MicroTCA for Medical Applications

By Claudia Bestler, Dr. Stephan Rupp, Kontron

Within the shortest period of time, a huge choice of Advanced Mezzanine Cards has become available on the market. The most popular developments include multi-core processor boards of all mainstream processor families, as well as a diversity of interface boards. Among the typical applications are multi-processor systems for image processing. This makes Advanced Mezzanine Cards (short AMCs) very well suited for medical applications. MicroTCA represents an industry standard for using AMCs to build a new generation of high-performance multi-processor systems. The usage of AMCs provides flexibility and a ready supply of off-the-shelf components for scalable system designs with high functional density, high throughput and minimum latencies. This paper shows an overview of the technology and how to apply it to medical image processing.

Trends in Medical Imaging

Images represent an increasingly popular source of information in medical screening, diagnostic, and therapy. Familiar procedures like the X-Ray still picture for pulmonary screening or X-Ray life pictures in fluoroscopy went digital a long time ago. Computer Tomography provides three-dimensional images based on X-Rays, Ultrasound, PET (Positron Emission) or MR (Magnetic Resonance). The next level are life pictures with real-time requirements used in therapy (e.g. to support surgical procedures) or moving three-dimensional images (like 4D ultrasound baby scans). Image processing is taking care of image generation and the rendering of images. While some applications can be handled off-line (by batch processing), in other cases the images need to be ready soon after the exposure or even under real-time conditions. But the need for processing power does not end there.

Processing can also provide assistance in the analysis of images and related patient information in screening and diagnostics. Among those tasks are pattern recognition, rendering of organs, volumetric analysis, the comparison of multiple image types, and the processing of related patient information from data bases. If devices for mobile (outpatient) and stationary patient monitoring are increasing in the future, they also provide valuably sources of information to be used in combination with images.

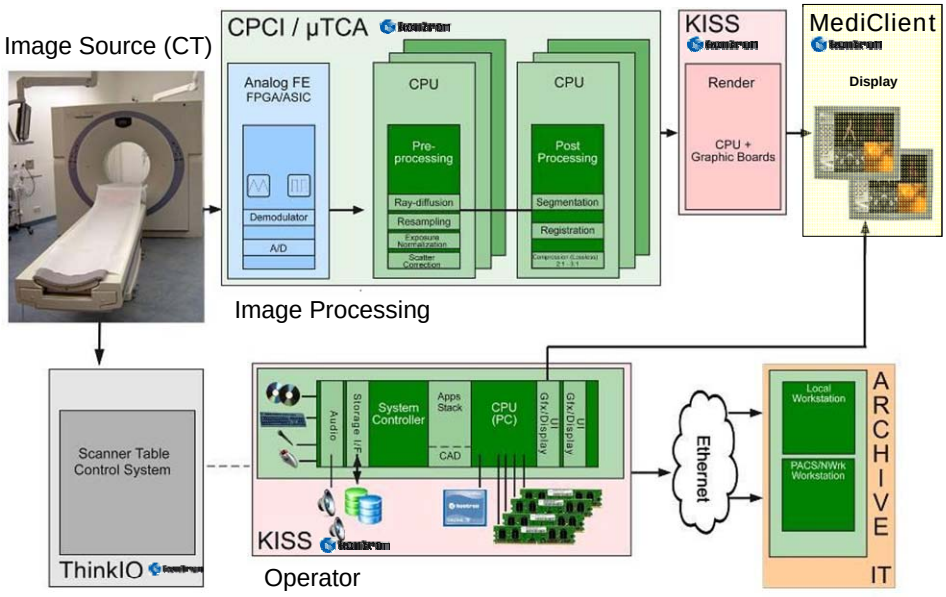

Figure 1 Image Processing System

Modern multi-core processors significantly facilitate the processing of image sources. In some cases, it allows the implementation of new methods which have been out of scope before for lack of computing power. Image processing frequently requires many cores, usually on multiple processor blades; either tightly coupled over high-speed busses (e.g. SRIO) or loosely coupled over network protocols (10GbE, GbE). Image rendering requires high-performance graphics boards, which interconnect over high-speed PCI-Express lanes. Figure 1 shows an example of an image processing systems as used for computer tomography.

On the upper left, the patient is positioned on a table, which moves through the image scanner. In tomography, the canner rotates around the patient to generate information, which is later processed to generate slices (axial sections) in a three-dimensional still picture of the inside. The reconstruction of this image is the main task of image processing. The first step is the digitization of the sensor data, which takes place in an I/O module in the first box of the block "Image Processing" in Figure 1.

The I/O modules then hands the raw image data to multiple processor boards over a high-speed connection using a popular protocol (such as SRIO or Ethernet). The next processing step (pre-processing) reconstructs the three-dimensional image from the raw image data. Cleaning out sensor artifacts, calibration and geometrical alignments also are typical pre-processing tasks. A subsequent step (post-processing) allows interpreting the 3D information in different ways (volume rendering, surface rendering or use of image segmentation to show concealed structures like vessels hidden by bones). Finally, the image is shown at a display and archived. Other components of the system are the control of the patient table and image source, as well as the operator position.

Benefits of MicroTCA for Image Processing

An image processing system as shown in Figure 1 can be implemented by a stack of servers, or by industry standard server blades such as CompactPCI or MicroTCA. In comparison to stacks of isolated servers, systems based on server blades provide the benefit of a higher density of computing and the tight coupling of processors over the backplane. In CompactPCI 2.16 and MicroTCA, the Ethernet network infrastructure is built into the system. Another benefit is the much lower power consumption by using embedded processor technology and power management. It is expected, that energy saving solutions will get much more attention in the future.

CompactPCI and MicroTCA are industry standards which are supported by a multitude of suppliers with a broad range of products. Commercial-off-the-shelf components significantly allow to shorten time-to-market and to save development costs. Originating from the PICMG, CompactPCI and MicroTCA provide in-built regulatory compliance and are aligned with further standards and trends in computer technology.

In comparison with traditional parallel bus architectures, MicroTCA supports a serial backplane and supports a variety of high-speed busses (such as SRIO, 1GbE, 10 GbE, PCI-Express), which also may be operated in parallel. This makes MicroTCA very well suited for tightly coupled multi-processor systems. At the same time, MicroTCA provides the smallest form factor on the market with a wide range of scalability. An entry level configuration may accommodate 4 to 6 processor blades (AMCs), a mid-range system up to 12 AMCs, a high-end system up to 24 AMCs with option to migrate to a fully featured AdvancedTCA system, while maintaining the same systems architecture and management.

System Components and Standards

AMCs originate from AdvancedTCA (ATCA), in industrial standard of PICMG. From a technical point of view, AdvancedTCA uses serial interfaces on the backplane (instead of the traditional parallel bus architectures), and systematically build in a management concept for all hardware components of a system based in IPMI (Intelligent Platform Management Interface). Multiple serial busses may be used at the same time within a chassis. The compatibility of boards on the bus is checked by an electronic keying, which is part of the in-built system management. Other management functions are power management and the provisioning of external interfaces for monitoring and control.

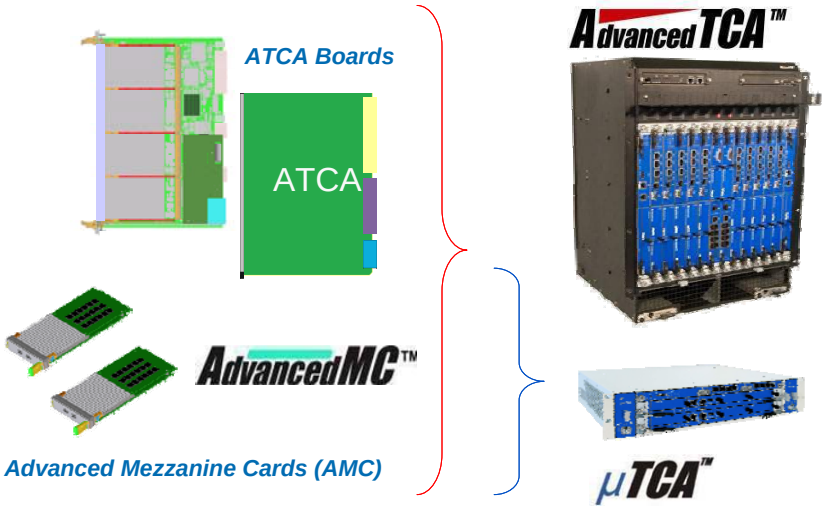

Figure 2 AdvancedMCs as system components in ATCA and MicroTCA

What are AdvancedTCA and MircoTCA systems made of? The major components are chassis of different form factors, ATCA blades and AMCs, as shown in Figure 2. The chassis contain one or two shelf managers, which are in charge of maintaining the proper operation of the hardware incl. management of power and fans. ATCA provides a high functional density by: (1) ATCA boards (8U x 160 mm x 1,2“) with a power budget of 200 Watts (2) the specification of AMCs as smallest field replaceable units on ATCA blades.

The standards from PICMG for AdvancedMC are structured in a similar way as the ATCA-Standard: AMC.0 represents the basic specification; further documents specify different transport systems on the bus (AMC.1 PCI-Express, AMC.2 Gigabit Ethernet, AMC.3 Storage Interfaces SAS/SATA, AMC.4 Serial Rapid IO). Figure 2 shows a fully equipped ATCA system on the right upper corner with 2 slots for extra AMCs on each blade, and a MicroTCA system for 8 AMCs below.

The MicroTCA standard handles the case of building systems from AMCs directly placed on a back plane. The base specification MicroTCA.0 has been released in July 2006 and describes systems with up to 12 AMCs including the functions of system management, power management and switching of Ethernet and other transport systems. The AMC specification remains untouched by MicroTCA. However, new components are specified to handle the extra functions (management and switching).

AMC Form Factors and Interfaces

AMC boards are 180.6 mm deep and 73.8 mm wide (in horizontal position) for Single modules, and 148.8 mm wide for Double modules. Depending on the pitch of the front plate, AMCs are specified in three standard sizes: Full-Size, Mid-Size and Compact, as shown in Figure 3.

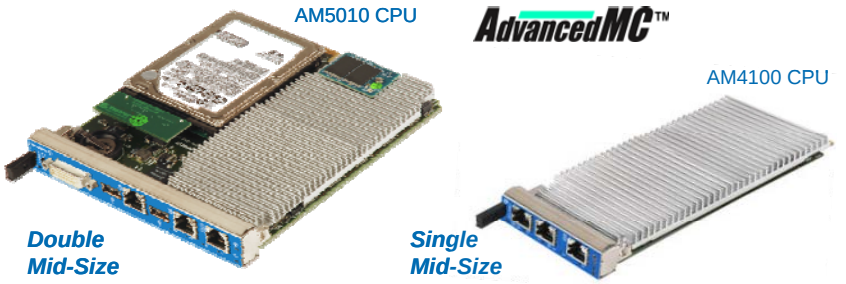

AMC	Height	Width
Single Full-Size	6 HP	73,8
Single Mid-Size	4 HP	mm
Single Compact	3 HP	
Double Full-Size	6 HP	148,8
Double Mid-Size	4 HP	mm
Double Compact	3 HP	

Figure 3 AdvancedMC Form Factors

The most popular form factor today is the Mid-Size, which allows implementing multi-core processors with a thermal budget of typically below 40 Watts. Depending on the form factor, up to 60 Watts are specified for Single modules, and up to 80 Watts for Double modules. The Single format is well suited for server types of processors without graphics and hard disks. It allows an extremely high density of processing. The Double format allows implementing complete single board computers incl. graphics and hard drives. Graphics, hard drives and a wide range of interface boards are also available on AMC.

The small form factor aligns very well with the next generations of embedded multi-core processors, which are of particular interest for image processing. As a practical case, Figure 4 shows a 12-core processor on AMC. The processor contains 12 MIPS64 cores, each operating at 600 MHz, which allows keeping total power consumption below 35 Watts (which could hardly be done by a corresponding 7.2 GHz engine). Processors in this range can process data with I/O speeds above 10 Gigabits per second, which makes such designs very well suited for the high-speed serial interfaces available on AMCs.

Irrespective of the type of function (processor, I/O, graphics etc.), each AMC holds a separate management entity, which is represented by a small micro controller which feeds from a separate power supply (3.3 Volts management power) and runs management interfaces into the system (IPMI – Intelligent Platform Management Interface). This separation allows the system manager to check the configuration, before the regular input feed (12 Volts) of the AMC is activated. At the same time, activation and deactivation of the input feed allows the replacement of AMCs on the running system (hot-swap capability).

- Cavium Octeon Plus 5650 Network Service Processor (12x MIPS64 R2 Cores @ 600 MHz each)
- 2 GB DDR2 800Mhz RAM with option for up to 4GB USB Flash drive
- SW configurable FAT pipe region Interface (10GbE, Quad GbE, PCIe x8)
- Standard Linux environment and toolchains available

Figure 4 Example – 12-Core Processor on AMC

AMCs hold card board connectors (similar to regular PCI-Express cards for PCs) on one or both sides of the board. With 85 contacts on each side, the total number of contacts is up to 170. For serial bus systems, a total of 42 contact pairs are available for differential signals. In total, an AMC supports 21 serial ports (each port representing one signal pair as receiver and one pair for transmission). Operating currently at clock speeds of 3.125 GHz and with 8/10 bit coding, each port can handle bi-directional data rates of 2.5 Gigabits per second over the backplane.

AMC-System Designs

When implementing AMC based systems based on MicroTCA, the design needs to reflect the fundamental requirements. Figure 5 shows a summary of present designs. The system on top represents a 2U server for 19" racks for 8 AMCs. This system is completely redundant (each function and component is duplicated) and thus fits requirements for high-availability systems in telecommunications. The systems below are compact designs for 4-6 AMCs (Single) and 6-12 AMCs (Single and Double).

OM5080

Telecommunications

Carrier-Grade System
(2U Server for 8 AMCs)

OM6040

OM6060

Industrial

Compact Systems
with AC Power (4 to 6 AMCs)

OM6062

OM6120

Cost optimized Systems
for Single & Double AMCs
with pluggable power supplies AC& DC
(6 to 12 AMCs)

Figure 5 AMC-Systems Designs

System management requires that each AMC connects to a manager, which handles the electronic keying and following a successful check of the configuration powers up the AMC. Concerning the AMC, this function is already part of the AMC specification and originates from using AMCs on AdvancedTCA blades. In absence of an AdvancedTCA environment, the MicroTCA specification provides the missing functionality in the following way: (1) The system configuration and the topology of the backplane is contained in a memory (EPROM) on the backplane, which is accessible from the system manager. (2) As new entities to be controlled by the system manager Cooling Units (CU) and Power Modules (PM) are defined. Power Modules represent a kind of intelligent power supply which can be addressed over IPMI and contains individual power switches for each AMC. (3) The specification of the system manager as MCH (MicroTCA Carrier Hub). The MCH is also represented with AMC form factor and can also provides external management interfaces for monitoring and control, e.g. over SNMP.

Figure 6 shows a summary of the management functions and different options for implementation. On top it shows the manager (MCH), which connects over an I2C bus to the following components:

- To each AMC by a high-level protocol on I2C (the IPMB-Protocol, IPMB meaning Intelligent Platform Management Bus),
- To the memory chip (EPROM) on the backplane which contains the system configuration over a low-level I2C interface,
- To Cooling Units (CU) and Power Modules (PM) over a high-level protocol (an additional IPM Bus, the IPMB-0).

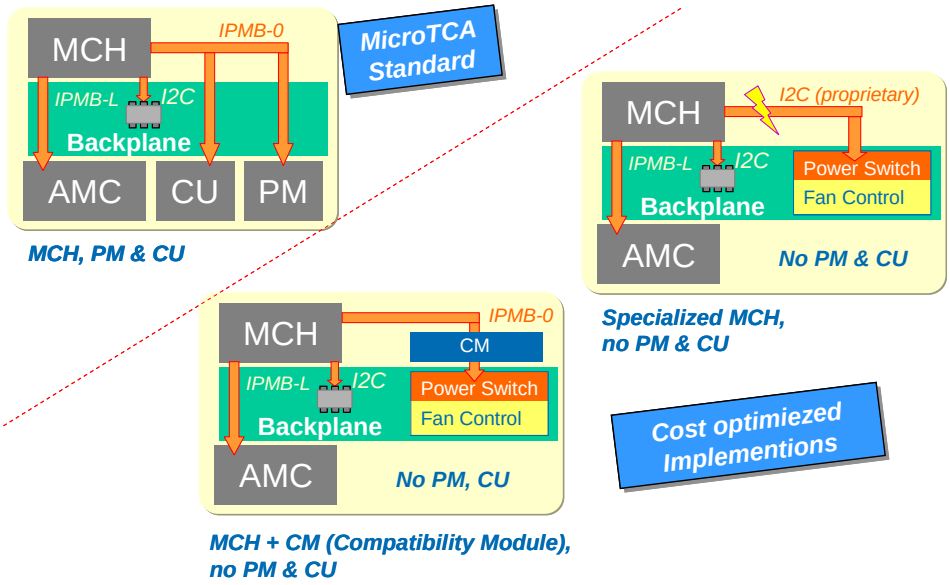

Figure 6 Management Functions and Implementation

Unfortunately the usage of specialized Power Modules and redundant power makes MicroTCA systems expensive and consumes space. For medical applications or in an industrial environment, the same results can be achieved in a much less expensive way. Instead of Power Modules, power switches and circuits for fan control may be placed on the backplane. Such a solution is entirely compliant with the AMC standards and does not require any restrictions concerning system management. As shown in Figure 6, the system manager (MCH) may control fans and power over an I2C connection. However, such a low-level interface is not in the MicroTCA standard, which has a high-level IPMI connection in this place. A low level I2C connection means a proprietary software extension in the MCH. Conformity concerning the MCH can be restored by placing a small micro controller on a socket on the backplane (the Compatibility Module or CM in Figure 6), which translates the high-level IPMB commands into proprietary low-level I2C commands. In very simple system designs, the Compatibility Module can also become the system manager.

Adding Capacity on the Serial Bus - PCIe, Gigabit Ethernet, SRIO and 10 GbE

In the much same way as PC cards on a mother board, processor AMCs may communicate with peripherals over PCI-Express (PCIe). The AMC standard defines PCIe clocks (Fabric Clock, respectively CLK3) on the AMC and backplane, as well as up to 4 PCIe lanes on the backplane for data transport with corresponding AMC ports. In the simplest case, the AMCs connect to each other over point-to-point connections on the backplane. Alternatively, a PCIe fabric switch may be placed on the MCH. In this latter case, the backplane connects the PCIe ports from each AMC to the MCH in a star topology. Figure 7 shows a MicroTCA system with such a backplane.

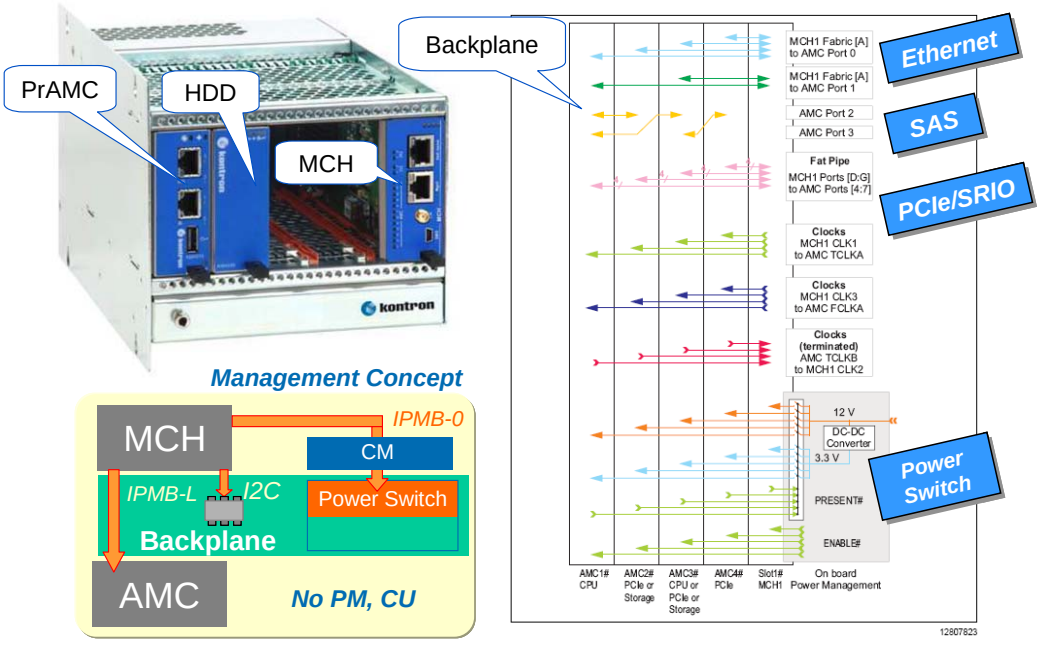

Figure 7 Example – OM6064 Systems Architecture

Other options to connect AMCs, which may be used in parallel, are Ethernet connections (1 GbE using one AMC port or 10 GbE using 4 AMC ports), or Serial Rapid IO (SRIO, also providing a speed of 10 Gigabits per second using 4 AMC ports). This extra capacity is usually needed for coupling processor AMCs e.g. for image processing. A corresponding Ethernet switch (respectively SRIO switch) can be placed on the MCH. The backplane then routes all corresponding AMC ports to the switch. The system shown in Figure provides 1 GbE connections to each AMC. In addition, a PCIe switch (or alternatively SRIO switch) can be placed on the MCH.

The backplane of an AMC system is not fixed by any standard. The AMC and MicroTCA standards provide conventions for the utilization of AMC ports and MCH ports. Which transport systems and which ports to be used are up to the systems designer. The backplane provides the corresponding mapping. An alternative system design to Figure 7 is shown in Figure 8. In this design, PCIe is routed directly from the processor AMC to 4 adjacent AMCs, which simplifies systems design and associated costs significantly: There is no need to a PCIe fabric switch and it simplifies the backplane (needs less layers on the board).

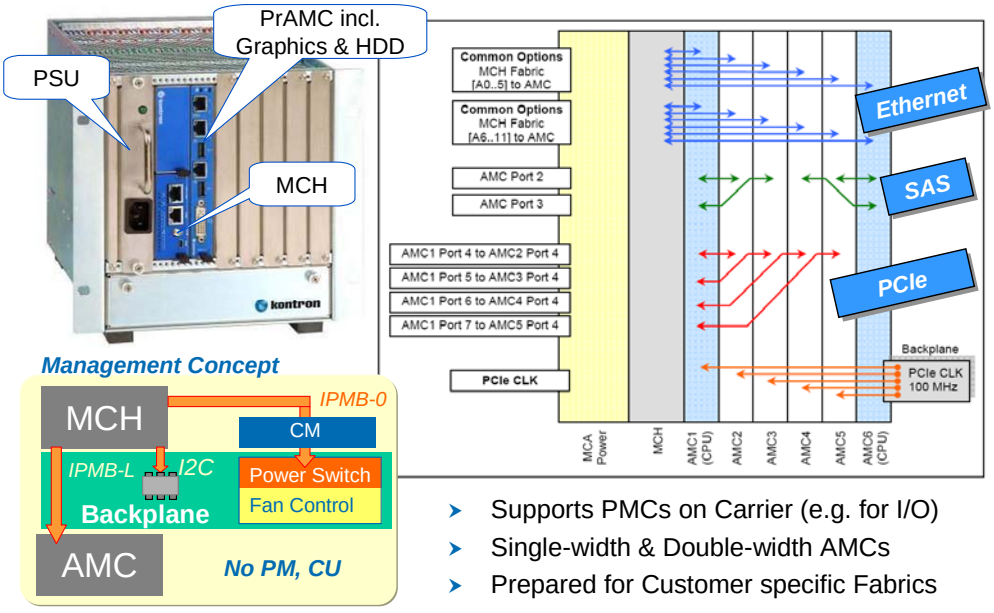

- Supports PMCs on Carrier (e.g. for I/O)
- Single-width & Double-width AMCs
- Prepared for Customer specific Fabrics

Figure 8 Example – OM6062 Systems Architecture

For AMCs, the conventions for port utilization are:

- If ports 0 und 1 are used, they should carry GbE. A basic MCH will contain a corresponding Ethernet switch and the backplane will route the AMC ports to this place. For redundant system designs, the system may contain dual MCHs with ports 0 routed to one MCH, and ports 1 routed to the other.
- If ports 2 or 3 are used, they should carry storage interfaces over SAS or SATA. Those interfaces are usually connected between processor AMCs and hard drives on AMC by point-to-point connections on the backplane
- Ports 4 to 7 may carry 1 to 4 lanes of PCIe. It is up to the system design to use point-to-point connections or a PCIe switch on the MCH.
- Ports 8 to 11 may carry another transport system for 10 Gigabits per second, frequently SRIO or 10 GbE. Usually, both systems are used in combination with a corresponding switch on the MCH. As one MCH can only carry one fabric switch, parallel usage of switches for ports 4-7 and 8-11 is only possible in systems with two MCHs.
- Ports 17-20 may carry another transport system for 10 Gigabits per second, frequently SRIO or 10 GbE. Again parallel usage of dual switched fabrics requires systems with dual MCHs. One option is the design with redundant fabrics, i.e. dual star backplane topologies for dual MCHs.

The standard handles AMC ports 4 till 7 as „Fat Pipes“, which may be used for different transport systems. In practice PCIe has become the predominant fabric in this place. Conventions for specific transport systems AMC ports facilitate interoperation between AMCs on the market. The trade-off is the re-usability of backplanes. With the system shown in Figure 7, any fabric on ports 4-7 is routed to the MCH. Depending on what AMCs present on those ports, the same system can be used for PCIe and SRIO. A convention, which places SRIO in the so called “Extended Fat Pipe” section e.g. at ports 8-11 would require a system with a different backplane.

One technical trend concerning MicroTCA systems surely is the flexibility concerning backplane designs. While the initial phase of lab prototypes in small quantities favors universal system designs, field deployment and quantities will drive cost improved designs with dedicated and application specific backplanes. Systems and components available on the market today provide everything which is needed to set up prototypes and low volume deployments from off-the-shelf components now. The modular approach assures a future proof and fast migration into larger quantities.