

New FMC Standard Drives Performance in Aerospace and Defense Signal Processing Systems

Jeremy Banks, Curtiss-Wright Controls Embedded Computing

New generations of FPGAs present developers with a level of processing performance and potential I/O bandwidth that cannot easily be matched by conventional CPU configurations. While many COTS solutions enable developers to readily make use of FPGAs for processing, the real challenge to an application is often measured in terms of I/O bandwidth, latency and connectivity. For example, military Electronic Counter Measures (ECM) applications require high bandwidth data input, processing and data output with minimum latency. The FPGA Mezzanine Card (FMC) (ANSI/VITA 57.1) directly addresses the challenges of FPGA I/O by solving the dual problem of how to maximize I/O bandwidth while still being able to change the I/O functionality.



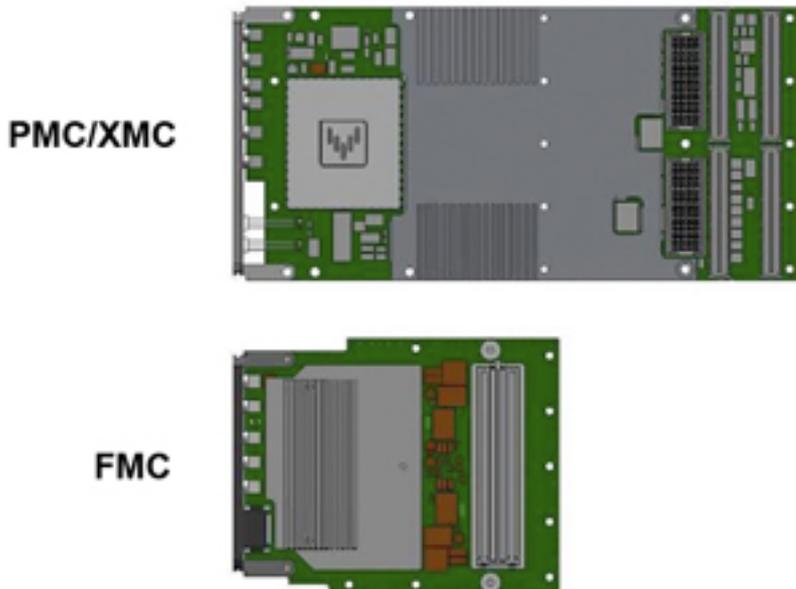
The FMC offers an elegant, simple solution because they only host I/O devices, such as ADCs, DACs or transceivers. FMC modules have no on-board processors or bus interfaces, such as PCI-X. Instead, FMC modules take advantage of the intrinsic I/O capability of FPGAs to separate the physical I/O functionality on the module from the FPGA board design of the module's host, while maintaining direct connectivity between the FPGA and the I/O interface.

Mezzanines for Rugged Computing

The most popular mezzanine format, for defense embedded computing are PMC which uses the PCI, and PCI-X. The newer XMC replaces PMC's parallel PCI or PCI-X bus with a serial interface, of which the most common protocol supported is PCI Express. Unfortunately, the throughput available from FPGAs is beyond the capabilities of PMC or XMC. FPGAs can be used to implement the necessary interfaces, so advantage can be made out of the direct coupling of processing performance and I/O bandwidth.

Summary of FMC Connectivity

The purpose of the FMC specification is to allow one or more FPGAs on a host card to connect directly with the I/O devices on the mezzanine module - just as if the device were on the host board. Busses like PCI-X are redundant and would get in the way of the FPGA and its I/O devices. This intimacy means the interface can be optimal and savings can be made in real estate, cost and power - and boost bandwidth and reduce latency.



An FMC is similar in height and width to a PMC, but almost half the length. The reduced width, compared with PMC or XMC enables up to three FMCs to be fitted to a 6U host. The FMC specification has a default stacking height of 10 mm, but permits a stacking height down to 8.5 mm, too, for low profile solutions. The majority of FMC host/carriers use VPX (3U and 6U), VXS and AMC formats but there are also PCI Express solutions such as the Xilinx ML605 Virtex-6 evaluation card.

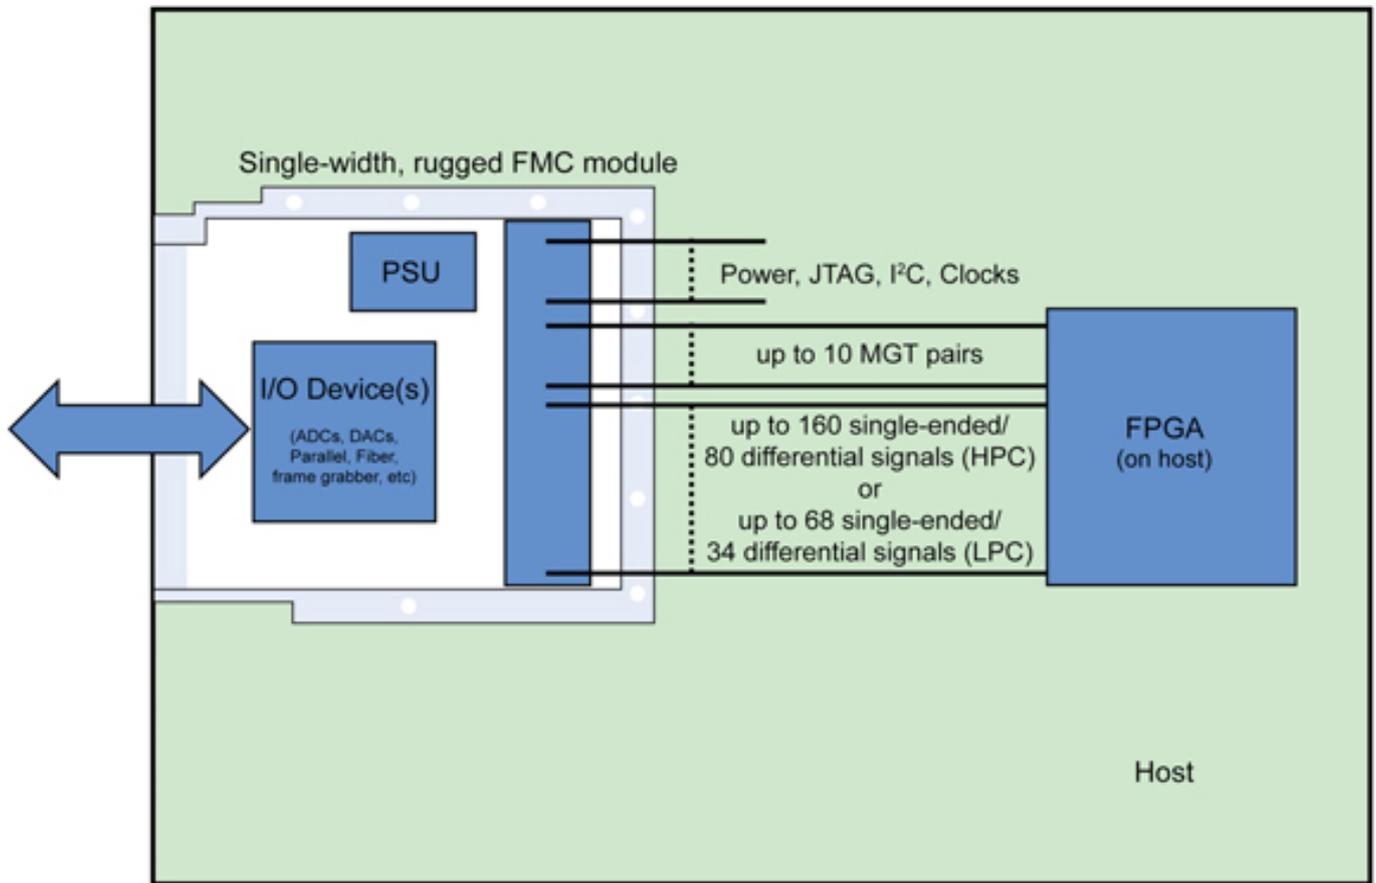
The FMC specification provides for a large number of differential connections, up to 80 pairs, (or 160 single-ended signals), to support one or more high speed parallel interfaces between the FPGA and I/O devices. There are also a number of serial connections (up to ten pairs) suitable for Multi-Gigabit Transceivers (MGTs) operating up to 10 Gb/s. FMC modules and hosts support two connector options; a Low Pin Count (LPC) 160-pin connector or High Pin Count (HPC) 400 pin connector. The majority of FMC solutions are likely to use the HPC variant. Although aimed at I/O, FMC can be used for any function that might connect to an FPGA including DSPs, memory or even another FPGA.

Connectivity

Connectivity for FMC modules is unusual in that the number of active connections is not defined; only the upper limit. This means that host carriers need not provide the same number of FPGA signals as another host. To fully populate an HPC solution may require a large FPGA, so reduced pin-out offers cost sensitivity. This is something to be aware of, but the specification defines the signals populate the LPC or HPC connector at a given position and add to the connector in a given sequence such that if two hosts provide x signal, they will use the same connector pins.

Power Supply

When it comes to power supply requirements, the FMC specification has a neat trick: the host detects what the FMC's power should be and the host provides it. This is achieved through the host interrogating the FMC's E2PROM and an adjustable power supply. The benefit to the FMC is a simplified power requirement thereby freeing up valuable real estate for more I/O.



Usable PCB Real Estate

Although around half the PWB area of an XMC, the FMC can sometimes achieve greater I/O functionality, most notably for rugged applications. If the solution requires a large FPGA and if the XMC module complies with the VITA 20 specification, there are restrictions as to where the FPGA can be located. In turn, this may limit the available area to fit the I/O devices. Consider an actual example with a pair of designs using the same I/O devices for a rugged application; one using an XMC format card and one using an FMC format card.

Because the rugged XMC specification requires an area across the middle of the board to mate up with a host stiffening bar (which doubles as a primary thermal interface conduction-cooled variants), a large FPGA (for example 35 mm x 35 mm) invariably needs to be fitted to the area of the circuit board closest to the front panel and just where the design would want to fit the I/O devices. The useful space in which to fit the I/O devices is perhaps a quarter of the overall real estate of the XMC and not very efficient. In comparison, the FMC even though it is around half the size of the XMC, has a far greater real estate area for the I/O devices. In this example, the FMC is able to support two ADCs for two 3GS/s channels compared the single channel of the XMC. Of course an XMC using a smaller FPGA, or not restricted

by the rugged XMC specification may not be affected to such an extent - provided it still has a sufficient number of I/O connections to the devices. An FMC may be smaller, but it may still be able to support greater functionality than its larger XMC equivalent.

The choice of which mezzanine format is best for rugged embedded computing solutions will ultimately be down to issues such as application details, perception of risk, development timeline and personal preference. The baseline for choosing which mezzanine is most suitable for certain applications is how it compares with a monolithic board i.e. a single PWB with all functionality onboard. A monolithic card usually provides the best technical solution because it does not have the restrictions imposed by segmenting the design such as number of connector I/O pins to the mezzanine.

Jeremy Banks is a product marketing manager at Curtiss-Wright Controls Embedded Computing. For more information, contact Curtiss-Wright Controls Embedded Computing, 9540 Vassar Ave., Chatsworth, CA 91311; (818) 998-0070; www.cwcembedded.com [1]

Source URL (retrieved on 04/26/2015 - 12:19pm):

<http://www.ecnmag.com/articles/2010/04/new-fmc-standard-drives-performance-aerospace-and-defense-signal-processing-systems>

Links:

[1] <http://www.cwcembedded.com>