White Paper

How to Reduce Cellular Backhaul Transport Costs While Improving 2G and 3G Network Operating Efficiencies

Contents

Introduction

Backhaul is one of the major contributors to the high costs of building out and running a mobile network. This is true whether the mobile operator owns the transport links or leases lines from a local carrier. It also applies if the carrier is deploying a dedicated transport segment for 3G or planning to integrate several generations into the same platform.

Industry consensus indicates that transport equipment accounts for 25% of the costs of building out a private cellular backhaul infrastructure whereas, in the case of leasing services, transport outlays vary between 40-60% of the total cost of renting lines, with backhaul contributing 75% of this. This is compounded by the fact that the backhaul access network spans the entire coverage area. Considering the huge impact that backhaul has on operating expenses (Opex) and capital expenses (Capex), mobile operators would be well advised to carefully review their cellular backhaul strategies before making further network infrastructure investments. While prudence is always good counsel, it is particularly worth heeding during this era of rapid transition.

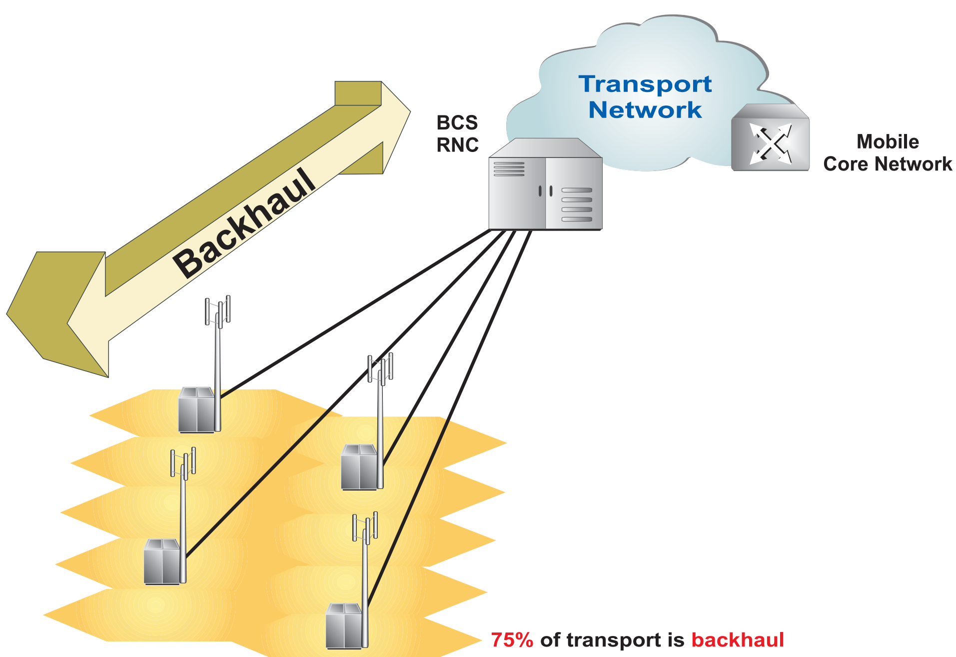

The Challenge – Efficient Backhaul

Complicating the equation, network operators must now be able to support simultaneously the divergent technological demands and applications of existing 2G/2.5G networks as well as those of newly emerging 3G operations. This migration from 2G circuit-switched networks (TDM) to 3G packet switched networks (ATM and eventually Gigabit Ethernet, IP, and MPLS) raises new challenges. In particular, the mobile network operator must weigh the cost, suitability and availability of the access platform chosen to handle the expected increase in bandwidth capacity and be able to manage the complexities of a converged voice and data network.

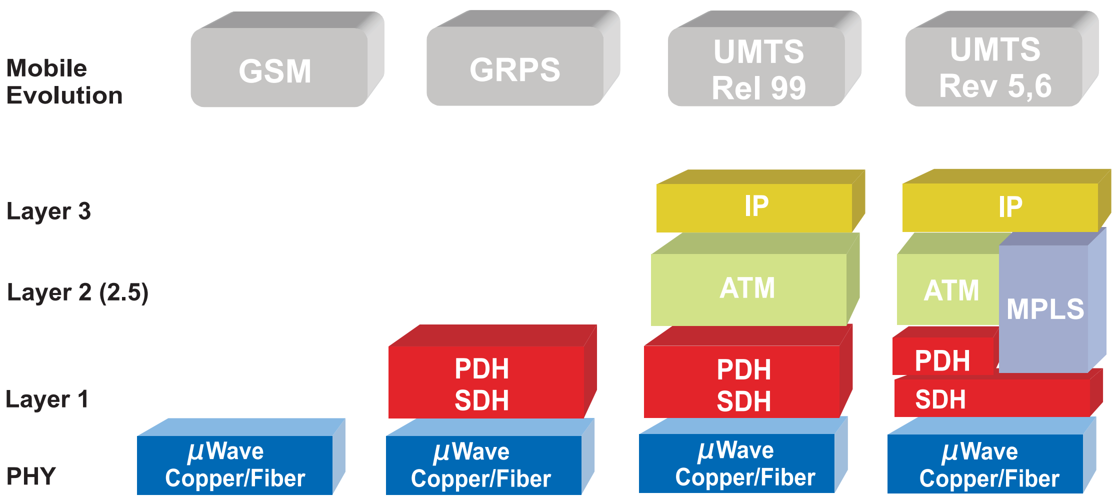

Technologies in the Mobile Transport Network

Efficient Backhaul Design

Before examining the requirements for efficient backhaul, it's necessary to identify the elements that make up the cellular backhaul network. Backhaul is typically defined as transmitting voice and data traffic from the radio cell site to an edge switch at the ingress/egress point of the backbone core network. Generally, this access network can be divided into two segments: from the Node B (in 3G UMTS networks) or Base Transceiver Station (in 2G GSM/TDMA networks) to an aggregation point; and from the aggregation point to the Radio Network Controller (3G) or Base Station Controller (2G). There is a third segment, between Mobile Switching Centers, which is not part of the access network but, rather, the core network, where the fundamentals of backhaul still apply.

Aggregation is an essential part of the network transport design because it allows for more efficient use of the transport bandwidth and simplifies network management. Generally, aggregation occurs where there is dense concentration of traffic from multiple radio cell sites, typically at the controller sites and the MSC. As 3G gathers steam, however, and new broadband wireless services increase, the quantity of backhaul links will rise exponentially, making upstream aggregation at the front end of the transport network very attractive.

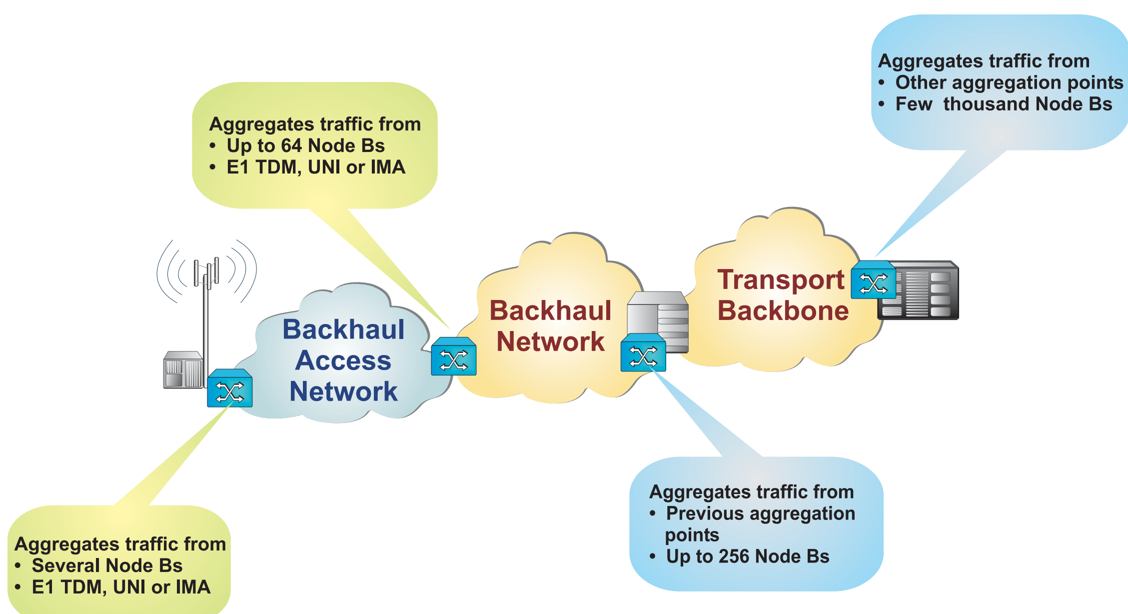

Aggregation Nodes in the Transport Network

Thus, there are sound business reasons for deploying aggregation devices to serve medium-sized or even smaller coverage areas. Locating a switch at a remote point-of-presence (POP) offers many economies, especially where link density is high and the price-per-port of the switching equipment justifies the investment in Capex.

Traffic aggregation is applicable in two instances:

1. Homogeneous traffic stream – Combining TDM links with Abis compression from several BTSs over a thicker pipe enables full use of the link (fanout) or, in the case of ATM, taking advantage of the inherent statistical multiplexing mechanism of the protocol to reduce bandwidth.

2. Heterogeneous traffic stream – Supporting different types of traffic from different mobile generations over the same transport link leverages the existing network infrastructure.

Any Access Network over Any Transport Network

Mobile operators are saddled with a bewildering choice of backhaul technologies and network interfaces as they try to anticipate which access infrastructure will best serve their current and future requirements. The easiest decision would be to build out parallel networks. This is not as straightforward as it would appear since the network equipment access technology, which is an embedded feature provided by the cellular vendor (as defined by industry standards), has to be considered as well as the available backhaul technology and related services. Using a dedicated transport network for successive mobile generations – PDH/SDH for 2G/2.5G; ATM for UMTS Release 99; and Ethernet/IP/MPLS for UMTS revision 5/6 – is not as efficient or potentially cost-effective as integrating diverse traffic streams over a single backhaul link. Having acknowledged that, neither transmitting TDM frames nor ATM cells over an Ethernet/IP/MPLS network, nor running Ethernet/IP packets over TDM in a cost-effective manner are trivial matters.

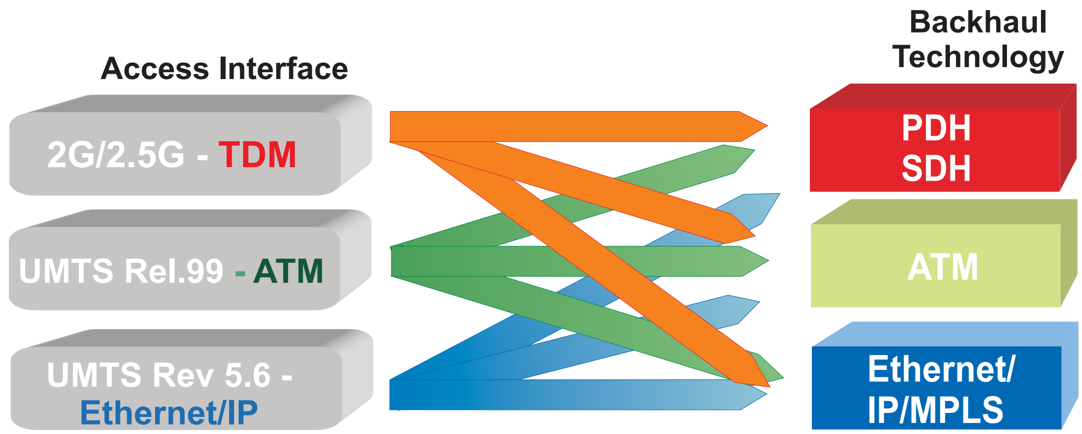

Any Access over Any Transport Infrastructure

Given the drawbacks stated above, does a converged backhaul access network solution exist that is technologically feasible, economically sound and readily available? To find out, we'll examine each backhaul technology and measure its suitability based on technological and commercial criteria.

Backhauling over PDH/SDH Transport

TDM-based networks – PDH and SDH/SONET – have long served as standard transport platforms for cellular traffic. PDH and SDH/SONET are optimized to handle bulk voice circuits with maximum uptime, minimal delay and guaranteed service continuity. Leveraging existing TDM-based networks for cellular backhaul affords a large degree of investment protection for operators of private networks. In addition to carrying legacy TDM streams, the TDM-based network is also suitable for ATM traffic. Thus the same TDM transport infrastructure may be used for migration to UMTS.

There are many advantages to settling on a TDM based infrastructure. The technology is time-tested, simple and deterministic – at any given time the operator knows what traffic is going over the links. SDH/SONET, in particular, offers incomparable reliability. TDM networks are ubiquitous, reaching even outlying rural areas. There the links tend to be under-populated, affording economical bandwidth sharing through aggregation.

Despite numerous benefits, TDM does have drawbacks. UMTS uses ATM technology, which is statistical in nature and has built-in service differentiation mechanisms. TDM is less flexible in its support of different types of traffic. Moreover, each circuit requires a dedicated timeslot, which in the long run could add to operating costs in comparison with ATM's statistical traffic handling capability. Furthermore, although the existing PDH or SDH/SONET equipment may easily support Node B's ATM interfaces, this is not the case for IP/Ethernet. Integrating new radio network equipment may require considerable investments.

These reservations aside, TDM-based networks can deliver significant cost reductions providing that the multiple traffic streams are aggregated wisely. This following illustration shows the aggregated traffic carried as VC-12 containers over the SDH transmission network. A TDM multiplexer merges several E1 trunks, bearing either ATM or TDM traffic, over a thicker pipe. There are break-even points where moving from separate E1 lines to a single STM-1 link becomes cost-effective. In these cases, the mobile service operator will shift from a PDH to an SDH transport network, either privately owned or leased.

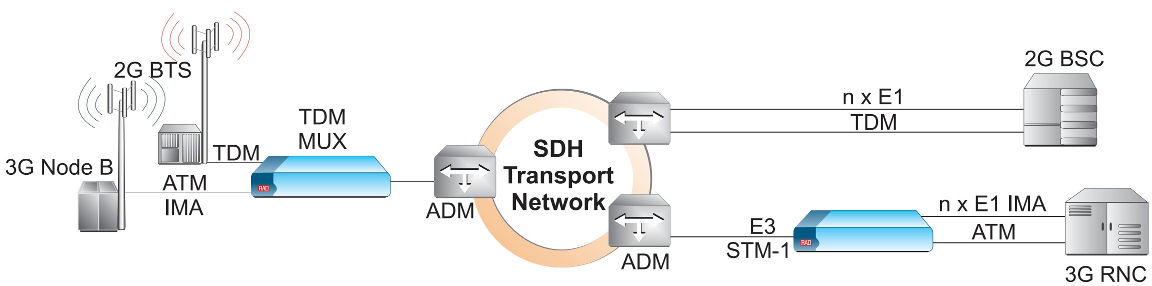

TDM aggregation for SDH Networks

The business case for using SDH instead of PDH is clear-cut. Multiplexing requires fewer switching ports and less expensive equipment. (As an example, an STM-1 port costs less than a 63 x E1 port card on a switch.) Also using higher speed modules allows for higher port density in the same device, further reducing Capex.

Aggregation of traffic invariably leads to lower operating expenses. This holds true even for outlying areas, accompanying lower density traffic flows.

Nevertheless, this approach has one main drawback and that is the requirement for E1 ports in the Radio Network Controller (RNC). This is because the inverse multiplexing over ATM traffic (IMA) is carried over the transport network in VC-12 containers and not in an STM-1 universal network interface (UNI) format as required by the RNC.

This limitation can be overcome by channelizing the ATM traffic over the SDH network so groups of ATM with IMA protocol are transported in VC-12 containers. The aggregated traffic (from multiple 2G and 3G radio sites) is transmitted in VC-12 containers across the SDH network to an ATM multiplexer, which terminates the channelized (n x VC-12) IMA cell streams and converts them to STM-1 ATM UNI (VC-4) for RNC connectivity. Basically, to benefit from using the same existing SDH infrastructure for both traffic types and to avoid having multiple E1 ports in the RNC, a dedicated device or module is required to make the conversion. What is gained by adding a dedicated device? This saves additional ports on the RNC or SDH ADM.

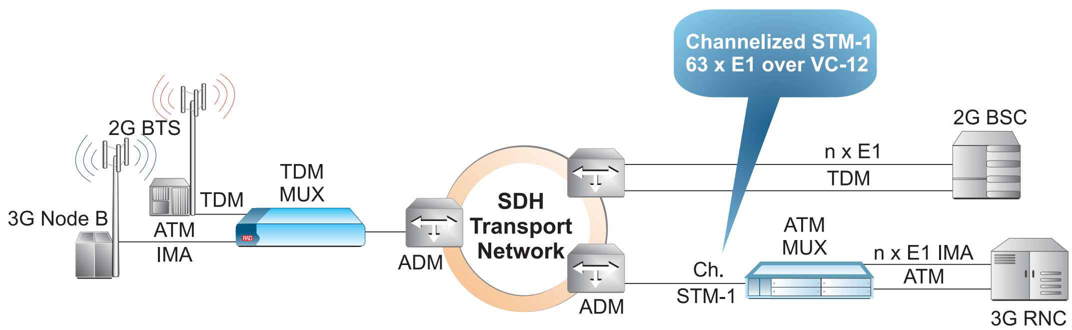

The Channelized Approach

Aggregation is also possible with a dedicated ATM device. The aggregated traffic (from multiple 2G and 3G radio sites) is carried in VC-4 containers over the SDH transmission network. Even though we are discussing a TDM-based transport platform and mixed traffic streams (ATM and TDM), the dedicated ATM aggregator enables the mobile network operator (MNO) to benefit from the ATM network's inherent QoS and traffic engineering mechanisms (OAM I.610) end-to-end (from aggregation point to aggregation point) across the transport network. This is a significant advantage since typically the transport network may be leased from another operator. There are other plusses as well. Throughput is higher with ATM (compare VC-4 bandwidth with n x VC-12) and by aggregating traffic at the remote site it is possible to apply statistical traffic handling for the variable rate flows.

Backhauling over ATM Transport

In the era of 3G, ATM is probably the ultimate cellular backhaul technology. Release 99 of the 3GPP standard denotes ATM as the de facto interface for bearing 3G traffic while GSM flows can be carried over AAL1. The protocol can cope with the multiple technologies of the various mobile generations, and can handle different traffic types with great efficiency. This is due to ATM's built-in QoS capabilities, offering traffic prioritization and service differentiation as well as traffic engineering capabilities. For example, ATM gives a mobile operator the ability to transport voice and video streaming with fixed resources using CBR while adaptive traffic flows like Internet surfing or messaging can be handled with VBR. Network administration is based on ATM's powerful (I.610) OAM mechanisms. All in all, mobile operators have greater ability to manage network operations and reduce their operating expenses. This extends to integration of 2G and 3G traffic flows, which is less problematic over a cell-based network than other networks. This also applies to supporting a future packet-based mobile access network since ATM can efficiently map Ethernet/IP over its core.

Wherever an ATM network is already in place, mobile operators stand to benefit immediately from using its powerful traffic handling features. Rolling out new infrastructure based on ATM, however, is a different matter since the high cost of switching equipment may offset the Opex savings.

Traffic aggregation also plays a role in reducing Opex and Capex in an ATM network. Dedicated aggregation devices, smaller in size and significantly lower in cost than ATM switches, can be located at remote POPs or even co-located at the radio sites themselves. This equipment can minimize the number of physical links and statistical multiplexing can contribute to reduce network-operating expenses. Furthermore, replacing low-speed ports with high-speed modules on ATM switches shaves Capex overhead. Finally, efficient, end-to-end management tools improve network performance and uptime.

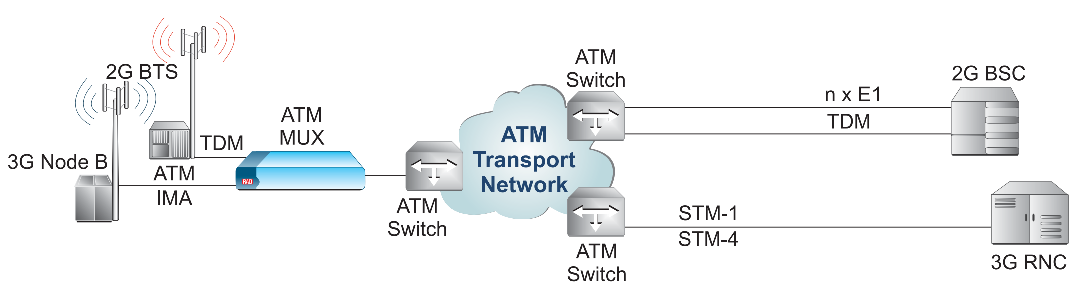

ATM Aggregation

Backhauling over DSL Links

ATM networks are often complemented in the Last Mile by DSL networks. That means that the access network is based on SHDSL and DSLAMs. Many service providers, particularly in Europe, are interested in using this popular and ubiquitous infrastructure to reduce Last Mile backhaul costs.

Modem-like intelligent integrated access devices (IADs) may be used to connect a GSM BTS to the DSLAM over a copper pair with a TDM IAD. Meanwhile the Node B can be linked to the DSLAM with an ATM IAD. Traffic from the DSLAM is tunneled over the ATM transport network. Both the TDM IAD and ATM IAD incorporate ATM quality of service features to monitor and control traffic up to the Last Mile.

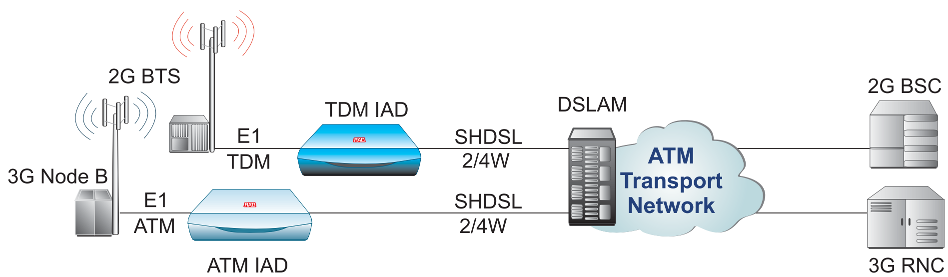

Backhaul over DSLAM Infrastructure

Leveraging the existing SHDSL infrastructure for GSM and UMTS backhaul conceivably can lead to lower access costs than traditional leased line services. The same infrastructure is also being used for Ethernet services to corporate and enterprise customers, which means it is suitable for Ethernet/IP interfaces and routing.

The bandwidth limitations of copper are also being overcome by a technology known as DSL bonding. This is based on adding copper capacity on demand and partitioning the traffic carried over those links. Thus the aggregated traffic from a BTS and Node B may be transported over multiple DSL links through the implementation of a bonding protocol like IMA or M2DSL. Alternatively, bonding could be implemented in a way that is transparent to the DSLAM and based on ML-PPP.

Backhauling over Packet Switched Networks (PSN) – Ethernet, IP/MPLS

What is the advantage of using a packet switched network for cellular backhaul? Mobile network planners looking for a comprehensive solution encompassing future developments will take a good look at the packet switched option. First, many mobile vendors are already in the process of designing radio access elements with Ethernet interfaces. Second, Metro Ethernet transport networks are mushrooming. Thus the trend is clearly toward a packetized network environment to handle 3G's data-intensive services. In this scenario, the main challenge will be to deal with legacy voice and data traffic. In other words, how to deliver services based on GSM TDM and UMTS ATM equipment over an Ethernet/IP backbone.

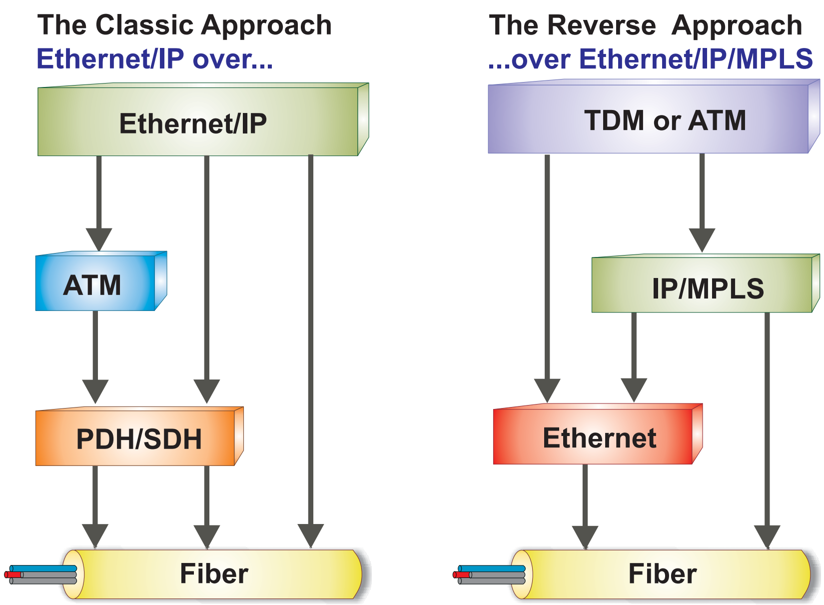

Handling Legacy Traffic

Transporting TDM over Ethernet/IP is made possible by implementing TDMoIP® (patented) and TDMoMPLS technologies. These protocols convey TDM circuits, mainly voice connections, over packet-based networks. In the case of cellular backhaul, these technologies apply nicely as they are transparent to the underlying traffic. Unlike VoIP, which requires translation of signaling, TDMoIP and TDMoMPLS provide a transport tunnel across the statistical packet network without distortion. These protocols are currently being standardized in the ITU, IETF, MEF and MPLS/Frame Relay Alliance.

TDMoIP Technology

TDMoIP is a technology patented by RAD (U.S. patent 6,731,649) for extending TDM circuits transparently across Ethernet/IP/MPLS networks. A TDMoIP gateway encapsulates segments of the constant rate TDM in AAL1 cells, which are carried in UDP/IP packets. A second TDM gateway at the destination regenerates the clock, strips the UDP/IP headers and delivers a synchronous bit stream. This solution also provides an adaptive clock for clock transfer across the entire network

Of course, TDM traffic, which is of a synchronous nature, also requires transparency of clock regeneration. Using TDMoIP technology, the mobile network operator can aggregate traffic from GSM base stations and connect them over the packet switched network to the BSC. The operator benefits from cost reductions associated with Ethernet/IP transport as well from the implementation of a future-proof infrastructure. This kind of solution is successfully being used to transmit mobile communications over packet-based satellite and microwave media.

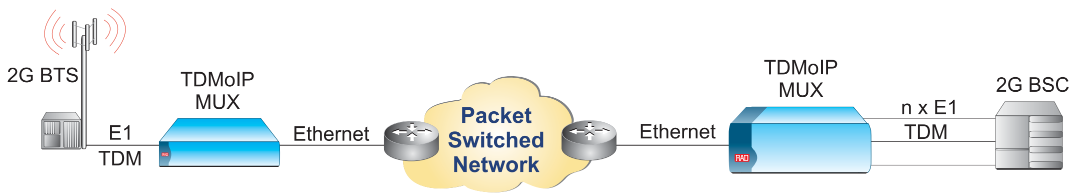

2G/2.5G Traffic over PSN

No solution to handling legacy radio traffic is complete without addressing UMTS Rel.99 traffic as well. The approach to this is similar to TDM. The ATM cells are transported over the packet-switched network using a dedicated ATMoPSN device that converts ATM cells into packets. Once again, a central device located at the edge of the packet-based transport network recovers the ATM cells in their original format and feeds them to the RNC.

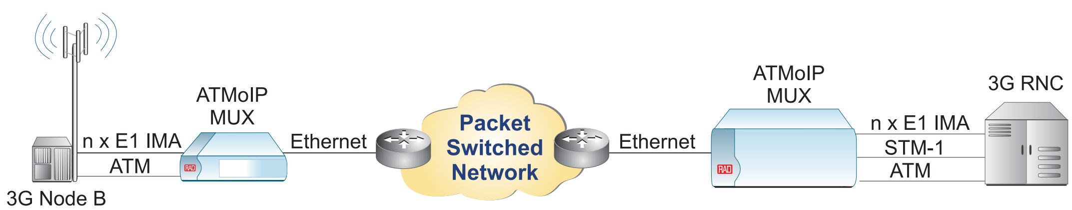

3G ATM Traffic over PSN

Implementation of ATM over PSN is based on tunneling the ATM cells according to emerging IETF or ITU standards for a pseudo-wire link over a connectionless packet-based network. [The IETF standard describes the way cells are to be mapped to packets.] Proposed options cover many-to-one (n x VCC/VPC to a single pseudo-wire) connectivity or one-to-one (a single VCC/VPC to a single pseudo-wire) connectivity. The PSN packet can hold either a single cell-in-frame or multiple cells-in-frame. Mapping of the ATM cells into packet flows is in accordance with the protocol's respective addressing method: Layer 2 (MAC or MAC and VLAN addressing); Layer 2.5 (MPLS addressing); Layer 3 (IP addressing).

The same procedure applies to QoS. To differentiate between flows and assure QoS in a connectionless network, the ATM traffic class is mapped to VLANs, in the case of Ethernet; a dedicated shim header or using EXP bits in MPLS; and IP ToS/DHCP in an IP backhaul network.

Once the mapping or tunneling issue is resolved, the next and perhaps more challenging task is to synchronize all network elements and distribute the clock throughout the transport network. This is critical for maintaining proper synchronous operation of the mobile network.

Several clock options are available:

1. Use an external reference clock based on a GPS receiver. Several mobile operators already implement this solution at aggregation nodes to avoid relying on the transport network for reference distribution. In such cases, the clock is distributed to the BTSs and Node Bs by the TDM/ATMoIP multiplexer located in the vicinity of the radio sites.

2. Provide each radio station with a dedicated PDH link for the sole purpose of reference distribution. In such a case, the reference would be from an accurate source – normally an external source or higher-level SDH transport. This dedicated clock may then be distributed by the multiplexer.

3. The third option is to distribute the clock via the PSN. Although this solution presents certain technological challenges, it would be the ideal one as long as the synchronization requirements are met.

Clock Distribution over the PSN

Recovering the clock from the PSN means that the centrally located device, which normally connects the BNC and RNC to the packet-based transport network, will distribute a dedicated reference clock together with the other traffic. To use the same stream for all locations, the transport network must support multicasting. At the remote sites, the local TDM/ATMoIP multiplexer would recover the clock by implementing TDMoIP technology or by using ATM nulling or adopting some other standard method.

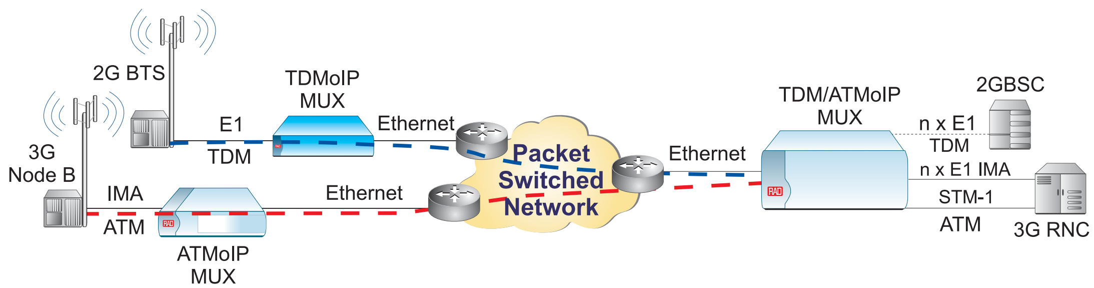

Clock Distribution over the PSN

IP-based Future

In the not so distant future, base station equipment will comport with rev.5 and rev.6, dictating an IP-based access interface. Then the issue will arise as to how to run this traffic over "legacy" networks. Solutions currently exist for supporting MLPPP over PDH, virtual concatenation with LCAS for SDH networks and AAL5 in the case of ATM.

Conclusion

There is no widely accepted blueprint or off-the-shelf solution for designing today's cellular backhaul network. A prudent, forward looking cellular backhaul strategy must take into account current voice requirements and limited data traffic while minimizing network operating expenses or capital equipment outlays. Yet it must also be able to factor in future broadband services, which will eventually encompass multimedia imaging, broadcast video and video conferencing. Despite the challenges, the mobile network operator has plenty of options available:

- Leasing or owning the transport network
- Leveraging the 2G infrastructure
- Building out backhaul from scratch
- Using copper, fiber or microwave
- Providing end-to-end QoS over the transport network
- Maintaining 2G and 3G simultaneously

The key to success is in choosing a suite of network components with the flexibility and price/performance characteristics that enable the mobile network operator to pursue all available options and implement those which best address present and future technological and commercial concerns.

Corporate Headquarters
RAD Data Communications Ltd.
24 Raoul Wallenberg Street
Tel Aviv 69719, Israel
Tel: 972-3-6458181
Fax: 972-3-6498250
email: market@rad.com
www.rad.com

U.S. Headquarters
RAD Data Communications Inc.
900 Corporate Drive
Mahwah, NJ 07430 USA
Tel: (201) 529-1100,
Toll Free: 1-800-444-7234
Fax: (201) 529-5777
email: market@radusa.com
www.radusa.com

RAD
data communications
Unique Access Solutions

Version 11/2004