

Clocking in: Real-world MicroTCA needs close clocking/fabric interaction

By Will Chu

In this article, Will briefly surveys MicroTCA architecture and then focuses on four different example systems, with a particular emphasis on the integration of clocks and fabrics in MicroTCA. The example systems are:

- General purpose computing and software development platform
- WiMAX basestation
- 3G/4G basestation
- Quadruple/triple play IP Multimedia System (IMS) platform

MicroTCA architecture overview

The MicroTCA specification provides system designers with a very flexible architecture. In developing a number of working MicroTCA systems at CorEdge, we found certain applications frequently demanded tight coupling of the MicroTCA clocks and network fabrics, yet the need for this close interaction is often misunderstood.

MicroTCA systems typically consist of the following elements: AdvancedMCs, MicroTCA Carrier Hubs (MCH), MicroTCA

Power Modules (PM), and a MicroTCA chassis that comprises a backplane, connectors, and cooling units. The AdvancedMCs deliver specific applications while the remaining elements serve as general purpose infrastructure. The MCH serves as the logical and physical hub for delivering IPMI management (the *Carrier* in MCH) plus networking of common option and fat pipe fabrics and clocking (the *Hub* in MCH). The PM accepts various AC or DC inputs and provides power conversion, distribution, management control, and safety functions. The MicroTCA chassis houses and cools the AdvancedMCs, MCH, and PM. Figure 1 shows the PICMG MicroTCA RC1.0RC2 block diagram.

By intelligently pairing specific types of AdvancedMCs with MCHs that support specific networking protocols and clock types, developers can leverage this generic MicroTCA platform to deploy a wide array of applications. Table 1 summarizes a number of possible AdvancedMC/MCH pairings used for different applications.

For general purpose computing/software development platform MicroTCA systems, there are a number of implementation

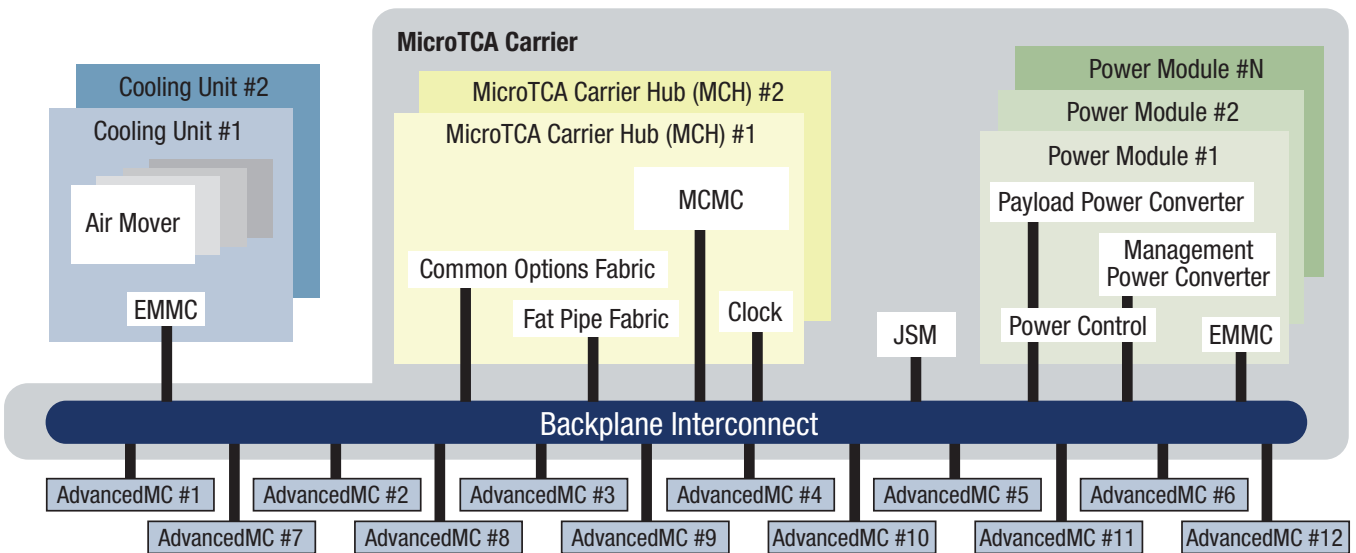


Figure 1

Application	AdvancedMCs	MCH Fabrics	MCH Clocks
General Purpose Computing/Software Development	Processor • Storage • I/O • VGA	1 GbE • SATA/SAS • PCI Express	100 MHz
WiMAX	WiMAX • Processor • Storage	1 GbE	1 PPS/30.72 MHz
3G/4G	DSP • Processor • I/O	1 GbE • Serial RapidIO	8 kHz/19.44 MHz
Triple/Quadruple Play (Voice, Video, Data, Wireless)	10 Gigabit Ethernet Processor • 10 Gigabit Ethernet I/O	1 GbE • 10 Gigabit Ethernet	8 kHz/19.44 MHz

Table 1

PRODUCT GUIDE

MicroTCA

options (Table 2). For our first implementation, we used processor, storage, I/O, and VGA AdvancedMCs. A MicroTCA general purpose computing/software development platform using a CorEdge Networks PicoTCA 2UE test and development platform is shown in Figure 2. In this case, the MCH is supplying a 1 GbE fabric between the processor AdvancedMC, I/O AdvancedMC, and the outside world. The MCH supplies a 100 MHz PCI Express oscillator to serve as the master clock domain between the processor and VGA AdvancedMCs. A direct connection on the backplane achieves the PCI Express link between the processor and VGA AdvanceMCs. The MCH supplies a SATA/SAS switch fabric between the processor and storage AdvancedMCs. It is

important to note that AdvancedMCs using PCI Express requires a single 100 MHz clock source to enable any PCI Express links. This is a very tight coupling indeed, and highlights the importance of having good clocking support in the MCH.

In future implementations (Table 3), an MCH PCI Express fabric could tie the processor, I/O, and VGA AdvancedMCs into a single PCI Express domain with the MCH 100 MHz PCI Express oscillator.

For a MicroTCA WiMAX basestation (Table 4), the key requirement is support for a 1 PPS clock typically sourced from an external GPS antenna. The MCH generates a 30.72 MHz clock locally that is synchronized to the 1 PPS external clock source and distributes those clocks to the WiMAX AdvancedMCs. Typically the base 1 GbE fabric provided by the MCH has enough bandwidth for the entire system. Some implementations of MicroTCA WiMAX basestations under consideration use a Serial RapidIO fabric in place of the 1 GbE fabric. In this case, an MCH SRIO fabric is required.

For a MicroTCA 3G/4G basestation (Table 5), the key requirement is support for a low latency Serial RapidIO (SRIO) fabric for data transport, 1 GbE fabric for management traffic, and




Figure 2

Modules	MicroTCA Carrier Hub/AdvancedMC Backplane Interfaces				
	CLK1	CLK3	Port 0	Port 2	Ports 4-7
Processor AdvancedMC		MCH CLK1	1 GbE link to MCH	MCH SATA/SAS Fabric	Direct connect to VGA
Storage AdvancedMC		MCH CLK1		MCH SATA/SAS Fabric	
VGA AdvancedMC		MCH CLK1			Direct connect to Processor
I/O AdvancedMC			1 GbE link to MCH		
MCH	Local PCI Express 100 MHz clock source to AdvancedMC CLK3		AdvancedMC 1 GbE Base Channel A	AdvancedMC SATA/SAS Fabric Base Channel B	

Table 2

Modules	MicroTCA Carrier Hub/AdvancedMC Backplane Interfaces				
	CLK1	CLK3	Port 0	Port 2	Ports 4-7
Processor AdvancedMC		MCH CLK1	1 GbE link to MCH	MCH SATA/SAS Fabric	MCH PCI Express Fabric
Storage AdvancedMC		MCH CLK1		MCH SATA/SAS Fabric	
VGA AdvancedMC		MCH CLK1			MCH PCI Express Fabric
I/O AdvancedMC		MCH CLK1	1 GbE link to MCH		MCH PCI Express Fabric
MCH	Local PCI Express 100 MHz clock source to AdvancedMC CLK3		AdvancedMC 1 GbE Base Channel A	AdvancedMC SATA/SAS Fabric Base Channel B	AdvancedMC PCI Express Fabric Base Channel D, E, F, G

Table 3

Modules	MicroTCA Carrier Hub/AdvancedMC Backplane Interfaces				
	CLK1	CLK2	CLK3	Port 0	Port 2
WiMAX AdvancedMC		MCH CLK2	MCH CLK1	1 GbE link to MCH	
Processor AdvancedMC				1 GbE link to MCH	Direct connect to processor
Storage AdvancedMC					Direct connect to processor
MCH	Local 30.72 MHz clock to AdvancedMC CLK3	External 1 PPS GPS clock from face plate to AdvancedMC CLK2		1 GbE Base Channel A with 1 GbE uplinks	

Table 4

potentially 8 kHz/19.44 MHz telco synchronization clocks sourced from an external master system clock. The MCH provides the 1 GbE and SRIO fabric and acts as a clock hub that distributes the incoming telco clocks to the AdvancedMCs.

For a MicroTCA quadruple/triple play IMS platform (Table 6), the key requirement is support for 10 Gigabit Ethernet with advanced flow control and congestion management support throughout the system. IMS networks must handle a massive amount of traffic with precisely controlled quality of service levels. With advanced flow control and congestion management support, both the data and management traffic can be carried across the 10 GbE fabric. This approach obviates the need for additional fabrics, which reduces overall system costs. In an all-Ethernet system, the requirement for synchronization clocks may also be relaxed because advanced flow control and congestion management capabilities coupled with a well-designed packet buffering scheme can provide a user experience and network performance similar to a fully synchronized network.

Implications for MCH design

In looking at these MicroTCA systems examples, it is clear that once you have selected your target applications, you select the appropriate AdvancedMCs and then select an MCH that supports both the application and AdvancedMCs. In a perfect world, a single MCH would support the entire spectrum of potential applications that could be addressed by the MicroTCA architecture. However, in the real world that is impossible. Different applications require different protocols, different bandwidths, and different clocking schemes. One size MCH cannot fit the complete range of requirements. For example, fabric options include:

- 1 GbE as a general purpose fabric for data and management traffic
- PCI Express fabric for processor AdvancedMCs to other peripheral AdvancedMCs (I/O, VGA, others)
- SATA/SAS fabric for processor to storage AdvancedMC communications
- Serial RapidIO fabric for Processor and DSP AdvancedMCs
- 10 GbE fabric for next-generation processor and I/O AdvancedMCs that will use a 10 GbE backplane interface
- Future fabrics

Designers also face a number of clocking choices:

- 100 MHz clock for PCI Express applications
- 1 PPS and 30.72 MHz clock for WiMAX applications
- 8 kHz/19.44 MHz clock for telco applications
- Custom clocking requirements

To handle this broad range of requirements, CorEdge has taken advantage of the multitongue connector architecture in the MicroTCA spec to create a series of modular replaceable clock cards for MCH tongue 2 and different fabric MCH cards for different protocols. This has allowed us to support a wide range of customer application needs, which in turn will be critical for jumpstarting the launching of the MicroTCA market.

Figure 3 shows a base MCH with PCI Express/SATA switch module and multitongue plug.





Figure 3

Conclusion

The key take-away for companies wanting to develop complete solutions using MicroTCA is to give careful consideration upfront to the interplay between clocks and fabrics in MicroTCA systems. With proper upfront thinking, MicroTCA applications can be developed relatively efficiently and, in many cases, using many off-the-shelf components including AdvancedMCs and MCHs. 



Will Chu is the president of CorEdge Networks, a developer of AdvancedTCA, MicroTCA, AdvancedMC, and IPMI products.

To learn more, contact Will at:

CorEdge Networks, Inc.
 50 Commonwealth Avenue, Suite 504
 Boston, MA 02116
 Tel: 617-267-5205
 E-mail: info@coredenetworks.com
 Website: www.coredgenetworks.com

Modules	MicroTCA Carrier Hub/AdvancedMC Backplane Interfaces			
	CLK1	CLK2	Port 0	Ports 4-7
DSP AdvancedMC	MCH CLK1	MCH CLK2	1 GbE link to MCH	MCH SRIO Fabric
Processor AdvancedMC	MCH CLK1	MCH CLK2	1 GbE link to MCH	MCH SRIO Fabric
I/O AdvancedMC	MCH CLK1	MCH CLK2	1 GbE link to MCH	MCH SRIO Fabric
MCH	External 8 kHz clock to AdvancedMC CLK1	External 19.44 MHz clock CLK2	1 GbE Base Channel A with 1 GbE uplinks	AdvancedMC SRIO Fabric Base Channel D, E, F, G w/ SRIO uplinks

Table 5

Modules	MicroTCA Carrier Hub/AdvancedMC Backplane Interfaces	
	CLK	Ports 4-7
10 GbE Processor AdvancedMC with advanced flow control and congestion management support	N/A	MCH 10 GbE Fabric
10 GbE I/O AdvancedMC with advanced flow control and congestion management support	N/A	MCH 10 GbE Fabric
MCH	N/A	AdvancedMC 10 GbE Fabric Base Channel D, E, F, G with 10 GbE uplinks

Table 6