


» Whitepaper «



Network Technology for Embedded Computing

Authors:

Christoph Neumann
Stephan Rupp

Network Technology for Embedded Computing

Networks are gaining importance in all disciplines of embedded computing: At the high-end, multi-processor systems need high-speed networks for interconnectivity and feeding in data. The scope of this paper are major trends in Ethernet network technology and how network technology, both hardware and software, is implemented in high-end embedded systems. As practical examples, embedded form factors CompactPCI and MicroTCA are shown at 10 Gigabit/s speed for switching and packet processing. The paper concludes with a practical look into the components of a switching platform for embedded products.

CONTENTS

Network Technology for Embedded Computing	1
Summary	2
Contents	2
Trends in Ethernet Network Technology	3
Embedded Switching Products	3
Embedded Product Lines	4
10G for CompactPCI and MicroTCA	4
Packet Processing	6
Platforms for Network Technology	6
About Kontron	8

Trends in Ethernet Network Technology

The evolution of network technology to higher densities is most visible when having a look at interface connectors for 10 Gigabit/s speed (see figure 1): over just five to six years, connector size and power consumption of fiber-optic or copper transceiver modules have been consistently shrinking (from XENPAK to XFP and SFP+). Compact high-speed transceivers easily find their way into embedded products, such as XFP and now SFP+ for 10GbE connectivity over optical fiber or copper.

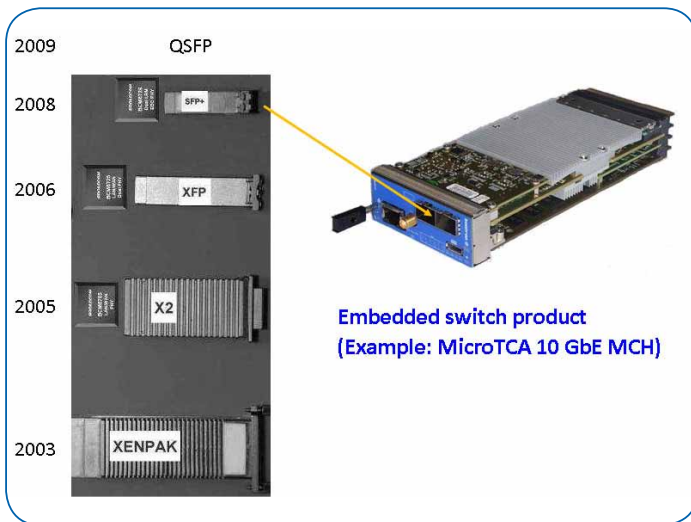


Figure 1: Evolution of high-speed switching to small form factors and high densities

Further down the road are 40 GbE technologies. Backplane speeds will also increase. Embedded form factors such as MicroTCA, AdvancedTCA and VPX currently enable 10 GbE over multiple 2.5 G lanes using the XAUI interface. With higher serial speeds, future backplanes will carry 10 GbE on a single port and 40G on the 4 lanes (today used for 10G XAUI). This type of extensibility makes the current embedded form factors future proof.

Higher speeds according to IEEE 802ap (10GBase-KR) have now been approved as standard. It will allow 10Gbit/s speeds to be carried over one single lane at 6.25 GHz frequency (with 8/10 coding). As usual, the interface will be backward compatible with XAUI and 1GbE. The standard does not end there: the IEEE P802.3ba 40GBase-KR4 task force laying the foundation for 40 Gbit/s over multiple lanes (4-lane XAUI infrastructure) and external interfaces over QSFP optical transceivers, respectively copper interconnects with twin-ax cabling over SFP connectors.

In communication networks, Ethernet is becoming a universal transmission standard for both end-to-end connectivity between end-users, as well as for aggregation

of traffic. The corresponding technologies and features are summarized as Carrier Ethernet. They include synchronous Ethernet, IEEE 1588 precision time protocols, ITU G.8261 and G.8262 synchronization standards for circuit emulation over Ethernet, as well as ITU G.8031 and G.81302 Ethernet protection switching for ring topologies.

At network layer, IPv6 is waiting for introduction and becoming a checkmark feature in many communication projects. It is becoming essential for any future proof product. For interworking with core networks, MPLS (multi-protocol label switching) on the Ethernet switch is gaining attention in Carrier Ethernet installations. Among other future requirements are end-to-end traffic management and priority based flow control to support Storage Area Networks and video enabled networks.

Some of those carrier types of requirements will find their way into industrial applications, such as managed layer 2 and 3 (which is state-of-the-art in communications today), and support of ring redundancy protocols incl. fast recovery. Other requirements for controls in the industrial environment are real-time-requirements with guaranteed latencies.

Embedded Switching Products

Implementing network technology in embedded computer products needs a closer look at the features required in hardware and software. The type of switching needed to provide high-end feature sets is always a combination of switching silicon, a switch controller and software on the switch controller. There is a choice of switching silicon on the market, as well as controlling software for configuration and management.

Figure 2 shows the development cycle for an embedded switching product. Development is driven by the customer feature requests and the target hardware platform definition, i.e. the chosen embedded form factor (such as CompactPCI, ATCA, VPX, MicroTCA or other). Requirements need to be mapped to suitable switching silicon on the market, as well as a matching software package containing the right feature set. Most importantly, the resulting embedded switching product needs to be maintained in terms of release upgrades. Depending on customer requests, feature sets are summarized into products with up to 2 software release updates per year.

The feature sets concerning silicon and software for switching products are complex by the nature of their application. Ethernet based networks are universally applied in the industries and the feature sets are diverse. Also, Ethernet based networks are on a constant path of evolution, which is the reason for their world-wide success through all sorts of applications, all speeds, distances and the most diverse kinds of physical interfaces, including backplanes.

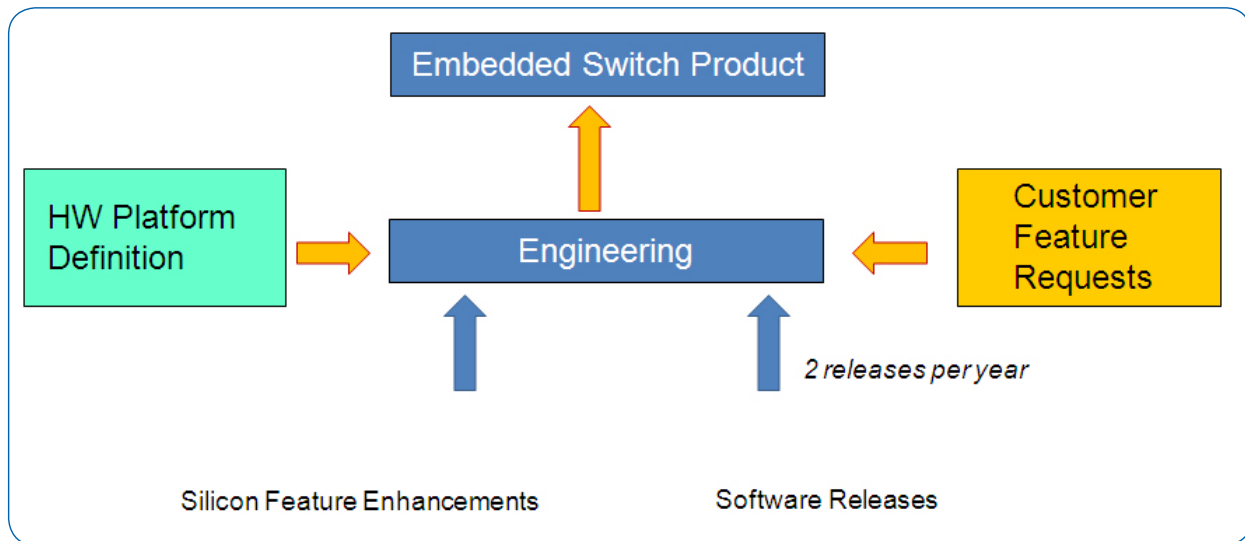


Figure 2: Product Development Cycle for Embedded Switching

Among the typical feature sets for switching silicon for embedded applications are the use of VLANs to manage traffic, multicasts to efficiently distribute traffic streams to processor blades, support functions for Layer 3 (network layer), or support of traffic classes with different priorities and traffic policies (QoS, Quality of Service). The use of spanning trees protocols (RSTP) allows repairing broken links inside a system and thus providing high fault resilience and high-availability. Depending on the area of application, extra feature sets apply, such as ring redundancy protocols in Ethernet ring topologies. Ring topologies are used inside of vehicles, vessels or communication networks. They need much less cabling than the star-type of networks applied in IT type of installations, while providing a redundant path. Special feature sets also apply for Next Generation Telecommunication networks, in particular for Carrier Ethernet in Metropolitan Networks or servers for IPTV.

Embedded Product Lines

There are three PICMG standards of embedded form factors, which incorporate network technology: AdvancedTCA, CompactPCI and MicroTCA. While CompactPCI uses parallel bus architectures, AdvancedTCA and MicroTCA have been moving to serial backplanes. PICMG 2.16 defines Gigabit Ethernet for CompactPCI on the backplane. Higher speeds are possible using front cabling.

The serial bus architecture used in AdvancedTCA and MicroTCA allows a multitude of high-speed fabrics, among them multiple 1 GbE or multiple 10 GbE links. Alternatively to Ethernet, local high-speed protocols such as Serial Rapid IO or PCIeExpress can be used to interconnect processor blades. On VME, the VITA group currently investigates the use of serial busses and GbE under the Acronym VPX. While many applications will never need backplane speeds beyond 1 Gigabit/s, speed matters for high-end applications such as in communications, military and aerospace, as well as image processing in industrial and medical environments.

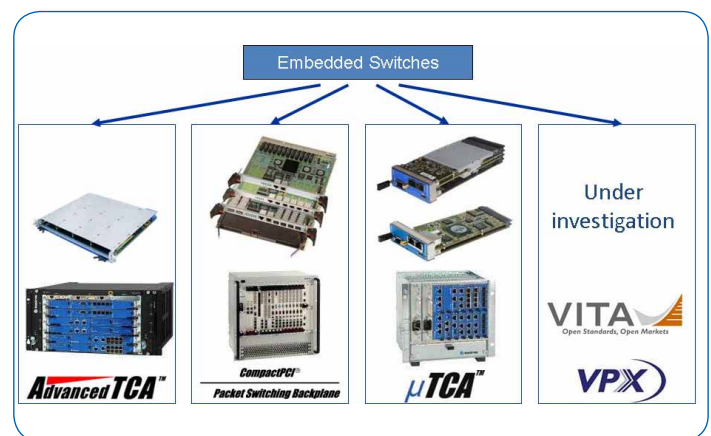


Figure 3: Embedded Product Lines

For a supplier providing embedded form factors, building blocks from one form factor may be reused on other form factors. This holds for the latest processor technology, as well as for the latest switching technology. This way, high-end processors on ATCA can also be made available in CompactPCI or MicroTCA. Switch designs can move between CompactPCI, MicroTCA and AdvancedTCA. In terms of functional density, the most challenging form factor is MicroTCA. Providing designs for MicroTCA with 10GbE switches and multi-core CPUs represents a skill that all other form factors benefit from. Depending on the standard, Ethernet in embedded product lines is applied in different ways.

10G for CompactPCI and MicroTCA

CompactPCI according to PICMG 2.16 supports Ethernet on the backplane at 1GbE speed. This speed may be sufficient in many multiple-processor environments. Matching Ethernet switches in 6U CompactPCI complete the system and provide L2 and L3 managed features such as IP-multicast. Optionally, the Ethernet switches provide extra ports on the front for 1GbE copper links or 10GbE copper or optical links.

The ports on the front may be used as uplinks to feed in traffic, and also to provide extra GbE connectivity to processor blades with traffic demands over 1 Gbits/s.

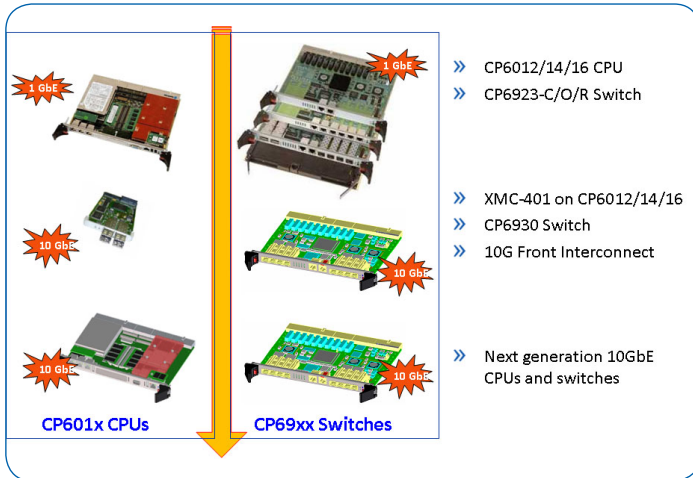


Figure 4: CPUs and Network – Matching performance in CompactPCI

Depending on the application and performance of the processor blade, traffic throughput beyond multiple GbE links can be covered by 10 GbE front cabling. This needs a 10 GbE interface on the processor blade (to be provided as upgrade by an XMC), as well as a 10 GbE switch with sufficient front ports. Such products are currently under development and will move CompactPCI products into the domain of speeds beyond 1 GbE. Such developments are standard compliant according to PICMG 2.16 and provide a migration path to any standard compliant system currently deployed.

CompactPCI designs are available for rugged applications in different packages: (1) commercial air cooled (the standard version), (2) Rugged Air Cooled (with enlarged heat sinks according to VITA47-EAC3/6 and E2 qualification), and (3) for Conduction cooled (with conduction cooled frame and wedge locks according to VITA47-ECC4 and E2 qualification).

In MicroTCA, the serial backplane already provides multiple Gigabit speeds over a variety of fabrics. On the 2U form factor of MicroTCA (i.e. single-width AMCs), processing power tends to be smaller than 6U CompactPCI, so multiple 1GbE speeds are sufficient for inter-processor-communication. Alternatively, PCIeExpress is used to connect peripherals.

Serial Rapid IO provides a lean, low level programming interface for high-speed interconnections between processors and for high-speed IO cards such as camera interfaces and image sources. Serial Rapid IO (SRIO) currently operates at speeds of up to 10 Gbits/s with low latencies. The programming interface may seem less comfortable to network-level programmers (memory mapped IO versus IP sockets). However, both types of programming interfaces apply for different areas of applications: SRIO in particular for any programming which is close to the

hardware, e.g. for safety critical applications or below the level of an operating system.

In MicroTCA, speeds at 10 Gbit/s have been the domain of dedicated multi-core packet processors, such as the Cavium Octeon. With multi-core processors becoming faster, the next generation of x86 and Freescale processor AMCs will be able to fill 10 GbE pipes over network interfaces. Also, there is a 4U form factor in MicroTCA (double-wide AMCs), which enables the implementation of processors with higher thermal design power. The new generation of MCHs meets those demands on high-end 10 GbE switch, i.e. advanced feature sets at full speed.

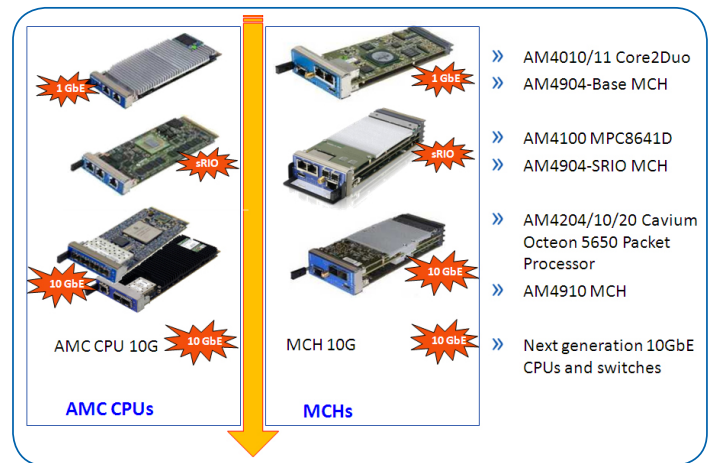


Figure 5: Highest density – CPUs and Switches in MicroTCA

For example, the Kontron AM4910 is a fully MicroTCA complaint MCH with high-end network features. It provides 1 GbE base interfaces to 12 AMCs, 10 GbE fabric interfaces to 12 AMCs and 10 GbE MCH interlink for redundant configurations. The AM4910 provides two 10 GbE uplinks on the front via SFP+. The modular concept allows configuration according to demand, with options such as a Telco Clock PLL module. The enterprise class, non-blocking Broadcom 10G switch provides the same advanced feature set as available in our AdvancedTCA and CompactPCI products.

Switching features include Static link aggregation (IEEE 802.3ad) on any port combination, classic and rapid spanning tree algorithms (IEEE 802.1D, IEEE 802.1w), Quality Of Service on all ports (IEEE 802.1p), Full Duplex operation and flow control on all ports (IEEE 802.3x), Static MAC filtering, Port Authentication (IEEE 802.1X), Layer 2 multicast services using GARP/GMRP (IEEE 802.1p), VLAN support including VLAN tagging (IEEE 802.3ac), dynamic VLAN registration with GARP/GVPR (IEEE 802.1Q) and Protocol based VLANs (IEEE 802.1v), Double VLAN tagging, Port Mirroring, IGMP snooping, CoS (Class of Service), DiffServ (Differentiated Services) and ACL (Access Control Lists). Switch management is provided via CLI (command-line interface), Telnet and SNMP, both out-of band (front panel FE) or in-band over the fabric interfaces. Additionally, a web based GUI is supplied for fast and easy configuration,

debugging and lab use. The AM4910 is prepared for upgrading to more L3 features on customer request. Using a common switch technology platform allows to leverage the synergies of a comprehensive, well tested and stable core platform across different form factors.

Packet Processing

In a 10 GbE based MicroTCA system, switches can be combined with 10 GbE packet processors, such as the Kontron AM4204 (4 GbE uplinks on front) or AM4220 (10 GbE uplinks on front). Ethernet switches operate on Layer 2, even if they are able to analyze Layer 3 information. They cannot conceal processor blades (hosts) in the system by terminating the network layer protocol (IP layer). Switches also cannot handle state based decisions such as handling sequence numbers or session IDs of packet streams. Such operations, which are required by virtual IP addresses, for encryption off-load or load balancing, are ideally suited for a packet processor. Other types of applications are firewalls, routers, deep packet inspection e.g. for virus scans, lawful interception or customization of packet streams.

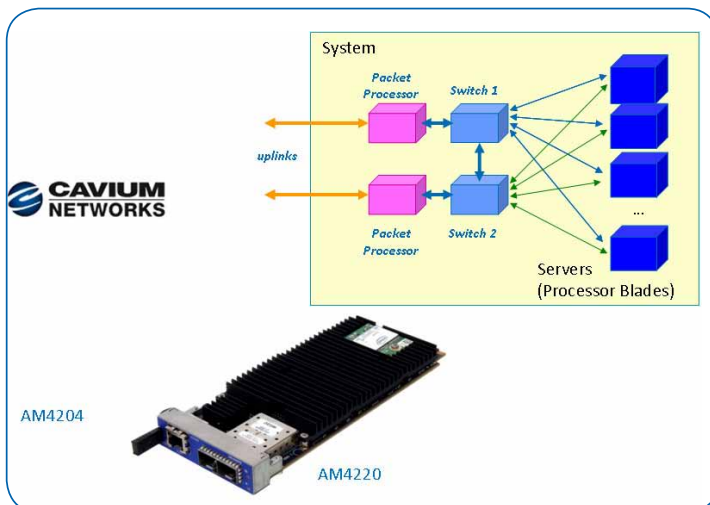


Figure 6: Packet processing at Network Layer

The AM42xx series of processor AMCs are using a 12-core MIPS64 processor (Cavium Octeon Plus 5650), which is especially designed for packet handling at high speeds. It is using 10 GbE interfaces at the AMC ports (as well as PCIe when used in combination with a host CPU on an ATCA blade). It supports a Linux operating system (such as WindRiver PNE) and optional software suits for packet processing, depending on the type of application.

With CPUs becoming faster, new designs in CompactPCI and MicroTCA are now ready to support speeds to 10 Gbit/s with enterprise class of network features. At system level, management capabilities are becoming increasingly relevant for monitoring and remote administration. The new generation of MCHs provide full system management functionality on board, incl. SNMP features, as well as advanced debugging and tracing features for mTCA environments.

The same management capabilities can be added to other embedded product lines, such as CompactPCI, VPX (currently under investigation) or small form factor conduction cooled systems. Technologies for switch administration such as CLIs, SNMP or Web-based management interfaces are of general use in systems administration and provide ease of use for any type of application. Key competence in CPU technology, network and management represents the right technological base for any project based on standard systems or customized designs.

Platforms for Network Technology

Which components are required in a platform for network products? Which hardware and software to chose for switching and management? As example, the components of the Kontron platform are shown. Kontron designs switches and network technology in-house. Providing network technology under the same roof enables strong synergies between engineering teams in terms of standards support for form factors such as CompactPCI, ATCA, VPX and MicroTCA, as well as for customized form factors, both Hardware and Software.

Kontron is using its own IPMI core implementation in the AdvancedTCA, CompactPCI and MicroTCA product lines. For switching hardware, Kontron has chosen the market leading silicon from Broadcom, the Strata XGS family, because the roadmap ensures regular feature updates, pin-compatible successors and because of the strong commitment of Broadcom to Ethernet at future higher speeds of 40G and beyond. For performance tests and regression testing, Kontron supports a comprehensive test set-up in its labs, including 10G network testers.

As software, the Fastpath Enterprise from Broadcom (formerly LVL7) is used, running on a professional WindRiver Linux environment on PowerPC processors as switch controllers from Freescale or AMCC. Depending on demand, there is a range of switch controllers deployed, starting from PPC 405GPr at 400 MHz, PPC 405Ex at 600 MHz to the Freescale MPC8547PQIII at 1 GHz. Current products are based on WindRiver 2.0 release, new products will move to Release 3.0. Typically, the operating system is not visible for customers, switches are sold as "black boxes" or closed environments. On request BSPs (board support packages) can be provided to enable customers to add their own applications.

Given the multitude of features and options, switching software tends to be tedious to test. Hence another major requirement is the strong support for software validation on the target silicon. The Broadcom FASTPATH SW package covers this requirement pretty well: there are more than 3500 test cases available.

Given the multitude of options and features, the configuration of switches also tends to be tedious for customers. A familiar environment such as the Cisco-like CLI (command line interface) which is provided with the



Figure 7: Network Technologies under one Roof

switches facilitates customer configurations. Also, the switches provide an SNMP interface for integration into management systems.

Feature updates need special attention to match customer needs with state-of-the-art switching. Kontron products today are based on Fastpath Release 5.x and Releases 6.x. The roadmaps provide an excellent view about features in future releases for customer projects. Examples for advanced switching features are L2/L3 multicast customer features (e.g. support for meshes of IPMG snooping enabled switches). Multicast is relevant in multi-processor systems for efficient distribution of traffic. It is also relevant for IPTV applications. Other custom features are customizable packet memories, the support of dynamic provisioning (slot based, system based, or network based), as well as platform specific adaptations and commands.

About Kontron

Kontron designs and manufactures standards-based and custom embedded and communications solutions for OEMs, systems integrators, and application providers in a variety of markets. Kontron engineering and manufacturing facilities, located throughout Europe, Americas, and Asia-Pacific, work together with streamlined global sales and support services to help customers reduce their time-to-market and gain a competitive advantage. Kontron's diverse product portfolio includes: boards and mezzanines, Computer-on-Modules, HMIs and displays, systems, and custom capabilities.

Kontron is a Premier member of the Intel® Embedded and Communications Alliance.

For half-a-decade now, Kontron has been named a VDC *Platinum Embedded Board Vendor*. Based entirely on user feedback, industry professionals evaluate vendors on over 45 non-product related criteria. Kontron is only one of two companies to receive the Platinum award 5-years running.

Kontron is listed on the German TecDAX stock exchange under the symbol „KBC“.

For more information, please visit: www.kontron.com

CORPORATE OFFICES

Europe, Middle East & Africa

Oskar-von-Miller-Str. 1
85386 Eching/Munich
Germany
Tel.: +49 (0)8165/ 77 777
Fax: +49 (0)8165/ 77 219
info@kontron.com

North America

14118 Stowe Drive
Poway, CA 92064-7147
USA
Tel.: +1 888 294 4558
Fax: +1 858 677 0898
info@us.kontron.com

Asia Pacific

17 Building,Block #1,ABP.
188 Southern West 4th Ring Road
Beijing 100070, P.R.China
Tel.: + 86 10 63751188
Fax: + 86 10 83682438
info@kontron.cn

