

The FINAL Pre-Processor Module (PPM) for the ATLAS Level-1 Calorimeter Trigger

The ATLAS group at KIP Heidelberg

September.2006, P.Hanke

This document shall serve as a “base” for the combined “Final Design Review / Production Readiness Review (FDR/PRR)” of the Pre-Processor Module. The review is scheduled to take place on 12.October 2006 at the Kirchhoff-Institut für Physik, Universität Heidelberg.

Introductory Remarks.

This document relates directly to the “Interim Design Review (IDR)” for the final pre-series of PPMs (“16 PPMs in a PPr-Crate”) of December 2005.

It is “incremental” in the sense, that items identified in that review are taken up. The implemented solutions are described here. The “system aspects” mentioned in the IDR (e.g. full-crate operation) are covered as well.

The first prototypes of PPMs came into operation in spring 2004. FOUR printed circuit boards were produced. A first one was equipped in-house with components and, finally, with daughter-boards to make up a functional Pre-Processor Module. The other three PCBs were assembled at a company to establish the out-sourced procedure for the full production of 160 PPMs later on. Also, these 3 PPMs came into operation during the second half of 2004. A decisive test took place at the “last” period of CERN test-beam in summer 2004. Important conclusions could be drawn from this “slice-test”, which involved the whole chain from signal-sources on detectors all the way to the “Central Trigger Processor” and to Data-Acquisition.

Many periods of detailed testing in laboratories followed, to clear up performance issues. Several modifications on details of implementation were done on the 4 prototypes including a complete “re-design” of an important daughter-board the “LVDS Cable Driver (LCD)”. The prototypes (PPM_1.0) in their present form fulfill all the requirements. This has been shown in the latest test-sessions (e.g. tests on readout at RAL, real-time signal recording from calorimeters at ATLAS-USA15).

All conclusions drawn and summarized in the “Interim Design Review” of Dec.2005 were worked into a re-design of the Pre-Processor Module (PPM_2.0). TWENTY PPMs were built since then; 16 of those new PPMs form the “crate-assembly” under test for this “Final Design/ Production Readiness Review”. It should be noted, that the crate-assembly constitutes 1/8 of eight identical crates required for ATLAS.

Related documentation.

One more document is of direct relevance for the “Final Design Review” procedure:

- “PPM modifications for a FINAL prototype”; Write-up for the “Interim Design Review”, KIP Heidelberg, 13.Dec.05.

URL: <http://www.kip.uni-heidelberg.de/atlas/L1/DISCUSS/InterimDRevPPM.pdf>

The document gives the “starting points” for this FDR/PRR. It lists all modifications/improvements identified in the IDR, which were also implemented as “updates” (wired “fixes”, added “SMD-submodules”) on the four boards of PPM_1.0.

The present state of the PPM is described in:

- “The Pre-Processor Module (PPM) for the ATLAS Level-1 Calorimeter Trigger”, KIP Heidelberg, Feb.2006

URL: http://www.kip.uni-heidelberg.de/atlas/L1/PP_Module/PPMod_Wrup.pdf

This document is intended to become a comprehensive description / user’s guide / programmer’s guide of the Pre-Processor Module. The document replaces all earlier versions of the PPM-specification and updates thereof. The “programming section” (VME register model, firmware) will be actualised after the review during the production-phase of the modules. References to more detailed documentation is given in its bibliography-section.

The following write-up is aimed at newcomers (not the reviewers), who want to get an overview of the Pre-Processor (PPr) system. It describes the ATLAS calorimetry briefly, the analog signal output and the calorimeter division for trigger-purposes. A simple (historical) allocation of signals to PPr-channels is used in:

- “The Pre-Processor System in the ATLAS Level-1 Calorimeter Trigger”, KIP Heidelberg, Nov.2005

URL: http://www.kip.uni-heidelberg.de/atlas/L1/PP_System/PPSys_Wrup.pdf

New collaborators intending to work on the Pre-Processor (or Level-1 trigger generally) can get more familiar with more details of the Pre-Processor Module by reading:

- “Channel Mapping on the Pre-Processor Module (PPM)”, KIP Heidelberg, Feb.2006

URL: http://www.kip.uni-heidelberg.de/atlas/L1/PP_Module/ppm_channel_mapping.pdf

The transition to the “expert-level” requires intimate knowledge of the system as implemented in ATLAS. The Pre-Processor is a strictly “channel-oriented” system. This means, that all signal-processing quantities can be set for each channel individually. Thus, other “architectural aspects” can prescribe, which signal-group (in a connector) is allocated to which module. ATLAS calorimeters and the trigger system have a certain architecture (azimuthal quadrants). The exact signal tracing is compiled in:

- “Level-1 Calorimeter Trigger: Cable Mappings and Crate Layouts from Analogue Inputs to Processors”, L1Calo group, Apr.2006, ATL-DA-ES-0036

URL: <http://www.hep.ph.qmul.ac.uk/l1calo/doc/out/CableMappings.pdf>

The FINAL design of the PPM.

It is the purpose of this final design review, to summarize and to check the completeness of the “updates/modifications” identified in the IDR. In addition, the complete functioning of a significant subset (1/8) of the Pre-Processor system has to be demonstrated. The collection of issues following shall show, that a viable system is at hand, which can be multiplied (“production readiness”) and installed at ATLAS.

The figure below shows the new board-implementation of the PP-Module. As matter of fact, the new design (PPM_2.0) can hardly be distinguished from the previous version on the level of such a picture.

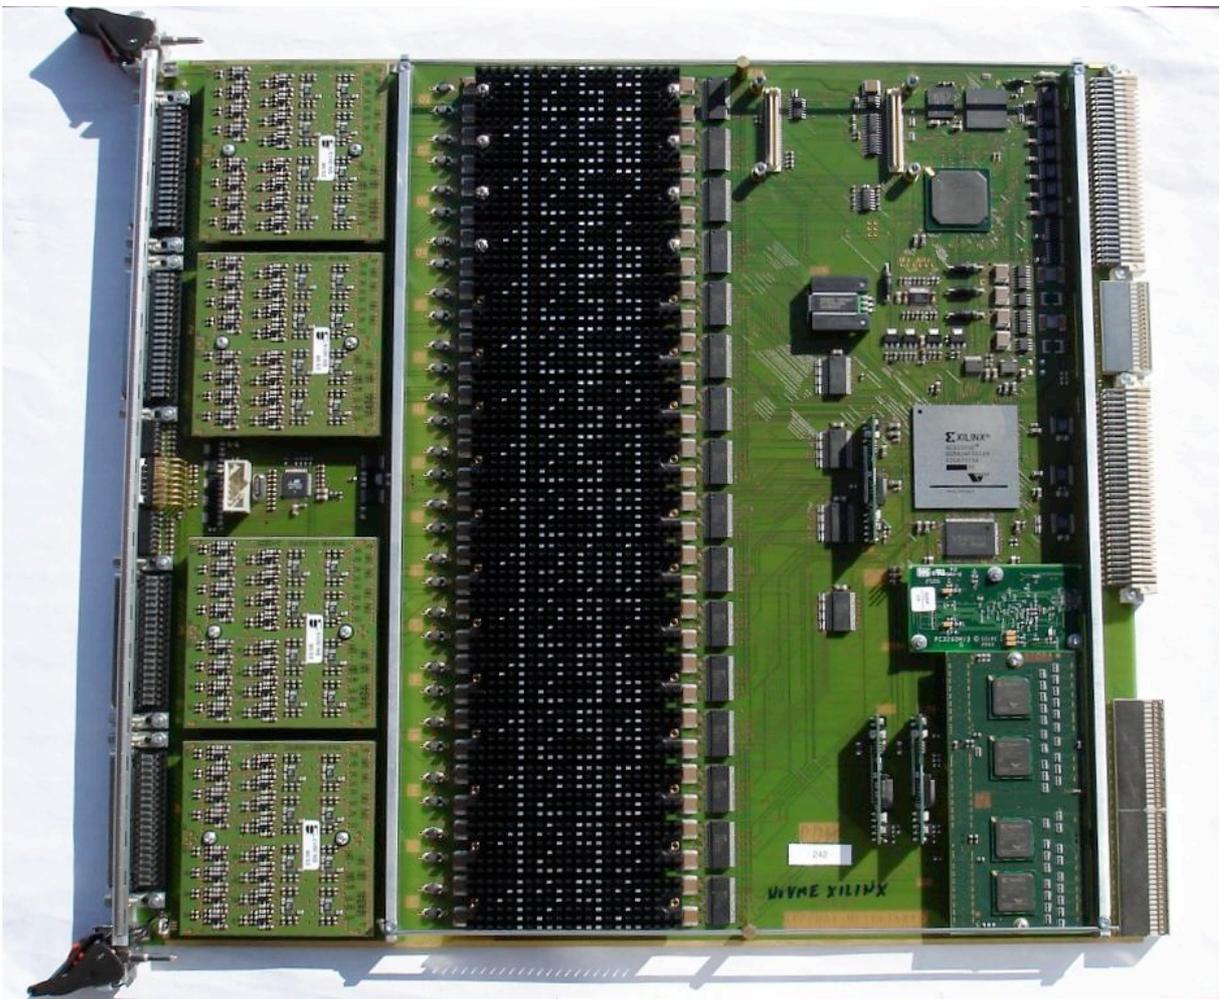


Figure 1 : The PPM and its periphery.

Easily visible are the many sub-components, which populate the motherboard:

1. Four “Analog Input boards (ANIN)” at left.
2. Sixteen “MultiChip Modules (MCMs)” as “column of heat-sinks” in the center.
3. One “Lvds Cable Driver (LCD)” at bottom-right.
4. One “TTC Decoder (TTCdec)” directly above the LCD.
5. One CAN-module at top-right, not yet inserted.
6. One “Readout Merger FPGA (Rem_FPGA)” at center-right (Xilinx).

Another “macroscopic” change is the addition of mechanical “stiffening” for board-insertion/ extraction: the vertical stiffening bars are connected by two horizontal rods (top and bottom), which are supported midway. These rods add significant longitudinal strength to the board, when insertion forces are applied.

Even though, the PPM_1.0 proved to be fully functional with “on-board updates”, all those solutions found were worked into the design/ layout of the new version PPM_2.0. The full set of schematics for PPM_2.0 is available at:

URL: http://www.kip.uni-heidelberg.de/atlas/L1/PP_Module/PPM20_schematics.pdf

The list of issues from the IDR and their implementation.

The following sections list all the items of concern. To impose structure on the list, we choose to group the issues according to functionality. The primary issue in the trigger system is the decision-making real-time data path.

Since a complex, highly integrated system as ours does not permit checks by diagnostic tools (oscilloscopes...), an equally important area is the readout data path.

Issues related to the module as a whole as well as issues concerning the assembly in an entire crate are treated in the last section dealing with the PPM’s infrastructure.

Finally, “production readiness” also requires evidence, that a system assembly can safely be operated. These points are addressed in the last part concerned with “System Aspects”.

The Real-time Signal Path

- Cross talk on ANIN.

Current implementation: The reason for 4-5 % cross-talk on some channels was localised at a quad line-receiver OpAmp. where other PCB routing ran parallel to the input of channel #4 for some millimeters. The solution involved re-routing. To avoid similar problems in a new layout, we decided to use single OpAmps for each analog signal arriving. The second and final ANIN_2 shows no observable cross-talk. The set of schematics for ANIN_2.0 is available at:

URL: http://www.kip.uni-heidelberg.de/atlas/L1/PP_Module/AnIn20_schematics.pdf

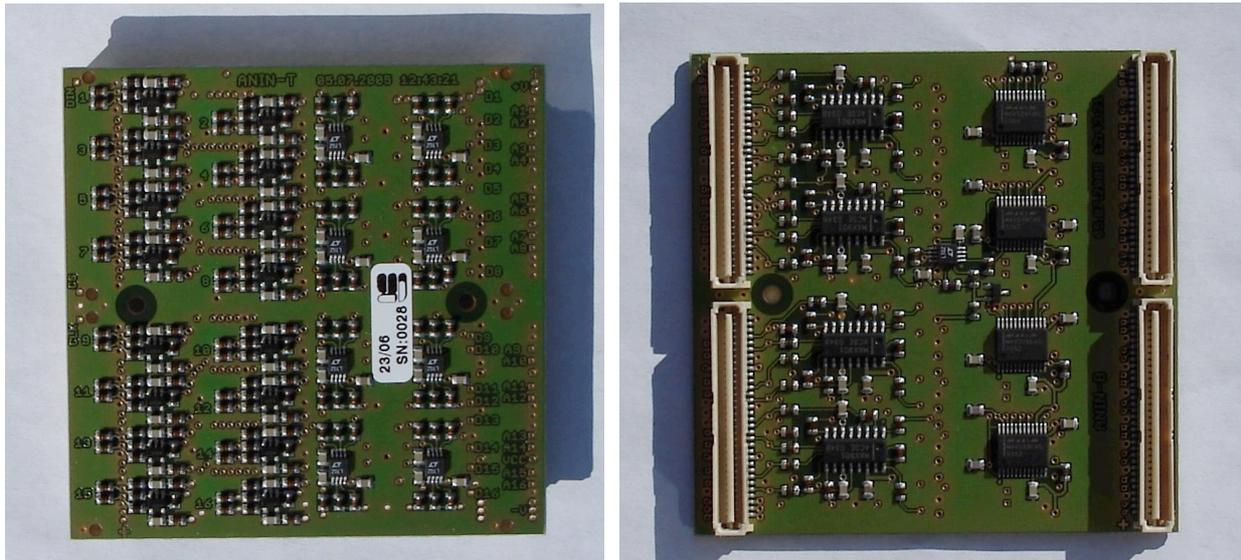


Figure 2 : AnIn daughterboard (left: Top with 16 single OpAmps. right: Bottom)

- ANIN changes on DAC-ranges for signal-offset.

Current implementation: The signal-conditioning is determined by the dimensioning of passive components around the crucial Op.-Amplifier in each channel (see LT1813 in the schematics of the AnIn). The re-layout contains a few passive components, whose values are slightly changed, i.e. “adapted”. This was done to make sure, that all signal-offsets can be adjusted with some safety-margin left for eventualities.

The figure below shows the linearity of the signal-transfer at the entry of the PPM. A differential signal from the calorimeters is transformed into a uni-polar signal for digitisation by the AnIn-board.

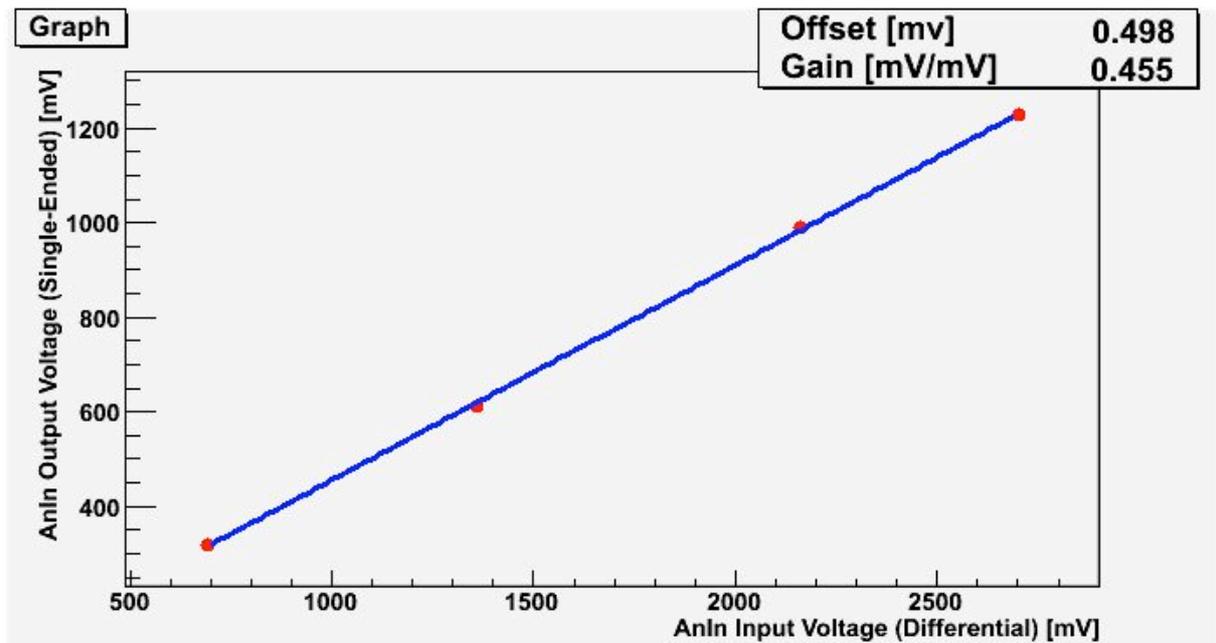


Figure 3 : Linearity of analog signal-transfer in the PPM input to the MCM.

The blue line is a fit to the red points measured. Output of the AnIn-board has been measured with AC-coupling. Therefore, DC-offsets are suppressed.

Gain of Op-Amplifier to map the signal into the FADC window: The feed-back is set to give Gain=0.45 on the signal (a measurement says: Gain=0.455). There is an additional (appr. 20%) signal-attenuation coming from the combination of a “filter-resistor” on the MCM-“balcony” and the input-resistance of the FADC (default-filter-RC is 47 Ohm, 150 pF).

DC-Offset to shift the signal to the FADC window: The FADC reference-voltage is 2.5 Volt (a measurement says: V-center=2.49V). This is the center-voltage around which the FADC digitises. The FADC starts to digitise at a Voltage of 1.92V.

DAC-range to shift a bipolar signal: An absent analog input-signal and DAC=0 yields 1.65V at the AnIn-output pin. One DAC-step corresponds to 2.06 FADC-counts (measured). DAC=255 (full range) would yield 2.25V at the AnIn-output. Hence, an analog signal-input can be shifted by an equivalent of 220 FADC-counts “downwards” and by equivalent of 305 FADC-counts “upwards”. The figure below attempts to illustrate the effect of the DAC.

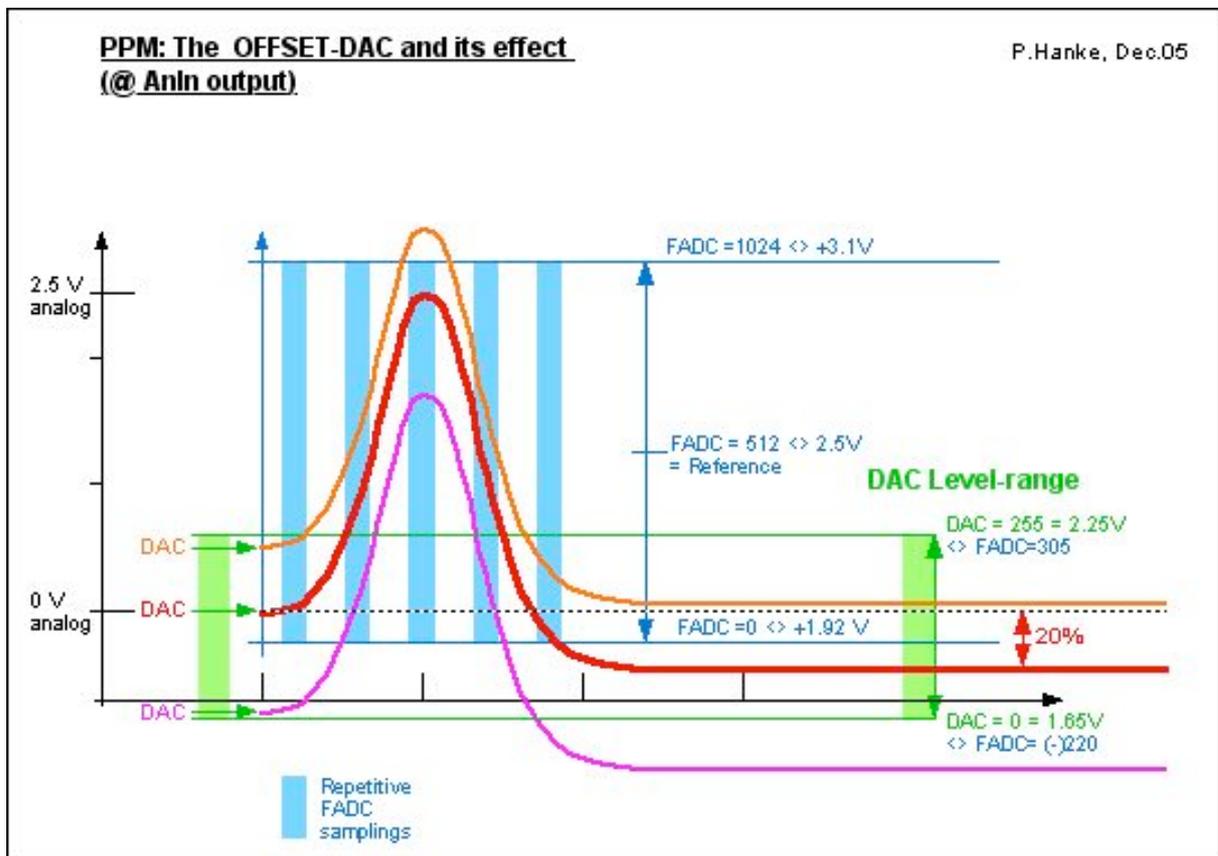


Figure 4 : Effect of DAC-setting on an analog Input-Signal (e.g. from LAr Calorimeter).

The final transformation to the digital scale of the FADC is shown in the figure below.

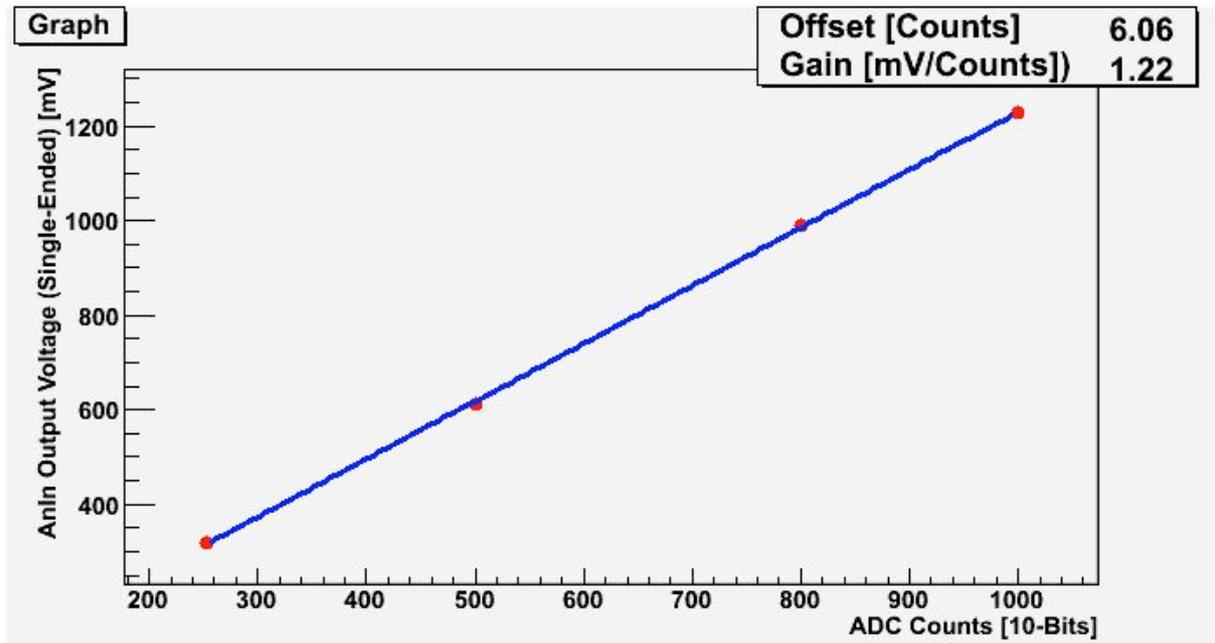


Figure 5 : Linearity of signal-digitisation (FADC on the MCM incl. MCM input-filter).

- Changes around the MCM.

There were absolutely **no changes** to this specific “daughter-board”. Most of our development-effort went with highest priority into the test and design-verification of this essential part of the Pre-Processor. The device was “PRRed” in January 2005. Production is finished meanwhile. There is a sufficient number of MCMs available to cover the needs of ATLAS including a generous supply of “spares” (appr. 50%).

1. Distribution of TTC clock on the PPM for “real-time” pipelining.

Current implementation: The first PPM had unwanted distortions on LHC clock-signals (40.08 MHz) being fanned out through buffers to the MCMs. Part of the layout in the “digital” PPM area was done by an “auto-router” including the clock-lines. The problems were largely cured by a “series termination” at the source . This removes “reflections” and/or “under-/over-shoots”. These terminations are included in the new layout. In addition, the “clock-tree” has been routed with great care manually.

2. I²C clock for serial data.

Current implementation: The clock of the I²C bus was implemented following the “standard” (bi-directional lines, open drain/collector with pull-up). Reliable setting of Phos4-parameters required improvement of the I²C clock-signal distributed to the MCMs (better rise-time of clock; removal of reflections). This is achieved by:

- higher drive-strength using a “push-pull” driver from the ReM_FPGA. This is possible, because the bus is used in “write-direction” only (to the Phos4).
- series resistors on the clock-distribution from the ReM_FPGA, which dampens the reflections. It works also as moderate “pull-up” (see figure below).

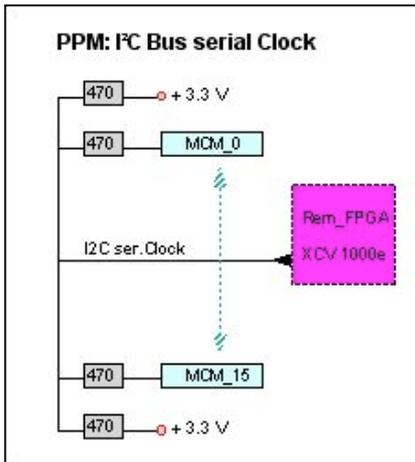


Figure 6 : Termination for I²C clock.

3. Phos4 initialisation.

Current implementation: Occasionally, the Phos4 chips could not be accessed after “power -ON”. This was tracked to some residual voltage (appr. 0.7 V from 5V supply) on the 3.3V at the Phos4. Hence, the 3.3V supply was not ramped up from “0” by the on-board “Hot-Swap” controller. This prevented I²C-access to the Phos4. Insertion of an active “pull-down” (transistor) on 3.3V allows clean initialisation of the Phos4s on the MCMs.

The figure below (colour recommended !) demonstrates the “delay-scanning” capability by stepping the digitisation-strobe in 1 nsec steps across a time-fixed analog signal using the Phos4. There are 24 pulse-shapes (24 Phos4 positions) superimposed in each histogram. Each pulse is read out from the “scrolling” real-time memory holding the FADC values in the PPrASIC.

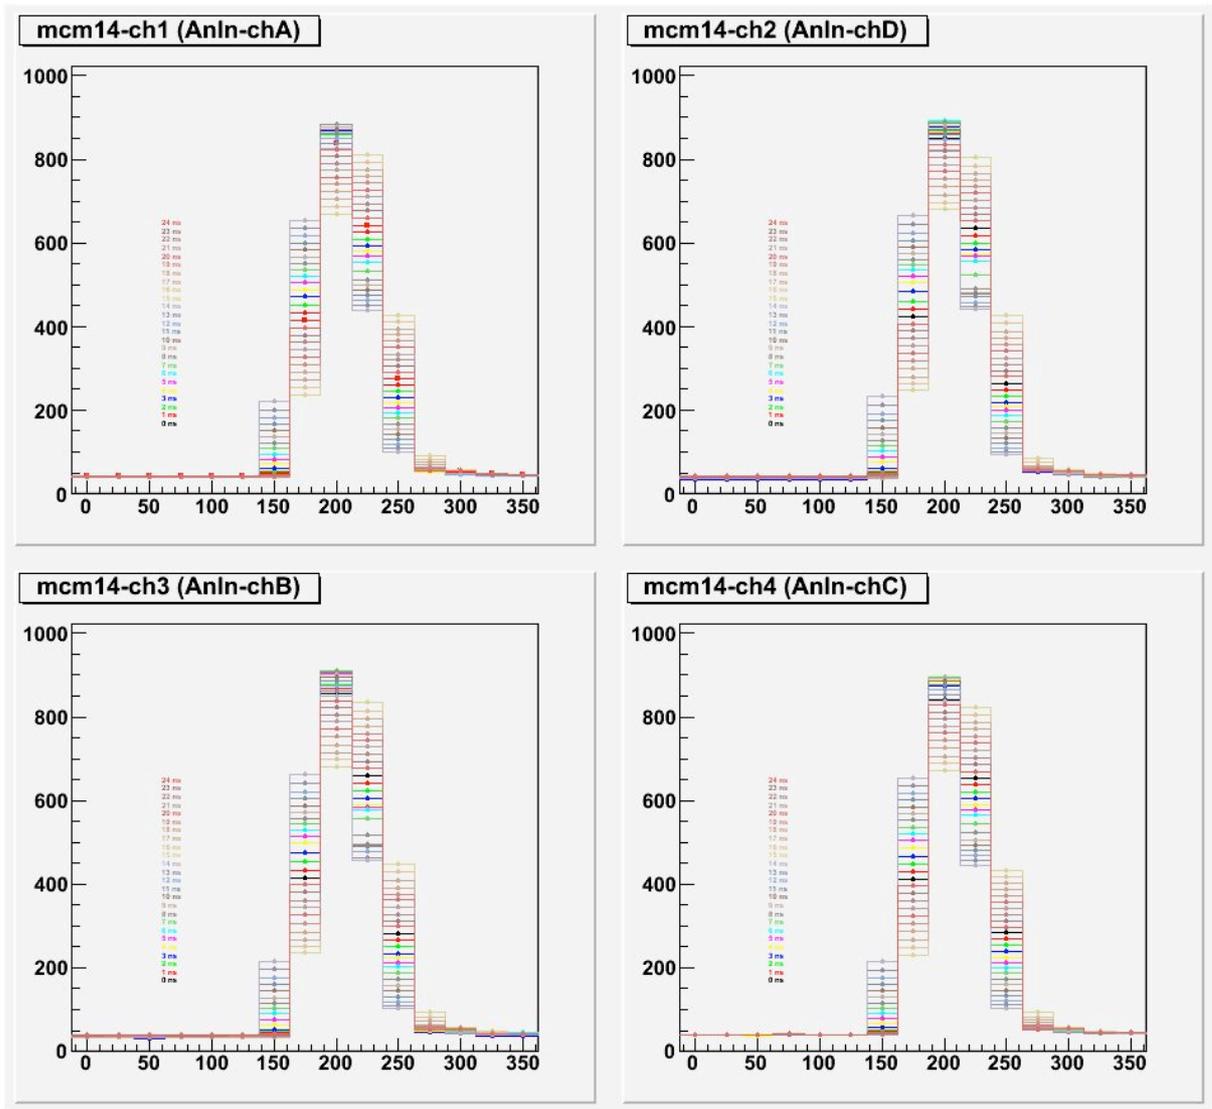


Figure 7 : Scanning an analog signal by 1 nsec steps of the Phos4.

- Pre-compensation on PPM for “MCM to LCD” signal transfer.

Current implementation: The PPM has high-speed serial LVDS links across the printed-circuit board, which run at 480 MBaud over “strip-lines”. The track-length can reach 30 cm. Investigations showed, that the differential LVDS-signal streams suffer from attenuation leading to reduced “opening of the data-eye”. Another problem was identified as “shift of the zero-line”, when the data-content is unbalanced (highly unequal number of “0” and “1” in the serial bit-stream).

The problems are cured by the introduction of “pre-emphasis” on the board directly behind the LVDS-sources on the MCMs. Furthermore, a “pull-up” on the negative signal-branch keeps the data-eye well centered at zero. Both features are implemented in the design of PPM_2.0.

- Re-design of LCD.

Current implementation: Similar problems, as described above, were present after the “LVDS repeater” called “LVDS Cable Driver”. This is the daughterboard driving the long (12

to 15 m) cable-lines to the subsequent Processors (Cluster Processor, Jet-Energy Processor). In fact, some effects observed on the board were “magnified” after the cables. LVDS cross-talk was also observed due to non-optimal positioning of the “pre-emphasis” components.

The two FPGAs implemented were found to be “over-loaded” driving the required number of outputs. These findings led to a complete re-design of the fanout and cable-driving daughterboard (LCD).

A new LCD_2.0 was developed with **four** FPGAs and modified usage of FPGA-internal banks of LVDS buffers. The passive RC-network for Pre-Emphasis of each signal was placed in the PCB layout as close as possible to LVDS-drivers on the FPGAs. All this resulted in very good transmission quality for these 480 MBaud data. Bit-Error Rates (BER) were measured to be so small, that safe operation of the Trigger is achieved, i.e. $BER \leq 1/10^{*}15$ guarantees a sufficiently small contribution of “fake triggers” in ATLAS.

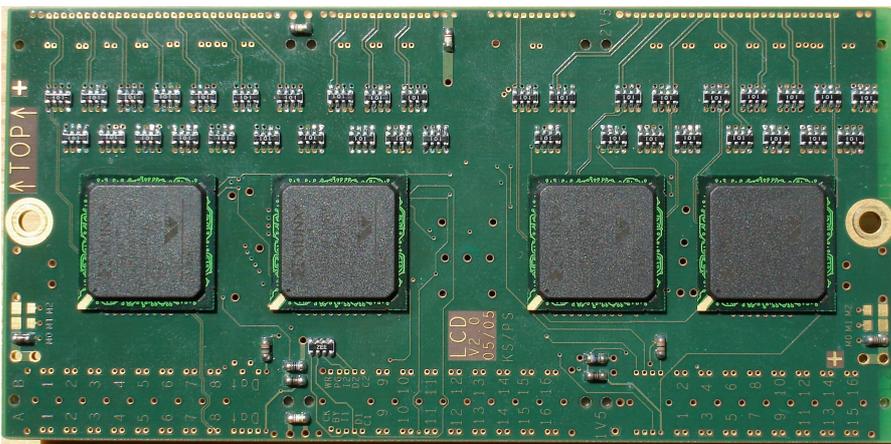


Figure 8 : LCD 2.0(top) with 4 FPGAs and Pre-Emphasis circuits (Rs) at outputs.

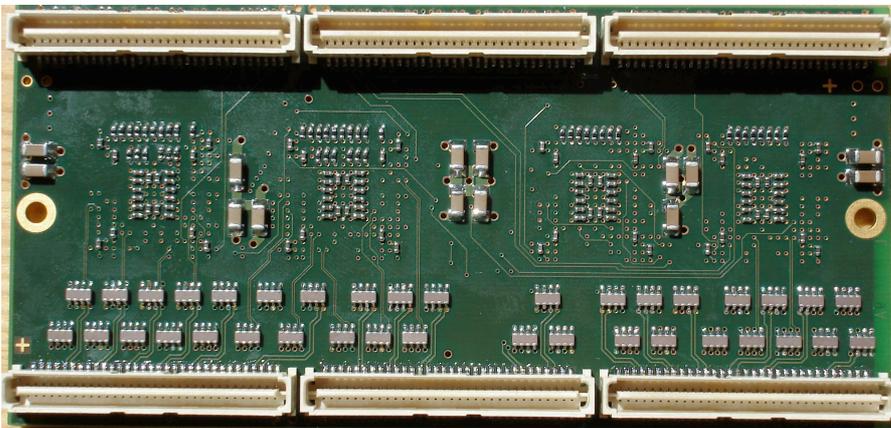


Figure 9 : LCD 2.0(bot) with Pre-Emphasis circuits (Cs) at outputs.

The combined improvements are shown in the following figure comparing the first design-approach to the current implementation.

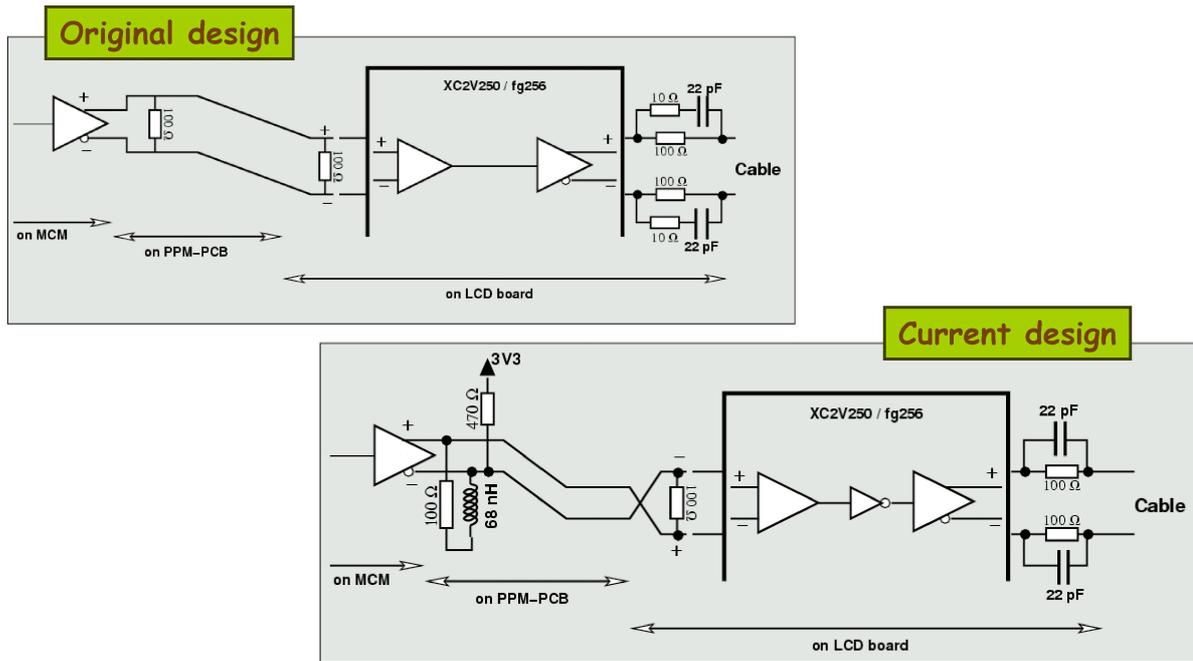


Figure 10 : The entire “pre-emphasis chain” on the PPM.

Note the LC pre-emphasis directly after the LVDS-serialiser on the MCM, the “pull-up” and the “inversion”, which is re-inverted in the driving FPGA. The “inversion” improves the data-eye by using the sharper high-low transition in chips.

The set of schematics for LCD_2.0 is available at:

URL: http://www.kip.uni-heidelberg.de/atlas/L1/PP_Module/LCD20_schematics.pdf

LCD_2.0 was re-developed already for PPM_1.0. Many test-sessions at various laboratories have shown the viability of the new design, which is proven by results as mentioned above. The daughterboards for PPM_2.0 are absolutely identical.

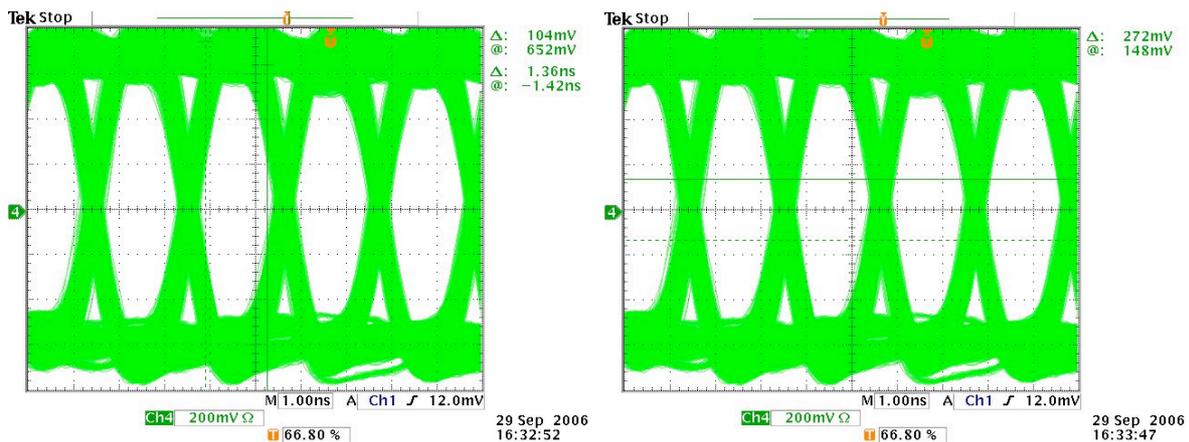


Figure 11 : LVDS data (left “eye-width” \gg 1 ns, right: “eye-height” \gg + - 100mV)

The figures above are shown as “re-verification” of the LVDS output quality. Note the “opening” of the data-eye after a 15m AMP twin-pair cable: the width is 1.2 nsec at + - 150 mV amplitude. Both quantities lie well above the requirements of a LVDS receiver.

The Readout Path.

- Firmware (i.e. loading of bit-files to devices on PPM).

Current implementation: The CPLD, handling VME traffic, had very limited resources left over to serve the Flash-RAM loader. Demands have increased due to the LCD_2.0 upgrade (4 FPGAs). The VME prototyping, however, was well tested, hence should not be changed in a new version of the PPM. For this reason the design of PPM_2.0 required a SECOND CPLD to handle Flash-RAM loading. This leaves the VME CPLD literally untouched. The new, “widened” scheme is shown in the **Appendix** (PPM_VME block-diagram - preliminary).

The Flash-RAM has a capacity of 8 MBytes, which can be allocated as bit-file store in the following way:

Six versions (production/ debug) of ReM_FPGA code can be stored (with a max. size of 1 MB each).

The code for each of the four “routing / fanout” FPGAs on the LCD-board with a max. size of 250 kB each.

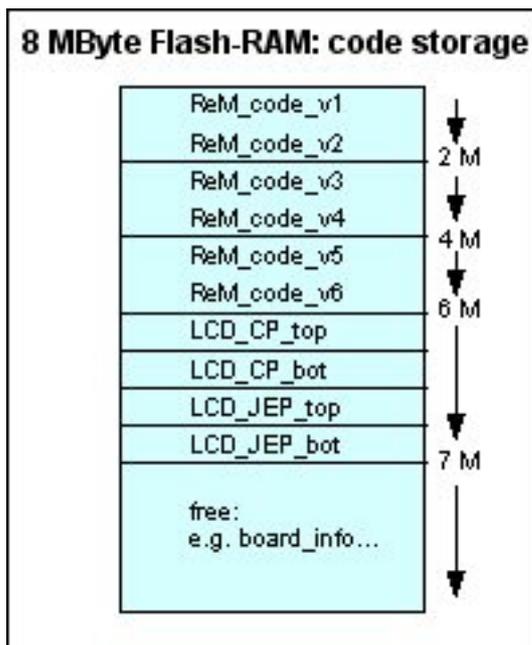


Figure 12 : Space allocation for “firmware” in Flash-RAM.

- Firmware-code

is not subject in this review, but has been developed much further during tests. The firmware-code has now reached a status of “stability”. NOTE, that there are no limitations due to availability of resources on the implemented XCV_1000E. Readout at full speed and specified max. data-volume have been verified using PPM_1.0.

- SRAM usage outside Rem FPGA

Current implementation: The PPM_1.0 held a Static RAM for fast data-storage at the “disposal” of the ReM_FPGA. Its size was 1 MByte (256k * 32 bits). The same component became available with larger capacity. The pin-compatible version providing 4 MBytes is implemented in PPM_2.0. Storage of “hardware” information like “unbiased rate/ cell” or “unbiased energy-histogram/ cell” is possible without space-restriction.

- Optical G-Link (RAL-board in the rear of the PPr-Crate).

Current implementation: A “electrical” G-Link connection was used in conjunction with PPM_1.0 to transmit event readout data to the 6u-RODs. It was decided later in the L1-Calo community to use the G-Links with optical data-transmission. Hence, another version of the “Rear G-Link Transmission Module (RGTM)” was built. The new RGTM-O (“Optical”) is “standard” from now on for sending Readout data to the new 9u “ReadOut Drivers (RODs)”.

POWER and CLOCK come from the PPM board by pin-through. Hence, switching ON/OFF the RGTM-O is controlled by the PPM in the same slot. Also, the G-Link status is available on the PPM for evaluation.

The RGTM-O has been successfully tested on the “back” of PPM_1.0 and with a 9u-ROD. As there are no changes in ReM_FPGA firmware nor on the parallel data-output of PPM_2.0, the readout functionality is only going to be RE-verified with the new PPM.

The PPM’s infrastructure, VME, TTC etc.

- “Status LEDs” at the front-panel.

Current implementation: The PPM_1.0 has some of built-in front-panel LEDs activated. For all of the LEDs, a “status” was attributed to make visual control for PPM_2.0 easier. The LED-attribution is shown in the figure below. Experience from tests has initiated this LED-allocation, which tries to follow the “colour-convention” once established for the Level-1 system.

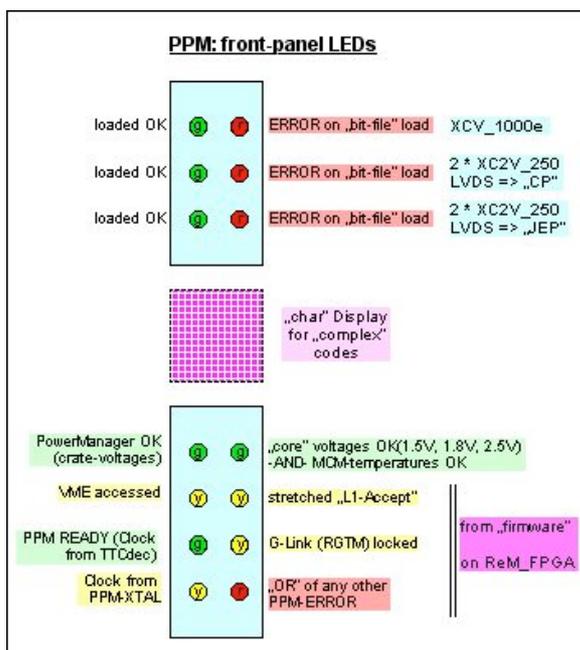


Figure 13 : The “illuminated” PPM front-panel (LEDs), if visible “behind cables”.

- VME-CPLD (see also: “firmware loading”).

Current implementation: The CPLD-code for VME-communication is well debugged on the existing PPM_1.0, hence shall remain “untouched” in its functionality. A second CPLD-device assumes the task of “Flash-RAM administration” (see above).

- TTCdec daughterboard.

Current implementation:

1. Meanwhile (for PPM_2.0), a new configurable TTCdec card has been designed, which can be used anywhere in the Level-1 Calorimeter Trigger system. It provides a “jumper-field”, which allows either to include a PLL in the clock-output or to use the clock-output without the PLL.

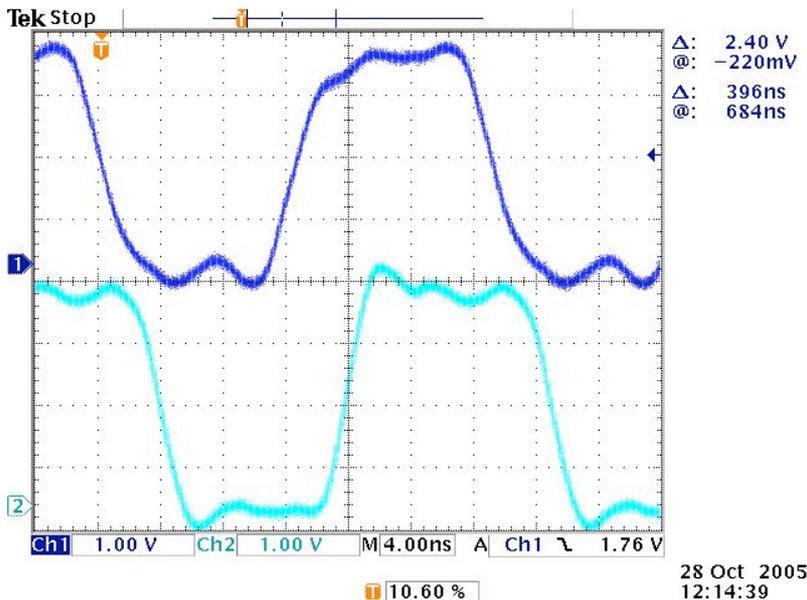


Figure 14 : TTCdec clock – with PLL (upper trace) – without PLL (lower trace)

2. Mounting holes on the new TTCdec-daughterboard are changed compared to the original version. The new PPM layout adapts to the mirrored hole-positions to have proper mechanical fixation also for the new TTCdec board.

Distribution of the TTC-protocol (as well as CAN and extended addressing for crates) across the crate-backplane is described below under “System Aspects”.

- CAN-Controller.

Current implementation: The Level-1 community has agreed on a certain micro-controller to be used on all modules. A newly developed daughterboard holds the Fujitsu MB90F594. The schematics and the physical layout of the daughterboard on the PPM is shown in the **Appendix**. The CAN controller is accessible from VME for local testing (e.g. MCM temperatures). A “minimal” interface consisting of 8-bit data (read) and 4 control-lines (write) to/from the VME-side is sufficient.

The ATLAS CAN-Bus infrastructure is carried to each PPM-slot on a special backplane plugged onto the P0 of VME64xP. PPM_2.0 connects the CAN daughterboard to the respective pins.

A prototype of the CAN daughtermodule is assembled and has been successfully powered on the PPM. A test with CAN-software might be carried out at CERN.

- “Hot-Swap” controller.

Current implementation: “Live-insertion” of a second PPM into a crate showed an undesired effect. When the “hot-swap” controller ramps up the inserted PPM, an in-rush of current results in a “dip” followed by a “spike” on the 3.3V of the crate. The amplitude of the dip/ spike is only small, but sufficiently large to trigger the controller of the first PPM into “switching off”.

The controller’s sensitivity is +- 10% (i.e. +- 330mV) and “fast” reacting. It is decreased by “flattening” the spike (passive filter on the sense-line). Now, the current-surge on the 3.3V does not lead anymore to “switching-off” a neighbour-PPM in the crate.

Another measure taken is a “delayed” switch-ON of the bank of LVDS-buffers. These buffers interface the control/readout of MCMs to the board (see picture of PPM: on the right of MCMs). Thus, the current leads no more to a dip/spike on 3.3V at all.

- Use of “Hot-Swap”.

Current implementation: The originally intended use of “hot-swap” implied only, that power at the PPM-board is locally switched-off for exchange of the module. It does NOT ascertain quiescent VME-bus traffic. Consequently, a reboot of a crate-controller or other non-intended actions could occur when a module is powered down.

The undesired occurrence is eliminated by powering the bus-defining CPLD from a source, which goes OFF last (5V instead of 3.3V). This prevents “undefined states” on VME bus-lines leading to undesired actions in the crate.

- Miscellaneous: PCB, assembly....

Manufacturing of printed-circuit boards (366 mm * 399.5 mm) has proven to be non-problematic. The provider has delivered electrically tested boards of very good quality. This was the case for the first batch (4+ pieces for PPM_1.0) as well as for the second batch (20+ pieces for PPM_2.0).

Assembly of PPMs: The assembling company must watch the alignment of “Mezzanine” connectors connecting the daughterboards. Those connectors have, unfortunately, no mechanical alignment pins. Hence, alignment is optically controlled at automatic placement onto the solder-pads.

Apart from this “special care”, the assembly followed the company’s standard procedure including optical QA control before delivery. On the bench-tests of 20 assembled PPMs, only some faults showed up. While BGA-soldering of 1.27mm pitch FPGAs showed no problems, some “shorts” were found under the 0.5 mm pitch MCM-sockets. Those connectors were replaced at KIP. It should be noted, that this action requires extraordinary skill from the performing person.

System Aspects.

- Power-Supply to serve a fully-loaded crate.

A first run-up of a fully equipped PPr-crate showed unexpected behaviour. The arrangement of the crate-cage and the power-supply followed closely recommendations from the manufacturer (Wiener) and CERN-experts (EP-Pool): “the length of low-voltage power-leads should not exceed 0.5 m” (see figure below). The grouping of the cables is “as delivered”; only the mech. connection to the bolts on the PSU was done in setting up the crate.

When the crate was filled with modules to half-capacity (7 or 8 PPMs out of 16), a strong fluctuation on the current meter of the 3.3V supply was seen. Closer investigation revealed not only huge current-variations, but also a big voltage-oscillation on the crate supply-backplane (see “oscilloscope” in the figure below). Similar observations were made in a loaded “Cluster Processor” crate at Birmingham University.

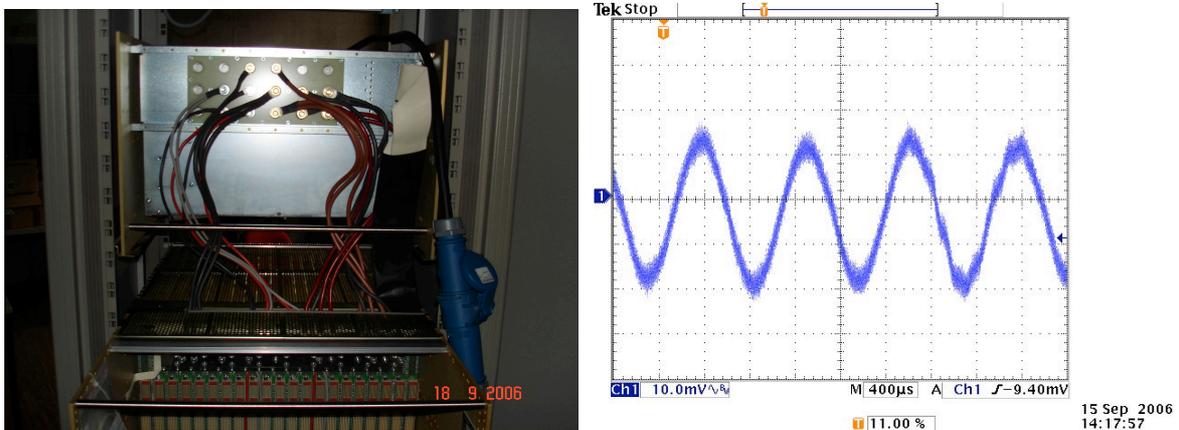


Figure 15 : The “original” supply-cabling: the 30 mV, ~1kHz oscillation on 3.3V.

The voltage fluctuation at appr. 1 kHz is far too large in amplitude (30 mV) compared to the requirement (< 10 mV). An estimate of a PPM’s capacity pointed to the order of milli-Farad ! Combined with an inductance of micro-Henry, a circuit resonating at appr. 1 kHz is plausible.

The inductance is provided by the separation of the “forward” current on 3.3V (multiple brown leads on the right in photo above) and the “return” current (multiple black leads on the left in photo above). Prove is given by the fact, that the frequency can be changed by “tuning” the inductance, i.e. changing the “coil-area” by moving the brown and black leads towards/ apart from each other.

The solution requires minimisation/ elimination of the inductance, i.e. rewire and tighten forward/ return to “zero” inductivity (see photo below). Inductance and combined module-capacity are not resonating any more in a disturbing frequency-interval (see “oscilloscope” in the figure below).

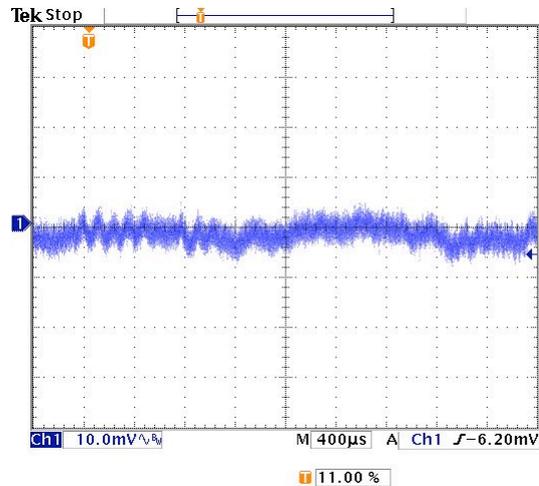


Figure 16 : The “zero-L” cabling: the <math><10\text{ mV}</math> ripple on the 3.3V supply.

- Current on 3.3 Volt under “heavy-duty” data-pattern.

Starting PPM operation by sending so-called “stress-patterns” continuously from “playback” memories down the RT data-paths of a PPM on all 64 channels simultaneously, challenges the on-board power controller by a fast change of current drawn . The result is an “overload” condition followed by a “switch-off”.

Such a starting condition is un-realistic. Nevertheless, re-programming the controller raising the current limit by 20% eliminates the problem,

- Total Consumption from the crate’s Power-Supply Unit (PSU).

Filled up with 16 PPMs and 16 RGTM in the rear the crate consumption amounts to:

- appr. 175 A on +3.3 Volt,
- appr. 150 A on +5.0 Volt.

This lies well within the capability of the PSU; it actually demands little over 50%. However, the current on 3.3V will still rise somewhat, when real-time LVDS outputs are terminated. Note, that the 5.0V current includes the G-Links on the “Rear G-link Transmission Module (RGTM)”.

- Temperature of MCMs.

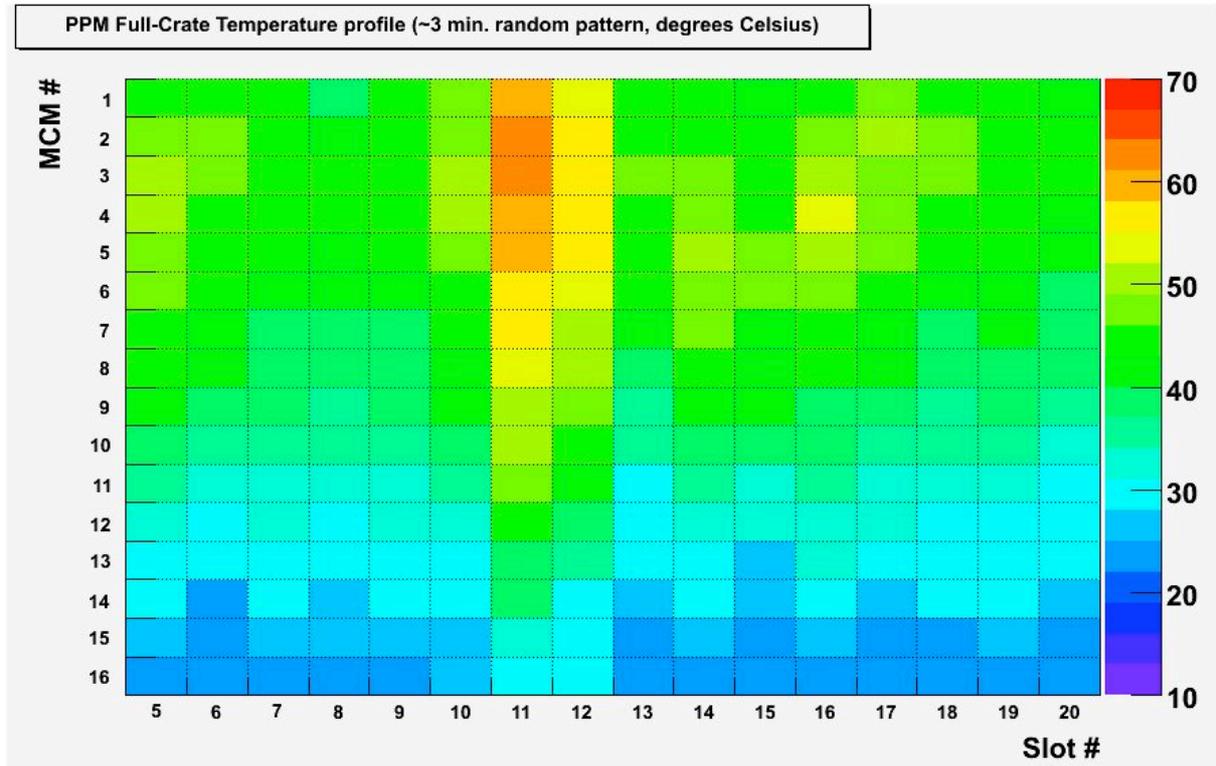
A major concern about system viability was the efficient cooling of densely packed, high-power electronics like our PPM arrangement. MCMs have been heated up 120 °C and been found fully operational afterwards.

The crate-environment at KIP has a “final crate” with its integrated cooling-unit, but no water-air heat-exchanger underneath (as in ATLAS-USA15). Hence, the set-up represents a “worst case” scenario using ambient air only.

Yet, operating all PPMs does not at all reach critical temperature values. Each, PPrASIC on a MCM is equipped with a temperature measuring diode of $\pm 2\text{ }^{\circ}\text{C}$ accuracy. The readings

are plotted in the figure below. The measured temperatures reach from 30 °C on the “bottom” MCMs next to the air-entry to appr. 60 °C for the “top” MCMs at the air-exit.

The temperature distributions for all MCMs in the PPr-Crate are shown below. The figures show the MCMs as seen from the front of the crate: slot# gives the position of the modules, MCM# shows the installed “column” (see PPM photo). Cooling air enters at the bottom of the crate, which matches with “cooler” temperature there.



Note that slot-numbers 11, 12 are “under-privileged” for cooling. The TWO “round” fans have a “dead zone” there. The “air-buffer” of 1u above the actual fans does obviously not redistribute ideally.

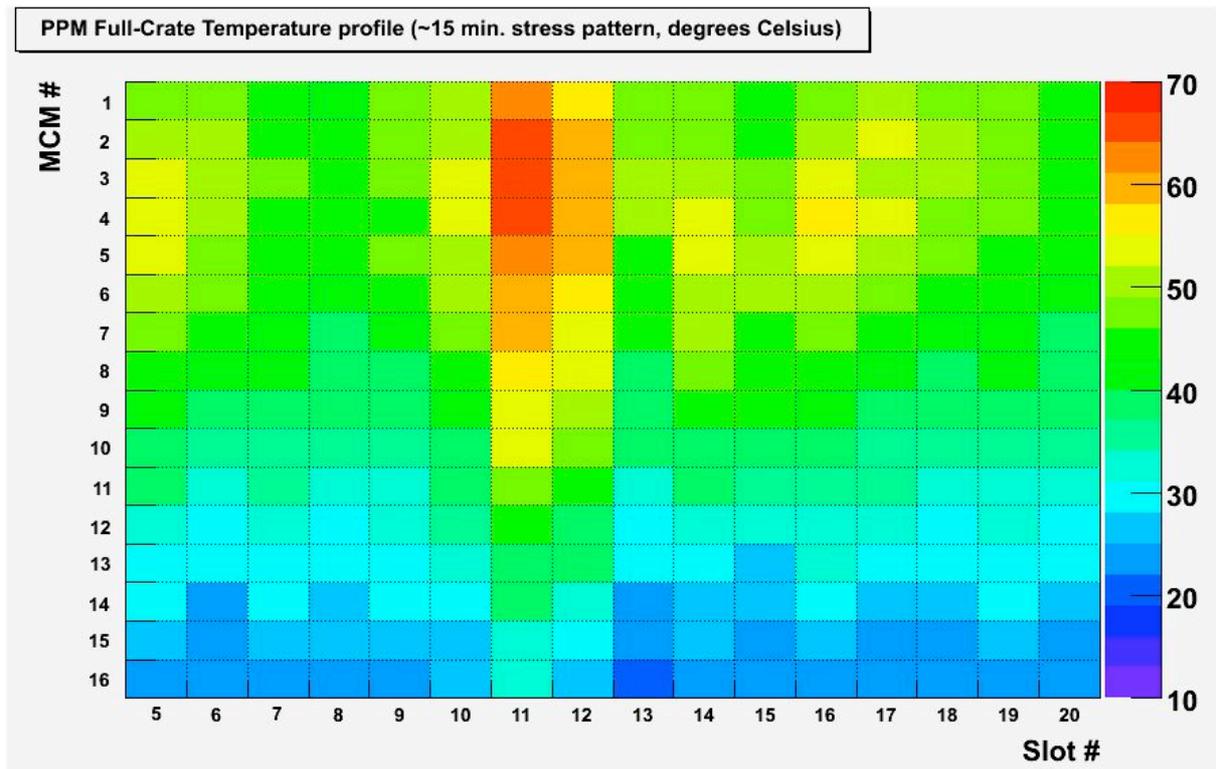


Figure 17 : MCM-temperatures in a powered crate.

A worst case demonstration is performed by loading all “playback” memories of the 1024 PP-channels installed with so-called “stress” data-patterns. Even then, the “underprivileged area” does heat at most into the 60s °C. Note, that MCMs have been exposed to temperatures up to 120 °C and worked afterwards. Note as well, that the crate at KIP only pushes ambient air into the gaps between modules. The installation in ATLAS-USA15 has water-cooled heat-exchangers under the crate.

- TTC-distribution across backplane to 16 PPMs.

RAL has produced a printed-circuit board to distribute the TTC-protocol including the LHC-Clock (as well as CAN-Bus and extended geographical addressing) from a “Trigger Control Module (TCM)” in slot#21 to the PPr-stations (slot#5 to slot#20). The same need exists in Level-1 ROD crates, hence the “common” implementation.

The PPr application requires different connectors, because access for additional services to the PPMs is wanted. The PPr-version is equipped with “long pin” versions of the “Compact-PCI” connectors protruding through the PCB. This allows insertion of a “servicing plug-on” for e.g. re-loading resident firmware on CPLDs.

A tool for insertion/extraction of this backplane-spanning device has been developed. It shall allow safe placement/removal of the TTC-distribution board.

- Strain-relief on “massive” LVDS cabling.

Each of the PPMs carries 16 connectors with LVDS real-time output (10 for CP, 6 for JEP). Each of those connectors holds 4 twin-pair cables bundled. Hence, the rear of the crate must accommodate 256 cables of the kind. More importantly, the cables hang for up to 2 m vertically in the before they reach a supporting horizontal cable-tray.

It is necessary to hold the cables mechanically. The strain from their own weight cannot be taken by the “flimsy” Compact-PCI connector in the backplane. The arrangement shown in the photo below has been designed to ensure a viable installation for the years of ATLAS data-taking.

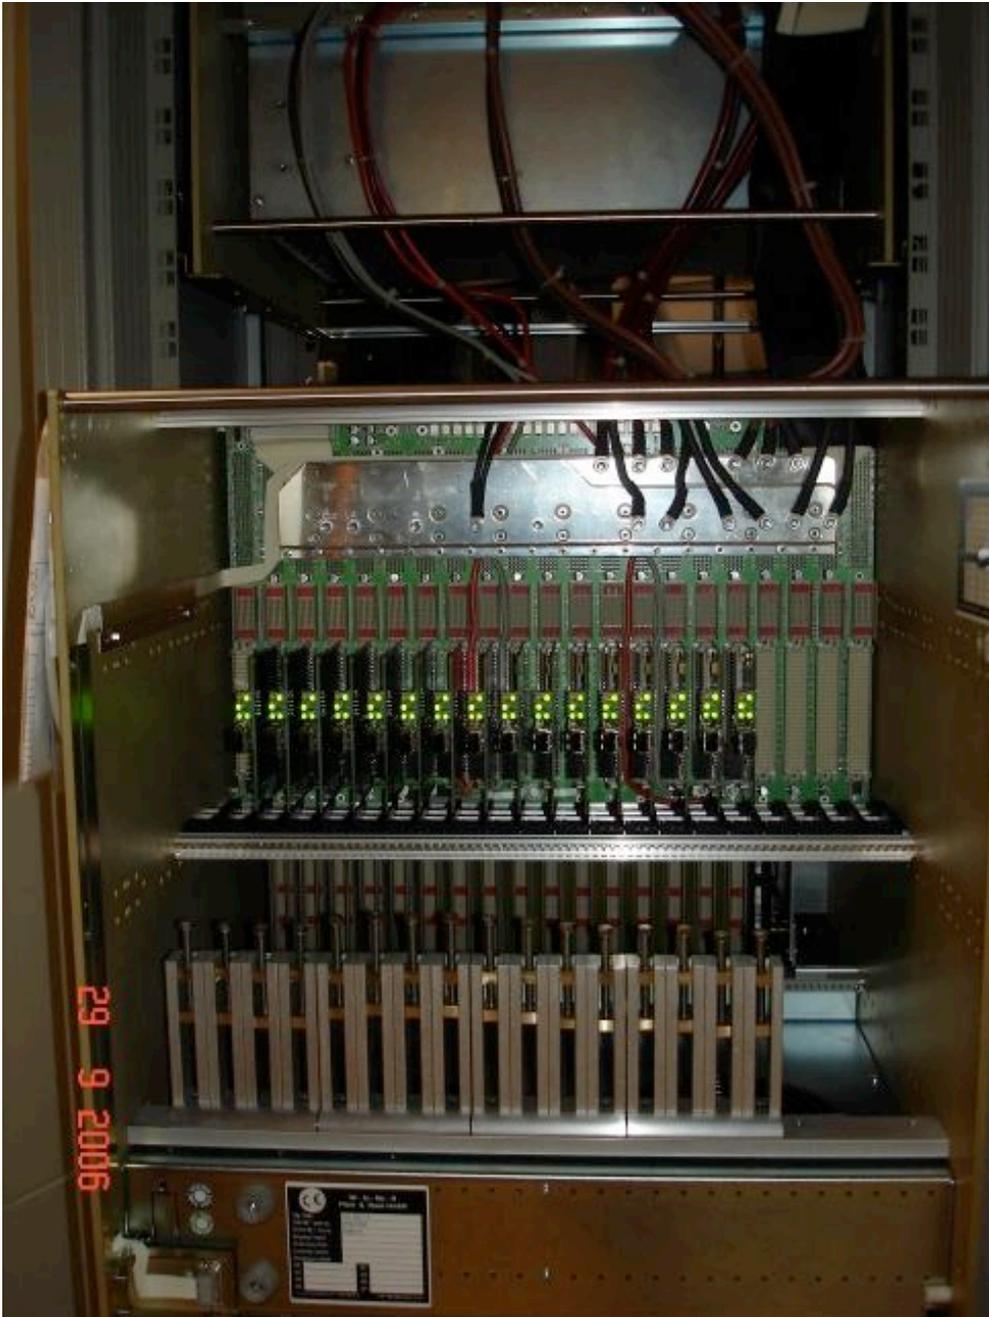


Figure 18 : The rear of a PPr-crate: rewired power, RGTMs, “strain-relief” for LVDS.

- Phos4: 256 of them in a crate.

The Phos4 as a component has proven to be problematic in all stages of the Pre-Processor development. The missing facility of an “external reset” led to problems and “work-arounds” with the hopeful conclusions that the “problem is solved”. It started with testing of the manufactured chips. It continued with testing devices, where the chip was implemented. Examples are our MCM or the “Mustard” and “BOC” of the ATLAS Pixel Detector . (see: http://sctpixel.home.cern.ch/sctpixel/PhonePresentations/20020807/PHOS4_Fix1.pdf)

For the first time a system was configured, which employs a large number of these chips: the Pre-Processor crate holds 16 Phos4s on each PPM amounting up to 256 Phos4s “running-up” simultaneously.

And here it is, where the “initialisation” problems re-occur on a small number of all those chips in a badly-reproducible way. Quantitatively, there are any number ranging from 3,4 or 5 up to 20 or 30 out of the total 256 Phos4s, which do not start proper operation.

Improvements implemented after the IDR (see also further up in this document) are necessary, but still do not guarantee that all the Phos4 can be initialised reliably. The device contains feed-back for a Delay-Lock-Loop to regulate the delay-time /step to 1 nsec accuracy. Initialisation starts with powering the chip followed by ascertainment of a reference clock (40 MHz XTAL). The actual initialisation is finished, when a “dummy”, negative-edge on the I²C bus-clock arrives.

Few of the chips do not “start”, i.e. the feed-back is latching either at the slow extreme (appr. 1.6 nsec/step) or at the fast extreme (appr.0.6 nsec/step).

Very time-consuming investigations resulted in two observations:

1. If “Switch-off, switch-back-on “ is too short in time, the “Power-Off” due to some storage-capacity is not apparent to the Phos4 chip, i.e. the “latch-up” is still there, not reset as wanted.
2. The “Reference clock” MUST be “symmetric”, i.e. the two half-periods must have identical length within a few percent. Our reference clock shows a 60/40 ratio, which is not tolerated by some Phos4s. Already the XTAL-output is slightly asymmetric (55/45). This is, unfortunately, magnified in the buffer/fanout tree following, where slightly off-centred switching-thresholds bring the final ratio to 60/40.

The possible solution, which can be implemented with very little change to the PCB-layout:

- Prevent “immediate” re-powering of the PPM after “Power-Off”.
- Produce a “symmetric” Reference Clock for initialisation of the Phos4.

“Immediate” re-powering.

The on-board controller allows introduction of a delay (few seconds), which lets remnant charges/ voltages disappear. No change to the layout is required.

The “symmetric” reference clock.

This modification is more complicated. There is a possibility to use a DLL in the Rem_FPGA, which produces a 40 MHz clock with a good duty-cycle (50/50). The aforementioned effect of the buffer/fanout does not deteriorate the symmetry enough to dissatisfy Phos4 requirements. This has been tested on all 20 PPMs.

Hence, it is suggested to implement a scheme as described in the figure below, where both versions are compared. The first buffer is used as “sub-grouped” device, i.e. the “output enable” for the groups are used separately.

One group passes the XTAL permanently. There are still “customers”, which need this clock (e.g. CPLDs),

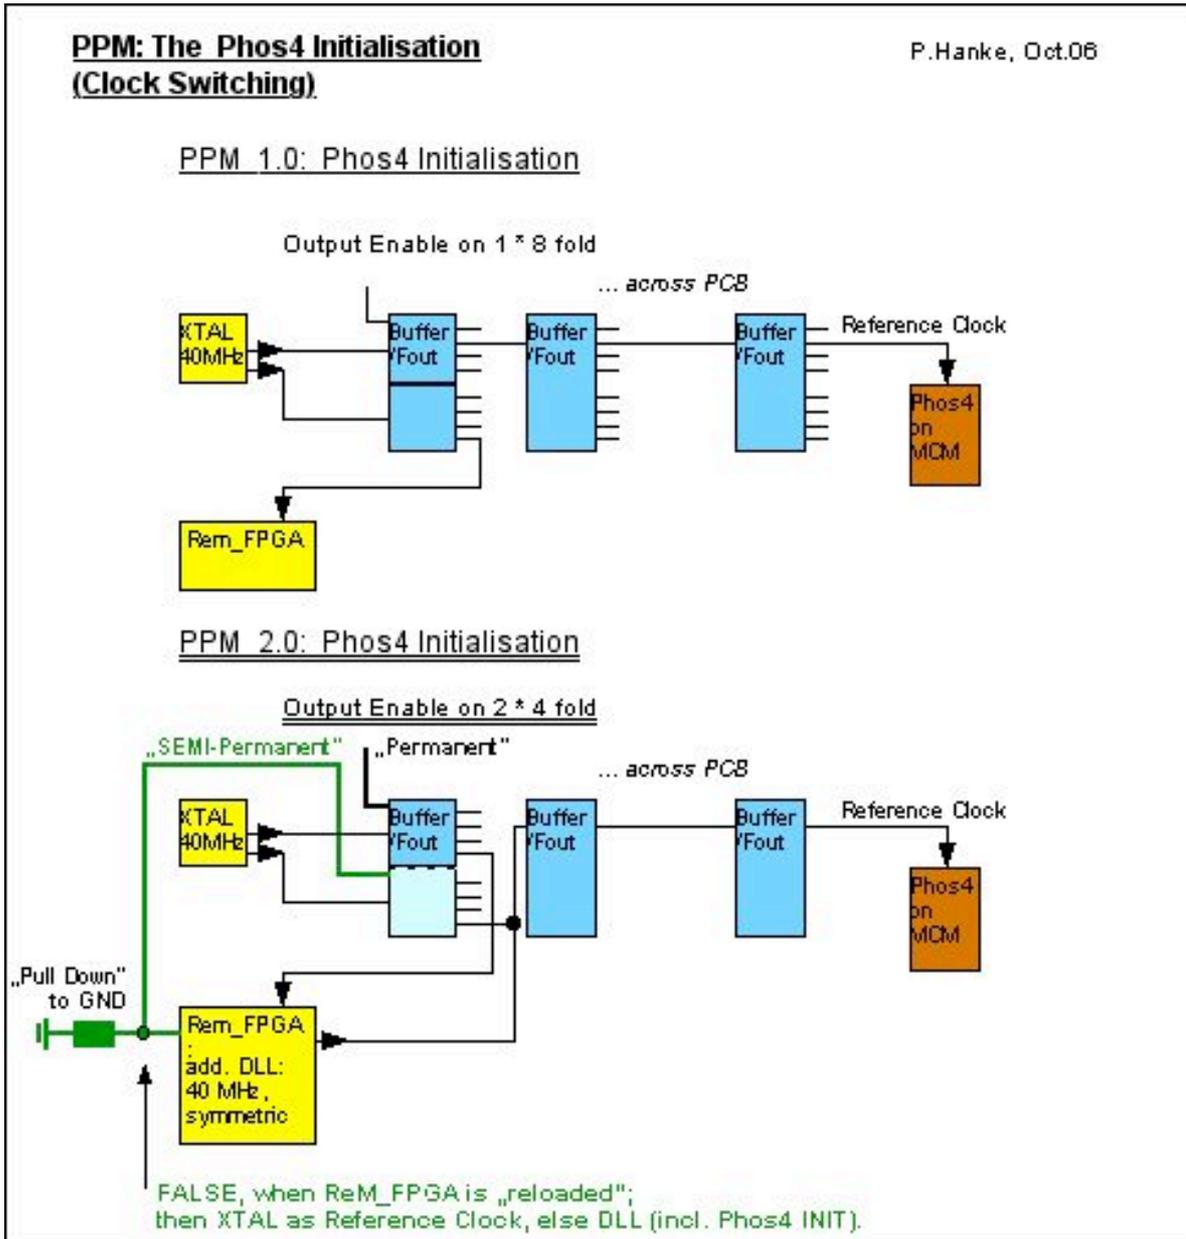


Figure 19 : Phos4 operation (and INITIALISATION) with “symmetric” Clock.

The other group’s outputs are controlled by the ReM_FPGA. It passes the “DDL-ed” clock down the tree after running up of the PPM. The Phos4s are initialised with this clock. However, a “reload” of firm-ware in mid-operation would stop the reference clock to Phos4s. This is an unwanted situation just like a drop-out of the TTC, which will also have to “bridged”. In this case the “permanent” XTAL from the first group keeps the Phos4 going.

To show the effect of non-initialised Phos4s, the same histograms as above have been produced. If the DLL feed-back is not in action (“stuck”), then the step-size of the Phos4 is either 0.6 nsec (“fast” operation) or 1.6 nsec (“slow” operation). In either case, the stepping of 0.6 nsec (or 1.6 ns) *24 does not cover the 25 nsec LHC clock-cycle. Consequently, certain parts of the pulse are never reached (or attributed to the wrong clock cycle).

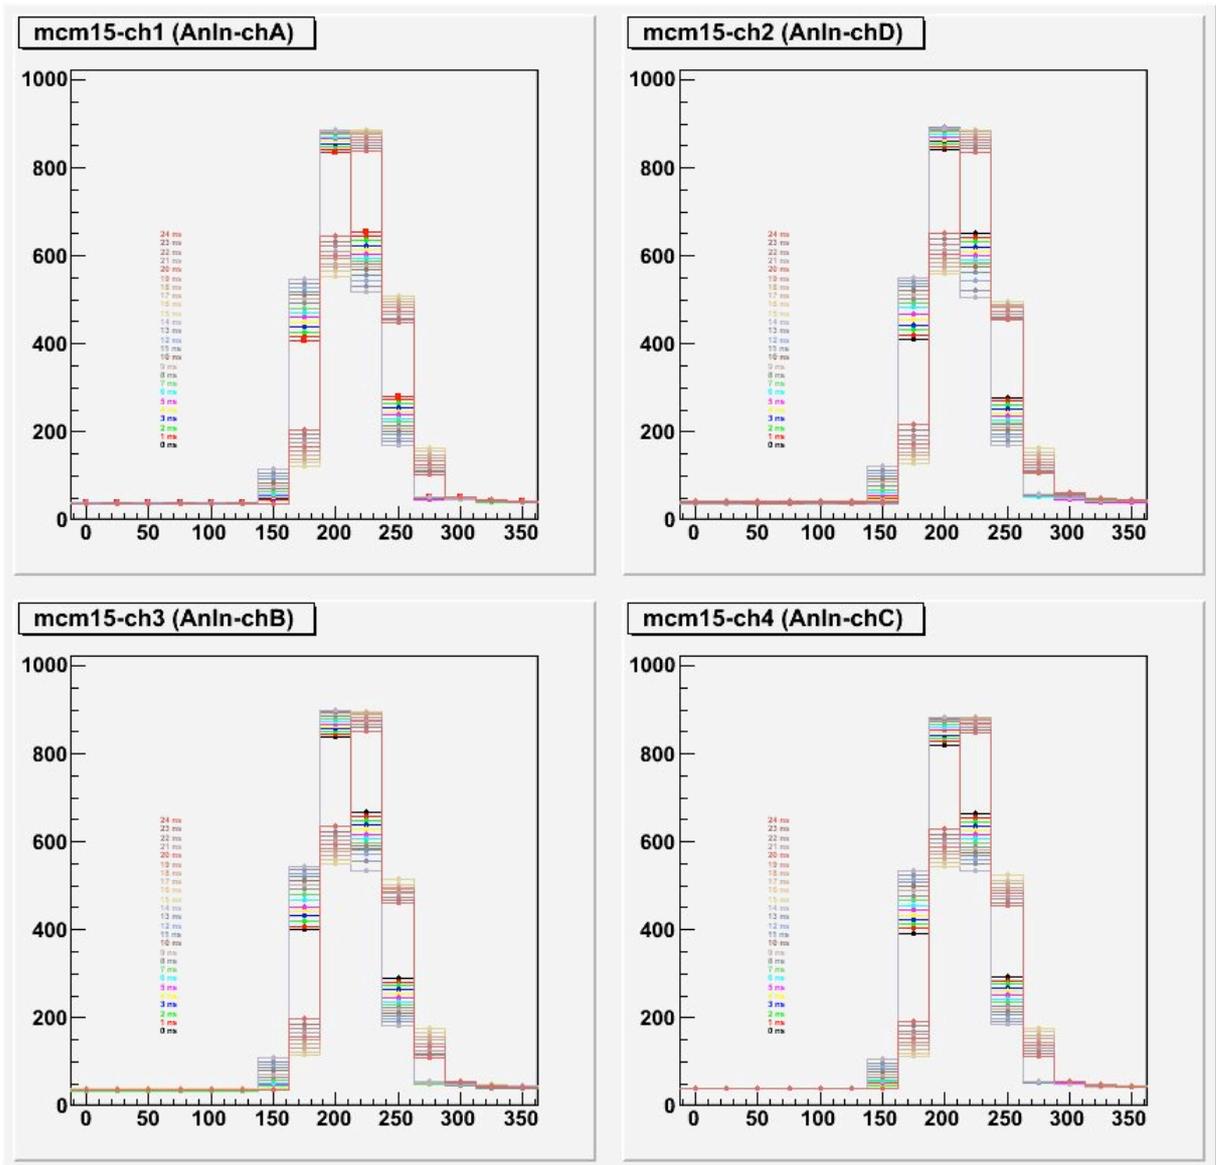


Figure 20 : Scanning an analog signal with a “STUCK” Phos4.

There is a “sacrifice” to the new scheme outlined above. Originally, the two buffer-groups were meant to serve the “real-time” path and the “readout” path separately. This was motivated by the “option” to use a different frequency (e.g. 60 MHz or other) for clocking the “readout” separately for higher bandwidth. This option is no longer available now. On the other hand, why should a sender (PPM) provide something, that is not acceptable to the receiver (ROD).

Summary, Schedule, Production Testing.

- Summary.

KIP Heidelberg has built a “pre-series” of 20 PPMs with all the improvements implemented, which were collected in the “Interim Design Review”. 16 of those PPMs are used to test operation of a full PPr-crate.

Layout work, starting from the reviewed schematics, was done with outmost care and to a very large extent “by hand”. The PCB-part, where most work has gone in, was the “rear part” of the PPM, i.e. the “digital” side behind the MCMs.

The “pre-series” of 16 (+ 4 = 20) FINAL PPMs is operational. They make up one final PPr-crate. It is the first of the required 8 identical crates in ATLAS-USA15. Tests of full-crate operation shows the viability of the system for running in the experiment for many years of ATLAS-DAQ to come.



Figure 21 : The fully-assembled PPr-crate with a TCM, 16 PPMs and a 9u VME-CPU.

- “Full Production” planning.

Production of PPMs (160) will take place straight after this review still during the last quarter of 2006 starting with the printed-circuit boards followed by assembly. In parallel, the daughterboards will be produced following the same pattern.

Scheduling of all production steps is determined by time-intervals given to us by the respective manufacturers. There are very limited, if any at all, possibilities to accelerate production:

- ⇒ gain is marginal (1-3 weeks), but additional cost is increasingly prohibitive to have PCBs earlier than the standard way.
- ⇒ neither the PCB maker nor the assembler can be changed (or another engaged in parallel) at this stage of the project.
- ⇒ both steps (PCB, assembly) have a risk of delay due to high demand imposed by “big customers”. KIP has tried to minimise the risk by early announcement (“reservation”), but such things are at most promised never guaranteed.

- “Full Production” testing.

KIP Heidelberg has infrastructure for production-testing of fully equipped PPMs. A VME-crate serving as “test-rig” is available at our laboratory. This crate’s PSU is also equipped with an additional 48Volt power-brick to run a 9u ROD in the test-environment as well.

Hence, the test-rig consists of:

9u VME-crate with CPU (set-up, control, spy readout) and a TCM for TTC-distribution.

6u VME-crate for TTC generation.

64-fold analog Input to feed a PPM.

Multiplexer to check 4 LVDS channels into LVDS-Receiver in 6u VME.

eventually, 9U ROD to check G-Link output.

The running of the test-rig requires software-overhead, which must be completed. Many software blocks are already put together into a package similar to the well-known MCM test-package.

- ATLAS Check-out.

1. The “updated” FOUR existing PPM-prototypes do their “job” properly. They cover the needs for calorimeter signal check-out during 2006. They will be used together with dedicated software (multi-step running with VME, i.e. ROD-DAQ readout) to complete the check-out work on calorimeter signals.

2. Two PPr-crates fully equipped (“rewired” PSU, TTC- distribution backplane, LVDS backplanes, LVDS cable supports etc.) along with 16 PPMs can be shipped to CERN after the FDR/PRR.

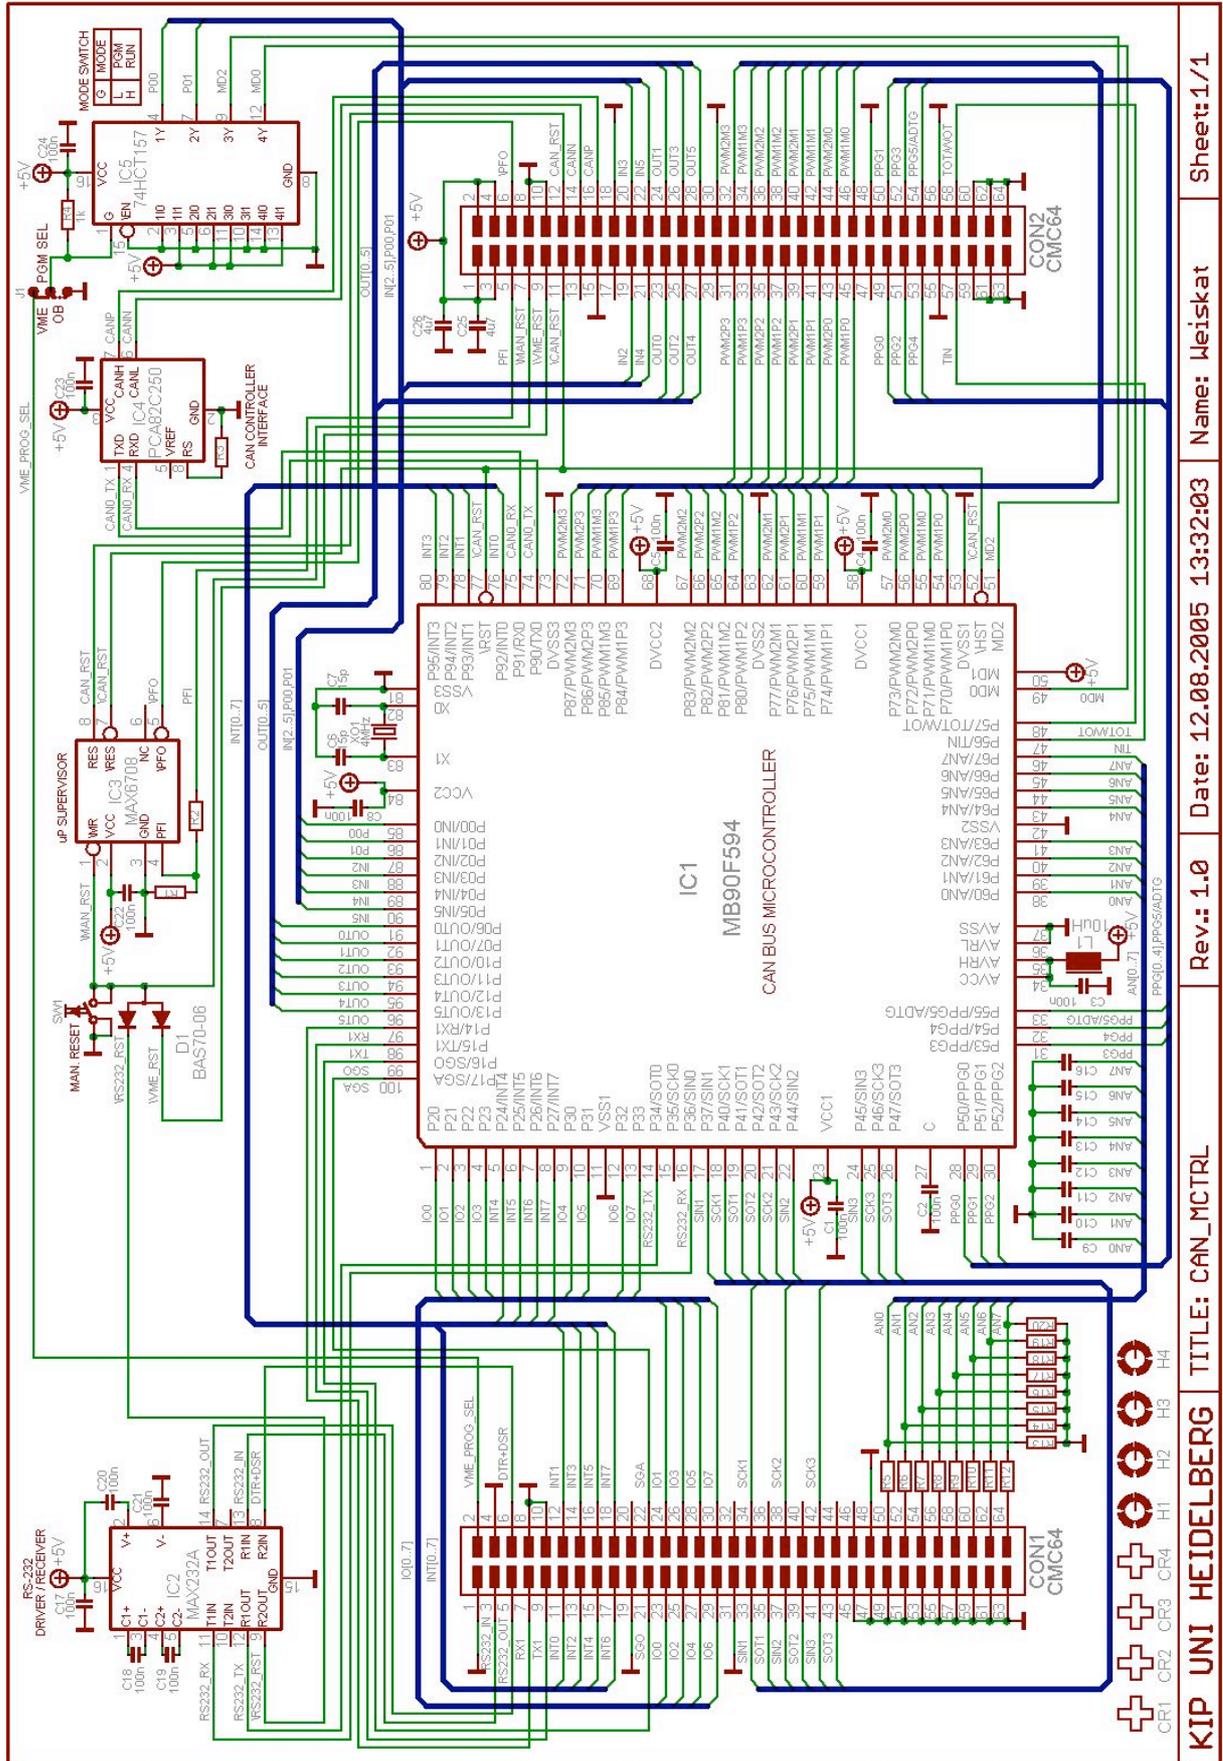
3. Crates will be equipped with infrastructure (backplanes, mechanics) and loaded with PPMs at KIP Heidelberg as modules come out of the testing-procedure. Only fully assembled and power-tested crates will be taken to CERN for installation in ATLAS-USA15.

CONTENT

The FINAL Pre-Processor Module (PPM) -----	1
for the ATLAS Level-1 Calorimeter Trigger -----	1
Introductory Remarks.	1
Related documentation.....	2
The FINAL design of the PPM.	3
The list of issues from the IDR and their implementation.....	4
The Real-time Signal Path	4
• Cross talk on ANIN.	4
• ANIN changes on DAC-ranges for signal-offset.	5
• Changes around the MCM.	7
• Pre-compensation on PPM for “MCM to LCD” signal transfer.	9
• Re-design of LCD.	9
The Readout Path.	12
• Firmware (i.e. loading of bit-files to devices on PPM).	12
• Firmware-code	12
• SRAM usage outside Rem_FPGA	12
• Optical G-Link (RAL-board in the rear of the PPr-Crate).	13
The PPM’s infrastructure, VME, TTC etc.	13
• “Status LEDs” at the front-panel.	13
• VME-CPLD (see also: “firmware loading”).	14
• TTCdec daughterboard.	14
• CAN-Controller.	14
• “Hot-Swap” controller.	15
• Use of “Hot-Swap”.	15
• Miscellaneous: PCB, assembly....	15
System Aspects.	15
• Power-Supply to serve a fully-loaded crate.	16
• Current on 3.3 Volt under “heavy-duty” data-pattern.	17
• Total Consumption from the crate’s Power-Supply Unit (PSU).	17
• Temperature of MCMs.	17
• TTC-distribution across backplane to 16 PPMs.	19
• Strain-relief on “massive” LVDS cabling.	19
• Phos4: 256 of them in a crate.	20
Summary, Schedule, Production Testing.....	23
• Summary.	23
• “Full Production”planning,	24
• “Full Production” testing.	25
• ATLAS Check-out.	25
Appendix (PPM-VME, CAN).....	28

Figure 1 : The PPM and its periphery. -----	3
Figure 2 : AnIn daughterboard (left: Top with 16 single OpAmps, right: Bottom)-----	5
Figure 3 : Linearity of analog signal-transfer in the PPM input to the MCM.-----	5
Figure 4 : Effect of DAC-setting on an analog Input-Signal (e.g. from LAr Calorimeter). ---	6
Figure 5 : Linearity of signal-digitisation (FADC on the MCM incl. MCM input-filter).-----	7
Figure 6 : Termination for I ² C clock.-----	8
Figure 7 : Scanning an analog signal by 1 nsec steps of the Phos4.-----	9
Figure 8 : LCD_2.0(top) with 4 FPGAs and Pre-Emphasis circuits (Rs) at outputs.-----	10
Figure 9 : LCD_2.0(bot) with Pre-Emphasis circuits (Cs) at outputs. -----	10
Figure 10 : The entire “pre-emphasis chain” on the PPM. -----	11
Figure 11 : LVDS data (left “eye-width”>>1 ns, right: “eye-height”>> +- 100mV)-----	11
Figure 12 : Space allocation for “firmware” in Flash-RAM. -----	12
Figure 13 : The “illuminated” PPM front-panel (LEDs), if visible “behind cables”.-----	14
Figure 14 : TTCdec clock – with PLL (upper trace) – wihtout PLL (lower trace)-----	14
Figure 15 : The “original” supply-cabling: the 30 mV, ~1kHz oscillation on 3.3V. -----	16
Figure 16 : The “zero-L” cabling: the <10 mV ripple on the 3.3V supply.-----	17
Figure 17 : MCM-temperatures in a powered crate. -----	19
Figure 18 : The rear af a PPr-crate: rewired power, RGTMs, “strain-relief” for LVDS. ---	20
Figure 19 : Phos4 operation (and INITIALISATION) with “symmetric” Clock. -----	22
Figure 20 : Scanning an analog signal with a “STUCK” Phos4. -----	23
Figure 21 : The fully-assembled PPr-crate with a TCM, 16 PPMs and a 9u VME-CPU. ----	24

• Schematics of CAN-Controller.



- Dimensions of CAN-Controller.

