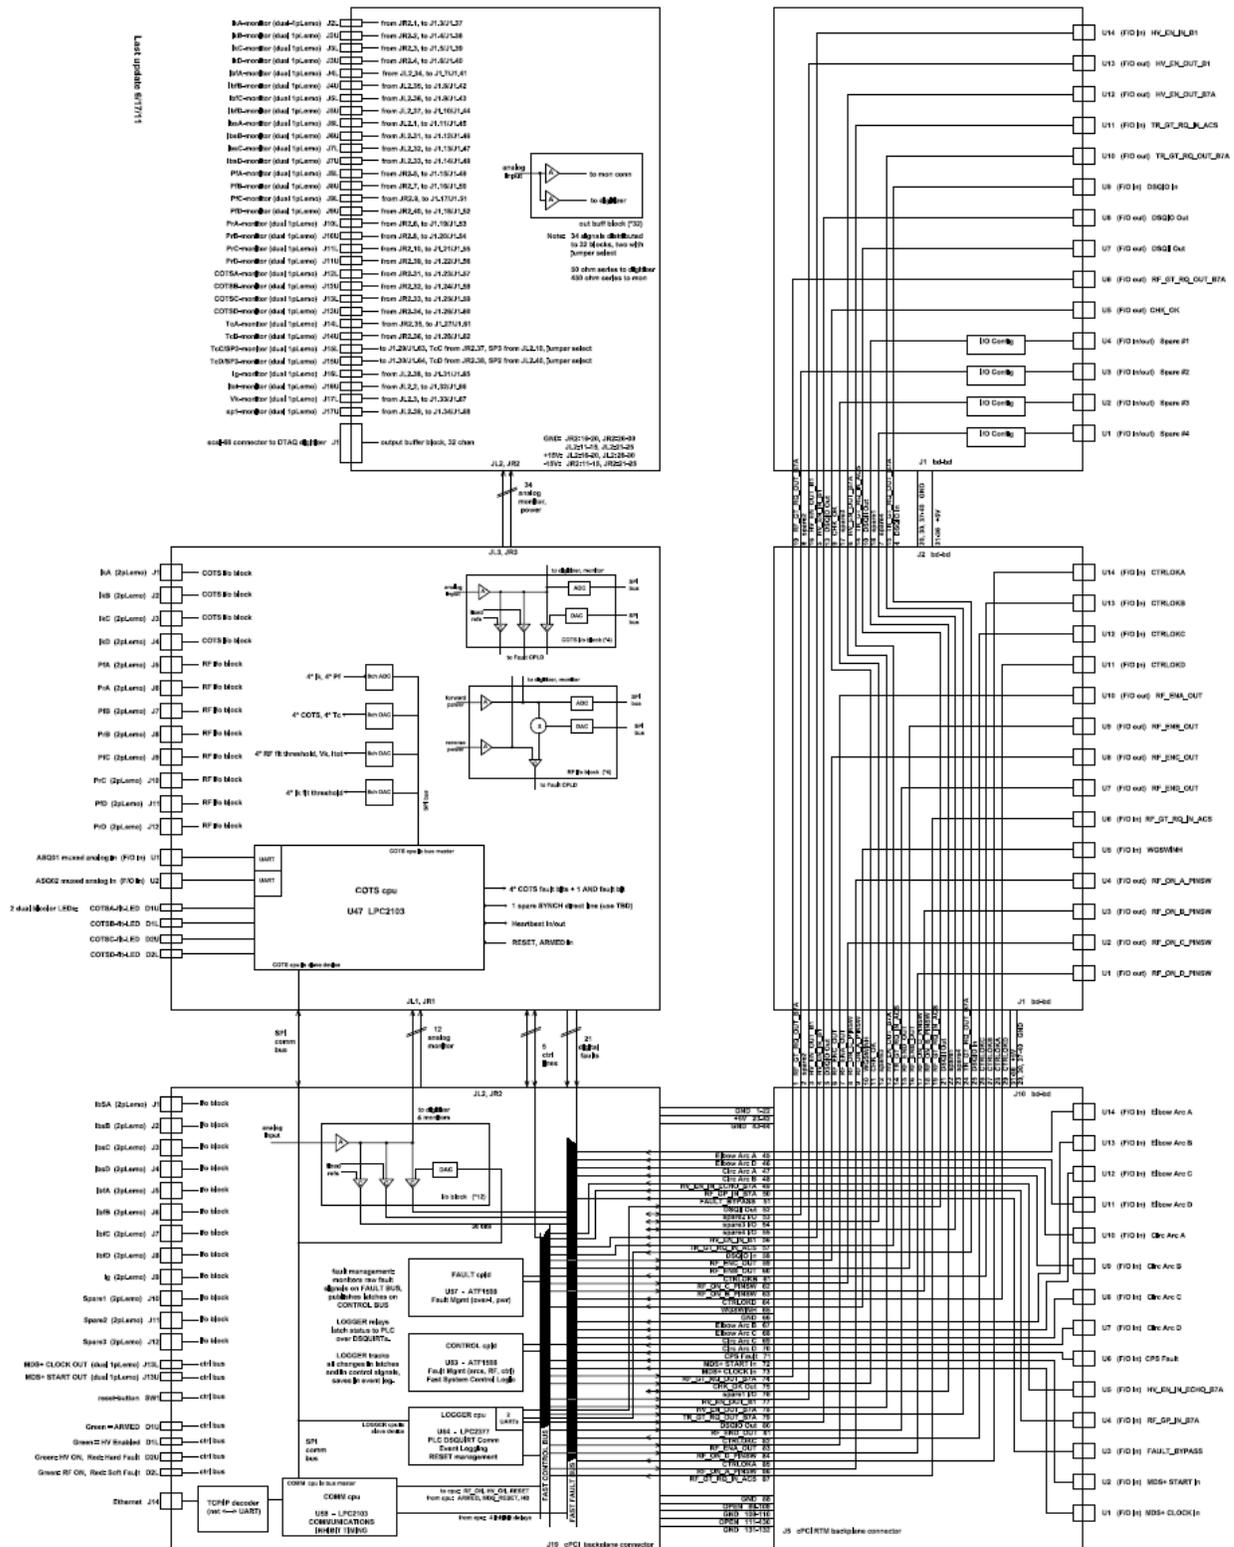


Figure 7. Block Diagram of the 6-board Core TPS architecture.

2.1.1 Collector Over-Temperature System (COTS)

The COTS – collector over-temperature system – is a digital model for the klystron collector temperature. It is necessary because the klystron vendor, CPI, notified MIT rather late in the game that the klystrons were in fact not sufficiently well cooled to accept full pulse-length unspent beam on the collector without incurring catastrophic damage. As the klystrons are quite expensive, this was deemed a serious matter. The klystrons are electronically vacuum diodes – high voltage is placed on the cathode, and a current is extracted towards the collector (according to the Child-Langmuir $V^{3/2}$ dependence, or perveance, typical for vacuum diodes). At full voltage (and thus current), this beam will present the full HV power on the collector surface, which must be removed by cooling water with neither copper vaporization nor water boiling. In fact, the klystron is an RF amplifier, and RF energy is emitted through the output coupler in proportion to the RF input. At reasonable efficiency, there is considerably less power presented to the collector than that emitted by the cathode, thus the cooling load is lessened.

There are several types of faults that the klystron must be protected against. Any internal arcing within the vacuum path must be responded to by immediately removing the HV power, which results in a shutdown of the HV. Lesser faults, however, such as air breakdown in the waveguides resulting in waveguide arcs, can be dealt with more easily by simply shutting off the RF input to the klystron. This is preferable, since we can wait a short time for the fault to clear (tens of ms), then re-establish the RF input to the klystron and continue operation after only a very brief interruption. Shutting down the RF, however, increases the energy in the spent beam, thus increases the energy load by the cooling.

At present, the MIT team must limit the RF pulse length to the worst case conditions – where RF is off for the entire shot, and the collector temperature is nevertheless safe. This is unsatisfactory, as longer pulsing is desired for physics experiments, and the cooling load is actually considerably less in practice since the RF is generally on. A collector temperature measurement would provide a simple answer, as we could interlock and shut down HV if the temperature approached the danger level – yet this is not possible, as the temperature is worst on deeply buried inside surfaces of the klystron. The difference between safe and disastrous are in fact contained within the time for water to transit from these deeply buried surfaces to the nearest available port for a thermal sensor – so we cannot safely measure downstream water temperature and shut off in time.

The COTS was thus designed as a predictor of the collector temperature. We gather all the inputs for power *into* the klystron: beam voltage and beam current, and calculate this load. We also gather all the inputs for power removal *from* the klystron: water temperature and flow, and RF power out. The difference must result in a net rise in the collector temperature – the rate of rise being determined by the heat capacity of the collector with its water load. Additionally, we can calculate the additional safety factor from the elevated boiling point of the water by measuring the water pressure as well. All of these signals are presented to the TPS in realtime, and the model maintained in realtime --- at all times, in fact, regardless of whether or not a shot is in progress. There are no adjustable parameters for the key power balance, and only the heat capacity of the tube is relative unknown, but that plays only into the dynamics, and is rather easy to match by comparing temperature decay rates after a pulse once the system is installed.

As designed, we dedicated an ARM7 microprocessor just to accommodate the COTS calculations for the four klystrons serviced by a single TPS unit. This microprocessor talks to

two analog multiplexed remote I/O units (one which provides beam voltage, the other which provides 16 channels of water temperatures, pressures, and flowrates). It also directly digitizes the discrete beam currents and RF diode signals for each of the four klystrons. Gathering all this information, it updates the model for the klystron internal temperature at 1 msec intervals. MDSplus settable parameters, as well as fault trip points, are maintained in this processor – and it generates dedicated digital fault channels for each klystron to signify that the model temperature has exceeded the allowed setpoint. These fault bits are treated like any of the analog path hardwired comparator outputs – fed to the programmable logic matrix for sub-microsecond fault response, with removal of HV permissives as a response.

COTS Implementation

The COTS algorithm is coded in ARMASM to run in realtime on an NXP (Phillips) LPC2103 ARM7 microprocessor. No operating system or high level compiler is used, thus increasing speed and reliability (at the cost of coding effort.) This is a relatively fast 32 bit micro, but has no floating point math skills. Some pre-processing has been done in coding, and is necessary in converting from MDS+ setpoints to shot-loadtime setpoints in order to streamline the calculations. No divide or transcendental math are performed in realtime – all such calculations are pre-fit to polynomial expansions over the range of interest, allowing exclusive use of hardware instructions for multiply, accumulate, and shift.

The COTS calculation is performed on the 2nd (middle) of three boards in the front TPS assembly (Figure 8). There is one TPS assembly on each cart, thus each one services four klystrons. The beam current transducers and the RF diodes are processed on this same board for proximity to the COTS usage.

Each of the beam current (cathode current, klystron current) signals is generated by a FW Bell RS-100, with a raw signal calibration of 40 mV/A. It is oriented for positive signal under normal operations.

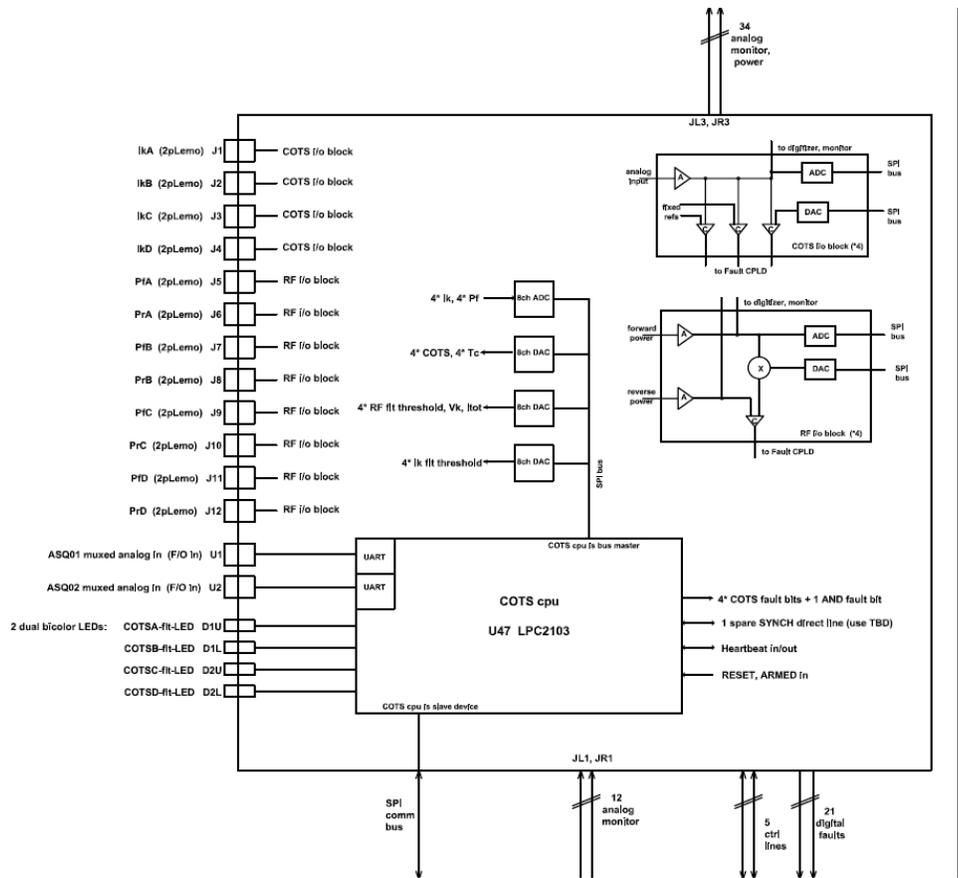


Figure 8. Expansion of middle-front board from Figure 7. The COTS model dominates this board, and the analog signals serviced on this board are the ones needed for the COTS calculation.

The analog processing section of the COTS board performs a clean-up (transient glitch protection, 1MHz *RC* filter), then buffers this through a 25 MHz op-amp (LT1358) with a differential gain of 5x. (The signal is treated differentially on input from a two-pin LEMO connector, although it is realized that the signal is technically single-ended). This buffered signal is fed to three comparators for fault detection (one to detect backwards-oriented installation, one with a pot-settable trip-point which is always active, and one with an MDS+ settable trip-point which is inhibited for a brief inhibit delay after the HV is turned on). The signal is then fed to several analog points of interest: 1) buffered and passed to a dTacq digitizer connector, 2) buffered and passed to a single pin LEMO testpoint, 3) routed to a local 8-channel 12 bit multiplexed digitizer. This last is digitized with a 12 bit scaling of 0-25A. The COTS microprocessor polls each ADC channel at 1 kHz.

Each of the RF diode pairs (forward, reverse), are provided by loop directional couplers with a default orientation for a negative signal. The directional couplers have a 50 dB attenuation, and the signal paths have an additional 15-20 dB attenuation. Typical signal amplitudes are about -2V.

Each detector input is again cleaned up (transient protection, 1 MHz *RC*), then buffered through a -2.5x gain differential amplifier, producing a positive signal with typical amplitude of about +5V. Each is similarly sent to the dTacq digitizer and to a single pin LEMO test point, and a ratiometric fault sensor is applied to each forward / reflected pair. This entails using an analog multiplier (AD633) with an MDS+ settable multiplication factor between 0.0 and 1.0 to attenuate the forward signal (and applying a 0.5V positive offset as well), and then faulting (via comparator) if the reflected power exceeds this calibrated fraction of the forward signal.

In addition, the forward signals (but not the reflected signals) are locally digitized by the same 8 channel 12 bit digitizer as the beam current. The scaling of this is such that the 12-bit range corresponds to 2.46V of the raw diode signal. This is also polled by the cpu at 1 kHz.

The main loop of the COTS code polls this digitizer at 8 kHz, with interleaved sampling of the four beam currents and forward powers ultimately sampled at 1 kHz each. The COTS calculation is called at 4 kHz intervals, with interleaving between the four klystron channels – and each klystron COTS calculation is done as soon as its relevant signals are acquired; thus the COTS calculation for each klystron is updated at 1 kHz. Benchmarking shows that the entire calculation update for a single klystron is performed in 26 μ sec.

The TPS system makes extensive use of ASQUIRT and DSQUIRT modules, which are very similar to the in-house AFOL and DFOL links that Alcator C-Mod already uses. They differ in that they are fully digital, with complex messages passing via F/O both ways. These messages include dip-switch settable addressing, checksums, firmware version numbers, and other failsafes to help prevent erroneous operation of equipment. The ASQUIRT01 units transmit and receive 0-10V signals at 12 bit resolution, with two channels updated at 10 kHz. The ASQUIRT02 units transmit and receive either 0-10V or 0-20mA signals at 12 bit resolution, with 16 channels updated at 1 kHz.

One ASQUIRT01 unit placed in the LH HV control area of the Alcator C-Mod cell transmits the power supply voltage on channel 1, with a 12-bit range of 0-60 kV. The total supply current is on channel 2, with a 12-bit range of 0-500 A. The ASQUIRT01 units are made with four parallel F/O outputs (but only a single F/O input). This ASQUIRT01 unit feeds to all 3 (4 after the upgrade) LH klystron carts. There is no corresponding ASQUIRT01 mate to receive this

signal – the F/O receiver is on the COTS board, and the signal is decoded by a “virtual ASQUIRT01” routine within the COTS cpu. The ASQUIRT01 timing is controlled by the sending unit, and the receiving decoder is polled, thus is asynchronous with the COTS calculations. When complete and valid message packets are received, the “outputs” (buffers in SRAM) and status flags are updated. Timeout on receive, receive errors, or other faults revert the outputs to values of zero, and error flags are set.

Four ASQUIRT02 units are used at each cart. One normal ASQUIRT02 pair monitors the RTD temperature sensors on the klystrons, with these signals received only at the PLC station, and dealt with normally there. One more ASQUIRT02 – modified to have dual parallel F/O outputs (but no F/O input to copper output section) monitors fourteen signals associated with the klystron cooling on that cart. One fiber goes to a second ASQUIRT02 receiver at the PLC station, while the other goes to a “virtual ASQUIRT02” receiver on the COTS board, again decoded by the COTS cpu. This is treated similarly – the reception is polled and timing is controlled by the sender, thus asynchronous with the COTS calculation. Valid received packets are decoded and the messages are updated (packets are actually transmitted at 8 kHz, with two channels sent per packet). Any fault or timeout on reception zeros all received signals and sets fault bits. The channels include outlet water temperatures, pressures, and flows for each klystron, as well as temperature and pressure at the inlet manifold.

Since both ASQUIRT0x units operate asynchronously, updating the “outputs” – which are really just buffer memory in the cpu – whenever valid data is received, this data is treated as a simple lookup table for the COTS calculations, supplying the most recently received data.

The COTS system will produce a calculated model temperature for the water in the klystron collector cooling loop for each tube. The goal is to treat this as a realtime signal, and to latch a fault if it exceeds a given threshold. For verification purposes, we also convert this model in realtime to an analog signal and ship it out to the dTacq digitizer for acquisition by MDS+. These signals are updated at 1 kHz (as the COTS calculation is performed), and scaled with a 0-10V range corresponding to a 0-250C temperature. The analog outputs are generated by a muxed 12-bit DAC. (These analog signals are also available as single-pin LEMO testpoints on the TPS unit).

Additional analog outputs are provided for convenience in benchmarking performance of the system. We echo the two ASQUIRT01 signals (beam voltage, total current) to 0-10V signals with the same ranges (0-60kV, 0-500A respectively), updating these as the ASQUIRT01 packets are received, and sending to both dTacq digitizer pins and LEMO testpoints.

We also echo the *measured* klystron outlet water temperatures which are received from the ASQUIRT02 transmission. These are similarly sent to dTacq pins and LEMO test points. These are rescaled to have the same range as the COTS temperature signal – 0-10V range corresponds to 0-250C, and are updated as ASQUIRT02 packets are received.

COTS Calculation

The COTS calculation is quite simple. We determine $Power\ In - Power\ Out$, and divide by the heat capacity of the klystron collector cooling system to get the dT/dt , i.e. the rate of change of temperature. The power in is simply the beam voltage multiplied by the beam current. The power out is the sum of the forward RF power as determined by a calibration curve on the forward RF diode, and the power removed by the cooling water. The latter is determined by the

difference between the COTS temperature and the measured water inlet temperature, and multiplied by the water flow rate. Thus, for a timer interval dt :

$$dT = \frac{I_{beam}V_{beam} - P_{rf} - Q(T_{COTS} - T_{in})}{k} dt .$$

The first part of the calculation requires us to calibrate the forward power diode signal. The C-Mod team¹ maintains calibration measurements, fit to a cubic with a calibration for the 50dB attenuator used. The form of these calibrations is defined as:

$$P = \frac{(a_1V^3 + a_2V^2 + a_3V + a_4) * 10^{(-a_5/10)}}{10^6} ,$$

where V is the raw diode voltage in Volts, *and which is negative*, and the power P is given in kW. The coefficients were provided by calibrations done by the MIT team.

Note that this can be reduced somewhat, as the above calibration does not really represent five independent constants, but only four. If we fold the attenuator calibration into the other constants we get a simple cubic expansion. Furthermore, our forward diode signal is digitized with a 12 bit (0-4095) range of 0-2.456 V, and we must remember that we have removed the sign flip in the process. Rolling this all together, we now get something with a more tractable form:

$$P = b_1N^3 + b_2N^2 + b_3N + b_4 .$$

N is now the number of raw counts of the digitized RF signal. (We will keep the unusual ordering of the coefficients, vs. the standard nomenclature where the subscript matches the exponent, to match Atma's calibration documents.)

Applying a few tricks to make this work well for 32 bit unsigned math, we can rewrite this as:

$$P = -c_1\left(\frac{N}{2^6}\right)^3 + c_2\left(\frac{N}{2^4}\right)^2 - c_3N + c_4 ,$$

with N being the raw RF diode counts over a 0-4095 range, and the power P given in mW, the calculation and all intermediate steps maintain maximum accuracy while staying within a 32 bit range. All four of the calibration coefficients are now defined as strictly positive, and the sign is maintained in the form of the calculation.

We will load the coefficients c_1 through c_4 for each klystron via MDS+ tables. These can be loaded prior to each shot, or written into flash memory within the cpu to become power-up default values – confirmatory writing prior to each shot is of course recommended.

The conversion from Atma's calibrations to these numbers is as follows:

$$c_1 \equiv a_1 * 10^{(-a_5/10)} * \left(599.75 * 10^{-6} \mu\text{V}/\text{count}\right)^3 * 2^{18} , \text{ typical} = \text{default} = 226,$$

and c_1 will be constrained to be a positive integer $< 2^{12}$.

¹ Personal communication, Atma Kanojia

$$c_2 \equiv a_2 * 10^{\left(\frac{-a_5}{10}\right)} * \left(599.75 * 10^{-6} \mu\text{V}/\text{count}\right)^2 * 2^8, \text{ typical} = \text{default} = 5400,$$

and c_2 will be constrained to be a positive integer $< 2^{14}$.

The temperature model is maintained in a similar fashion – we roll the intermediate coefficients up into single unsigned integers, with care to maintain the range of all intermediate terms as well as the final result. We repeat the COTS calculation at 1 kHz, so the value of dt is fixed. The form of the calculation is:

$$\Delta T = k_1 N_V N_I + k_2 N_Q N_{Tin} - k_3 P_{cal} - k_4 N_Q N_{COTS} .$$

We have already discussed most of these terms (N_V is the raw counts of the ASQUIRT01 beam voltage measurement, N_I is the raw counts of the locally digitized beam current, N_Q is the raw counts of the ASQUIRT02 flow rate, N_{Tin} is the raw counts of the ASQUIRT02 inlet temperature, P_{cal} is the previously calibrated RF power, N_{COTS} is the model outlet water temperature, and ΔT is the incremental correction to T_{COTS}).

The calculation is performed as a chain multiply-accumulate-shift with 64 bit accuracy. The additional accuracy is necessary in that very small temperature changes must be allowed for the small (1 ms) time intervals, particularly at the tail end of the exponential decay as the tube cools back down towards the inlet temperature. The four constants, k_1 through k_4 are actually a factor of 2^{32} oversized, so that the extra accuracy is on the correct (LSB) end of the calculation, and so that we may simply harvest the top half of the result at the end of each calculation. (64 bit, or two word, multiply and multiply-accumulate instructions are also primitives for the ARM7 processor.)

This calculation is continuously iterated at 1 kHz, maintaining the COTS model temperature. Clearly, with no signal for beam current or RF power, the COTS temperature should always converge quickly to the inlet water temperature. Note that we maintain no synchronicity with the HV or RF controls which are separately maintained elsewhere within the TPS hardware – this calculation proceeds blindly and continuously as long as the TPS is powered, with no knowledge of whether or not a shot is taking place.

COTS Fault Threshold

The previous calculation generates a model for the outlet water temperature, thus making a calculated proxy for the actual internal water temperature (which we cannot measure). We would like to set a trip-point threshold which we use to shutdown the HV for the remainder of the shot – one high enough to allow the largest possible window of run parameters (particularly “picket-fence” operation with reasonable duty cycle of RF low or off) – but also within safe bounds to avoid damage to the klystron.

Literature from CPI has given us guidance on the water temperature at the skin layer adjacent to the copper surface in the collector – typically 25% higher than the bulk water temperature (in deg C), thus we need to depress the fault threshold below 100C to account for this. Pressurization of the water, on the other hand, elevates the boiling point, which thus elevates the fault threshold. Finally, prudence dictates that we allow an additional safety margin from the actual boiling point.

In equation form, this yields:

$$T_{fault} = T_{boil}(P) * k - \Delta T_{margin}$$

We've already discussed how the pressure is measured in the suite of signals received over the ASQUIRT02. This pressure measurement can be used to first calculate the boiling point of the water at the elevated pressure. Reducing steam table curves to a cubic fit, and using our usual technique for multiply-accumulate-shift with single 32-bit word internal accuracy, we use the massaged formula:

$$T_{boil} = \frac{\left[6312 \left(\frac{N_p}{2^7} \right)^3 - 18580 \left(\frac{N_p}{2^5} \right)^2 + 45371 N_p + 55246848 \right]}{2^{15}}$$

Where N_p is the raw pressure measurement (in counts) with the 819 count floor subtracted off (i.e. 0-3276 counts = 0-232 psig), and the boiling temperature is in units of 4095 counts = 250 C. All of these coefficients are fixed. This fit gives better than 1% error over much of the range 100-200C, with the worst error (about 3C) right at 0 psig.

The factor k has a default value of 0.8, and will be maintained by MDS+, either loaded every shot or overwritten in flash memory for future power-on defaults. It is entered as an integer, to be divided by 4096, e.g. the default value is actually 3277. (The MDS+ value loaded must be a positive integer < 4096).

The factor ΔT_{margin} is also maintained by MDS+. It is recommended that this have a minimum value of 5C, due to the errors in the boiling point approximation, and prudence dictates that it actually be considerably larger. The default will be 15C. This is entered as an integer on the 250C = 4095 count scale, thus the default is 246.

This fault temperature threshold is updated in realtime on every cycle that the COTS temperature is updated, and used for determining whether or not a fault condition has occurred.

COTS Fault Logic

The COTS cpu produces five fault (status) bits. There is one bit for each klystron (COTSAOK, COTSBOK, COTSCOK, COTSDOK) which give the status of the COTS temperature model relative to the fault threshold. If the fault condition is tripped for one of these, it will LATCH off, and will not reset until the main cpu on the TPS has issued a RESET command, *and* the over-temperature condition no longer exists.

These four bits are also hardwired directly to four bicolor LEDs on the front of the COTS board, and piped to the layer1 board of the TPS, where they are monitored by both the EVENT LOGGER and by the DSQUIRT communications which echo them to the PLC.

In addition, there is a single master fault bit of opposite polarity, COTSFLT. It is latched on if any of the COTSFLT latches go on, or if the local power supply monitor (this board only) shows a fault. It also will only reset on a master RESET with all fault conditions clear. The COTSFLT bit is transmitted to the layer1 board of the TPS, where in addition to the EVENT LOGGER and DSQUIRT communications, it is used in the CPLD logic matrix to generate a "hard fault" which takes down all HV immediately, and for the duration of the shot.

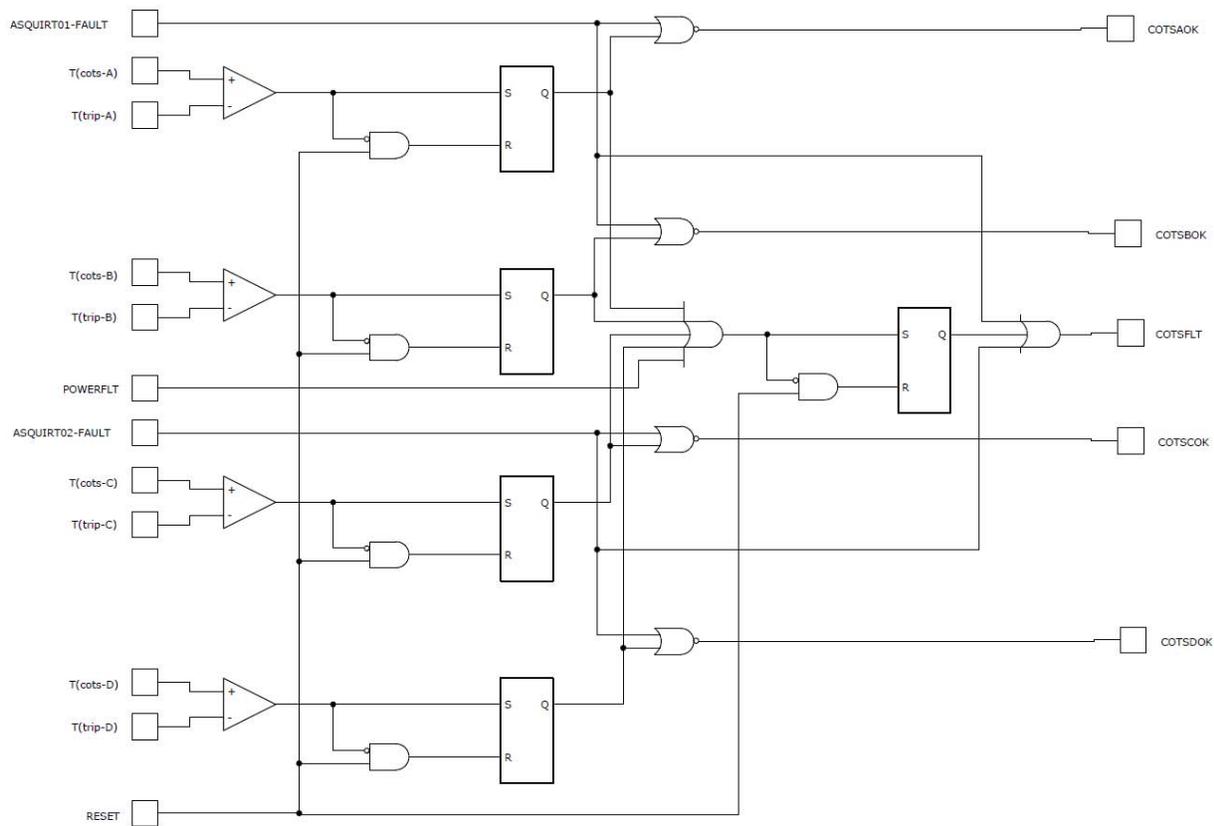


Figure 9. Logic diagram for the COTS fault determination.

We maintain a different logic for the faults on the ASQUIRT receiver routines. These will fault out if the receivers timeout without a valid packet, or if addresses or checksums don't match, or if there is a format error in the received packet. If either of these receivers fault out, we express the COTSFLT, but without latching. Similarly, we turn the COTSAOK and COTSBOK bits false without latching if the ASQUIRT01 faults out, and we turn the COTSCOK and COTSDOK bits false without latching if the ASQUIRT02 faults out.

This serves the following purpose:

- we generate a COTSFLT which inhibits all HV if there is any ASQUIRT failure, which is as it should be, as the COTS model temperature is no longer valid,
- we have an unlatched LED indicator at the front panel to indicate the validity of the communications link with the ASQUIRT01 and ASQUIRT02 senders – these will flicker red to green as the fibers are plugged in and the communications synchronize.

COTS Testing and Verification

For tests, we load in all the defaults mentioned above *except* for the fault threshold coefficients. For those, we assume a k multiplier of 1.0, and a margin of zero. We will thus begin with the fault trip point equal to the boiling point calculation.

In addition, we bypass the ASQUIRT02 readings, and superimpose fixed numbers for the following:

- $T_{in} = 20C$
- $Q = 38 \text{ gpm}$
- $P = 0 \text{ psig}$ (thus the fault trip should be 100C)
- $T_{out} = 40C$

Similarly, we bypass the ASQUIRT01 readings, and superimpose fixed numbers for

- $V_{beam} = 50 \text{ kV}$

We will apply no signal to the forward RF signal. Our only variable for this first test will then be the beam current. We will apply a signal from an arbitrary waveform generator to the I_{beam} input for channel A, and simultaneously monitor the analog COTS signal. We will also temporarily suspend the latching nature of the COTS fault bit, and echo the COTSA fault out to a spare pin on the cpu, which we also monitor.

In Figure 10, we see a signal on channel 4 (green) applied to the I_{beam} input of the COTS board for channel A. It is a pulse 700 mV in height for 50 ms. This corresponds to 17.5 A, or at 50 kV, to 875 kW of input power.

The steady state cooled temperature has no free parameters at all - only the transients depend on our assumptions for the heat capacity. Applying the steady state relation that $V \cdot I = Q \Delta T$, we quickly calculate an expected temperature rise between inlet and outlet of 88 degC.

The purple trace (channel 3) is the COTS model temperature for channel A, which rises from 800 mV to 4.20 V. At 10V = 250C, this translates as rising from a base of 20C to a steady state of 105C, for a rise of 85C. Check (close enough).

The blue trace, channel 2, is monitoring the temporary test point which echoes the (unlatched) \COTSAOK bit. Careful inspection shows that this is acting correctly – as a comparator for the COTS temperature relative to a reference of 100C. Check.

If we capture the same data to a spreadsheet and run the transient model, we can verify the dynamics of the rising and falling exponentials. In this case, we have the default coefficients derived from a heat capacity of 100 J/degC. The green and purple traces in Figure 11 are the

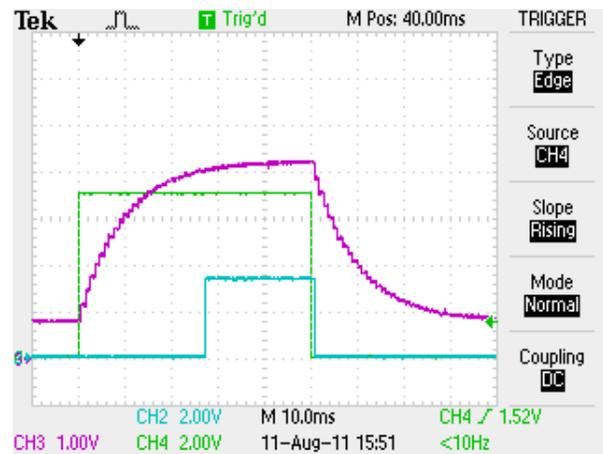


Figure 10. Test of COTS – the green trace is an emulation of 17.5A of beam current for 50 ms. The purple trace is the real-time COTS temperature model, and the blue trace is the overtemperature fault signal.

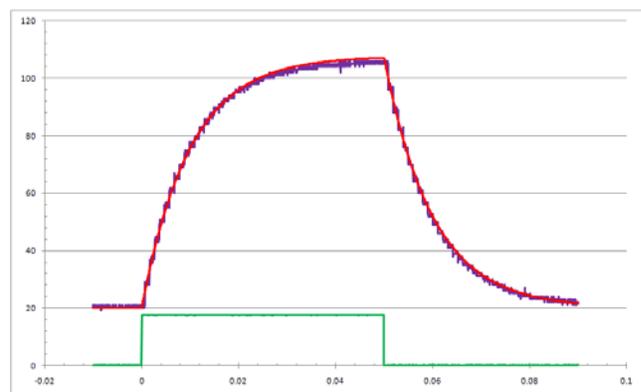


Figure 11. Test of COTS – same data as above, compared to a separate spreadsheet calculation for dynamics.

same (now in Amps and degC, respectively), while we have added a red trace calculated from first principles with the above numbers. Check.

Other contributions to the calculation are checked by, one at a time, removing the debug code which overwrites the input values from the source, and applying a signal to the appropriate input (in the case of the ASQUIRT0x modules, to the inputs on the remote sending unit).

- For a 10 psig signal (overwritten), the boiling point should rise to 116C. At $k=1.0$ and zero margin, we now observe the fault threshold at 116C
- Entering $k=3277$, for a factor of 0.8, we should observe the fault threshold at 93C. Check.
- Leaving $k=0.8$, we raise the margin to our 15C (246), which should drop the fault threshold to 78C. Check.
- Leaving $k=0.8$ and margin of 15C, we eliminate the overwrite of the measured pressure and use the ASQUIRT02 value. At 5 mA, corresponding to 14.5 psi, the boiling point should be elevated to 121C. That should raise the threshold to 82C. 81C observed. Check. At 15 mA, corresponding to 160 psi, the boiling point should be elevated to 187C. That should raise the threshold to 135C. 132C observed. Check.

We now leave zero signal into the pressure input, which should leave the boiling point at 100C (this however is the point of worst case error in the cubic fit, and it instead is about 103C). For k and margin as set, the fault trip point should now remain fixed at 67C. Check.

- We now remove the debug overwrites for all five of the temperature readings over ASQUIRT02. Supplying a signal to the input temperature channel, we now correctly see the baseline for the COTS model temperature shift up and down. At 8mA, corresponding to 25.3C, we correctly see a 25.3C baseline.
- Next, we remove the debug overwrites for all four of the flow readings over ASQUIRT02, and apply a signal to the remote ASQUIRT02. This, unsurprisingly, is the most profound change. At 18mA signal, which corresponds to 87 gpm cooling, a 100 ms 750 mV pulse into the I_{beam} input brings the temperature to only 40.5C (from a 0C baseline). 750 mV corresponds to 18.75A, or 938 kW at our fixed 50 kV. At 87 gpm, the steady state power balance requires a 40.8 degree ΔT . Check. At the other end, we crank the cooling way down, for only a 5mA signal, corresponding to 6.8 gpm, we must crank the beam current way down so as not to saturate the COTS model at 250C. Similarly, it takes much longer to reach steady state. We fix on a 300 mV pulse for 300 msec. This corresponds to 7.5A of beam current, or 375kW. The steady state power balance requires 210 degrees of ΔT . From a 0C baseline, we observe 210C. Check.
- Next, we put back the debug overwrites for flow at 38 gpm, and we remove the overwrites for beam voltage from ASQUIRT01. Applying 5V to the ASQUIRT01 input corresponds to 30kV. Pulsing the I_{beam} input with 500 mV (12.5A), or 375kW should give us a 37.4C temperature rise at 38 gpm. We measure 37C. Applying 9V to the ASQUIRT01 input corresponds to 54kV. For the same beam current, we should have 675 kW, and a 67.3C temperature rise. We observe 68C. Check.

- Finally, we apply a *negative* pulse to the forward RF diode input. This has the curious and unphysical interpretation of RF power being extracted when there is no beam, thus magically cooling the klystron. In Figure 12, we are applying a steady 8mA to the ASQUIRT02 input temperature terminal, thus giving us an elevated baseline of about 25C. We apply a 300 ms negative pulse of 1.5V (the green trace), which should calibrate to be 85 kW of RF power out. At our fixed 38 gpm, this should give us a ΔT of 8.5C. We observe 8.4C. Check.

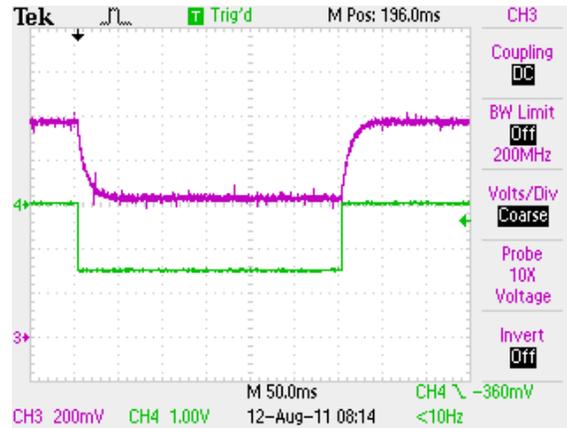


Figure 12. Test of COTS – a contrived test for pseudo cooling of the klystron via RF power extraction.

Additional Functions of COTS System

The COTS board is only part of the TPS assembly. The main TPS board serves the all-important function of monitoring all fault signals for conditions which require shutdown of RF, HV, or both. We process the RF and Beam Current signals on the COTS board because these signals are needed for local digitization for the COTS algorithm. In addition, these signals require fast comparators to sense fault levels, and the digital outputs are transmitted promptly back to the main board.

Additional functions that are provided by the COTS cpu, besides those discussed above, include:

- Controlling the multi-channel DACs which generate the MDS+ programmable fault-thresholds for beam current and RF reflected power ratio,
- Pass-thru of a master ‘heartbeat’ signal for watchdogging the TPS intelligencia,
- Communications (over a dedicated SPI port, which the COTS cpu is a slave to) with the master cpu on the main TPS board, and interpretation of commands passed by MDS+ (e.g. modification of setpoints, etc)
- Read/write of FLASH memory as needed, or as directed by MDS+ command.

2.1.2 Event Logger

The other novel feature in the TPS upgrade is an event logger. This is essentially a virtual logic analyzer, which tracks up to 96 channels of digital signals, and records an ‘event’ whenever ANY of them change flag. This event is stored in local memory, along with a timestamp referenced to the 1 MHz shot clock used by the MDSplus data system. Post-shot, MDSplus downloads this event log, and thus can create a table or logic-analyzer-like plot of all status and fault bits, thus observing the sequence of controls and faults in a complex cause-and-effect situation. Since the timebase is the same as the MDSplus system, this plot can be overlaid with any physics signal from the MDSplus system for reconciling events across the C-Mod experiment.

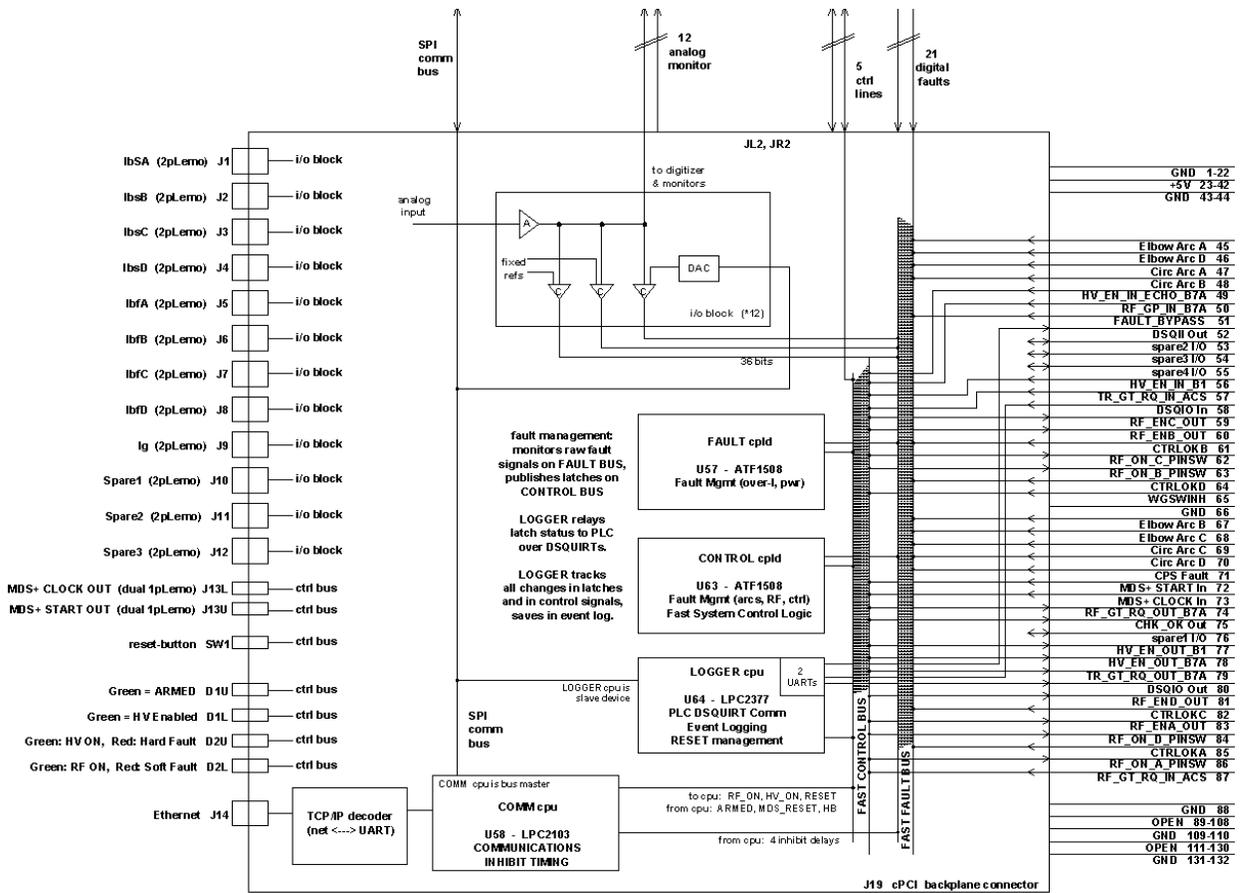


Figure 13. Expansion of main board from Figure 7. This board has 12 analog channels, plus the master/communications CPU, plus the Event Logger CPU. The Event Logger also serves as a dual DSQUIRT virtual transceiver for PLC communications.

This event logger is implemented on another dedicated ARM7 processor within the TPS on the main board (Figure 13). We have a 64 channel interface to the associated PLC, using two of the digital multiplexed DSQUIRT I/O units, and because so many of the fault bits and control signals are shared between the PLC I/O and the event logger, we perform both functions within this same ARM7. The high signal count requires that this be a very high pin-count chip – otherwise it is functionally very similar to the COTS cpu, sharing many architectural features.

Post-shot readout of the Event Logger allows comparison of complex fault-chains to compare to other tokamak activity, as well as to determine sequence of events (Figure 14).

2.1.3 ARM7 Programming

The coding of the ARM7 microprocessors was deemed to be highly sensitive and time-critical, and is thus not layered on an operating system or high-level language. All ARM7 coding is done in bare ARMASM, with full and absolute control over the microprocessor hardware. Interrupts are used for critical timing functions, and all chips are fully watchdogged against code errors which may hangup operation. In addition to conventional on-board

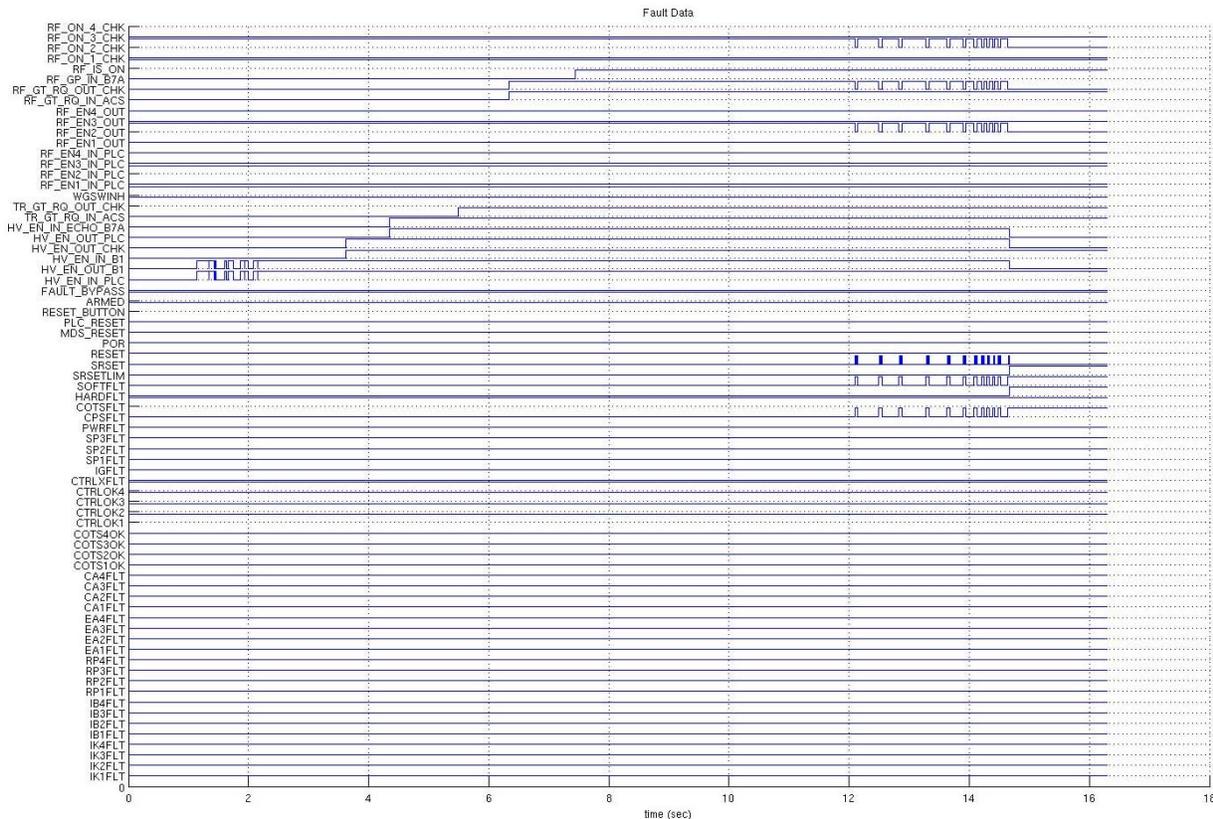


Figure 14. A readout of the Event Logger, post-shot, shows the time sequencing of all status and fault bits related to the TPS operation. The timebase is that used by the MDSplus data system. Events are defined as any change of any of the 96 bits – an event memory 2048 deep preserves a snapshot of all status/fault bits and a 32 bit timestamp on every event. (Readout plot provided by Atma Kanojia of the Alcator C-Mod team).

watchdogging of the cpus, we do an additional external “heartbeat” watchdog within the TPS system. This feature generates a 100Hz clock on one cpu, and passes it in turn through each of the other cpus and each of the programmable logic chips. In all cases, the “pass-through” is not interrupt driven, but is handled by polling at the highest level outer service loop – thus will fail at any hangup in the code. The output of the last stage is filtered, and two comparators detect if this filtered signal ever drifts off the middle towards either rail. The output of this fault-sense feeds a hardwired inhibit, the activation of which shuts down all HV and RF activity. Failure of any of the chips to service all features at the outermost service loop level thus results in a fail-safe shutdown of HV to protect equipment.

A third ARM7 microprocessor (on the main board) serves as a communications gateway from the TPS to the external Ethernet port. This processor is only “busy” with Ethernet between shots, and would otherwise be idle during shots. We thus keep it busy by assigning it the task of generating programmable inhibit delays, which we use to tune the sensitivity of the comparator faults with respect to leading and trailing edge of high voltage or RF commands.

While the COTS and MASTER cpus were demonstrated to operate with full functionality in default mode – running user code directly from onboard FLASH memory, we required more speed for the Event Logger. This program takes the additional step of copying itself from user

FLASH to SRAM on bootup, and transferring all operation (particularly interrupts) to run exclusively out of SRAM, resulting in about 3-4x speed improvement. The main loop, which services all polled functions, was redundantly programmed to repeatedly call the port polling function, keeping the latency to event detection down to $\sim 1 \mu\text{s}$.

2.1.4 Programmable Fault Logic

The internal architecture of the TPS breaks down into a conventional one used for modern radar controls. A network of programmable logic provides parallel and robust digital signal paths which ensure that HV and or RF are shut down promptly (sub-microsecond) in the event of a fault such as tube over-current or waveguide arc. At the same time, the use of programmable logic here allows for maximum flexibility and reprogrammability to accommodate specification changes.

This programmable digital fault logic is guided by a few state bits, which are indicative of the overall status of the system (i.e. idle, armed, debugging, etc). These state bits are set via slower, but higher level, controls, such as MDSplus commands or PLC commands. We thus have an interface which communicates with MDSplus, PLC logic, the ACS timing system, and the individual fault comparators and arc detection channels within the cart electronics, and performs the proper responses using all of the above as inputs as necessary. The TPS thus serves not just as its name implies (transmitter protection system), but also as a sequencer/interface for ross relays, HV supply turn-on, RF calibration/debugging, etc.

In Figure 15, we see a schematic representation of half of the programmable logic within the TPS CPLD array. This half is the crucial section for the above mentioned sequencing, and the logic shows the permissive structure which ensures that HV and RF are properly interlocked with all the safety and control features specified. (The other half of the programmable logic comprises a large repeating array of fault latches for many of the analog comparator channels – the net result being two fault bits which are treated as inputs to this page).

2.1.5 TPS Architecture Summary

The internal architecture thus consists of three interlocked ARM7 microprocessors, and a programmable logic array which is distributed between two Atmel CPLD array chips. The RTM contains 42 channels of Fiber-Optic I/O which feed directly to the logic array, and the front panel consists of 24 analog inputs for current transformers and RF diode inputs, 32 channels of analog monitor testpoints, and a port to directly connect 32 analog channels out to the dTACQ fast digitizer on a compact SCSI68 connector. The remainder of the front panel contains a few display LEDs, a manual reset pushbutton, F/O ports for the COTS analog mux links, and the Ethernet port.

A key additional task of this SBIR effort was to produce not just a specific solution for MIT's LH upgrade, but a general purpose transmitter control and protection system that could be easily installed for other applications – either in fusion, phased array radar, or other unanticipated needs. The design has discrete analog channels which are quite reconfigurable with respect to input termination, polarity, and initial gain, as well as to comparator filtering, thresholds, and logical dependence. The logical combination of fault signals is entirely contained within the programmable logic arrays, thus is easily reconfigurable for any foreseeable needs – and the



Figure 16. Five complete TPS front assemblies after testing and ready for shipment to MIT.

network resources within each of the programmable chips is considerably underutilized. Similarly, the possible uses of the onboard computational power are nearly limitless – they are quite interconnected, with spare pins allocated and unused for future expansion.

Essentially, we have designed a “kit” with a modular structure, containing a generous number of buffered analog inputs, fault comparators, and with all logical links reprogrammable for the application – easily adapted to nearly any fault-control use of similar scale.

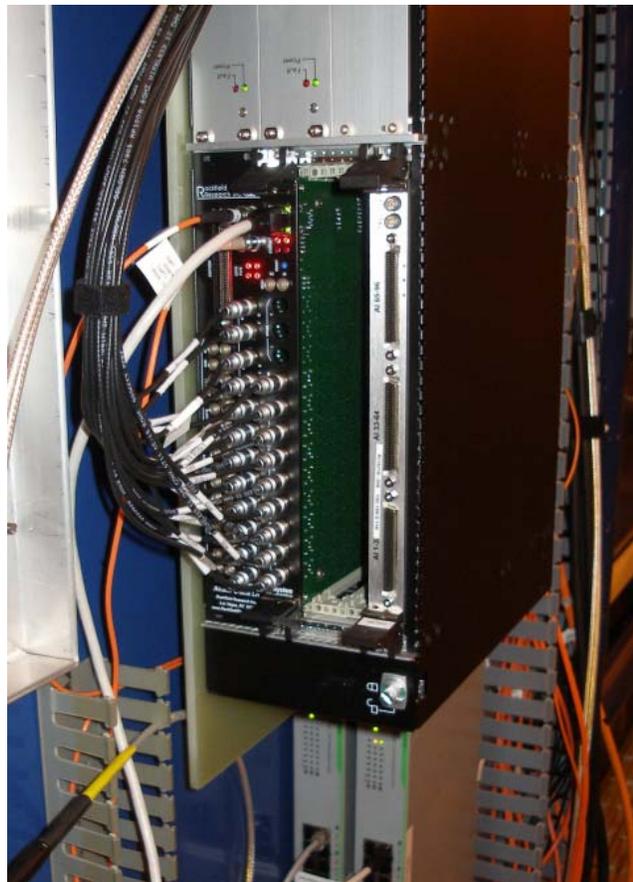


Figure 17. One of the TPS assemblies, installed in cPCI crate and installed on the side of klystron cart on the Alcator C-Mod LH facility.

3.0 Review of Auxiliary Interface Components of TPS

Refer to Figure 2 for each of the following auxiliary interfaces and interconnect relations to the rest of the LH transmitter controls.

3.1 Box 1 – HV Control Box (Figure 18, Figure 19)

This interface resides within the high voltage restricted access area of the cart, and controls the action of the Ross Relays (which apply HV to the klystrons). In addition, a number of auxiliary services are performed: filament control, airflow sense, body current fault testing, Ross Relay positional sensing, cart panel safety switches, etc. This box is designed to be a key permissive element in the hardwired safety interlocks of the high voltage system.

The design for this closely emulated the functionality of the legacy controls for simplicity in

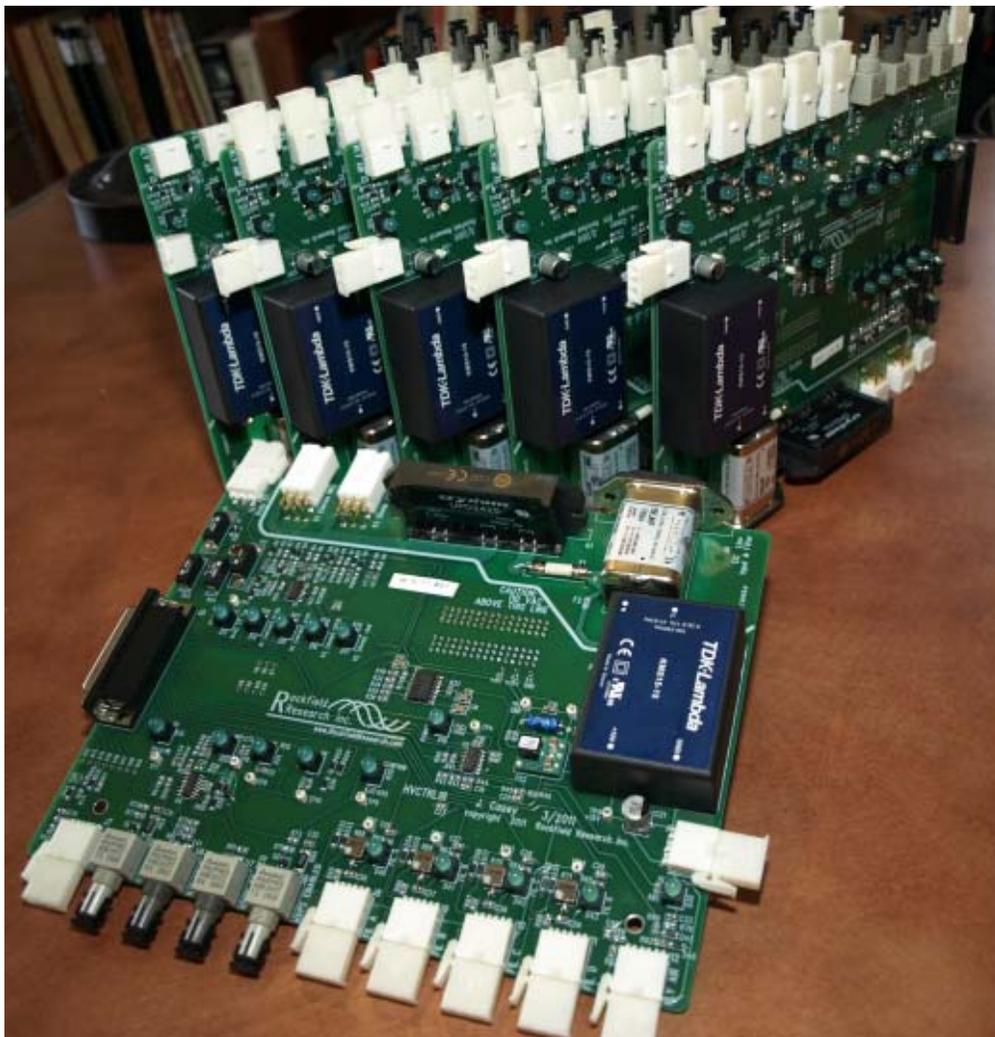


Figure 18. Six “box 1 HV control” units assembled and tested at Rockfield Research facility. Each of these boards interfaces to a DSQUIRT digital multiplexed fiber-optic I/O transceiver, and controls hardware in the HV area of the four-klystron cart.

wiring and interfacing. Cable runs external to the HV cart were largely eliminated by providing interface communications via fiber optics. Controls of the Ross Relays and filaments are managed by the PLC (under the responsibility of the MIT team), but the permissive reporting goes directly into the TPS as part of the interlock chain.

The block diagram specification and top-level design was performed by David Terry of the MIT team, and Rockfield Research did the detailed circuit design and all manufacturing, assembly, and testing under the supplementary Phase II program. One of these boxes is required for each cart, thus five were shipped to MIT (four carts plus one spare). One additional unit was retained at Rockfield Research for future support. The circuitry in each unit includes a DSQUIRT 16/16 channel bidirectional F/O transceiver, to provide non-galvanic connection to the PLC.



Figure 19. One of the “box 1 HV control” units is shown installed in the HV area of the cart, adjacent to the legacy filament power supplies, and underneath the klystron sockets.

3.2 Box 2 – Waveguide Switch Box (Figure 20)

This is a straightforward interface to operate the waveguide switches which alternate the klystron loading between the dummy load and the antenna. This box was the responsibility of the MIT team, although it did include one DSQUIRT 16/16 channel bidirectional F/O transceiver to connect to the PLC, which was provided by Rockfield Research under this effort.



Figure 20. Each cart has a “box 2 – Waveguide Switch Box”, shown here installed in the C-Mod Cell. Opto-22 interface modules on the left, and a custom interface board at right interface to a DSQUIRT I/O board, mounted underneath at right.

3.3 Box 3 – Elbow Arc Detector Interface (Figure 21, Figure 22)

This is also a relatively simple interface box. Most of the functionality is to power and sense the waveguide arc detectors and echo the signals to the TPS over dedicated fibers for minimum fault response latency. Additional interfaces to solenoid water flow switches, and waveguide pressure transducers are provided here for convenience, due to the physical proximity to those sensor areas.

This box also powers and interfaces to the PIN diode switches, which gate the RF to the klystrons.

The sheer number and size of the connectors required a complex physical setup, even if the circuitry was quite simple. The Elbow



Figure 21. Each cart requires one “Box 3 Elbow Arc Detector Interface”, which comprises two custom boards and two DSQUIRT fiber optic transceivers. The boards are physically constructed to “wrap around” the DSQUIRTS to fit a stacked assembly in a compact box

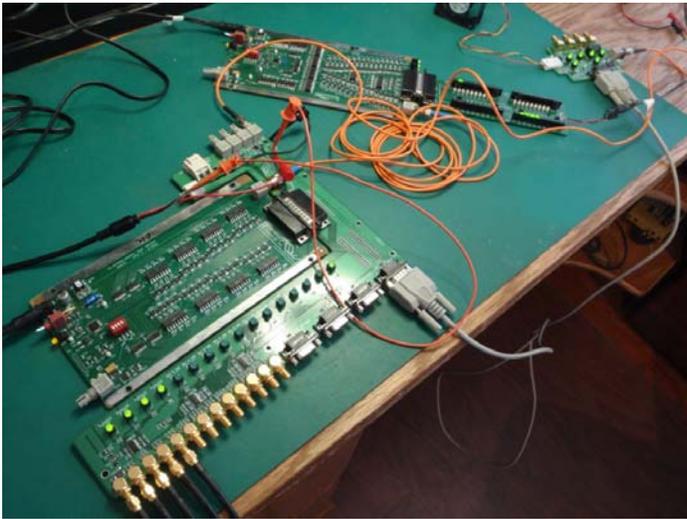


Figure 22. One of the Elbow Arc interface boards is shown on the testbench, plugged into its embedded DSQUIRT interface (above). One of the Elbow Arc units is shown mounted on the side of a cart, installed at MIT (right).



interface was built on two circuit boards which are stacked vertically to give two rows of detectors on opposite faces of a commercial enclosure. The boards are built to “wrap around” the embedded two DSQUIRT digital transceiver boards which link these controls to the PLC.

The LH system requires one Elbow interface for each cart, thus six were manufactured – four for install, one spare for MIT, and one to remain at Rockfield Research for support.

3.4 Box 4 – Analog Interface Box (COTS)

The water cooling system for the klystrons requires support of a number of transducers reporting water temperature, pressure and flow. In general, these report to the PLC for conventional analog logging and fault protection. For the upgrade TPS, we also report these signals directly to the TPS where they are used as inputs to the realtime collector overtemperature system. (Figure 23)

In Box 4, we power and interface these transducers, and bring the signals to a special version of the ASQUIRT 16 channel transceiver. This version omits the fiber-to-copper output section, but implements two parallel fiber outputs for a single 16 channel copper-to-fiber section. We also use the 0-



Figure 23. The analog interface box, Box 4, contains two ASQUIRT02 units to echo water cooling information to both PLC and TPS.

20mA option for the input scaling, so that the 4-20 mA analog outputs of the transducers interface easily to the ASQUIRT. The two output fibers separately feed the PLC and the TPS the same information on klystron cooling.

A separate additional ASQUIRT is used in each box for an additional number of RTD sensors that the newer CPI klystrons are outfitted with for improved diagnostics. These signals only feed to the PLC for monitoring and fault interlocking.

The Box 4 Analog Interface units – one per cart – were the responsibility of the MIT team. Each one used two ASQUIRT analog multiplexed transceivers, which were manufactured by Rockfield Research under this contract.

3.5 Box 5 – PIN Switch Diode Interface (Figure 24)

The klystrons are ultimately modulated via the input RF signal, which is gated via the PIN switches, mounted at the RF input coupler of the tube. These switches require DC power and a gating signal, which has been upgraded with the new TPS for galvanic isolation with a very simple fiber interface. The Box 5 PIN switch box performs these functions.

The fiber gating signal comes directly from the TPS for minimum latency, as the final end product of a long string of logic permissives and fault checks.

There is also an interface to the nearby Elbow Interface box, which supplies DC power, and reports back on power-ok and fan rotation status to the PLC via its embedded DSQUIRT interface.

One PIN switch box is required for each klystron. Nineteen were thus built – sixteen for four carts fully populated with klystrons, two spares to MIT, and one spare remaining at Rockfield Research for support. These were manufactured and supplied by Rockfield Research under this program.

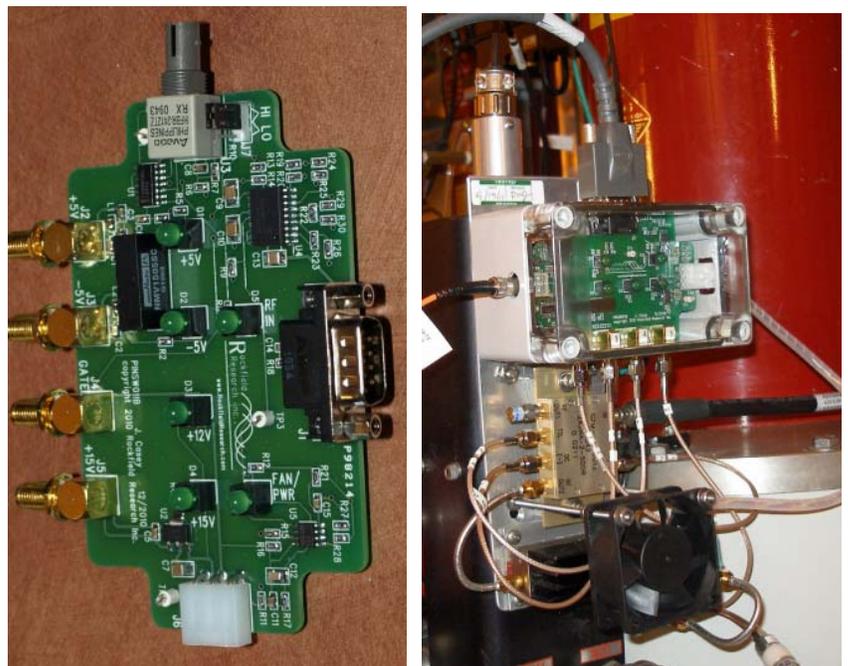


Figure 24. Box 5 is the interface to the PIN switch diodes – a very simple interface indeed. At left is the bare board, at right is the board installed in the protective case and mounted adjacent to the PIN Switch on the klystron body.

3.6 Boxes 6 & 8 – PLC Interface (Figure 25)

The “other end” of many of these isolated communications links are generally at either the PLC or the TPS central unit. While the TPS virtualizes the mating communications unit, the PLC does not and real physical mates are required to produce the proper signals on copper, which then route conventionally to the PLC for sensing.

Box 8 is a general interface holding six DSQUIRTs and two ASQUIRTs at the PLC. The fiber interfaces from these reach out to far flung remote sensing stations across noise and voltage barriers, isolating the signals and eliminating miles of cable in noisy and dangerous areas.

Each cart requires a fully populated Box 8. This was a group effort, with Rockfield Research providing the DSQUIRT and ASQUIRT units, and MIT housing them in a 19” rack mount frame, and wiring up the PLC cabling.

Box 6 is a special case – a single higher bandwidth two channel multiplexed F/O transceiver, which monitors the Beam Voltage monitor point at the HV controls, as well as the total HV supply current. Both of these channels are pumped out over four parallel fibers, one to each TPS for use in COTS calculations.

Since there is only one Box 6 in the entire LH system it is added to the Box 8 assembly for Cart 3 only



Figure 25. The Box 8 interface is primarily a rack mount housing for six DSQUIRT and two ASQUIRT modules....



....but with considerable wiring overhead, shown on the two trays between the Box 8 rack and the PLC.

3.7 Boxes 7

The nomenclature for Box 7 is somewhat confused, as it was broken into two separate interfaces during specification revisions. One of these is an interface to the pump controllers, and is relatively straightforward.

The other, Box 7A, is quite important. It is the control interface to the HV power supply, and it is also the “group control” for the suite of independent TPS controllers. We require that any RF fault on any one klystron immediately shutdown RF on all klystrons to prevent arcing on the antenna, and we also require that any hard fault which requires HV shutdown also immediately shut down the entire system. Box 7A gathers RF and HV “ok” signals from all TPS units and echos appropriate signals back to each of the TPS units – to confirm that group permission is given (or denied).

This box, by definition, must be reconfigurable to account for initial operation with the legacy 12 klystron system, but be easily upgraded to accommodate a 4th cart bringing the system to 16 klystrons at a later date. The TPS units have no such requirement – each controls 4 klystrons in an isolated universe.

Both the Box 7A and 7B assemblies were the responsibility of the MIT team.

3.8 Boxes 9, 10, 11

Boxes 9 and 10 do not exist, having been defined earlier in the specification, and later rolled into the previously described interfaces where convenient.

Box 11 is a small interface which includes the interlock keys and emergency stop button into the HV permissive chain associated with the HV controls (Box 1). It was the responsibility of the MIT team.

3.9 Two channel 10 kHz 12-bit Analog F/O Communication Link

In concept, this is a simple analog repeater – any signal presented at an input of one box is echoed at an output of the other box. Each box has two inputs and two outputs, making this a 2-channel bidirectional unit. The signals are transmitted over F/O link with 10 kHz update rates and 12 bit resolution. Shown in Figure 26, this was the first I/O unit designed.

The link accommodates configurable options for 20mA input or 10V input, as well as an option for single-pin LEMO rather than 2-pin LEMO input and output connectors. Two have been delivered to MIT – one at the power supply/modulator monitor point, and one for spare/debug. (The complementary receiver for these



Figure 26. The “ASQUIRT01” 2 channel 10 kHz analog F/O communications link, a sub-unit of the Alcator C-Mod TPS.

signals is virtual within the TPS).

Internally, the unit uses a 32-bit ARM micro-processor for both communication and internal logic, with UART interface to the F/O links, and local SPI interface to on-board digital-to-analog and analog-to-digital converters. The analog I/O sections use high-speed precision preamplifiers with 0.1% scaling resistors for absolute accuracy in signal reproduction at the far end of the link. The communications message includes an absolute tag for the starting byte of the message, a DIP-switch selectable address (to prevent activation if plugged into the wrong fiber), a code for the firmware revision, and a checksum for the entire message. Any failure (revision, address, broken fiber, broken wire, lost power at either end) results in a zero volt “safe” output. A watchdog reverts to “safe” if no signal is received over the fiber within a timeout period (a few msec). The unit is powered from a DC input – a commercial power “brick” with 6V output is supplied with the unit, but any DC voltage from 6V to 16V will serve – and installation at MIT will use the 15VDC local bus.

This unit is now a standalone product in the Rockfield Research offering, available as COTS for future program needs for the fusion community or elsewhere. (See <http://www.rockfieldresearch.com/products.html>).

3.10 Sixteen channel 1 kHz 12-bit Analog F/O Communication Link

The 16-channel twin of the fast analog link (Figure 27) was designed using similar circuit fragments. A slower version of the ADC and DAC chips with onboard 8-channel multiplexers were used for data conversion, the analog front ends were similar, but micropower preamps were used at lower speed. The inputs are designed for a default of 0 to 20mA, with optional stuff for 0 to +10V.

The system overhead is nearly identical – the microprocessor is the same, and the firmware shares many features. The 1 kHz update rate is achieved by transmitting an entire message at 8 kHz, with all i.d. and checksum features and data from two channels. The 16 channels are thus updated in sequence for a net 1 kHz update rate per channel.

Five units were delivered to MIT (one per cart plus one spare). Late in 2009, a realization that we needed to echo pump station signals to *both* TPS and PLC required fabrication of a modified version – with dual F/O outputs, one way communication, and no analog output. Four of these were fabricated to reside at the pumphouse – the existing transceiver units will receive one output at the PLC, while the other output is received in a virtual unit in the TPS (and used for COTS calculations in the collector temperature interlock).

This unit is also a standalone COTS product in the Rockfield Research offering, see above link.



Figure 27. The “ASQUIRT02” 16 channel 1 kHz analog F/O communications link, a sub-unit of the Alcator C-Mod TPS.

3.11 Digital F/O Communication Link

Digital communication links were designed as close cousins of the analog units. The design settled on 32-channel multiplexed digital units in three varieties: 32 channel copper-to-fiber, 32 channel fiber-to-copper, and 16 channel (each way) bidirectional. The input and output sections were designed identically to Allen-Bradley PLC modules # 1771-IBN/B (10-30V high true input) and # 1771-OBN/B (10-30V high true output). 15V logic interfaces are used, and a number of these interface directly to mating PLC modules. This matching to the Allen-Bradley design extends to pin-for-pin compatibility on the I/O, output fusing, as well as the isolation circuitry of each group of eight channels within each module.

The modules operate with a 5 kHz update rate, although the actual signal risetime is a bit slower than this due to the implicit filtering of the PLC circuitry used. Addressing, communications algorithms, checksums, and safeties are all similar to those of the analog links. A “witness” fiber output is present on all modules (of all three varieties) for debugging and verifying signals without disconnecting existing fiber highways.

Six units are installed at the PLC station, packaged in a single 3U height rack-mount chassis adjacent to the PLC (along with the 16 channel analog receivers for pump station signals). Other remote units are in extruded boxes similar to the analog units, and many units will not have hardware counterparts, instead talking to counterparts “virtualized” within the TPS. The inventory of digital I/O units is quite large, reflecting the complexity of the TPS controls. For four carts, we have manufactured and delivered to MIT:

- 30 bidirectional units (7 per cart plus 2 spares)
- 6 copper-to-fiber units (1 per cart plus 2 spares)
- 10 fiber-to-copper units (2 per cart plus 2 spares)

Custom test fixturing was developed to aid in qualifying the large number of units, and a full set of this fixturing was delivered to MIT as well.

All three varieties of DSQUIRT02 units are now standard COTS products from Rockfield Research, available for future program needs for the fusion community or elsewhere. (See <http://www.rockfieldresearch.com/products.html>).



Figure 28. The DSQUIRT02 digital I/O modules come in three varieties. At left is an “IO” module – 16 channel bidirectional, and next to it is the assembled board for same. Center right is the “OO” module – 32 channel fiber to copper. Far right is the “II” module – 32 channel copper to fiber. Both “II” and “OO” modules assemble into boxes identical to the “IO” assembly shown.



Figure 29. The entire suite of DSQUIRT02 boards – II, OO, and IO varieties – are shown after testing. Some are installed in the extruded aluminum chassis for standalone installation, some remain bare boards for embedding within other assemblies.