IPMI Support for FMC modules in MicroTCA.4 Systems

Dariusz Makowski
Patryk Nowak vel Nowakowski

FPGA Mezzanine Card Standards (2)

- Two standards available:
  - **FMC** ANSI/**VITA 57.1**-2019 (+Errata)
  - **FMC+** ANSI/**VITA 57.4**-2018 (+Errata)
- Modular expansion of FPGA systems
- High speed communication interfaces (up to 28 Gbps)
  - **Low/High Pin Count** (LPC/HPC) connectors for I/O flexibility (160/400 pins)
  - **High Serial Pin Count** (HSPC) connector (560 pins)
  - **High Serial Pin Count extension** (HSPCe) connector (80 pins)
- Mechanical dimensions allow for up to two **FMC modules on a single MTCA.4 AMC board**

Example of FMC module

VITA
Open Standards, Open Markets

FMC™

Support for FMC Module in MTCA.4 Systems

- Many existing **MicroTCA.4** designs utilizes **FMC modules**
- FMC I/O available on front panel of **AMC or RTM carriers**
- Identification of connected FMC module by reading FRU inside fixed-address I2C EEPROM
- FMC module managed and controlled by MMC on AMC carrier board

Support for FMC Module in MTCA.4 Systems

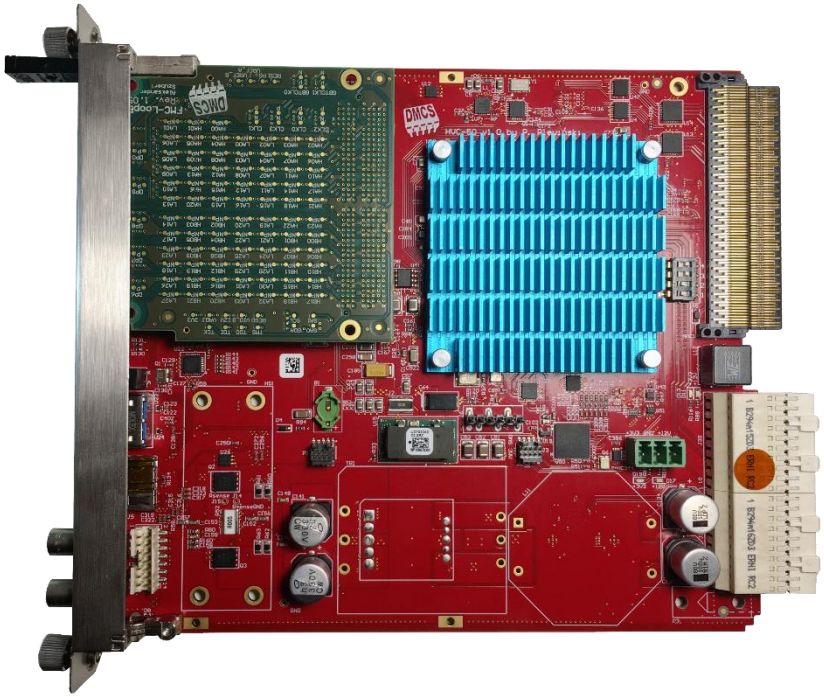

The HVC-50 board with single FMC module (Xilinx Zynq US+ XCZU4CG)

The MFMC carrier board with two FMC modules (Xilinx Artix XC7A200T)

Current Issues with FMC modules in MTCA.4 Systems

- **Lack of management compatibility of MicroTCA and FMC standards**
- **No thermal management – Risk of overheating FMC module**
  - FMC modules may dissipate significant amount of heat (more than 10 Watts)
  - FMC modules need thermometers, now included in VITA57.x standards
  - **FMC modules** should be involved in **temperature stabilisation process** in the whole **MicroTCA infrastructure**

- **No centralized management of FMC – Full IPMI support needed**
  - FMC managed only locally by AMC carrier board - MCH should have access to FMC FRU as well as sensors and initiate power on/off of the FMC module.
  - **FMC should be managed by MCH in similar manner as it is done currently with RTM**
  - **FPGA IO voltages** are automatically configured according to **FMC module requirements and carrier possibilities**

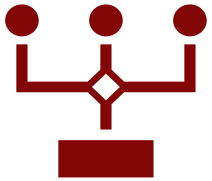

Powering up FMC must be done after checking if FMC is compatible with the carrier – DC load records need to be read from the FMC's FRU to chose right supply voltage

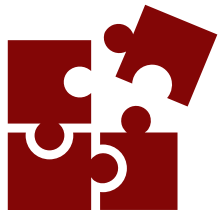

Proposed Assumptions

- Allow for **2 FMC modules on AMC** and **2 modules on RTM** (total 4 modules)
- Prefer a **simple solution** based on **I2C devices** (no programmable IPMI controller on FMC, no complex sensors)
- A single FMC module should support:
  - FRU EEPROM with Subtype 0 and Subtype 1 data
    - Configurable voltages for each FMC module
  - Temperature sensor: **max. 2 sensors** per FMC module
- MCH should provide support to detect 4 additional FRUs on AMC/RTM pair
  - Display FRU data (reference voltages, clock direction, etc.)
- Additional sensors implemented as dynamic IPMI sensors
- **Additional voltage sensors** – **do we need them? Other sensors?**

Possible IPMI Support for FMC modules on MTCA.4 systems

FMC modules with:

- EEPROM for FRU storage
- At least two temperature sensors at fixed I2C address are obligatory
- Other sensors if needed (avoid I2C address collisions)
- **FRU 0 is on AMC**
- **FRU 1 is on RTM**
- **FRU 2 and 3** are FMCs on AMC
- FRU 4 and 5 are FMCs on RTM (also on AMC possible)
- **Do we need FRU 4 and 5?**

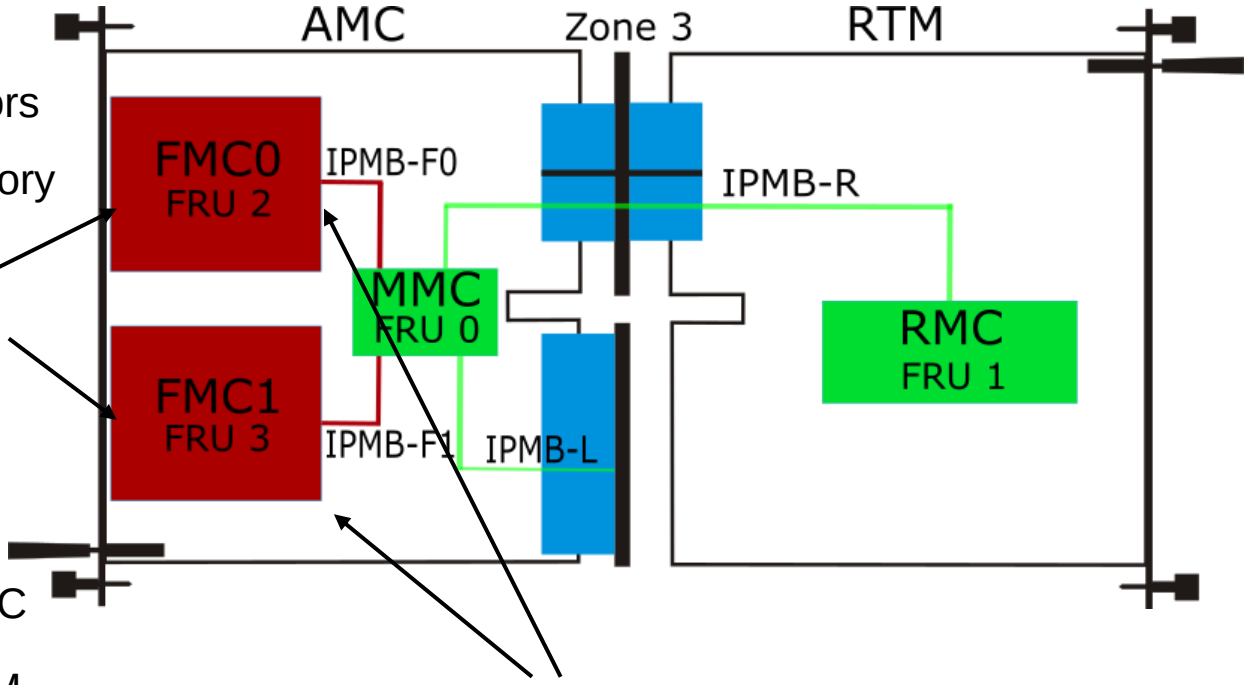

MMC ←→ FMC I2C bus or buses
(single bus should be enough, separate buses could give more freedom)

We prefer a simple solution

VITA 57.1 – Important Signals

IPMB-F Bus:

SCL – System Management I2C serial clock. This signal provides a clock reference to the FMC+ module from the carrier card for a two-wire serial management bus.

SDA - System Management I2C serial data. This signal provides a data line for a two-wire serial management bus.

GA[0..1] - These signals provide geographical addresses of the module(s) and are used for I2C channel

PRSNT_M2C_L - Module present signal. The ANSI/VITA-57.1 FMC standard defines this signal for a carrier card to successfully detect that an IO mezzanine module is installed.

CLK_DIR – Used to determine whether the mezzanine module or the carrier card is the driver for CLK[2..3].

PG_C2M – Power Good Carrier Card. This signal asserts high by the carrier card when power supplies VADJ, 12P0V, 3P3V are within tolerance.

PG_M2C – Power Good Mezzanine. This signal asserts high by the mezzanine module when power supplies VIO_B_M2C, VREF_A_M2C, VREF_B_M2C are within tolerance.

IPMI Support

The FMC mezzanine module provides hardware definition information that is read by an external controller using **IPMI commands and I2C serial bus transactions**. The mezzanine module supports the I2C link on pins SDA and SCL and optionally supports the base **IPMI commands defined in the PICMG 2.9** specification.

- **FRU records** according to Platform Management FRU Information Storage Definition V1.0
- **EEPROM shall include** the BOARD_INFO area for identification of the mezzanine module
- Other data provided via **FMC-specific MultiRecords** (type 0xFA): **Subtype 0 and Subtype 1.**

FRU EEPROM Address

- Standard defines I2C addresses for EEPROM depending on GA [1/0] addresses (4 possibilities)
- Now double byte addressing also possible (32 kB)
- MicroTCA.4 needs suport for 32 kB

Table 10: I2C Address Decoding 2Kb and 32Kb EEPROM

	GA[0..1]	I2C Address for serial 2Kb EEPROM	I2C Addresses for Optional Devices
2Kb EEPROM	00	0b101 0000	0bxxx xx00
	01	0b101 0001	0bxxx xx01
	10	0b101 0010	0bxxx xx10
	11	0b101 0011	0bxxx xx11
	GA[0..1]	**I2C Address for serial 32Kb EEPROM**	**I2C Addresses for Optional Devices**
32Kb EEPROM	00	0b101 0100	0bxxx xx00
	01	0b101 0101	0bxxx xx01
	10	0b101 0110	0bxxx xx10
	11	0b101 0111	0bxxx xx11

IPMI Support – Record Subtype 0

Provides informative data

- Module size
- P1 and P2 connectors types
- Clock direction
- P1/P2 signals number
- Max clk frequency

Field	Byte Offset	Bit Location	Length	Description
Subtype	0	7:4	4 bits	0 for main definition type
Version	0	3:0	4 bits	0 for current version
Module Size	1	7:6	2 bits	0b00 = single width, 0b01 = double width
P1 Connector size	1	5:4	2 bits	0b00 = LPC, 0b01 = HPC
P2 Connector size	1	3:2	2 bits	0b00 = LPC, 0b01 = HPC, 0b11 = not fitted
Clock Direction	1	1	1 bit	Defines direction of CLKx_BIDIR 0b0 = from Mezzanine to Carrier 0b1 = from Carrier to Mezzanine
reserved	1	0	1 bit	= 0b0
P1 Bank A Number Signals	2	7:0	8 bits	= number needed
P1 Bank B Number Signals	3	7:0	8 bits	= number needed
P2 Bank A Number Signals	4	7:0	8 bits	= number needed
P2 Bank B Number Signals	5	7:0	8 bits	= number needed
P1 GBT Number Transceivers	6	7:4	4 bits	= number needed
P2 GBT Number Transceivers	6	3:0	4 bits	= number needed
Max Clock for TCK	7	7:0	8 bits	= clock in MHz

IPMI Support – DC Output/DC Load Records

Provides **more useful** data

- **Supply voltages:**
  - VIO_B_M2C
  - VREF_A_M2C
  - VREF_B_M2C
- **Power loads**
  - VADJ
  - 3P3V
  - 12P0V
- **VADJ voltage ranges**
  - minimum
  - maximum
  - preferred

DC Output Record	Connector	Record Type	Output Number	Description
VADJ	P1	LOAD	0	Mandatory
3P3V	P1	LOAD	1	Mandatory
12P0V	P1	LOAD	2	Mandatory
VIO_B_M2C	P1	OUTPUT	3	Mandatory
VREF_A_M2C	P1	OUTPUT	4	Mandatory
VREF_B_M2C	P1	OUTPUT	5	Mandatory
VADJ	P2	LOAD	6	Mandatory for double width with P2 fitted, n/a for Single Width
3P3V	P2	LOAD	7	Mandatory for double width with P2 fitted, n/a for Single Width
12P0V	P2	LOAD	8	Mandatory for double width with P2 fitted, n/a for Single Width
VIO_B_M2C	P2	OUTPUT	9	Mandatory for double width with P2 fitted, n/a for Single Width
VREF_A_M2C	P2	OUTPUT	10	Mandatory for double width with P2 fitted, n/a for Single Width
VREF_B_M2C	P2	OUTPUT	11	Mandatory for double width with P2 fitted, n/a for Single Width

Temperature Sensors

Observation 5.23-2: The VITA 57.1 standard started with development cards, and as the market has grown and matured, more of the FMC sockets are being deployed on production systems. These systems depend on IPMI to enable easier system integration of boards from different vendors. In future revisions, this Standard could more tightly specify how the I2C bus can be used. **The potential addition of temperature sensor documentation is an example of this issue.**

- **We could easily support up too 8 thermometers (2 per module)**
- **No strict definition of the temperature sensors addresses**

Observation 5.23-3: To fully support using the I2C bus in a system supporting IPMI, the mezzanine module would need to have a **Mezzanine Management Controller** that can communicate with a Carrier Management Controller.

- **We recommend a simple solution without MMC controller**

Temperature Sensors (2)

Currently, I2C addresses for thermometers can be specified only optionally in FRU via Subtype 1 record

The device string portion of this **MultiRecord subtype consists of 6-bit ASCII text** as defined in the ISD.
The string is divided into one or more I2C device records. Each device record consists of one or more address characters followed by one or more bytes of device name, example:
* FMC module had two temperature sensors at addresses 0b0011nnn and 0b0100nnn
!CTLR$%LM75 // **I2C Controller addrres 0**, **I2C controller name: CTRL**
// sensor 1 (address: 3), sensor 2 (address: 4)

Table 9: Subtype 1: I2C Device Definition (variable length and optional)

Field	Byte Offset	Bit Location	Length	Description
Subtype	0	7:4	4 bits	1 for I2C device definition subtype
Version	0	3:0	4 bits	0 for current version
Device String	1..N/8		N bits	Device address / name strings, see below

Temperature Sensors (3)

- Addresses of I2C thermometers could be specified in **MicroTCA.X** specification similarly as it is currently for EEPROM
- We need 2 address (A1/A0) for 4 FMC modules
- 1 address (A2) for sensor ID on FMC module (2 sensors max.)
- More possible using „Float" addressing (3 sensors)

Table 10: I2C Address Decoding 2Kb and 32Kb EEPROM

	GA[0..1]	I2C Address for serial 2Kb EEPROM	I2C Addresses for Optional Devices
2Kb EEPROM	00	0b101 0000	0bxxx xx00
	01	0b101 0001	0bxxx xx01
	10	0b101 0010	0bxxx xx10
	11	0b101 0011	0bxxx xx11
	GA[0..1]	**I2C Address for serial 32Kb EEPROM**	**I2C Addresses for Optional Devices**
32Kb EEPROM	00	0b101 0100	0bxxx xx00
	01	0b101 0101	0bxxx xx01
	10	0b101 0110	0bxxx xx10
	11	0b101 0111	0bxxx xx11

	GA[0..1]	I2C Address for thermometer #1	I2C Address for thermometer #2
I2C thermometer i.e. MAX6626 LM75 MCP980	00	0b100 1000	0b100 1100
	01	0b100 1001	0b100 1101
	10	0b100 1010	0b100 1110
	11	0b100 1011	0b100 1111

Temperature Sensors (4)

- TMP75 – standard I2C thermometer
- 8 thermometers possible on a single bus, 2 on FMC module
- **More is also possible, do we need?**
- Implemented as dynamic sensors

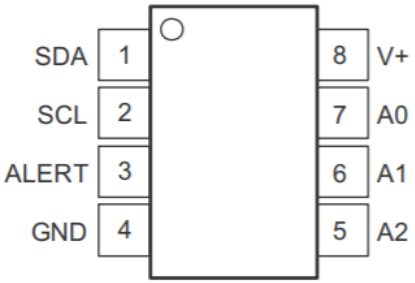

	GA[0..1]	I2C Address for thermometer #1	I2C Address for thermometer #2
I2C thermometer i.e. MAX6626 LM75 MCP980	00	0b100 1000	0b100 1100
	10	0b100 1001	0b100 1101
	01	0b100 1010	0b100 1110
	11	0b100 1011	0b100 1111

A2	A1	A0	SLAVE ADDRESS
0	0	0	1001000
0	0	1	1001001
0	1	0	1001010
0	1	1	1001011
1	0	0	1001100
1	0	1	1001101
1	1	0	1001110
1	1	1	1001111

VITA 57.4 (FMC+) Specification

Specifies new requirements:

- Extended FMC HPC connector HSPC (P1 and P2)
- Optional HSPCe Connector (P3 and P4)
- Same as ANSI/VITA-57.1 FMC, IPMI Support Section, **except IPMI subtype 0 record**

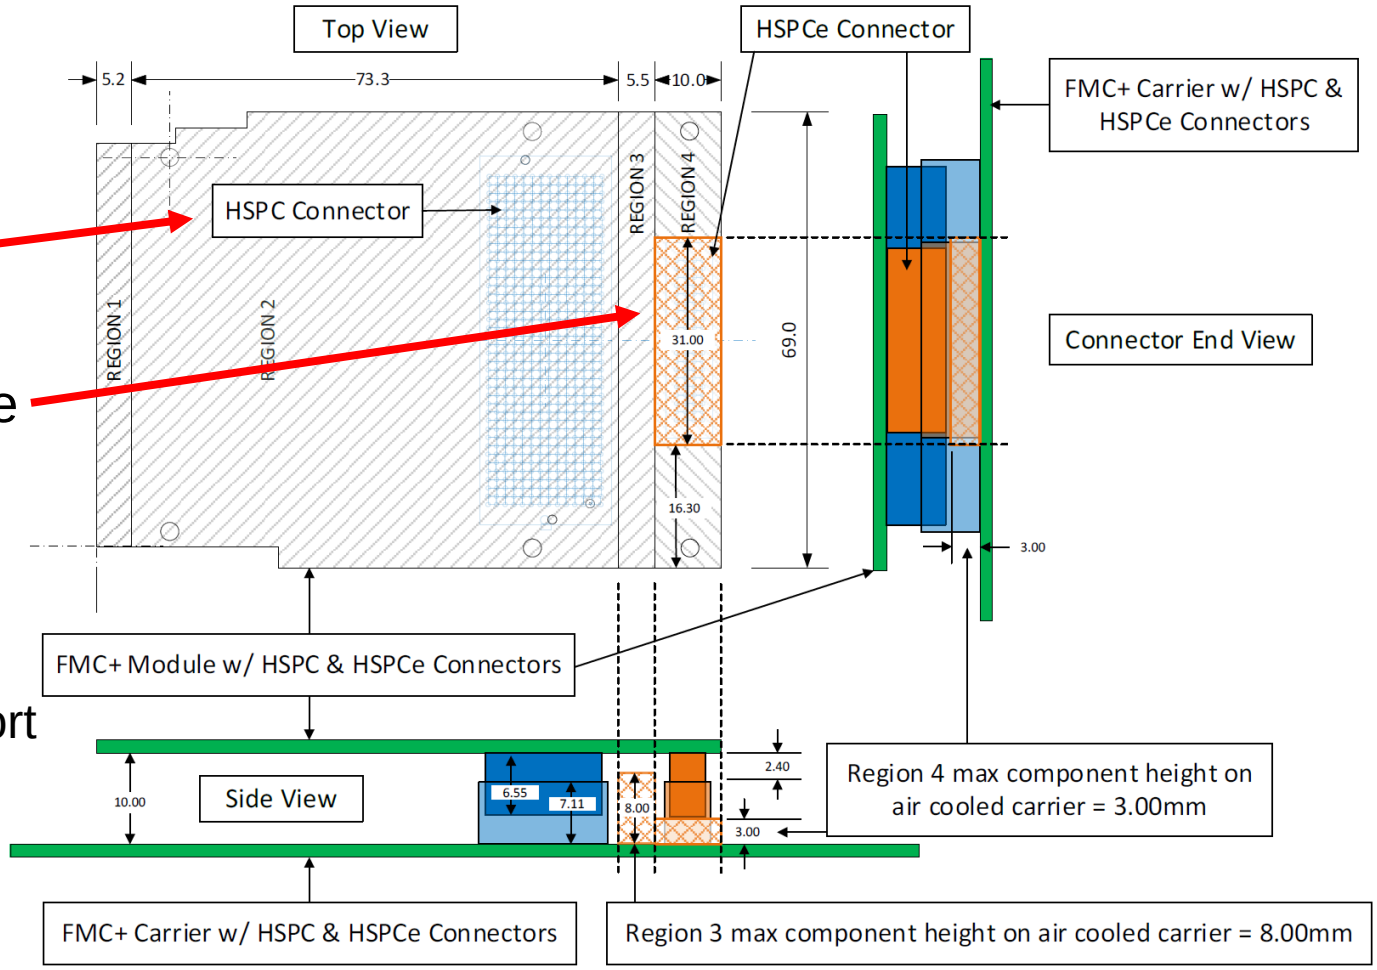

VITA 57.4 – Important Signals

SCL – System Management I2C serial clock. This signal provides a clock reference to the FMC+ module from the carrier card for a two-wire serial management bus.

SDA - System Management I2C serial data. This signal provides a data line for a two-wire serial management bus.

GA[0..1] - These signals provide geographical addresses of the module(s) and are used for I2C channel

PRSNT_M2C_L - Module present signal. The ANSI/VITA-57.1 FMC standard defines this signal for a carrier card to successfully detect that an IO mezzanine module is installed.

HSPC_PRSNT_M2C_L - Module present signal. This signal allows the carrier to determine whether an FMC+ module is present.

HSPCe_PRSNT_M2C_L – Module with HSPC and HSPCe connector present signal. This signal allows the carrier to determine whether an FMC+ module with optional HSPCe is present.

CLK_DIR – Used to determine whether the mezzanine module or the carrier card is the driver for CLK[2..3].

PG_C2M – Power Good Carrier Card. This signal asserts high by the carrier card when power supplies VADJ, 12P0V, 3P3V are within tolerance.

PG_M2C – Power Good Mezzanine. This signal asserts high by the mezzanine module when power supplies VIO_B_M2C, VREF_A_M2C, VREF_B_M2C are within tolerance.

VITA 57.4 – IPMI Subtype 0 Record

- New Subtype 0 Record proposed
- The same EEPROM addresses, size and type
- Same address of thermometers
- Same additional signals required (HSPC_PRSNT_M2C_L, HSPCe_PRSNT_M2C_L)

Field	Byte Offset	Bit Location	Length	Description
Subtype	0	7:4	4 bits	0b0000 for main definition type
Version	0	3:0	4 bits	0b0001 for FMC+ (see Observations in Section 5.3.1)
Module Size	1	7:6	2 bits	0b00 = single width, 0b01 = double width
P1/P3 Connector Size	1	5:4	2 bits	0b00 = LPC only, 0b01 = HPC only, 0b10 = HSPC only, 0b11 = P1:HSPC + P3:HSPCe
P2/P4 Connector Size	1	3:1	3 bits	0b000 = LPC only, 0b001 = HPC only, 0b010 = HSPC only, 0b011 = P2:HSPC + P4:HSPCe, 0b111 = not fitted
Clock Direction	1	0	1 bit	Defines direction of CLKx_BIDIR 0b0 = from Mezzanine to Carrier 0b1 = from Carrier to Mezzanine
P1 Bank A Number Signals	2	7:0	8 bits	= number needed
P2 Bank A Number Signals	3	7:0	8 bits	= number needed
P1 Bank B Number Signals	4	5:0	6 bits	= number needed
P2 Bank B Number Signals[1:0]	4	7:6	2 bits	= number needed
P2 Bank B Number Signals[5:2]	5	3:0	4 bits	= number needed
P1 GBT Number Transceivers[3:0]	5	7:4	4 bits	= number needed
P1 GBT Number Transceivers[5:4]	6	1:0	2 bits	= number needed
P2 GBT Number Transceivers	6	7:2	6 bits	= number needed
Max Clock for TCK	7	7:0	8 bits	= clock in MHz

Possible IPMI Support for FMC modules on MTCA.4 systems

FMC FRU content

Board info and Product info:

- Information about board name, manufacturer, serial number, part number FRU_ID,.

```
---------------------------------------
Board Info Area              : at offs=264, len=21
Manufacturer(04)             : DMCS
Board Name(06)               : FMC_LB
Serial Number(04)            : 0000
Part Number(04)              : 0000
FRU file ID(03)              : LB_F
---------------------------------------
Product Info Area            : at offs=285, len=90
Manufacturer(04)             : DMCS
Product Name(06)             : FMC_LB
Product Number(02)           : 00
Part Version(25)             : 000000000000000000000000
Product Serial Number(25)    : 000000000000000000000000
Asset Tag(25)                : ________________________
FRU file ID(03)              : LB_F
Customer Info-0 (0)          :
---------------------------------------
```

Possible IPMI Support for FMC modules on MTCA.4 systems

FMC FRU content:

MultiRecord Area (single-width module)

- **DC Load/Output Records (Min, Nominal, Max voltage):**

  VADJ (Most important), 3P3V**,** 12P0V, VIO_B_M2C, VREF_A_M2C, VREF_B_M2C

- **FMC unique info:**

  Module width, Connector type, number of used I/Os per bank, number of GBT Transceivers, Max clock for TCK

```
Multirecord Area:
Record Checksum OK
Record Checksum OK
FMC unique record
---subtype 7
---version 3
---modulesize Single-width
---P1ConnSize LPC
---P2ConnSize None
---P1Bank_A_NumberSignals 21
---P1Bank_B_NumberSignals 0
---P2Bank_A_NumberSignals 0
---P2Bank_B_NumberSignals 0
---P1_GBT_Number 0x0
---P2_GBT_Number 0x0
---MaxClockForTCK 10 MHz
```

```
Multirecord Area:
Record Checksum OK
Record Checksum OK
DC LOAD record:
---Output P1_VADJ
---Nominal Voltage 2.5V
---Max Neg Voltage 2.4V
---Min Neg Voltage 2.6V
---Minimal Current 10mA
---Maximal Current 100mA

Multirecord Area:
Record Checksum OK
Record Checksum OK
DC LOAD record:
---Output P1_3P3V
---Nominal Voltage 0.0V
---Max Neg Voltage 0.0V
---Min Neg Voltage 0.0V
---Minimal Current 0mA
---Maximal Current 0mA

Multirecord Area:
Record Checksum OK
Record Checksum OK
DC LOAD record:
---Output P1_12P0V
---Nominal Voltage 12.0V
---Max Neg Voltage 11.0V
---Min Neg Voltage 13.0V
---Minimal Current 10mA
---Maximal Current 500mA
```

First Implementation in Hardware

- Simple test implementation of IPMI support for FMC module on HVC-50 AMC board
- Special MMC firmware tricks MCH into thinking that connected FMC module is RTM

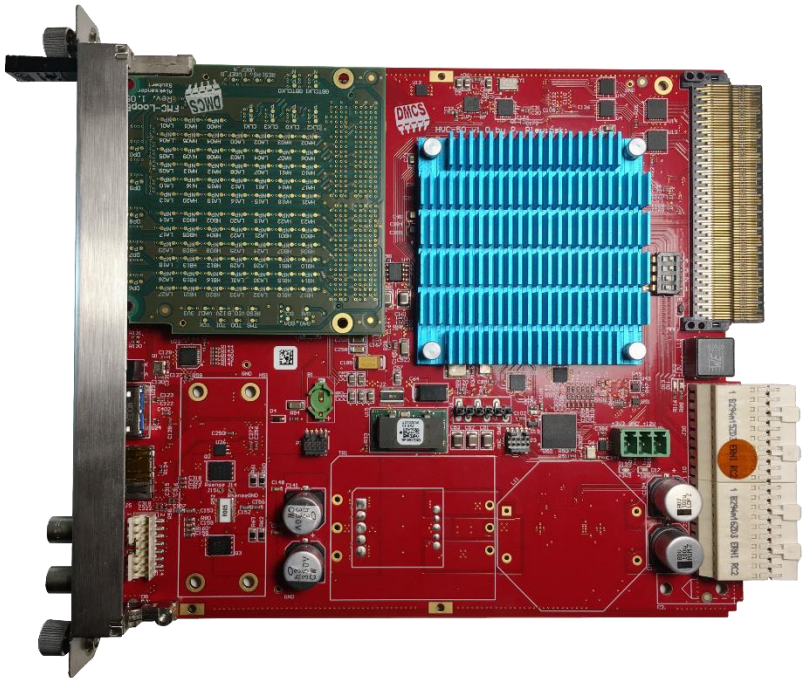

HVC-50 - FMC carrier used during tests

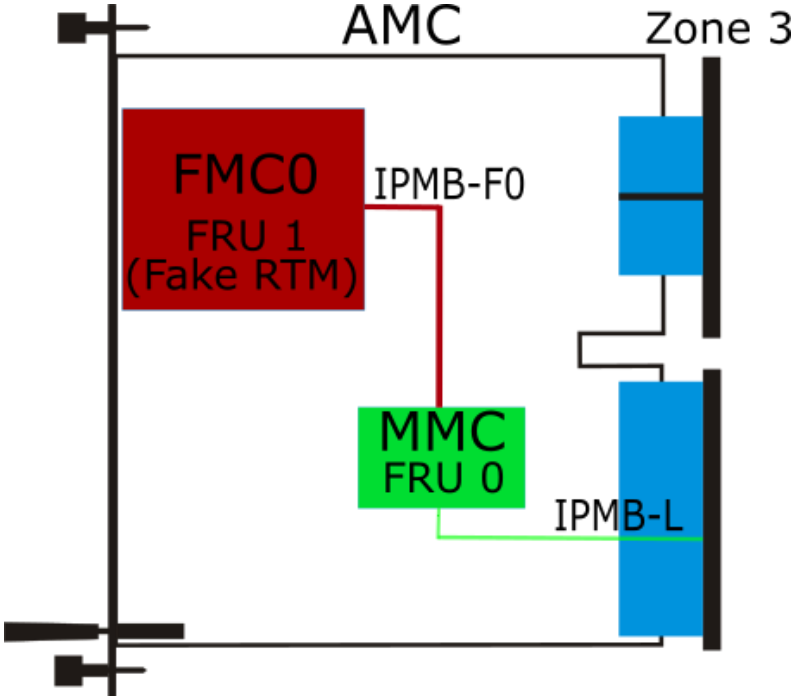

Scheme of connections

First Implementation in Hardware

- MMC provides FRU and sensors from FMC to MCH
- MCH reads FMC FRU and displays created sensors
- System can read temperature of FMC module and enable fans if it will be required

FMC sensors visible as one entity

FMC Temperature sensor

Voltage sensor

```
nat>show_sensorinfo 92
Sensor Information for FRU 92 / RTM3
=============================================================
  #   SDRType  Sensor Entity Inst  Value    State    Name
-------------------------------------------------------------
  -   FDevLoc          0xc0  0x63                   BASIC FMC
 32   Full     0xf2    0xc0  0x63  0xa0             IPMI FMC TEST
 33   Compact  0x0b    0xc0  0x63  0x00             0x00  UID:08000000000
 34   Full     Temp    0xc0  0x63  26.0 C     ok    FMC TEMP
 35   Full     Voltage 0xc0  0x63  0.3648 V   ok    FMC_VIOB
 36   Full     Voltage 0xc0  0x63  3.2528 V   ok    FMC_3V3
 37   Full     Voltage 0xc0  0x63  1.7632 V   ok    FMC_VADJ
 38   Compact  0xf0    0xc0  0x63  0x02             HS 092 RTM3
-------------------------------------------------------------
```

FMC Sensor Information in MCH

Summary

- Support for 4 FMC modules on a single AMC/RTM pair is feasible
- Update of MTCA.4.1 including the latest VITA 57.X revisions is required
- Single FRU and up to 2-3 sensors could be easily obtained
- Support from MCH is need for additional FRUs: 2-3, and
- 4-5 if we decide (but full MTCA.4 chassis could have 48 FRUs more)

Thank you for your attention