DCC Memo

Jianhui Gu

Department of Physics

The Ohio State University

Dec. 22, 2005

Chapter 1: Introduction

The DCC design starts in 2002. The main purpose of DCC is to collect data from several DDUs and send to SLINK-64 CMS main DAQ (data concentration). In the EMU system, there are 478 chambers (ME1/1, ME1/2, ME1/3, ME2/1, ME2/2, ME3/1, ME3/2 and ME4/1), each with one DAQMB. The DDU is designed to receive thirteen DAQMB data (fifteen if there would be ME4/2). And the CMS main DAQ allocated eight SLINK64 for EMU according to EMU data rate (2KB/event, 100KHz → 200MB/s). This means that 4.5 DDU goes to one Slink64 on average.

DCC version1 (project D785G) has been used for beam test in 2004 and cosmic ray test early 2005. This board is partially loaded with two (out of five) input FPGAs, six FIFOs (out of 10). The board requires TTC connection to get CMS clock, L1A and bus control. The VME function for TTC is not routed to the custom backplane. Only one board is produced for this version. Right now, it is sitting in cern building 904. The version 1 is shown in lower left picture.

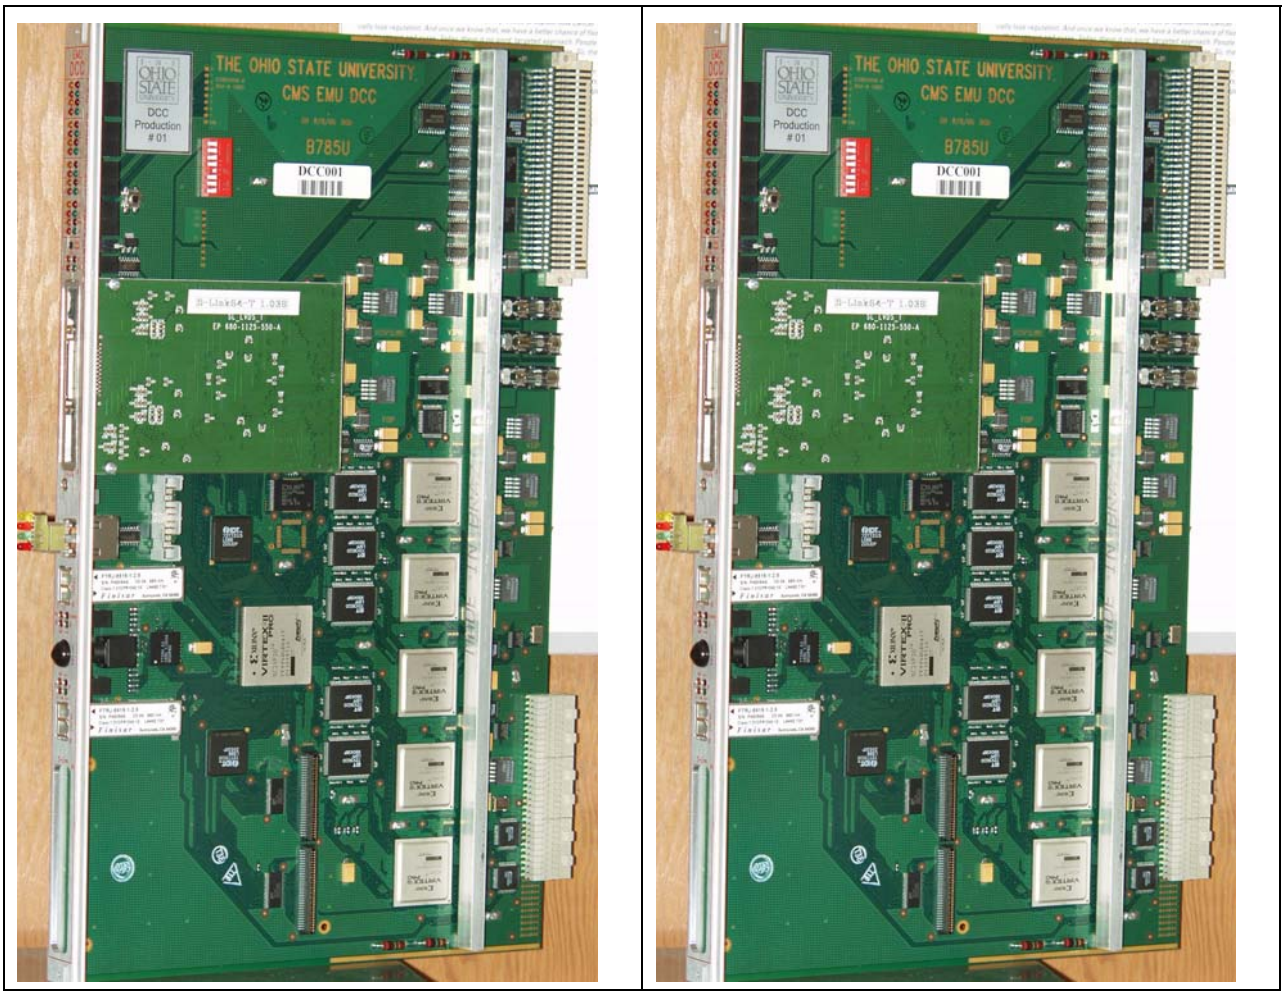

DCC version2 (project D785L) has slight improvement over the version1. This version fixed the clock driver problem. The LVDS clock driver used in version 1 can not driver capacitive load (10pf). This version used 2.5V PECL driver for 156.25MHz clock distribution. On the backplane, design out the National Semiconductor's DS90LV040 bus LVDS driver. This is the first fully loaded DCC. This board was used for 2004 beam test at H2, but destroyed by re-flowing the main FPGA. As no urgent need to use this card, the board is not repaired, but with the TTC connection, it can still function to supply clock and command for the DDU in the crate. Only one board is produced for this version. Right now, this board is also sitting in cern building 904. The board is shown in upper right picture.

DCC version 3 (project D785S) is the pre-production version. This version has major change from the previous version. The DDR input FIFOs are replaced by the conventional 36-bit wide fifos, because the high bit error rate on the DDR FIFO. This caused the re-design of FPGA firmware and PCB FIFO bus architecture. The bus is designed to run 160MHz. The input FIFOs are divided to both top and bottom sides. The package is also changed from ball-grid to flat pined. One board is stuffed for this version. One board is stuffed, and used in cern cosmic ray challenge. It behaves very stable, and reliable, works well. Though there are some fixes on the VME backplane grounding as the connector pin clearance was set too big that the grounds are disconnected at the connector. The problem is fixed by soldering some ground straps, and fully fixed in the production version. The board is used in CERN point 5 test. The board is shown in lower left picture.

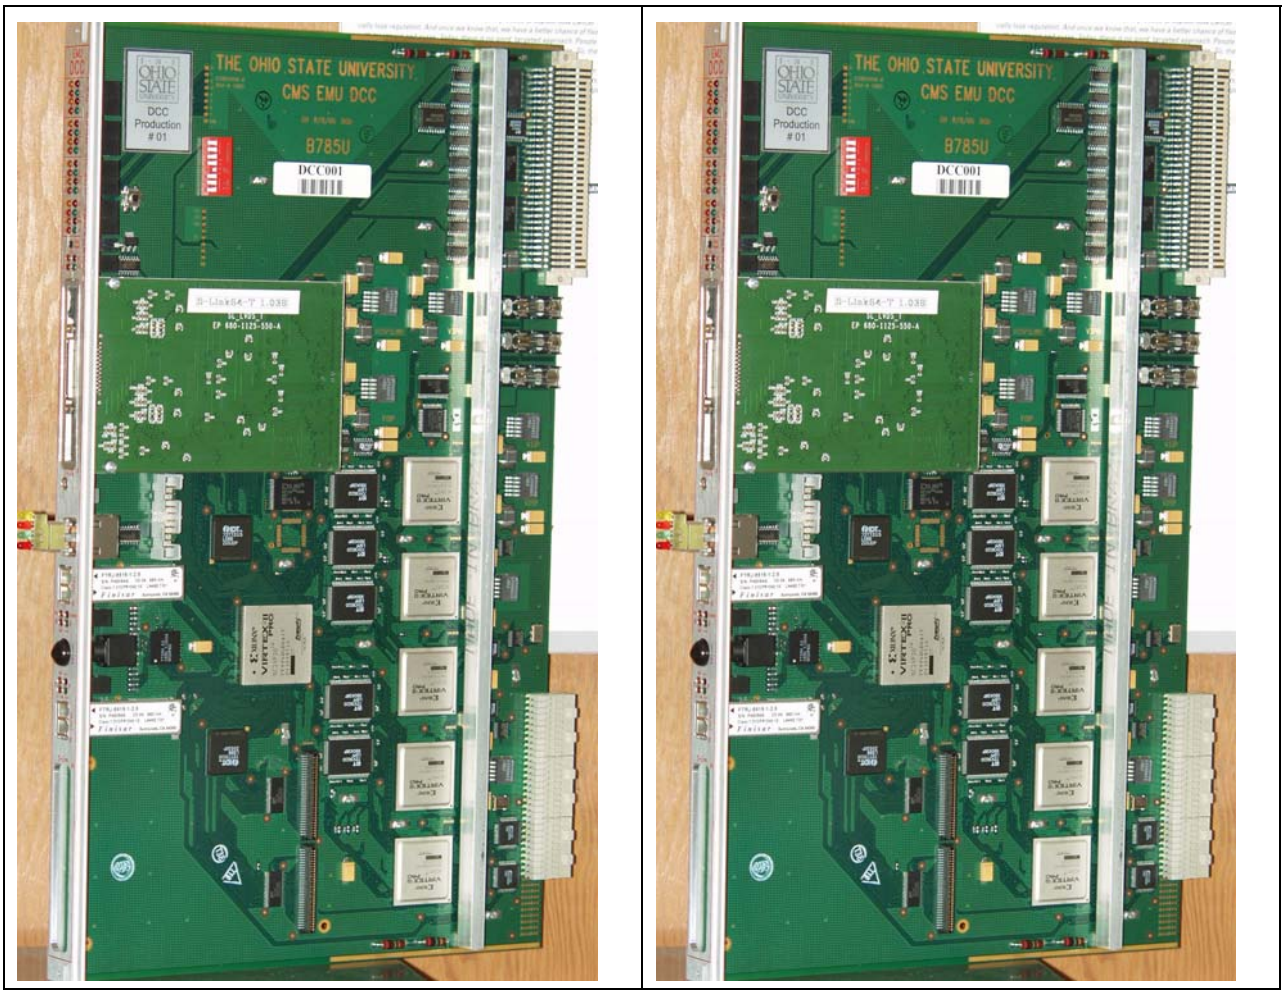

Production DCC (project D785U), or DCC version 4 has minor change from pre-production board. The two versions share the same firmware design. The production version has enhanced LED displays. There are minor PCB layout changes to facilitate the automatic soldering. Totally, ten production board are being made, and two of them do not have TTCrx, and four input FIFOs. These two boards are dedicated to the slot 17 and will be used for high data rate case only. All the other eight boards can be used either in slot 8 or slot 17. Even the pre-production version can be used too, except that the LED front panel display is different. The production DCC will be used at cern test from next Februrary. The board is shown in upper right picture.

Chapter 2: Design strategy

2.1 General description:

Here is the picture of the production DCC with annotations:

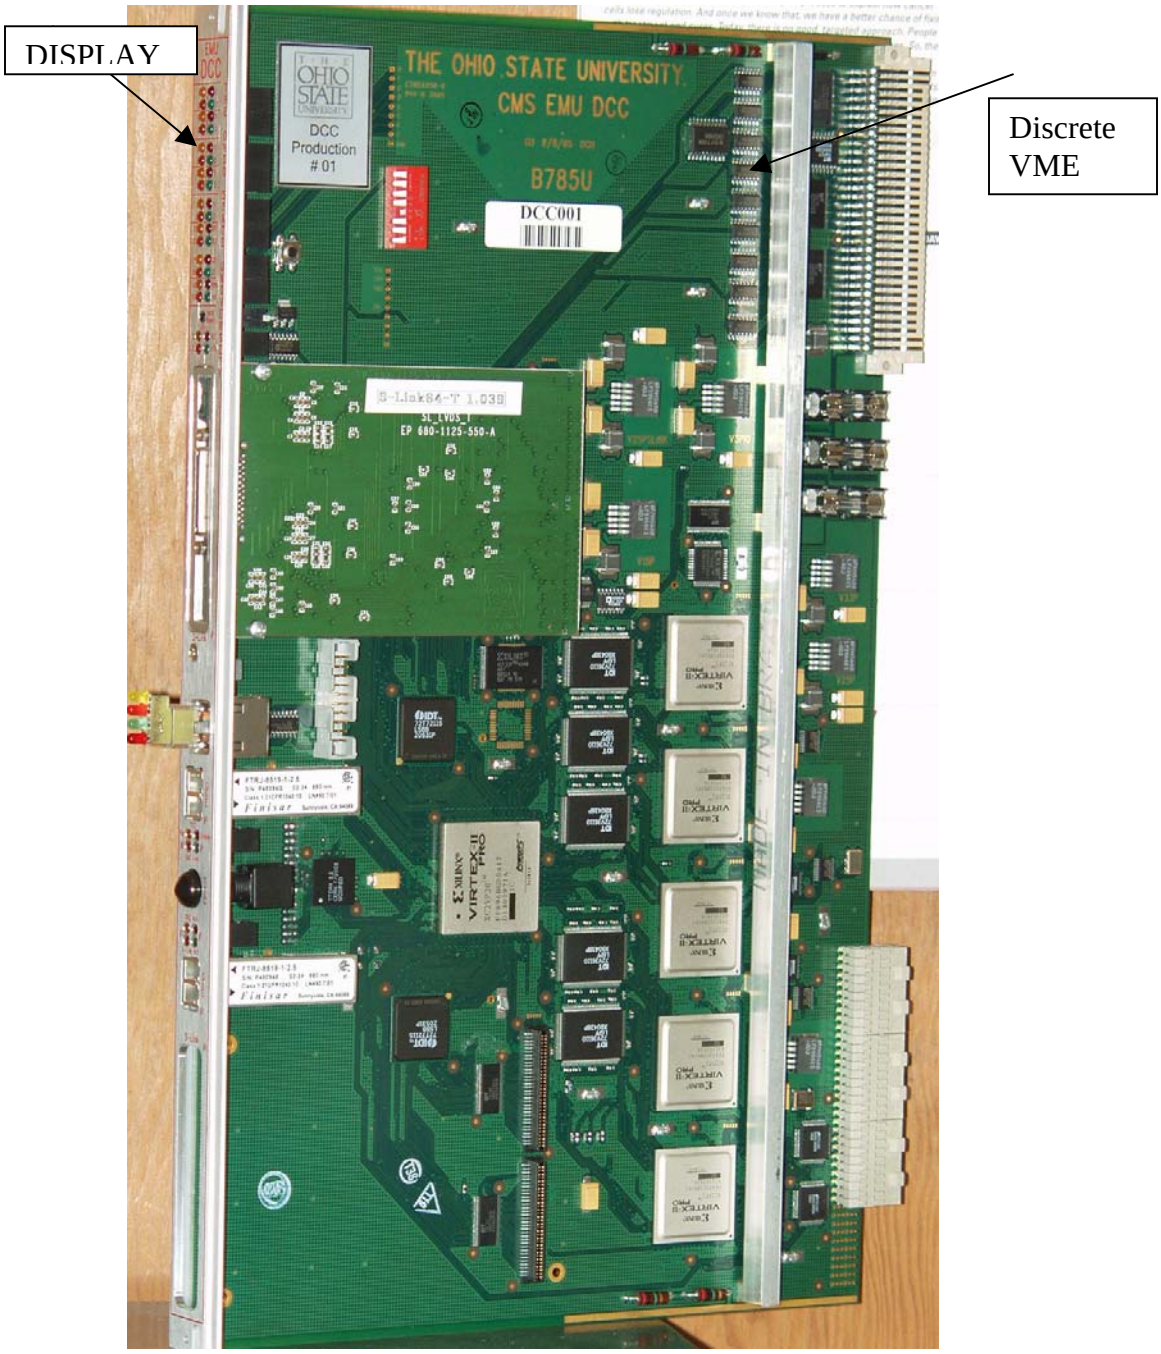

There are five input FPGAs (XC2VP2), each has four MGT, which are configured at 3.125Gbps. Two MGT (half of the FPGA) are connected to one DDU to received data at 625Mbyte/s. These are connected through high speed connector and custom P3 backplane. Each input FPGA is virtually two separate parts for two DDUs. There are ten input FIFOs, each is dedicated to one DDU. The ten FIFOs are equally divided into two high-speed read buses. The FIFOs have individual read/write control. The main FPGA consists of the VME slow control functions (JTAG, register access), TTC fast control and two identical data read out functions. The DCC can also fan out the TTC commands to the FED crate. It automatically switches to VME generated TTC command if the TTC is not plugged in. The slink FIFO are used to buffer the data before the data goes to the SLINK64. Two SLINK64 are designed on the DCC. Two gigabit Ethernet spy paths to local DAQ are also designed. One FMM is used for fast feedback to trigger control in case of problem, like rate being too high, out of synchronization, and some other error states. The TTCrx receives CMS fast control and fans out to the custom backplane for the FED crate.

2.2: Architecture:

As a heritage from DMB design and the popularity, VME form factor is selected for the board. The input and output space requirement on the front panel helps to select 9U wide form factor. But the DDU data flow is from front (DMB input) to the back (to custom backplane), and the DCC data flow is from back (custom backplane) to the front (SLINK), the board should be designed as short as possible, which I selected as 9U X 220mm. The boards look compact and no excessive long lines. Same as the DAQMB, the VME communication uses A24D16 only. The address is determined by the VME64X geographic address. For broad cast, 22 for DCCs, 28 for DDUs, and 30 for DDU and DCC.

On baseline design, EMU data rate is about 1.5Gbyte/s, and each SLINK is designed to be 200Mbyte/s. This means that overage 4.5 DDUs go to one SLINK. But the uncertainty is large. The DCC is so designed that it can take into account a large fluctuation of data rate, and uncertainty in cooperation with the FED crate design. On each DCC, there are two slink connection, and ten DDU inputs. For base line design, 9 of the ten DDU inputs are used, and two slinks are used. This gives the 4/5 DDU to one slink concentration. At very slow rate, the nine DDU inputs are used, and only one SLINK is enabled, this gives 9 DDU to on slink concentration. At higher rate, five of the ten DDUs are used, and two slinks are used, this gives 2/3 DDUs to one slink concentration. In this case, two DCCs are needed in one FED crate. The FRL deign can accommodate event lower data rate with two SLINKs plugging into one FRL, in this case, the expected slink rate is only 100Mbyte/s. The DDU can accommodate even higher data rate, with each DDU connects to one SLINK (on DDU). The one Mega Byte buffer between the main FPGA and the slink on DCC will weather out the large trigger rate fluctuation.

2.3: Speed considerations

2.3.1: Backplane: The maximum rocketIO data rate is 3.125 Gbps (156.25MHz) for xc2vp. The matching high speed connector is available from AMP. This also matches well with the input FIFO with maximum speed of 166MHz. On the input FPGA, a small buffer is necessary to equalize the two rocketIOs from each DDU. The input and out output data rate on input FPGA just match well. This gives a single DDU data rate of 625Mbyte/s.

2.3.2: On PCB: There are five fifos on each bus. The fifos are arranged to have two/three on top/bottom side to minimize the bus length. The read clock, which runs at 156.25MHz are divided into four sections with each clock drivers two or three loads, and terminated with 100//100 Ohm resister. The signals look fine.

2.3.3: SLINK interface speed: The slink fifo write clock is the same as the main FPGA processing clock, that is 156.25/2, about 78MHz. It does not make much sense to increase the FIFO write speed dramatically. At this speed, the data rate is about 625Mbyte/s, which is about three times of the SLINK data throuput. The SLINK fifo readout clock is the same as slink clock with some phase difference. That speed can vary with maximum at 200MHz, the speed limit of the slink FIFO. Right now, the speed is set to 62.5MHz. The 200Mbyte/s through put is equivalent to 25MHz with 100% occupancy.

2.4: Buffer size considerations:

There are two kinds of buffers, internal buffer and external buffers. On the input FPGA, the internal buffers (one block ram) are used to equalize the two rocketIO channels from the same DDU. This buffer size is determined by the rocketIO uncertainties, and 1Kbit is sufficient. The input FIFO are used as the main buffer for DDU data before the DDU fragments are assembled as a single event. This size is determined by the maximum DDU event size, which could be 20Kbyte*15=300Kbyte (assuming CFEB 16 time samples, and all the five CFEBs are fired, and there are TMB and ALCT data, and 15 DMBs per DDU). The 4Mbit IDT fifo are chosen for input FIFO. On the main FPGA, the gigabit spy channel data is buffered on the FPGA. The buffer size is determined by the difference between the data processing speed and the gigabit speed, and the event size (not the extreme event size). Right now, 16Kbyte are used. The SLINK fifo size is to used to even out the EMU event size fluctuation and rate fluctuation. Because there is only one fifo per channel, the cost is not a big factor, so the largest fifo is chosen for best performance, which is 1Mbyte.

2.5: Power supply:

The Power are supplied by the VME64X. The onboard 5V gets directly from VME 5V, the onboard +1.5V, +1.8V, +2.5V are regulated from VME +3.3V, the onboard 3VIO and +3.3V are regulated from +5V from VME. But the slink +3.3V is directly from VME, with filter

capacitors. For the DCC, the +3.3V on VME is about 5A, the +5V is about 1A. The total power consumption is about 25W. The slink card makes little difference.

2.6: PROM selection:

The input FPGAs share one prom XC18v02vq44. Each FPGA has three pins connected to the ground or power, to set the FPGA ID (1, 2, 3, 4 and 5). The main FPGA has two options, either using two XC18v04vq44 or one xcf32p-vo48. The production board will be loaded with xcf32p. In the Prom, there will be four revisions loaded in. These four revisions will be the same in design, but with different SLINK ID code. With the onboard two slink position, this will give eight slink ID numbers. When the slot 17 is used, that will another extra eight unique ID numbers. The DCC serial numbers will be saved in the input FPGA usercode. The FPGA usercode will keep the information of firmware version, and date code etc.

Chapter 3: Board operation

3.1 Before the board plugging in the crate:

Check the 8-bit toggle switch. Close to PCB is low, away from PCB is high. From top to bottom:

SW1. SW2 and SW3 are settings on the controller FPGA: Single L1A generation. One L1A is generated on SW1 from low to high. SW2, SW3: If both SW2 and SW3 are high, the L1A will be continuous with ~1.6ms between L1As. And the SW2 and SW3 are also used to control the logic analyzer port. Refer to the schematics for the details (page 10 and 20J). If all the SW1, SW2 and SW3 are high, the DCM on the controller FPGA will be reset.

SW4 and SW5 are passed to INFPGAs from controller FPGA. There are four combinations:

SW4	SW5	Function
Low	Low	The INFPGA generate data on L1A and feed to the INFIFO
High	Low	Normal operation mode
Low	High	Disable the propagation of SLINK backpressure to DDU
High	High	Serial loop back on the custom backplane RocketIO

SW6, SW7 and SW8 are used to control the xilinx platform PROM:

SW6	SW7	SW8	Revision selection
Low	Low	Low	Revision 0
High	Low	Low	Revision 1
Low	High	Low	Revision 2
High	High	Low	Revision 3
Low/High	Low/High	High	External revision selection disabled, depending on the PROM setting, default revision is selected

For now, the SlinkID is determined by the prom revision number. The CSC is assigned 750-757 (in decimal) for SLINKID.

Revision	Rev0	Rev1	Rev2	Rev3
SLINK_ID	750/751 (2EE/2EF)	752/753 (2F0/2F1)	754/755 (2F2/2F3)	756/757 (2F4/2F5)

The 0/1 is controlled by the slink position on the DCC, top 0, lower 1. If the slot 17 DCC were used, the extra slinkID bit comes from the geographic address bit 4.

3.2: After the board plugging in the crate:

The primary DCC goes to slot 8 of FED crate, the optional secondary DCC goes to slot 17.

Push-button switch on the front panel will hard-reset all the DCCs, DDUs and the TTCrx.

Push-button switch on the DCC will sync-reset the DCC only.

All the TTC commands will be fanned out to the FED crate. The TTC 0x04 is TTC hard-reset, which the DCC will translate and generate an 400ns pulse to hard-reset DDUs, but not the DCC. The DCC FPGAs can be hard-reset by sending command to the PROMs to pulse the PG pins.

After power up, the following procedure is necessary for the FED crate to work properly:

a. set the appropriate available DDUs.

DDU slot	3	4	5	6	7	9	10	11	12	13
setting	001/0	004/0	010/0	040/1	100/1	200/1	080/1	020/1	008/0	002/0

DDU slot	14	15	16	18	19	20
Setting	004/0	010/0	040/1	080/1	020/1	008/0

Here the /0 and /1 mean the slink 0, or slink 1. If there are multiple DDUs in the crate, the setting bits should be 'OR'ed together. For example, slot 4 and 12, the FIFO_in_use setting should be: '004'|'008' = '00C'.

b. Set the TTC sync-reset, which is TTC command 03. The FED crate will be ready.

Chapter 4: Front panel diagnostic display

4.1: Top part of the LED meaning:

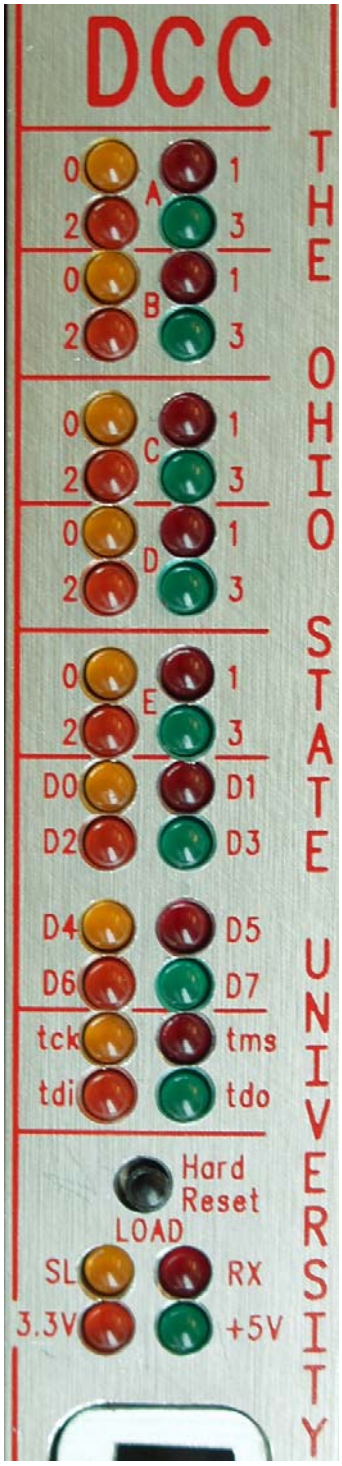

A0: ON: Slot 3 DDU rocketIO linked

OFF: at least one of the two slot3 DDU rocketIO not linked

Flashing: Data activity from Slot 3 DDU

A1: OFF: DCC is ready for slot3 DDU data

ON: DCC is not ready for slot 3 DDU data

Flashing: DCC backpressure the DDU, not to send data

slot	3	4	5	6	7	9	10	11	12	13
Link/act	A0	B0	C0	D0	E0	E2	D2	C2	B2	A2
Ready/bp	A1	B1	C1	D1	E1	E3	D3	C3	B3	A3

D0, D1, D2, D3, D4, D5, D6 and D7 are L1A rate indication. The rate are updated at 5Hz.

D0	D1	D2	D3	D4	D5	D6	D7
EB1	EB0	DB5	DB4	DB3	DB2	DB1	DB0

The L1A rate is: $5 * DB[5:0] * 2^{(3*EB[1:0])}$. Special case, the D7 will flash once if there is only single L1A.

TCK, TMS, TDI and TDO is the activity LED for the connector JTAG to the INPROM. The TCK, TMS and TDI are enabled only if the external cable is plugged in. The TDO is always active.

LOAD SL: ON: The main control FPGA is programmed (DONE);

LOAD RX: ON: The input FPGAs are programmed (DONE);

3.3V: ON: The +3.3V power from VME backplane is OK, the LED will be OFF if the on-board fuse is blown;

+5V: ON: The +5V power from VME backplane is OK, the LED will be OFF if the on-board fuse is blown.

4.2: Middle part of the LED meaning

FMM plug in LEDs:

Yellow: out of Sync Orange red: Ready Green: Warning Red: Busy	The FMM four bits will be: OrangeRed:Red:Yellow:Green If both OrangeRed and Red are ON, error happened

Normally, the state is READY; when the slink FIFO is almost full, the overflow WARNING will be on; when the INFPGA backpressure the DDU, the BUSY will be ON; When the slink FIFO is full, or the INFIFO not ready, the OUT_of_Sync will be ON; If any fifo is full, or other error state, go to ERROR state, both Ready and Busy will be ON.

FMM connector's Green LED: ON: Slink0 card is plugged in (link ready); OFF: slink0 card not link, either no power or not plugged in; Flashing: The slink0 busy and backpressuring.

FMM connector's Yellow LED: Same meaning, but for Slink1

Rdy: ON: TTC fiber connected, and the TTCrx is ready;

Enable: ON: the TTC data is enabled and fan out to the FED backplane.

GbE Link#2: ON: Fiber from gigabit Ethernet card to spy path GbE#2 is good;

GbE Link#1: ON: Fiber from gigabit Ethernet card to spy path GbE#1 is good. These two LEDs are driven by the optical transceivers directly.

GbE Act #2: Flashing: DCC is sending data to spy gigabit Ethernet #2, OFF: no data is sending out from DCC. ON: Continuously sending data;

GbE Act #1: Same meaning, but for Spy path gigabit Ethernet #1.

S-Link Act #2: Same meaning as GbEAct#2, but for the Slink #2 (lower Slink);

S-link Act #1: Same meaning as GbEAct#1, but for the Slink #1 (Upper Slink).