Size Matters – Applying the OM5080 2U Carrier Grade Platform

By Dr. Stephan Rupp, Claudia Bestler, Kontron

ATCA and MicroTCA have been specially designed for communication solutions in public and private networks. The pre-dominant feature in such applications is high availability, which is achieved by a redundant systems design without single point of failure. However, not every carrier grade system is large and contains a high number of processor blades. In small systems, the price for redundant network hubs and redundant shelf management becomes a major cost driver. The OM5080 represents a small carrier grade platform, which is designed to minimize the cost overhead associated with conventional ATCA or MicroTCA systems, while providing full redundancy and superior space efficiency. This paper shows areas of applications for the OM5080 in public networks and professional mobile radio networks. Among the applications are Media Servers, Signaling Gateways, CMTS-Cores, Radio Network Controllers and Tetra Switches.

Public Networks

Public communication networks have become an indispensable part of our life. Everywhere we use our mobile phones world wide, or connect or PCs and laptops to the Internet, there is access to a public network. The public network infrastructure collects traffic from millions of subscribers. Depending on the type of network, there are different technologies and infrastructures in place. Figure 2 shows a summary.

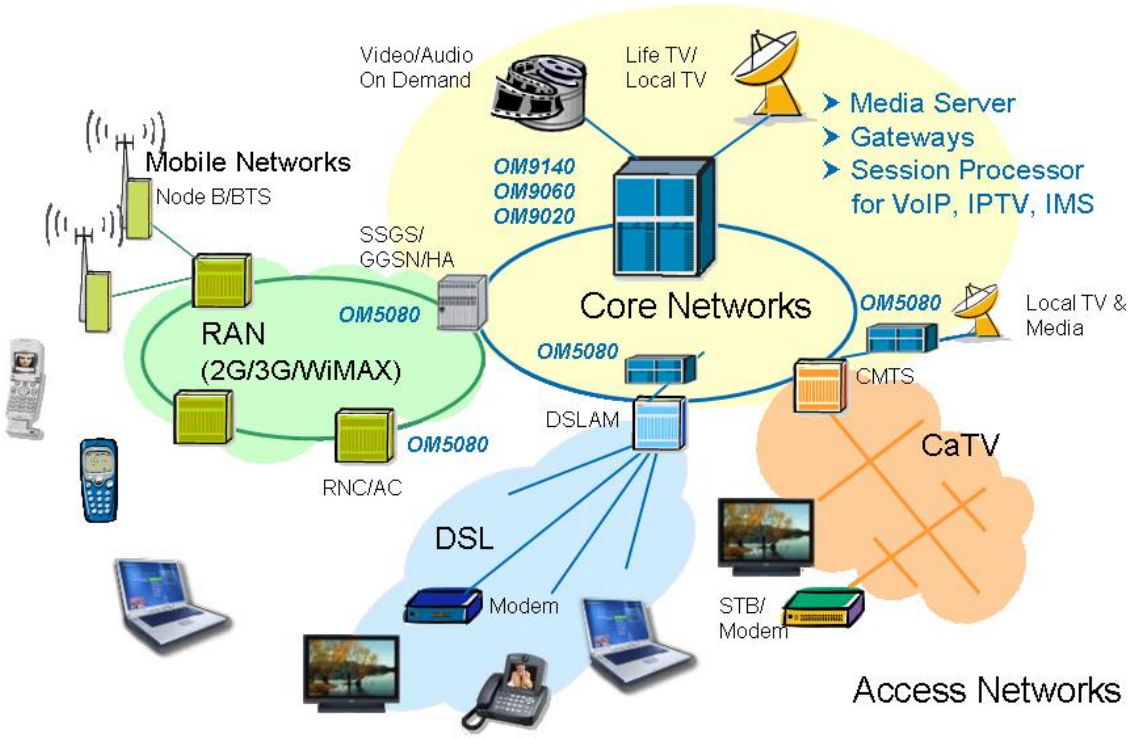

Figure 1: Public Communication Networks

Mobile phones connect over Radio Access Networks (RAN). Depending on the density of subscribers, a large number of base stations is in place to collect the traffic. Each base station handles a number of radio cells, which can accommodate a few hundreds of subscribers. Depending on the technology, there are 2nd generation mobile networks in place (predominantly GSM with close to 3 billion subscribers world wide), and its 3rd generation successor UMTS. Other popular radio technologies are WiMAX and the future 4th generation of GSM, the so called LTE (Long Term Evolution). The differences are basically in the radio interfaces, the network infrastructure is similar. Traffic from a number of base stations is collected and aggregated at base station controllers, a next level in the hierarchy represents a Radio Network Controller. The Radio Network Controller also handles authentication for network access fur subscribers, as well as the mobility management for mobile subscribers (so they can move in the network and still be reachable). The Radio Network Controller aggregates a high number of subscribers and represents the interconnection to the core network. Some figures on numbers of subscribers and corresponding traffic models in terms of transactions and throughput can be found in [1].

When we connect our Laptop or PC to the Internet over LAN or Wireless LAN, we frequently use a DSL connection. The DSL modem places the data channel on the telephone wire. Again, the network topology in the telephone access network is hierarchical. In some place, a matching entity to the DSL-Modem is needed, to extract the data channels from the telephone wire. The matching entity to the DSL modem is the so called DSLAM, Digital Subscriber Line Access Multiplexer. A DSLAM is either placed at the curb (e.g. VDSL networks), or at a Central Office location where tens of thousands of telephone lines terminate (ADSL). Behind the DSLAM, traffic is further aggregated them over a data network (usually a metropolitan network of optical fibres). Like the Radio Network Controller, the DSLAM is also responsible for authentication of network access: While the Radio Network Controller relies on access code on the SIM card, the DSLAM relies on the User ID and Password that have been entered into the DSL modem at the subscriber, as well as the ID of the telephone line. Hence, the DSAM also represents the logical point of interconnection between the telephone access network for the data channel and the Core Network.

Another popular network infrastructure are Cable TV networks (CaTV). CaTV networks cannot only accommodate analogue and digital TV channels, but also data channels for Internet access. In order to put the data channel in TV channels, a Cable TV modem or Set-Top-Box is needed. The CaTV network also has a hierarchical structure. In some place, a matching entity to the CaTV modem is needed in order to extract the data channel from the TV network. This entity is called CMTS, or Cable Modem Termination System. A CMTS interconnects the CaTV access network to a metropolitan network of optical fibre. Hence, a CMTS can aggregate thousands of data channels from subscribers over the CaTV network. Just like the Radio Network Controller and DSLAM, the CMTS is also responsible for authenticating subscribers for network access. In order to do this, it relies on keys and certificates on the cable modem (or set-to-box).

Access Networks aggregate traffic and authenticate subscribers. Core Networks provide the services to the subscribers, i.e. telephone service, SMS, E-Mail, Web, Media-Streams. Core Networks contain the servers to provide such services, i.e. engines to handle subscriber sessions, transfer or play-out data streams, and to store information. Figure 1 just symbolically shows one server for the local functions. Information made available at the core is of interest to all the subscriber in the network. The size of such servers is big and can consist of multiple shelfs with dozens of processor blades per shelf. Among the hardware platforms offered by Kontron in this place are the OM9140 and OM9060 (see Kontron web-page for details).

Regional or local information, such as regional TV channels, can be made available closer to the access networks. Also, servers may be operating in proxy mode and thus just contain and make available information which is most frequently requested. As small platforms, Kontron offers the OM9020 and OM5080. OM5080 represents the most compact carrier grade platform and still offers the versatility to adapt to a broad range of applications. Figure 1 just shows the principle network architecture. In practice, networks are further segmented into different domains of ownership, i.e. being operated and owned by different network operators. The size of a platform depends on the number of subscribers that one platform is intended to serve in a specific domain or a specific region.

Media Servers

While Media Servers in the core apply for information, which is of general interest to all subscribers within the network, small Media Servers apply to regional networks over Radio Access Networks, DSL or Cable TV. They feed in local information or operate as proxies to servers in the core (i.e. just store frequently demanded information and adapt this offer on demand).

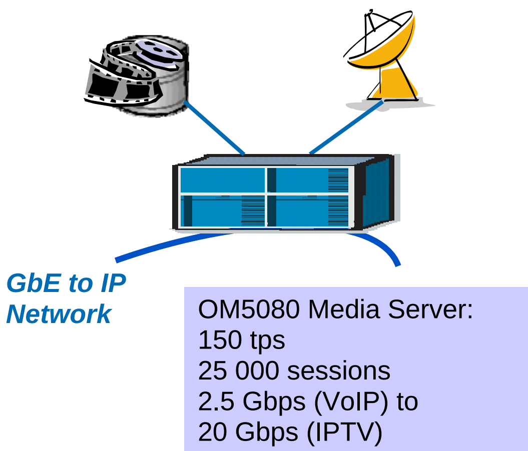

- Serves Regional Networks over RAN, DSL and CaTV
- Proxy/Web Server
- Media Processing: transcoding, content filter, content adaptation
- IPTV and streaming
- VoIP: SIP/RTP including Voice Quality Enhancement

Figure 2: Media Server Application

The demands on processing on a Media Server vary with the application. In the simplest case, the Media Server operates as a Web-Server and provides IPTV and streaming from local archives or from life streams. Frequently, the media processing: requires the trans-coding or media streams into a different format. Other applications are the filtering of content, or adaptation of the content stream according to the quality of the radio channel (in a Radio Access Network). Another case is in processing of VoIP data streams, for example for Voice Quality Enhancement.

A small Media Server may be benchmarked by numbers of transaction per second, number of simultaneous session, and the throughput or bit-stream. Such figures can be derived from the number of subscribers to be served by a traffic model (see [1] for some examples). Given the high performance of today's processors on a small form factor such as AMC, Media servers can be designed with high density in compact way. The OM5080 as 2U server allows accommodating 8 AMCs (with Multi-Core CPUs or DSPs). Despite the compact size, all functions are replicated and implemented on field replaceable units in order to maintain a continuous operation.

Signalling Gateway and Media Gateways

A Gateway interconnects different domains, i.e. domains of different network operators or domains of different technology (for example VoIP networks and traditional telephone networks). Communication networks typically differentiate between the control actions (such as setting up a call), and the user traffic (such as the voice connection as soon as the call is set up). This differentiation provides two different types of gateways: Signalling Gateways just handle control traffic (or signalling traffic) between network domains; Media Gateways handle the associated user traffic.

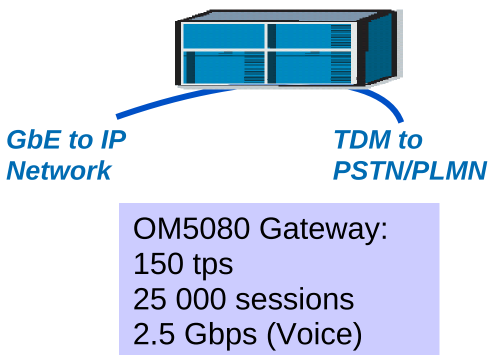

- Interconnects VoIP Networks with tradtional Telephone Networks
- Signalling Gateway: Control plane only (SIP to SS7 signalling)
- Media Gateway: Control Plane and Data Plane (additional IP and TDM voice traffic)

Figure 3: Gateway Application

A Signalling Gateway, which interconnects a VoIP Networks with a traditional telephone network needs two different types of interfaces: (1) for VoIP it needs to terminate the signalling over the SIP protocol (or other protocol used), (2) for the traditional telephone network it needs to terminate the so called SS7 signalling protocol. The Signalling Gateway needs to maintain the state of the transaction (e.g. ringing the phone of the called party), and translate between the protocols of each domain.

While IP networks (VoIP domain) interconnect over the Ethernet interfaces of the Gateway, the traditional network connects through dedicated interface and needs a corresponding interface card on AMC. The state engines can be implemented by processors (Multi-Core CPUs) on AMC. The type and number of processors depend on CPU performance and the desired performance of Signalling Gateway in terms of transactions per second (i.e. new sessions per second) and the number of simultaneous sessions. Those target figures can be derived from a traffic model. Again, each component of the Signalling Gateway needs to be replicated and field replaceable in order to maintain a continuous operation.

In order to call from a VoIP domain into a traditional network, it is not sufficient to handle the signalling traffic, which controls the transaction. As soon as the call is set-up, the data stream also needs to pass. The data stream comes in different formats in each domain: (1) an RTP-based packet stream in the VoIP domain, (2) a continuous bit-stream in the traditional telephone network. Also, the data streams follow specific kinds of coding in both domains. A Media Gateway converts the corresponding formats of the data stream in both domains. Depending on the type of application, the media processing may contain algorithms for Voice Quality Enhancements.

Trans-coding requires powerful Multi-Core CPUs and eventually DSPs, which are available on AMC form factor. The requirements again follow the target system performance in terns of transactions per second, number of parallel sessions, throughput, and implementation of specific codex. The system needs to be configured in a redundant way without single points of failure. This also requires the handover of the media stream, should one of the interfaces fail.

CMTS-Core Application

As the name indicates, a CMTS (Cable Modem Termination System) represents the opposite end of the Cable Modem at the subscriber. The cable modem essentially allows to place IP data channels into TV channels on the cable. The CMTS unpacks those channels and places them on the IP Core network. By definition, the CMTS is a type of gateway: it connects (1) video channels on the TV cable over special interface cards, (2) IP content over the regular Ethernet ports.

This nature of the interconnection has kept previous CMTS implementations outside the scope of COTS hardware platforms: the CaTV part with the corresponding radio frequency interfaces has dominated the design as a proprietary form factor. Also, it has rendered CMTS implementation rather inflexible to adapt to growing needs of providing more IP based data channels over CaTV. On order to better address flexibility, the standardization bodies came out with a modular approach to implement the CMTS: M-CMTS, see Figure 4.

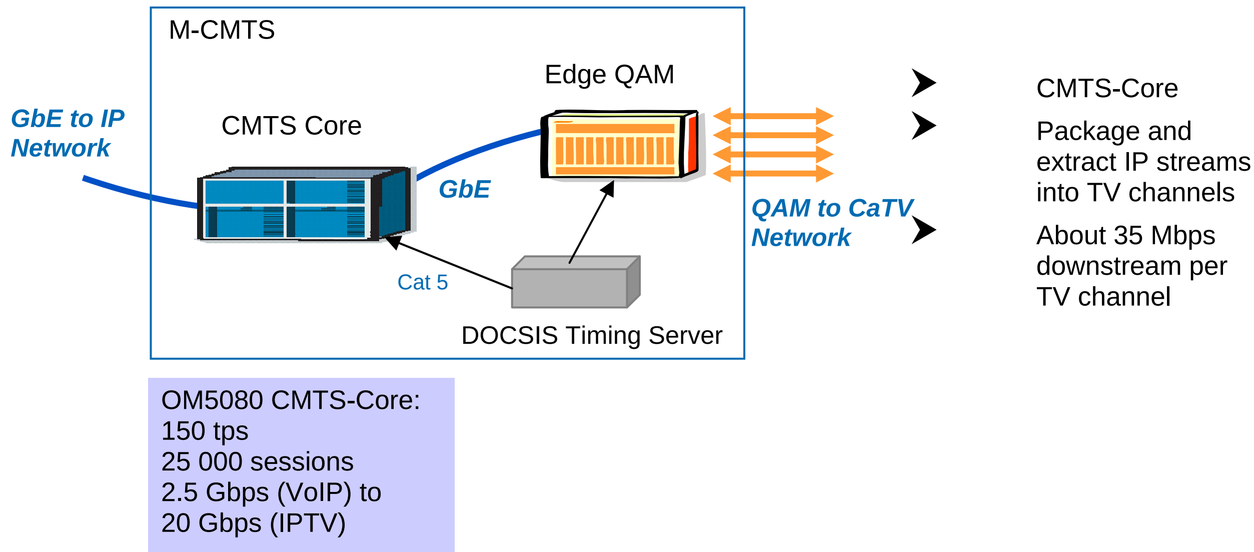

Figure 4: CMTS Application with CMTS-Core on OM5080

The point of M-CMTS is to reduce investment for provision of extra IP downlink channels and facilitate capacity upgrades. This is achieved by separating the CMTS-Core (Control, Authentication and IP Interfaces) from the radio frequency part (Edge-QAM, in phase 2 also to contain the upstream QAM). The separate RF part makes CMTS Core more interesting for ATCA and MicroTCA as carrier grade COTS solutions.

The CMTS-Core is in charge or packaging and extract IP streams (for VoIP, IPTV and Internet) into TV channels (at about 35 Mbps downstream per TV channel), to handle authentication and data sessions according to the respective standard Docsis 3.0 (or Euro DOCSIS 3.0 with different physical layer). The release 3 specification provides the bonding of multiple TV channels for data channels and thus significantly allows increasing bandwith. It also supports IPv6, which essentially provides a much bigger address space compared to IPv6, as well as IP multicast to support IPTV applications.

The OM5080 represents a nice fit for a compact, carrier grade CMTS-Core. In particular the 10 Gigabit version provides multiple 10 GbE uplinks and the high throughput between AMCs, which is essential for this kind of application.

RNC Application

In a mobile network such as GSM (2G) or UMTS (3G), subscribers are not attached to one specific line or connection, they are free to move. Mobility essentially covers two concepts (1) Roaming, i.e. subscribers are free to move and always reachable, as soon as the mobile phone is switched on, (2) Hand-over, i.e. staying connected during a session or call while moving from one radio cell to another. The management of mobility is an essential part of the Radio Network Controller (RNC). To do this, the RNC needs to access data bases where the location of a subscriber is stored, as well as access to the base station controllers to find out the exact location and to arrange allocation of radio resources.

Another function handled by the RNC is authentication of mobile subscribers in order to allow them to access the network. To perform authentication, the RNC needs access to the information stored on the SIM card of the subscriber, as well as to the matching part if the SIM card in the network (on a so called AAA-Server, where AAA stands for Authentication, Authorization and Accounting). Access to other network elements is handled by special network protocols. Figure 5 shows a summary of the network architecture.

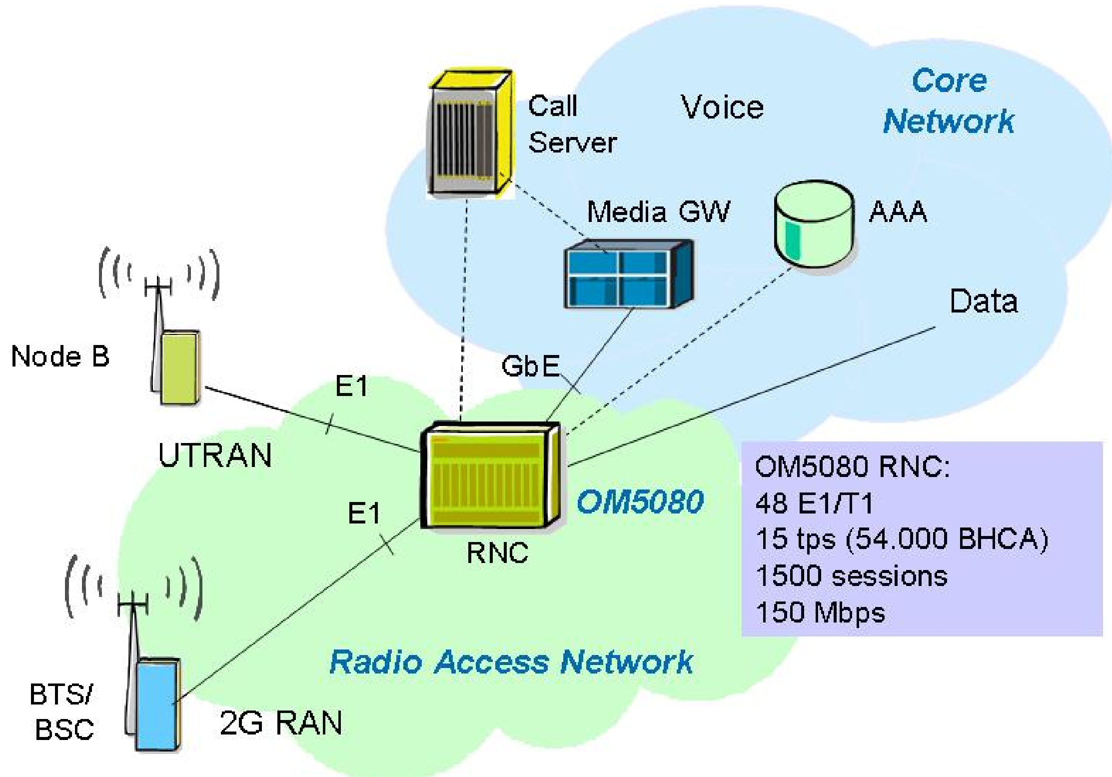

Figure 5: Radio Network Controller Application

The type of protocols varies with the representative technology. Protocols of 2G (GSM) differ from 3G (UMTS) protocols. Other technologies, such as WiMAX, use the same concepts (the RNC is part of the WiMAC Access Controller), but again different protocols for the same kind of functions. Also the type if interfaces are different. While all connections to components in the core network use IP based protocols and thus are handled over Ethernet ports on the RNC, the base stations connect over different interfaces: 2G base stations (or base station controllers) connect over E1 lines and 2G protocols, 3G base stations (Node B) connect over E1 or ATM links with higher bandwidth and 3G protocols, other implementations such as WiMAX use Ethernet connections with IP specific protocols.

Much like a Gateway, the Radio Network Controller needs to keep track of sessions including mobility management and authentication through the diverse protocols and implementations. It can connect 2G and 3G base stations. If traditional E1 or other implementations of TDM links are used, it needs corresponding interface cards or line cards on AMC. Other components are Multi-Core CPUs on AMC as main controllers and traffic controllers. Depending on the amount of trans-coding (e.g. from 2G to 3G), DSPs on AMC can be used. For handling of IP tunnels, Network Address translation and Core Network protocols, the usage of a packet processor on a dedicated Network Service Processor can be recommended.

An RNC implementation can be measured in terms of physical connections, transactions per second (including location updates), number of simultaneous sessions and traffic throughput. Those figures can be derived from a traffic model and assumptions on the number of subscribers the RNC covers. The operation of an RNC needs to be continuous, i.e. without single point of failure and with every component to be field replacable.

Professional Mobile Radio (Tetra/P25)

Professional Mobile Radio (or Private Mobile Radio) represents mobile communication networks for special applications for the police, fire brigades, rescue and emergency medical services. Other areas of application are airports, military, oil & gas and transportation, wherever public safety is of concern. Such networks need to operate independently of Public Communication Networks. They provide a slightly different set of features, such as point-to-multipoint communication, push-to-talk, closed user groups and coverage of large areas. Popular standards to implement such networks are Tetra (in Europe) and P25 (in North America). Figure 6 shows a summary of a Tetra- or PMR-network.

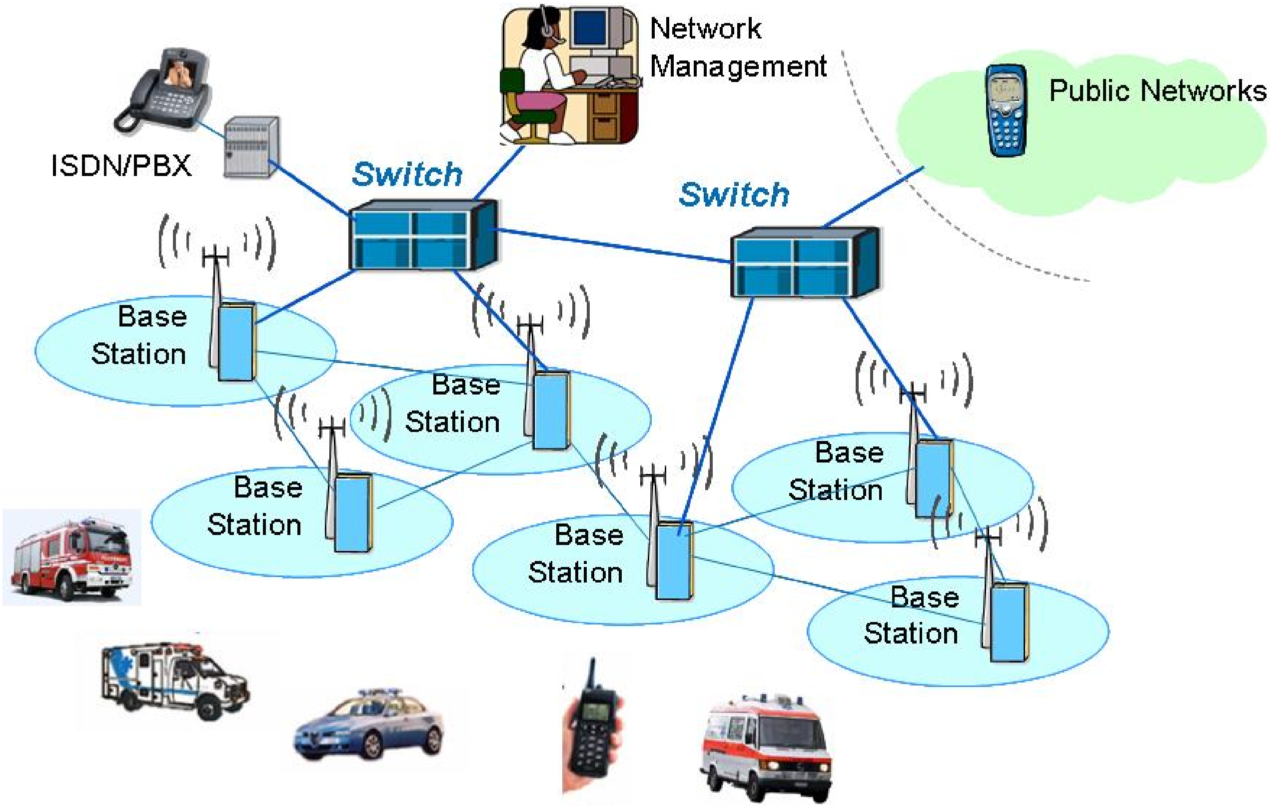

Figure 6: Tetra Switch Application

The PMR network represents a two-tier architecture: on tier 1, Base Stations interconnect subscribers and also interconnect to other base stations, tier 2 Switches control and interconnect base stations and other Switches. PMR Switches also interconnect to PBXs and to other networks, such as public mobile networks or public fixed line networks. PMR networks support a high degree of resilience against failures. Should the connection between a Base Station and the Switch break, the base station may fall back to operating in autonomous mode and perform essential switching functions.

The Switches handle the set-up of calls within the network including interconnection to external networks. Depending on the implementation, Switches are using IP type of interconnections. They also allow interconnection to conventional or PBXs. As autonomous network, the PMR also needs an administrational infrastructure for management of subscribers incl. user groups and access rights, as well as the administration of network resources for operation. Such functions are allocated in a Network Management-Centre, which also interconnects to the Switches over IP.

In comparison to public networks, the number of subscribers and density of subscribers per area is comparatively low. Thus PMR switches may be implemented on small systems. Nevertheless the operation of the network is highly critical and needs fault resilience in the switch, i.e. a carrier grade platform without single point of failure. The OM5080 represents an excellent choice to implement PMR switches. It can be configured in the most flexible way with interface cards to connect Base Stations, PBXs or external networks, and a choice of Multi-Core processors on AMC. The build-in uplink capacity provides connectivity over Ethernet for IP type of external systems. The architecture of the OM5080 allows seamless failover between any redundant components in operation.

OM5080 Carrier Grade System

The OM5080 consists of 2U ATCA chassis with AT840xM ATCA Carrier boards, which host the AMCs. Thus, the system is closest to the origin of the MicroTCA standard, which essentially translates the ATCA carrier concepts to an autonomous system with AMCs on a backplane. On order to provide the extra functionality beyond the ATCA carrier, the OM5080 contains a portable software extension for AT840x for Shelf Management Functions (e.g. advanced fan control). System management supports IPMI 1.5, SNMP V1, V2, V3, as well as HPI and is accessible inband over the Ethernet uplinks, as well as locally at the serial connectors on the carrier blades.

Figure 7: OM5080 Carrier Grade Platform

The OM5080 adds some features, which are part of ATCA, but missing in the MicroTCA standard: (1) Shelf Alarm Panel on Management Module, (2) Uplink Modules (1G and/or 10G). In MicroTCA, the MCH is intended to provide the uplink capacity. However, due to the AMC form factor, front space is limited for high uplink capacity. The use of Uplink Modules eliminates this bottleneck: The OM5080 provides 8x GbE uplinks (or alternatively 2x 10 GbE plus 4x GbE). In addition, the OM5080 provides space for two ATCA RTM for I/O extensions (optional) and dual power feeds and power distribution. Figure 9 shows the system architecture.

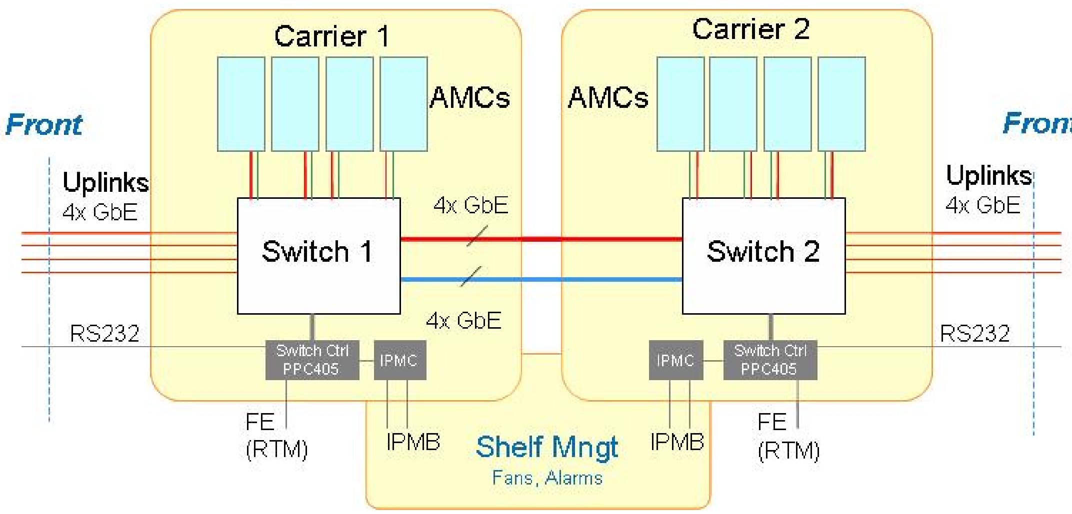

Figure 8: OM5080 Architecture – 1 GbE Version

From a conceptual point of view, the ATCA carrier represents an MCH plus PM. As network capacity, the 1GbE version of the OM5080 provides one uplink per AMC, which can be configured in transparent mode. The system represents a fully redundant configuration with 8x GbE interlink capacity between carriers. The carriers of the 1 GbE version also provide extra fabric capacity: 4 lanes of PCIe are available between AMCs and between Carriers (via update channel), there is SAS/SATA between AMCs and between Carriers, as well as Telco clocks. At the rear, the OM5080 provides two RTM slots which can be used with an RTM providing an external HDD and external SAS connection. The OM5080 provides power management and fan control as in ATCA systems. It also provides connectivity to telecom alarms, as well as connectivity to external management over SNMP.

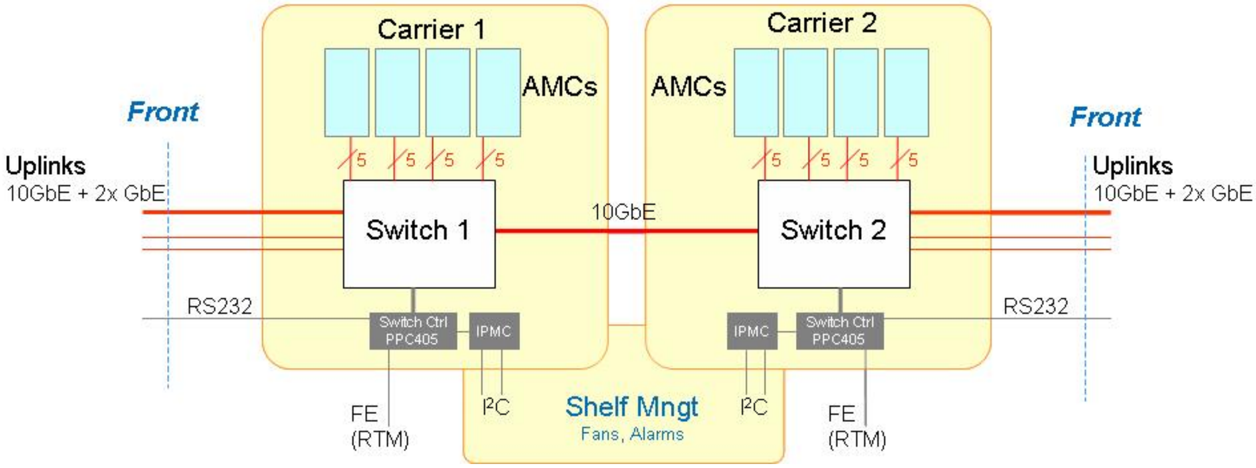

Figure 9: OM5080 Architecture – 10 GbE Version

The 10 GbE version of the OM5080 follows the same approach, but is based on the AT8402 ATCA carrier with 10 GbE fabric. Conceptually, the carrier again represents the MCH combined with the Power Module. The OM5080-10G also provides one uplink per AMC (can be configured in transparent mode) with fully redundant configuration with 10GbE interlink capacity between carriers. Each AMCs connects over 5x GbE ports to the switch. As additional fabric, the OM5080-10G provides fat pipe interconnects between paired slots (PCIe, sRIO), one 2.5Gbs APS link per AMCs between the Carriers (via update channel), SAS/SATA between AMCs and between Carriers via RTM as well as Telco clocks. At the rear, 2 RTM slots can be used for extra capacity.

The OM5080 applies for a variety of carrier grade applications, such as media servers for audio and video, signalling gateway in telecommunication networks and at the boundary of IP-based networks and traditional TDM based mobile and wireline networks (Time Division Multiplexed, rather than packet oriented). Other applications include radio network controllers, or Tetra/P25 switches. Due to its unique combination of high capacity for interfaces on a small form factor, the OM5080 is extremely well suited for such applications.

At the TDM part, E1/T1 links or STM-1/OCR-3 may be connected over I/O AMCs, at the IP part, the OM5080 provides high IP throughput over the Uplink Modules (8x GbE or 2x 10 GbE plus 4x GbE). For media conversion, DSPs may be included on AMC. As system controller for session set-up and tear-down, there is a choice of high-performance multi-core processors on AMC. Still, every component is replicated, i.e. each module may fail-over to a replicated module and traffic switched accordingly. Failing parts may be replaced on the running system (hop-swap capability). Figure 10 shows a sample configuration of the OM5080.

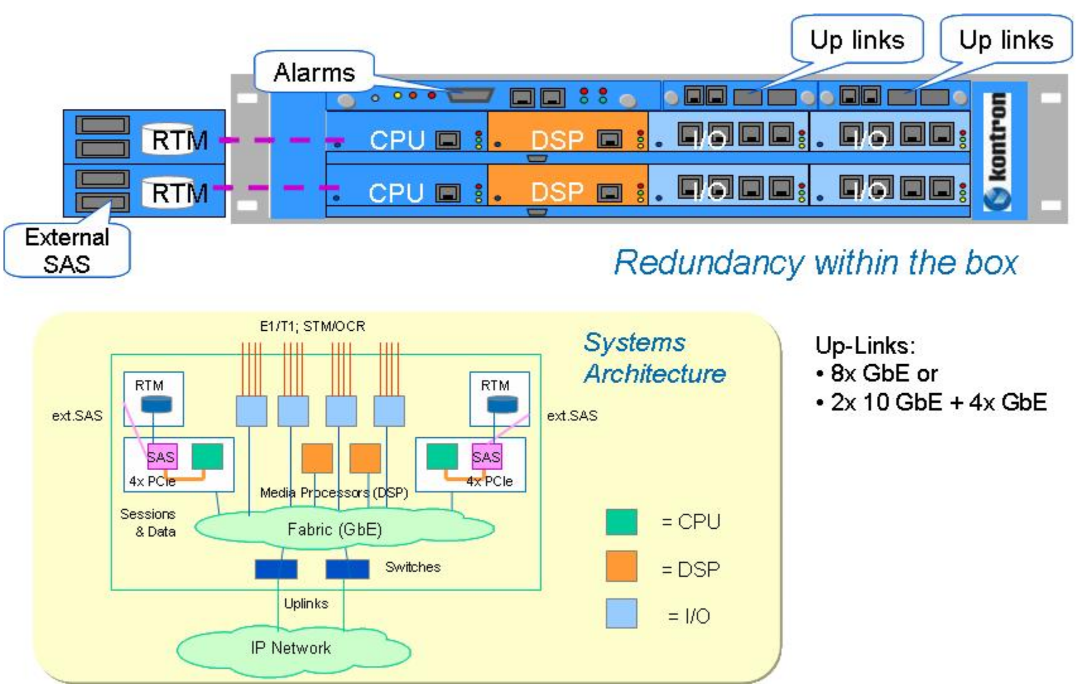

Figure 10: OM5080 Sample Configuration - Gateway, RNC and Tetra Switch

In the upper part, a schematic of the OM5080 is shown. Each of the AMC carriers contain one system controller (CPU on AMC), one media processor (DSP on AMC), and 2 I/O cards (e.g. each with 8x E1/T1 interfaces). With a total of 16 E1 ports connecting to Tetra/P25 stations, the system can handle 480 concurrent voice streams. The Uplink Modules provide IP-connectivity at the IP-domain of the network, e.g. towards a SIP based switch controller and media gateway. The systems architecture is shown in the diagram in the lower part of Figure 11. In this case, all modules connect over Ethernet, i.e. the I/O cards represent intelligent I/O (e.g. over iTDM). In this case the DSPs convert media streams between TDM and IP based protocols (such as iTDM and RTP). The CPUs handle the session set-up and tear down, and control the system at application level.

For high-availability, the OM5080 contains a pre-configured network with redundancy and fail-over at Layer 2, i.e. Ethernet with pre-configured VLANs and MSTP to repair broken links. At Network Layer and Application Layer, there is a choice of mechanisms to provide master and slave configurations with synchronisation at process level and failover. Among the building blocks are open source tools in the environment of the Linux ecosystem, as well as commercial tools with comfortable configuration tools and management support.

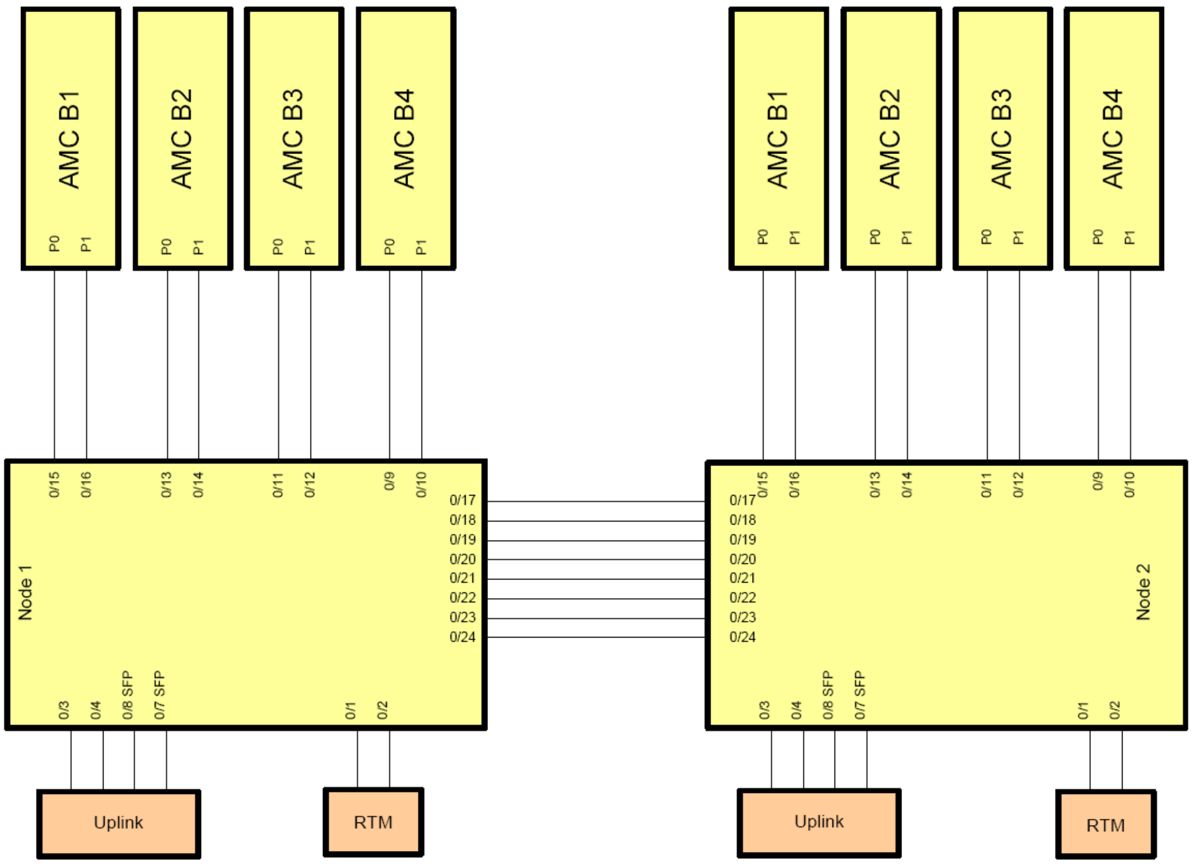

Figure 11: OM5080 Ethernet Configuration

The Ethernet configuration of the OM5080 is shown in Figure 11. The basic OM5080 contains two processor-AMCs (AM4010) plugged into Slot B2 in each of the carriers. The processor AMCs are ready-to-run with Linux OS installed and network ports including IP addresses pre-configured. In the default network configuration, the uplink-modules operate in separate VLANs, so they can be used in parallel. The inter-links between the carriers operate in another VLAN. This way MSTP can be used to automatically bypass broken links. The PCIe slots on each carrier are also pre-configured with Slot B2 (processor AMC) as upstream ports, all other ports in downstream. Over the update channel, PCIe can also extend to the PCIe switch of the second carrier blade.

References:

[1] Traffic Models (Spreadheet) for DSL, 2G/3G and CaTV Networks

[2] Media Server Design (White Paper), S. Rupp

[2] Open Modular Systems: Ready to run for customized Projects (White Paper), S. Rupp, C. Bestler