



MLPPP in the Evolving Radio Access Network

by *Networking and Multimedia Group*
Freescale Semiconductor, Inc.
East Kilbride, Scotland

This white paper looks at the usage of multilink point-to-point protocol (MLPPP) within the evolving radio access network (RAN). It looks at how quality of service (QoS) can be implemented in MLPPP-based network architectures to support the new range of services on offer to users.

Contents

1. What is MLPPP?	2
2. Why MLPPP in the RAN	2
3. Efficient Packet Transfer in the Base Station	3
4. Meeting QoS Requirements	4
5. Summary	6

1 What is MLPPP?

MLPPP can combine multiple links from different services onto one physical pipe, which can be made up of one or multiple MLPPP bundles. It is possible to add or remove links as required from a bundle providing bandwidth-on-demand capabilities. Different data types are transferred over all of the MLPPP links simultaneously, decreasing latency for packet transfer.

Without MLPPP, sending varying traffic types over T1/E1 can be done by using an individual low speed link for each traffic type. However, this can limit the performance of certain applications, and in many scenarios it does not make efficient use of the available bandwidth.

Link control protocol (LCP) is used to negotiate MLPPP between peers, with each system indicating that it supports multilink in the LCP configuration request. Once established, a unique value is assigned to the bundle between peers. This bundle can consist of multiple physical links of the same or varying types.

Figure 1 shows the packet format for MLPPP. The address and control field contains an 8-bit all-station address (default 0xFF) and an 8-bit control field (default 0x03). Multilink PPP fragments use a protocol ID of 0x003D, which is followed by a 4 byte MLPPP header. The MLPPP header contains beginning and end fields that signify whether the payload is part of a fragmented packet and a sequence number (SN) field for tracking the order of packets during transfer. The first four fields shown in Figure 1 are then protected using a frame check sequence (FCS).




Figure 1. MLPPP Packet Format

MLPPP can also be used with multiclass extensions that provide the ability to prioritize packets. When combined with fragmentation, this can ensure that large, low-priority packets do not delay higher priority packets. In this case, each class has its own incremental sequence number count to ensure that the order is maintained per class.

2 Why MLPPP in the RAN

3GPP release 5 saw a shift from ATM to IP as the bearer in the RAN. Ever since, a variety of backhaul implementations have appeared to meet the bandwidth and cost requirements of operators. Reusing existing infrastructure allows operators a path to IP without a complete overhaul of the access network.

In channelized TDM links, such as those used in 2G networks between the BTS and BSC, periods of low traffic mean the unused bandwidth on the link cannot be utilized by other users on the network. This connection is also limited in high density traffic, as it cannot be easily scaled to add more bandwidth. MLPPP allows multiple services to share the same bundled links, and the number of links in the bundle can be increased if more bandwidth is required. This is essential for handling the increasing numbers of voice users alongside advanced data services.

This bundling of T1/E1 links is not a new concept and was utilized in 3GPP release 99 networks by using inverse multiplexing over ATM (IMA) to transport ATM traffic over low speed links. Due to the increasing number of IP applications running across the RAN, the move to MLPPP from IMA reduces the overhead

involved in getting packets across the links, thus making use of the available bandwidth more efficiently. It should be noted that support for ATM is still a requirement for legacy systems and will be for a considerable time to come.

LTE standards will eventually mean Ethernet is the standard interface choice for the backhaul, but MLPPP provides an efficient and scalable solution in the interim for low density cells. However, for metro cell sites, the move to Ethernet is a requirement now.

3 Efficient Packet Transfer in the Base Station

With collocation a hot topic for operators, equipment manufacturers must support a number of legacy connections and provide interworking between the various protocols in the system. For example, [Figure 2](#) shows the network architecture for a base transceiver station (BTS) collocated with a release 99 NodeB. In this scenario, the TDM traffic from the BTS is interworked to the NodeB MLPPP interface in order to reach the core network. The NodeB must also handle interworking MLPPP traffic to the backplane (e.g. Gigabit Ethernet).

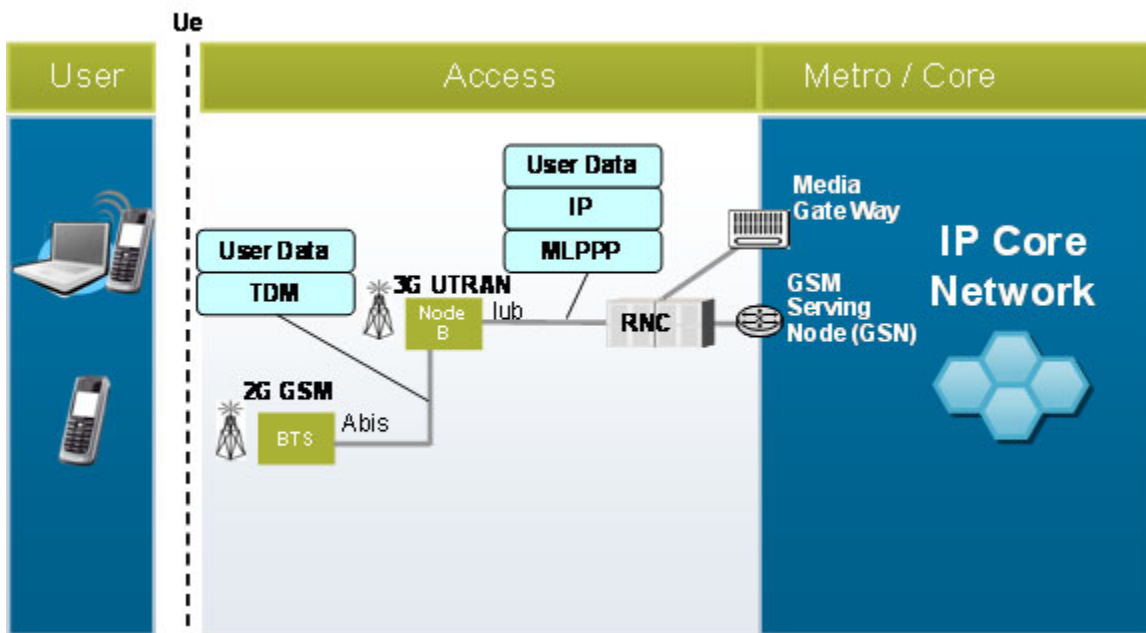


Figure 2. 2G BTS Collocated with 3G Release 99 NodeB

The base station interworking requirements can be efficiently handled by a communications processor, such as the MPC8569E from Freescale Semiconductor (see [Figure 3](#)). The system-on-chip (SoC) methodology coupled with 45 nm SOI results in a high level of integration. This allows board manufacturers to design complex systems with small foot prints, low power, and high performance.

The MPC8569E is compatible with the MLPPP standard (RFC1990) and the multiclass extensions (RFC2686). The QUICC Engine communications block on the device is capable of classifying traffic and automatically interworking to or from Ethernet interfaces based on multiple fields within the packet headers. This device also has a feature known as weighted fair queuing (WFQ) for class 0 that allows packets to be prioritized using WFQ on MLPPP transmit even if the system does not enable multiclass

extensions. Multiclass cannot always be used if the base station must operate in the network with peers that do not have this feature. WFQ for class 0 gives equipment manufacturers the benefit of WFQ without causing potential incompatibilities with peers.

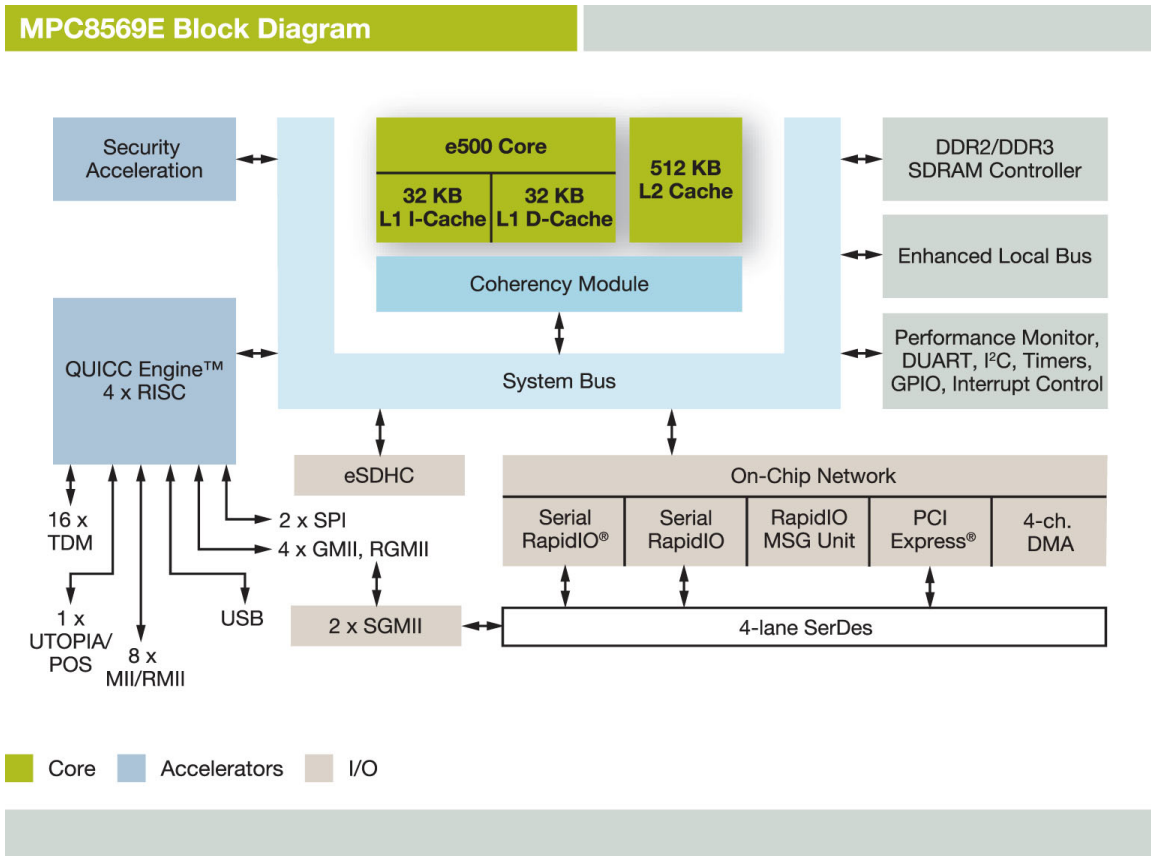


Figure 3. Freescale’s MPC8569 Communications Processor

Packet transfer can also be improved on the TDM links by ensuring that the number of idles/flags between adjacent frames is kept to a minimum. The MPC8569E provides a Tx bandwidth optimization feature that can significantly increase the utilization of the TDM link.

4 Meeting QoS Requirements

Due to the many services in use across mobile networks, traffic must be analyzed to meet the QoS requirements of a particular service or user. It is important to parse, classify, and prioritize traffic to meet QoE expectations because this can be a differentiating factor for operators.

Prior to release 5, backhaul networks used the in-built QoS features of ATM for traffic shaping or scheduling. As IP is a layer 3 protocol, on average frame lengths are much larger than for ATM. MLPPP fragmentation can be used to prevent short latency sensitive packets being delayed by larger lower priority packets. This technique splits a packet into fragments so that the higher priority traffic can be handled in between the lower priority fragments, which helps meet latency requirements.

Fragmentation and prioritization of packets may not be enough to meet QoS requirements in all scenarios. Various other QoS techniques, such as IP header compression (IPHC), can be applied to IP packets in the RAN to optimize bandwidth usage. Network control protocols (NCPs) are used to establish IPHC over MLPPP links. Freescale's MPC8569E processor provides IPv4/UDP and IP-only header compression options that can be handled automatically as traffic is interworked between Ethernet and MLPPP. Figure 4 shows an uncompressed IPv4/UDP header of 28 bytes, and Figure 5 shows a compressed header of 7 bytes when using an 8-bit context identifier (CID). If a 16-bit CID is required, the compressed header will be 1 byte longer. If the UDP checksum in the original header is equal to 0, the compressed packet can be as little as 4 bytes, assuming CID is 8-bit and the D bit equals 0. With such significant reductions in header size, it is clear why IPHC is used to utilize the available bandwidth more efficiently.

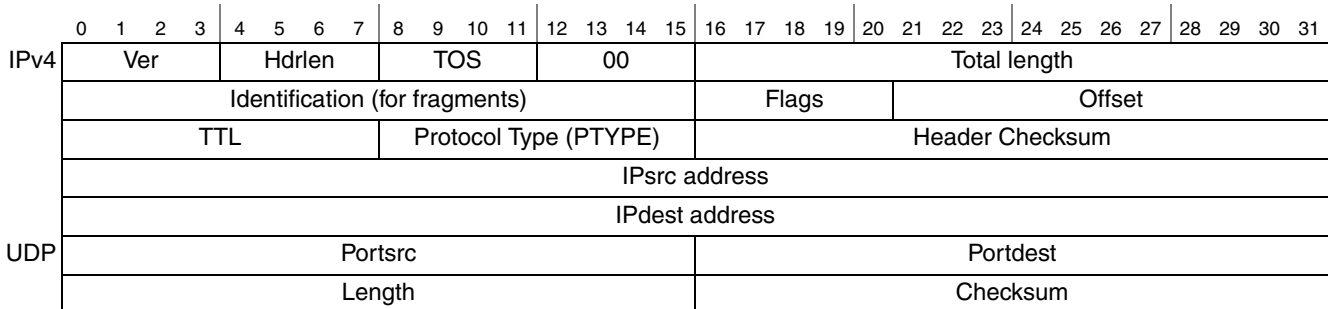


Figure 4. IPv4/UDP Header




Figure 5. IPv4/UDP Compressed Header with 8-bit CID

MLPPP also provides a configuration option to take advantage of header compression in the address and control, and the protocol ID fields of the MLPPP packet. Both these 2 byte fields can be reduced to 1 byte if agreed during negotiation.

Increasing use of mobile phones for web access is an important factor in the evolving network but voice must not be allowed to drop in quality because of changes in the network or its interaction with data services. With this in mind one of the issues that MLPPP solutions must address is the storing of fragments within the base station in order to meet the strict sequence number ordering. Even if packets from the peer are transmitted in order they can arrive at the base station out of order due to differential delay across the links and differences in packet sizes. This is okay until packet loss comes into play and can result in out of order packets waiting in the base station for too long to meet real-time latency requirements. The M algorithm for detecting packet loss over MLPPP links will only process out of order fragments when the minimum sequence number across all links is greater than the missing sequence number. In a base station with 8 T1/E1 links and low density traffic the time to detect fragment loss by the M algorithm could be significant enough to cause disturbance on a call. To limit the influence that packet loss could have on the latency of real-time packets the MPC8569E provides a mechanism for the base station to flush out of sequence packets immediately when a system defined time has expired. This time must be small enough to meet latency needs of real-time packets, but large enough to allow all packets the opportunity to arrive in a realistic time which means differential delay and the transmission time difference between large and small packets must be taken into account.

5 Summary

MLPPP is being used today to transport 3G (and 2G if colocation is used) traffic over T1/E1 lines. Using a combination of the features discussed, it is possible to ensure that QoS requirements for voice/video/data traffic are achieved. Operators who choose to migrate these base stations from TDM to native Ethernet will have an easier task because they already have an IP-based system in place.

6 Revision History

Table 1 provides a revision history for this white paper.

Table 1. Document Revision History

Rev. Number	Date	Substantive Change(s)
0	09/2010	Initial public release

How to Reach Us:

Home Page:

www.freescale.com

Web Support:

<http://www.freescale.com/support>

USA/Europe or Locations Not Listed:

Freescale Semiconductor, Inc.
Technical Information Center, EL516
2100 East Elliot Road
Tempe, Arizona 85284
1-800-521-6274 or
+1-480-768-2130
www.freescale.com/support

Europe, Middle East, and Africa:

Freescale Halbleiter Deutschland GmbH
Technical Information Center
Schatzbogen 7
81829 Muenchen, Germany
+44 1296 380 456 (English)
+46 8 52200080 (English)
+49 89 92103 559 (German)
+33 1 69 35 48 48 (French)
www.freescale.com/support

Japan:

Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku
Tokyo 153-0064
Japan
0120 191014 or
+81 3 5437 9125
support.japan@freescale.com

Asia/Pacific:

Freescale Semiconductor China Ltd.
Exchange Building 23F
No. 118 Jianguo Road
Chaoyang District
Beijing 100022
China
+86 10 5879 8000
support.asia@freescale.com

For Literature Requests Only:

Freescale Semiconductor
Literature Distribution Center
1-800 441-2447 or
+1-303-675-2140
Fax: +1-303-675-2150
LDCForFreescaleSemiconductor@hibbertgroup.com

Information in this document is provided solely to enable system and software implementers to use Freescale Semiconductor products. There are no express or implied copyright licenses granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document.

Freescale Semiconductor reserves the right to make changes without further notice to any products herein. Freescale Semiconductor makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Freescale Semiconductor assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters which may be provided in Freescale Semiconductor data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Freescale Semiconductor does not convey any license under its patent rights nor the rights of others. Freescale Semiconductor products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Freescale Semiconductor product could create a situation where personal injury or death may occur. Should Buyer purchase or use Freescale Semiconductor products for any such unintended or unauthorized application, Buyer shall indemnify and hold Freescale Semiconductor and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Freescale Semiconductor was negligent regarding the design or manufacture of the part.

Freescale, the Freescale logo, and PowerQUICC are trademarks of Freescale Semiconductor, Inc. Reg. U.S. Pat. & Tm. Off. QUICC Engine iis a trademark of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. The Power Architecture and Power.org word marks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.
© 2010 Freescale Semiconductor, Inc.

