

ATM Transport in Mobile Networks with Tellabs® 8600 Managed Edge System

MPLS is the latest technology for telecommunication networks that enables the transport of existing Layer 2 protocols like ATM and Frame Relay as well as new services carried on Ethernet or IP.

Executive Summary

Asynchronous Transfer Mode (ATM) is a transport protocol in 3G Code Division Multiple Access (CDMA) Radio Access Networks (RAN). ATM transport can be implemented with native ATM switches or with MPLS switch/ routers. MultiProtocol Label Switching (MPLS) is the latest technology for telecommunication networks that enables the transport of existing Layer 2 protocols like ATM and Frame Relay as well as new services carried on Ethernet or IP. This document shows 3G RAN deployment alternatives for MPLS switches, introduces ATM pseudowires which carry ATM via MPLS encapsulation and shows some delay calculation examples. These examples demonstrate that MPLS packet switches can carry real-time services like voice and video in 3G mobile RAN. Additionally, the statistical multiplexing feature enables cost-efficient data services in the same infrastructure.

Introduction

This white paper illustrates how the Tellabs® 8600 managed edge system can enable ATM transport. There is an increasing demand for cost-effective ATM transport solutions in 3G WCDMA Radio Access Networks, so it is shown as a specific application example. However, the same principles in terms of queuing delays and transport effectiveness apply to any ATM transport.

Figure 1 shows a typical mobile access network using leased lines and microwave radio. Both the 2G and 3G base stations are connected to microwave and leased line transport with E-1 interfaces. The BSC has E-1 or STM-1/VC-12 interfaces so it directly interfaces with the SDH or PDH infrastructure. Cross-connection equipment with 64 kbps level granularity improves transmission efficiency in the 2G network, if 2 Mbps circuits are only being partially filled.

A 3G WCDMA RNC is equipped with E-1 or STM-1/VC-4 ATM interfaces while 3G WCDMA base stations are usually equipped with E-1 ATM interfaces. Figure 2 illustrates the transmission protocol stacks at the 3G WCDMA base station and RNC. If 2 Mbps transport (E-1 or VC-12) is used, ATM cells are carried in 30 timeslots for 1920 kbps of bandwidth. ATM denotes Virtual Path (VP) and Virtual Channel (VC) switching layers. All user voice and data traffic is carried on the AAL2 adaptation layer.

There are a limited number of E-1 interfaces in the RNC, so STM-1 interfaces must be used to fully utilize the capacity of an RNC. Moving ATM cells from E-1 to STM-1 ATM

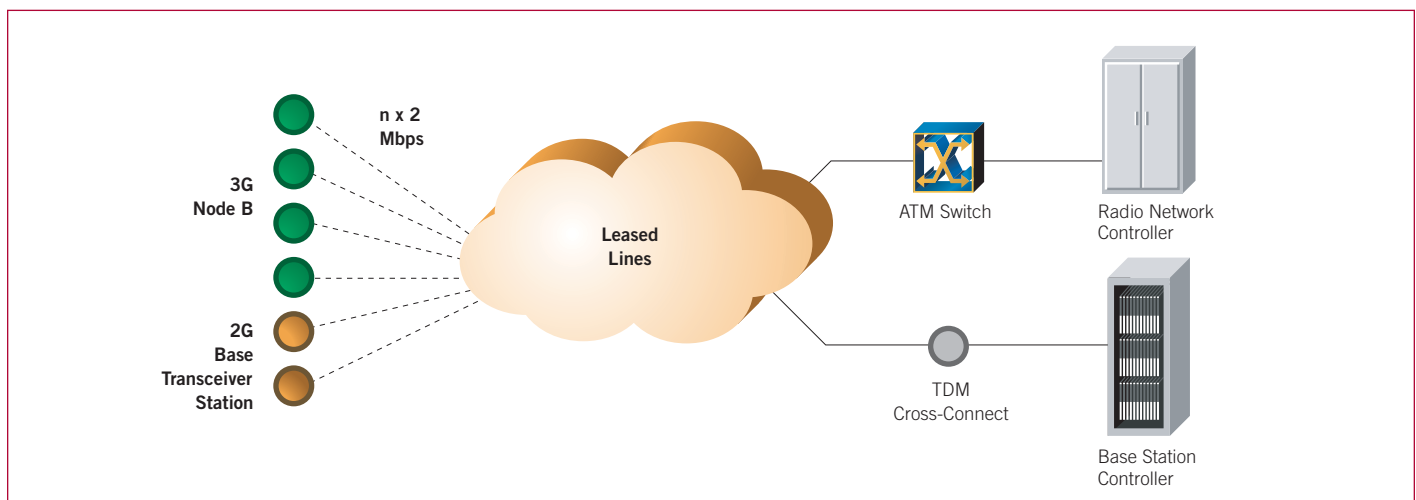


Figure 1. Leased line and microwave radio transport in 2G and 3G networks.

requires an ATM switch function that terminates E-1 and forwards ATM cells to a STM-1 interface. The ATM switch also improves the RNC port utilization as the traffic from partially utilized E-1 interfaces is forwarded to high-capacity STM-1 interfaces.




Figure 2. Transport protocol stacks in 3G RAN.

Figure 3 shows evolution scenarios from a leased-line network to a fiber-based infrastructure. The main driver for any of these deployment alternatives is the cost savings compared to traditional 2 Mbps leased-line infrastructure. The growing amount of 3G traffic makes leased-line costs especially prohibitive to commercially justify wireless data services. The scenarios assume that the last kilometers to base stations are typically implemented with microwave radios, free space optics or copper. In some cases, fiber can be available.

Case 1 illustrates the situation where all transport is based on SDH and a packet switch is deployed at the RNC site to consolidate traffic from E-1/VC-12 interfaces to STM-1 ATM. 2G and 3G base stations are connected with E-1 interfaces to SDH ADMs or cross-connects which all operate at VC-12 (2 Mbps level). This is a straightforward step from 2 Mbps leased-line infrastructure and requires minimal changes to existing transport network planning.

Case 2 suggests the deployment of packet switching equipment to hub sites closer to the base stations. One hub may aggregate from 2 to 50 3G base stations. The

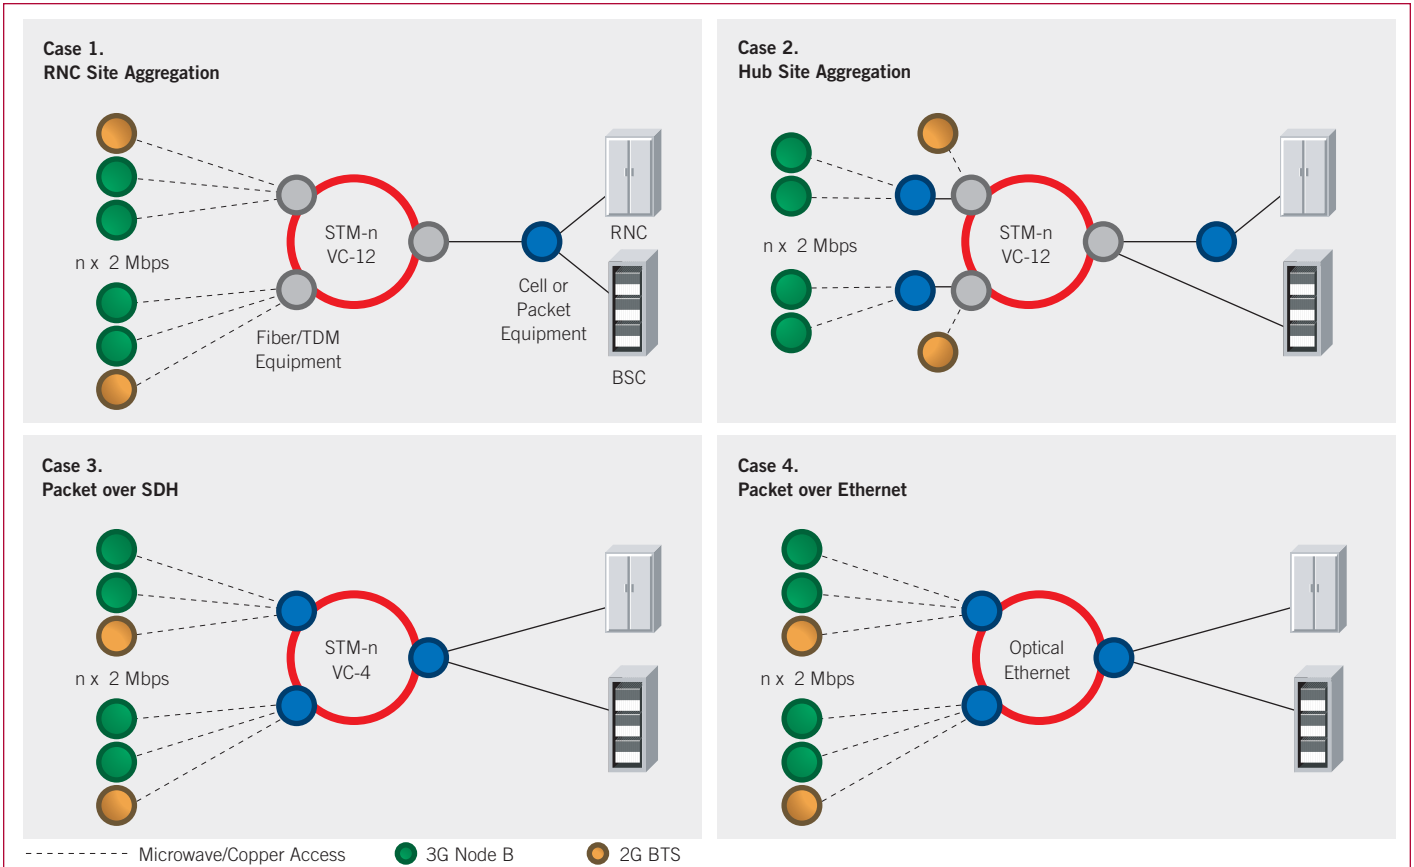


Figure 3. RAN evolution scenarios.

underlying SDH infrastructure is still based on VC-12 cross-connection which enables a direct connection of 2G base stations to SDH network elements. The benefit of this alternative is that a packet switch at the hub site can perform statistical multiplexing which improves the efficiency of the uplinks towards a RNC. The packet switch at the hub site can also create $n \times$ VC-12 (2 Mbps) bundles using ATM IMA or ML-PPP protocol. This means that more 3G base stations can be connected to the SDH infrastructure compared to Case 1. This will become an important issue to consider when High-Speed Downlink Packet Access (HSDPA) based data services are used in large volume.

Case 3 illustrates an alternative where both the 2G and 3G base stations are connected to packet switching equipment. This shows that 2G traffic is transported over the packet network using Time Division Multiplexing (TDM) circuit emulation that has similar characteristics to SDH-based transport. As both the 2G and 3G traffic is transported over the packet infrastructure, the transport layer can be simplified. The VC-12 layer is no longer required — all traffic can be transported at VC-4 layer. The main benefits of this architecture are:

- Lower cost VC-4 SDH equipment can be used
- Simpler SDH infrastructure, no VC-12 layer network planning
- All traffic is in common packet equipment
- Less expensive VC-4 POS interfaces in packet equipment
- Statistical multiplexing savings for data traffic
- Traffic protection at MPLS layer to enable differentiated protection schemes for voice and data

Case 4 shows a packet network model that utilizes Ethernet as a transport infrastructure. Ethernet can transport both ATM and IP, therefore, it fits both to 3G R99 and R5 RAN transport. The benefits of an Ethernet-based infrastructure are:

- Low cost
- Well suited to IP-based services

However, Ethernet alone is not a solution to a network that has service requirements from delay-critical real-time services to high-volume best effort data services. When Ethernet is combined with MPLS, we have an optimum solution:

- Fast (<50ms) network protection schemes
- Excellent network and service scalability
- Multiprotocol support: IP, ATM, Ethernet, Frame Relay and TDM can be carried on MPLS
- Good Quality of Service (QoS) support: performance of delay critical services can be guaranteed while over-provisioning best-effort services
- Admission control guarantees that the performance of existing service is not degraded while new services are provisioned

Figure 4 illustrates common transport of TDM and ATM over an MPLS network using pseudowire technology. MPLS network links can be SDH or Ethernet so this applies to both Case 3 and Case 4 in Figure 3.




Figure 4. TDM and ATM over MPLS pseudowire.

ATM and MPLS in 3G RAN

ATM switching in 3G RAN R99 can be implemented with ATM or MPLS-based equipment. This section discusses implementation alternatives. Figure 2 shows the minimum functionality that has to be implemented in every network: there has to be an element that switches ATM cells from E-1 to STM-1 ATM interface. This can be located at the same site with RNC or it can be in a remote location as illustrated in Figure 5. If ATM switching takes place in a remote location, VC-4 has to be used between the ATM switch and RNC. This resembles Case 3 in Figure 3. However, if an ATM switch is missing from the BSC/RNC site, the network can carry 3G traffic only and 2G traffic needs a separate VC-12 granularity SDH network. If the ATM switching takes place at the RNC site, common VC-12 SDH transport can be used for 2G and 3G. There needs to be a 4/1 SDH cross-connect at the RNC site that separates 2G traffic from 3G, as illustrated in Case 1 in Figure 3. The optimal physical interface between the ATM switch and the SDH cross-connect is in a STM-1/VC-12 because it gives the greatest 2 Mbps (E-1) density and lowest cost. The other alternative is to use physical E-1 interfaces in case the ATM switch does not support STM-1/VC-12 interface.

Next we will compare ATM and MPLS when deploying hub site and RNC site packet aggregation, as illustrated in Figure 6 and Figure 7. Figure 6 illustrates the case when using ATM only.

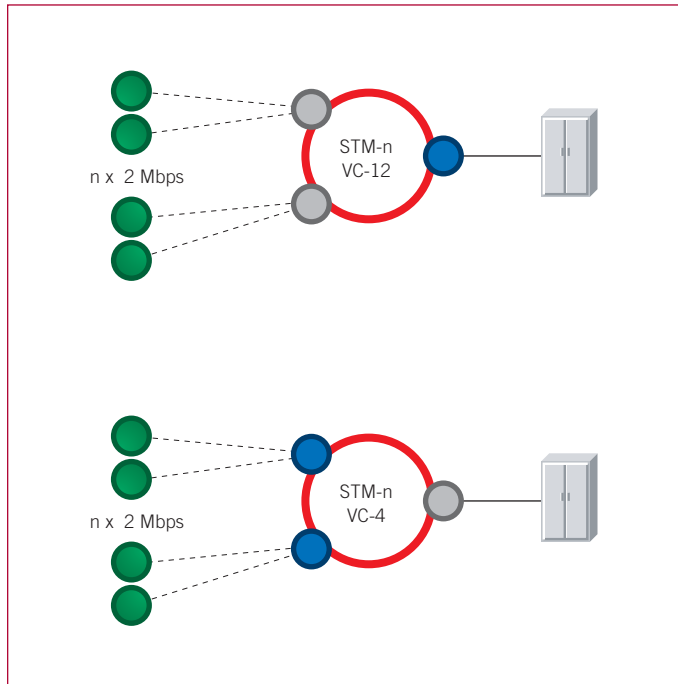


Figure 5. ATM switching at RNC site or remote hub site.

End-user traffic from a 3G base station is carried on one or several E-1s via Inverse Multiplexing over ATM (IMA) which enables a grouping of several E-1s into a logical n x 2 Mbps transmission pipe. IMA can also be used on VC-12, which is illustrated on the left of Figure 6. In this scenario, we utilize the existing VC-12 based SDH transport and want to optimize VC-12 usage by utilizing statistical multiplexing. On the right of Figure 6, a simplified SDH infrastructure is based on VC-4.

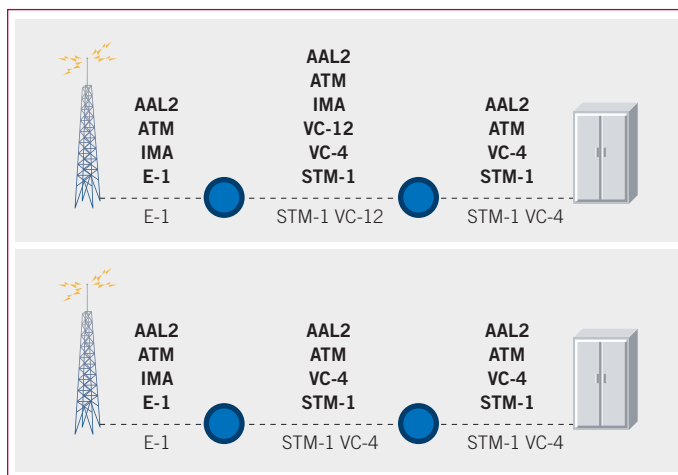


Figure 6. ATM transport and protocol stack alternatives.

Figure 7 shows some similar network structures, but this time ATM is carried on MPLS. Figure 7 shows three optical transport alternatives: VC-12 or VC-4 based SDH and Ethernet. The protocol stack between the 3G base station and the hub site is the same as before, but the protocol stack between the two packet switching nodes has changed because MPLS is added. Multi-link PPP (ML-PPP) has a similar function as IMA in Figure 6: it creates a bonding between several VC-12 forming a logical n x 2 Mbps transmission pipe. ATM cells are encapsulated with MPLS labels with pseudowire (PWE3) technology. The benefit of MPLS is that it enables QoS with admission control, provides good protection schemes. It also detaches ATM from underlying transmission which enables using either SDH or Ethernet as a transmission technology. MPLS is also future-proof, as it enables transport of Ethernet or IP which are the requirements of 3GPP R5.

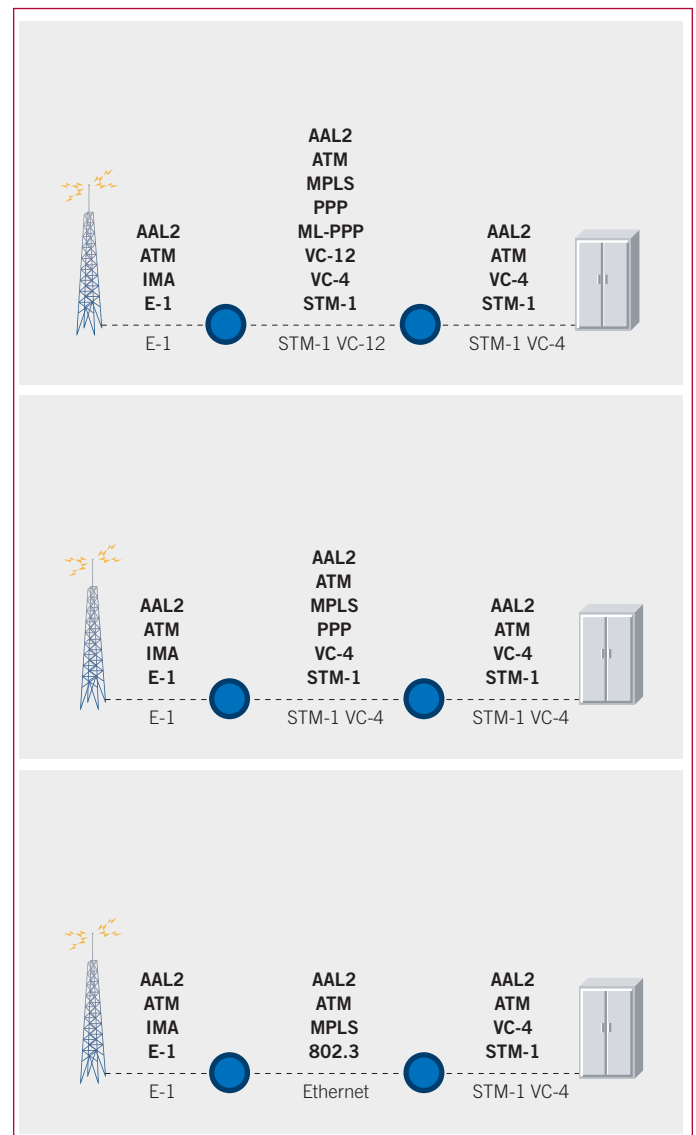


Figure 7. ATM over MPLS via pseudowires.

ATM Pseudowires

The Internet Engineering Task Force (IETF) PWE3 working group has defined a method to encapsulate ATM cells with MPLS labels. There are several alternatives for encapsulation and this section introduces two of them: n-to-one and one-to-one mode.

Figure 8 illustrates ATM cell treatment in n-to-one and one-to-one modes. The original UNI cell received from an ATM interface is on the left and the new formats for

MPLS encapsulation are on the right. N-to-one mode removes the HEC field (one byte) while one-to-one mode removes the whole cell header (five bytes). As n-to-one mode keeps VPI and VCI in the cell header, it is possible to encapsulate multiple ATM VPs or VCs into a single pseudowire. One-to-one mode enables mapping of a single VC into a pseudowire, and a mapping between VPI and VCI needs to be configured in the setup phase as this information is not carried in the cell itself. If a VP is mapped to pseudowire using one-to-one mode, VCI needs to be carried in the ATM cell header, so three bytes (GFC, VPI, PT, C and HEC fields) will be removed from an ATM cell in the encapsulation phase.

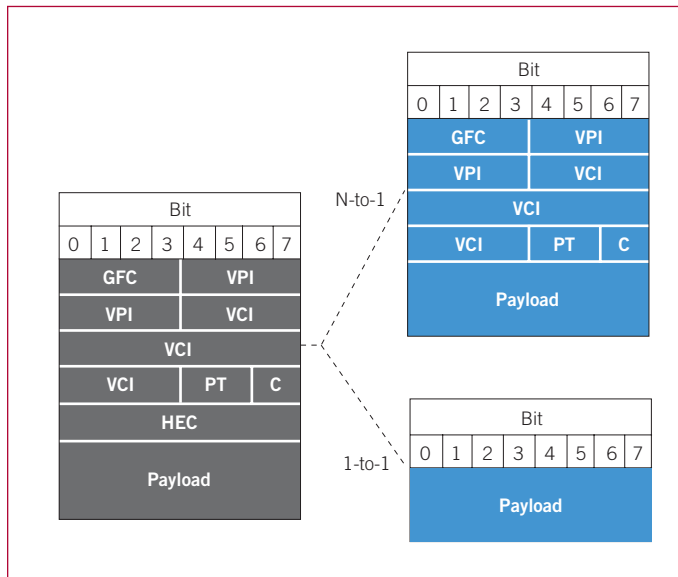


Figure 8. N-to-one and one-to-one ATM cell encapsulation modes.

Figure 9 illustrates n-to-one mode ATM cell encapsulation into MPLS pseudowire. Two MPLS labels are required: 1) PSN label defines the path over the MPLS network and 2) PW label (pseudowire) identifies the ATM connection. Pseudowire encapsulation defines also a third label, control word, which is optional in n-to-one mode. The ATM cell starts immediately after the PW label as the control word is not used in this case.

Figure 10 shows one-to-one mode ATM cell encapsulation. The PSN label and PW label are added in a similar way as in the n-to-one case and the control word needs to be added as well because it is mandatory in one-to-one mode. The ATM cell starts after the control word and in this case it is just the ATM payload as described in Figure 8.

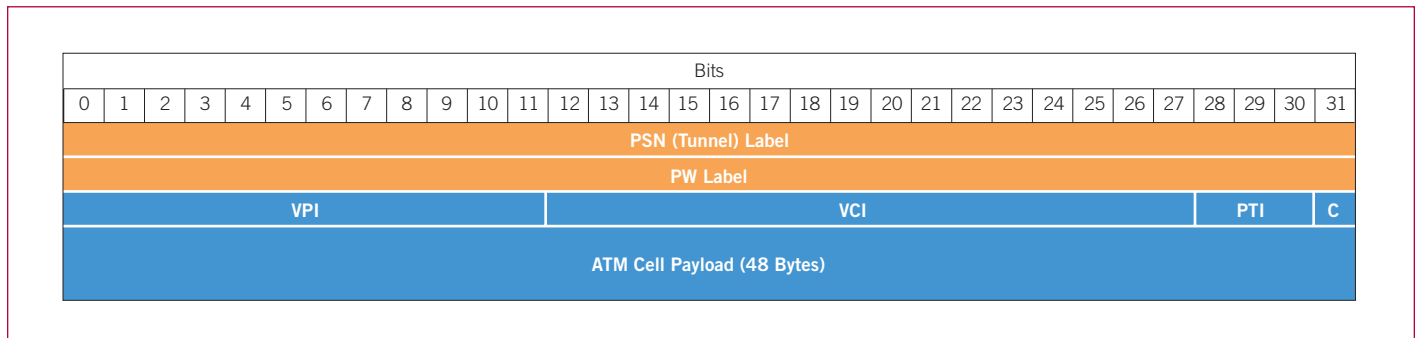


Figure 9. N-to-one encapsulation with two MPLS labels.

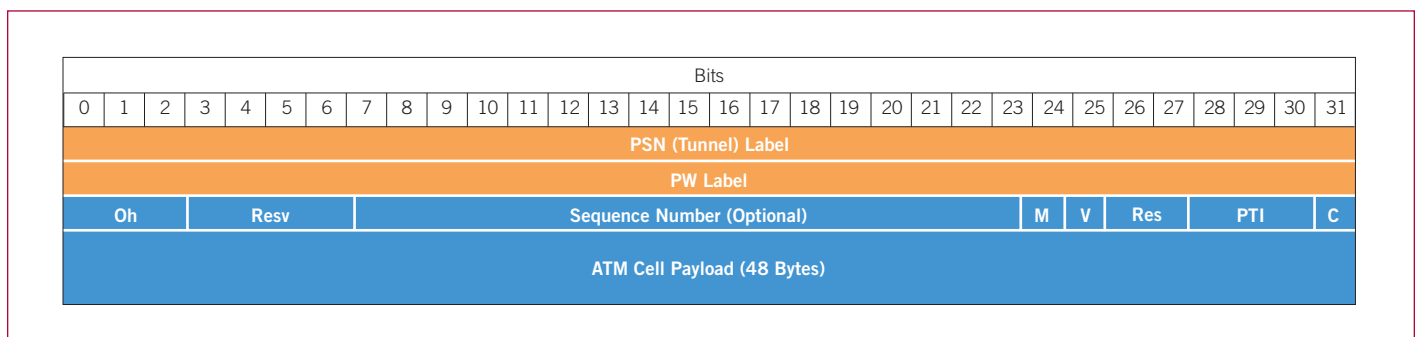


Figure 10. One-to-one encapsulation with three MPLS labels.

MPLS pseudowire encapsulation adds 8 bytes to the ATM cell if PSN and PW labels are added. One can argue that MPLS increases bandwidth requirements. However, bandwidth efficiency can be increased with cell concatenation which is illustrated in Figure 11.

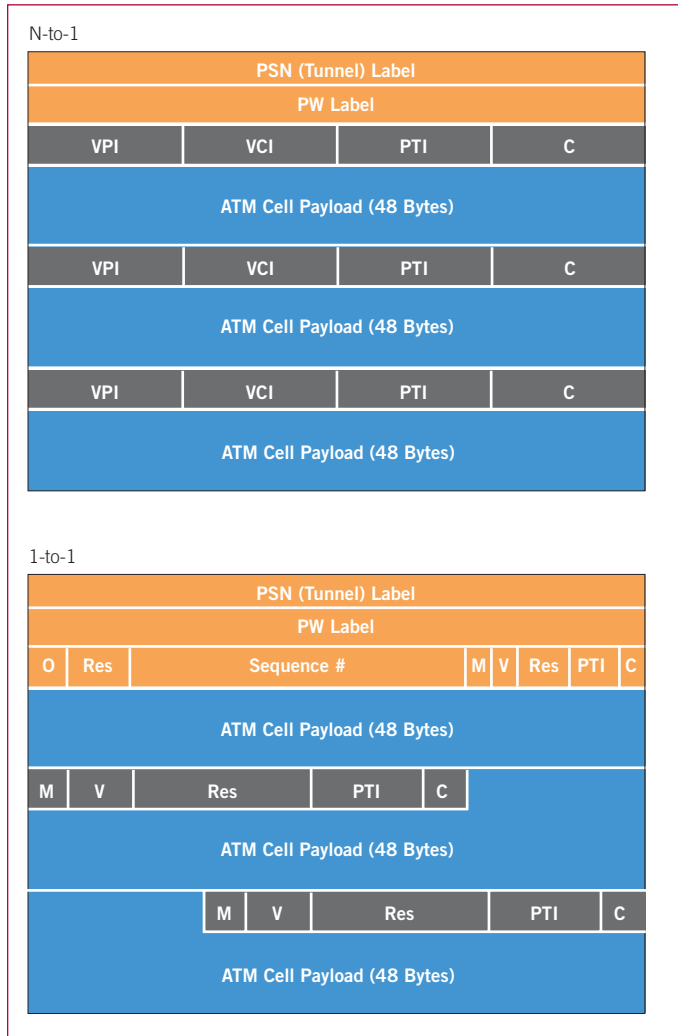


Figure 11. ATM cell concatenation in MPLS pseudowires.

Figure 12 illustrates the bandwidth efficiency of ATM pseudowires with cell concatenation. If a single cell is encapsulated with MPLS, more than 10% overhead is created when comparing against plain ATM transport. However, when cells are concatenated, the overhead of MPLS labels is soon compensated. In one-to-one mode the overhead is compensated with three concatenated ATM cells, and in n-to-one mode the overhead is compensated when eight cells are concatenated.

The downside of cell concatenation is that it adds delay and jitter to traffic as cells have to be buffered on the transmit side to create an MPLS packet from concatenated cells.




Figure 12. Bandwidth efficiency of ATM pseudowires.

The whole protocol stack requires consideration to fully understand the implications of MPLS in terms of bandwidth efficiency. The protocol stack alternatives were illustrated in Figure 6 and Figure 7. The most efficient way of transporting ATM is on a VC-4 SDH virtual container. The overhead caused by Multiplex Section Overhead (MSOH), Regenerator Section Overhead (RSOH) and VC-4 Path Overhead (POH) is 90 bytes out of 2430 bytes which is 3.7%. If ATM cells are carried on VC-12, the total SDH overhead increases to 22.2% and the overhead caused by IMA is one cell out of 128 cells, which is 0.78%. The protocol stacks for this comparison are shown in Figure 6 between the two packet nodes.

If ATM PWE3 is used between the two packet nodes as shown in Figure 7, the protocol stack becomes more complex. In this case, ML-PPP, PPP/HDLC and MPLS pseudowire overheads need to be taken into account as shown in Table 1. In n-to-one mode ATM cells are first encapsulated with two MPLS labels which add eight bytes. MPLS labels are then encapsulated with PPP protocol ID (two bytes) indicating MPLS payload. Multilink-PPP fragment control adds four or two bytes depending if long or short format is used. Finally, multilink-PPP is encapsulated with PPP in HDLC-like framing before transmitting it on a serial link. The recommended PPP/HDLC settings on high-speed links (RFC 1662) should use HDLC address and control without compression, PPP protocol in two byte mode and 16-bit FCS. Therefore, the total PPP/HDLC overhead is seven (1+1+1+2+2) bytes.

| Protocol | Bytes |
|-----------|-------|
| MPLS PWE3 | 8 |
| PPP | 2 |
| ML-PPP | 4 |
| PPP/HDLC | 7 |
| Total | 21 |

Table 1. Protocol overhead for ATM PWE3 with ML-PPP.

As this overhead is added per packet, cell concatenation improves bandwidth efficiency as shown in Figure 12. If the packet size is small (1-2 cells), PWE3 on an n x VC-12 links bundle consumes significantly more bandwidth than ATM IMA, as Figure 13 illustrates.




Figure 13. Overhead of ATM PWE3 compared to ATM IMA on n x VC-12.

Using VC-4 or Ethernet with ATM PWE3 is more efficient. Figure 14 shows an ATM PWE3 compared to an ATM on VC-12 IMA and an ATM on VC-4. ATM VC-4 is the most efficient mechanism if less than ten cells are concatenated in ATM PWE3. If more than ten cells are concatenated, ATM PWE3 encapsulation on VC-4 is almost equal while ATM PWE3 on Ethernet provides some bandwidth savings. Figure 14 also shows that ATM PWE3 on VC-4 or on Ethernet is more efficient than ATM IMA on n x VC-12, if two cells or more are concatenated.




Figure 14. Overhead of ATM PWE3 on Ethernet, VC-4 and ATM IMA on VC-12 compared to ATM on VC-4.

Delay Calculation

There are four delay components in a packet network:

- Serialization delay
- Signal propagation delay
- Forwarding delay
- Queuing delay

Serialization delay is dependent on the egress port bandwidth and the packet size. The calculation formula is $\text{delay} = \frac{\text{packet-size (bits)}}{\text{port-bandwidth (bits/second)}}$. This delay component becomes significant with lower-speed interfaces (E-1 or below) and large packet sizes. Signal propagation delay is the speed of signal in the physical medium, e.g. the speed of light in a fiber is 5 ms/1000 km. The forwarding delay is an internal delay by a packet node making a forwarding decision for a packet. This delay is <50 microseconds in modern packet switches. When ATM IMA or ML-PPP is used, internal processing of those protocols adds a few milliseconds of delay to the calculation. The last delay component is queuing delay, which is caused by statistical multiplexing and asynchronous arrival times of packets. This is the time a packet has to wait for service in a queue. There is some queuing delay, even for real-time packets because there is a chance that two or more packets arrive at once to the same queue. The more flows that enter the queue and the larger the packet size is, the longer the queuing delay gets.

Figure 15 illustrates a delay calculation scenario. The same network architecture has also been used in the protocol overhead calculation. A 3G base station (Node B) transmits ATM cells on single E-1 or multiple E-1s using ATM IMA. An ATM cell stream is received by a hub packet node which forwards them toward an RNC over STM-1 or Gigabit Ethernet (GigE). The serialization delay is experienced on all egress interfaces, in this case on $n \times E-1$ in Node B and on STM-1/GE in the packet node. If ATM cells are concatenated on the pseudowire, the delay increases. The delay decreases if the $n \times E-1$ IMA bundle is larger. Table 2 illustrates serialization delay at various packet and IMA bundle sizes. A single cell delay is always 0.221 ms. If eight or more cells are concatenated, the size of the IMA group should be at least $2 \times E-1$ to keep the serialization delay below 1 ms. The same calculation applies to the STM-1/VC-12 interface.

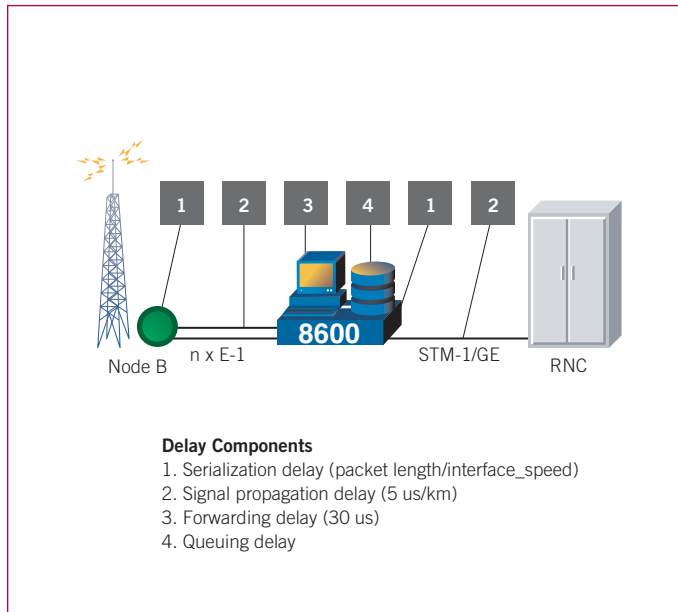


Figure 15. Delay calculation example.

| n x E-1/P12s | Concatenated Cells | | | | |
|--------------|--------------------|-------|-------|-------|-------|
| | 1 | 2 | 4 | 8 | 16 |
| 1 | 0.221 | 0.442 | 0.883 | 1.767 | 3.533 |
| 2 | | 0.221 | 0.442 | 0.883 | 1.767 |
| 4 | | | 0.221 | 0.442 | 0.883 |
| 8 | | | | 0.221 | 0.442 |
| 16 | | | | | 0.221 |

Table 2. Serialization delay on n x E-1 (ms).

Table 3 shows the serialization delay on an STM-1 VC-4 and Gigabit Ethernet. The delay ranges from four to 45 microseconds depending on the packet size (1-16 ATM cells). This calculation takes pseudowire overheads into account, which were illustrated in Figure 14. As can be seen, the serialization delay on the STM-1 and Gigabit Ethernet interface is very small so it becomes insignificant compared to other delay components.

| Cells | STM-1 | 1 Gbps |
|-------|-------|--------|
| 1 | 0.004 | 0.001 |
| 2 | 0.006 | 0.001 |
| 4 | 0.012 | 0.002 |
| 8 | 0.023 | 0.004 |
| 16 | 0.045 | 0.007 |

Table 3. Serialization delay (ms) on STM-1 and Gigabit Ethernet interface.




Figure 16. Queuing delay in Hub Site Switch.

Figure 16 shows the network assumption for a queuing delay calculation. The queuing delay is dependent on packet size, the number of flows and egress port bandwidth. The higher the number of flows and the larger the packet size, the more queuing delay there is. The reason for this is that packets from several flows may arrive at once into the queue and the scheduler will serve them one by one. The 10-5 quantile means that the probability for a packet to exceed this delay value is one to 100,000.

Table 4 illustrates queuing delay in the Hub Site Switch with GigE uplink. The GigE uplink utilization in this case is fairly low, 16.8% with 60 x E-1 interfaces, so the queuing delay is a few microseconds only. The packet size is one ATM cell.

| Interface | CBR delay (10 ⁻⁵ quantile) | Delay variation |
|-----------|---------------------------------------|-----------------|
| 4 x E-1 | 0.00184 ms | 0.00115 ms |
| 12 x E-1 | 0.00199 ms | 0.00130 ms |
| 20 x E-1 | 0.00226 ms | 0.00157 ms |
| 40 x E-1 | 0.00267 ms | 0.00198 ms |
| 60 x E-1 | 0.00312 ms | 0.00243 ms |

Table 4. CBR queuing delay and delay variation in Hub Site with GigE uplink.

When an STM-1 uplink is used, the CBR delay ranges from 7.9 to 31.6 microseconds in Table 5. The utilization of STM-1 link is 69.2% with 60 x E-1 interfaces offering CBR traffic so there is no overbooking in these cases.

| Interface | CBR delay (10 ⁻⁵ quantile) | Delay variation |
|-----------|---------------------------------------|-----------------|
| 4 x E-1 | 0.0079 ms | 0.0051 ms |
| 12 x E-1 | 0.0106 ms | 0.0078 ms |
| 20 x E-1 | 0.0135 ms | 0.0107 ms |
| 40 x E-1 | 0.0211 ms | 0.0183 ms |
| 60 x E-1 | 0.0316 ms | 0.0288 ms |

Table 5. CBR queuing delay and delay variation in Hub Site with STM-1 uplink.

Table 6 shows the delay of CBR traffic in the RNC Site Switch. The peak load and average load show as a percentage how much traffic arrives to a STM-1 uplink in the RNC Site Switch. When the percentage is above 100%, there is a overbooking situation as the peak load is more than the bandwidth of the STM-1 uplink.

| Peak load STM-1 (%) | Ave. load STM-1 (%) | CBR delay (ms) 10 ⁻⁵ quantile | Delay variation (ms) |
|---------------------|---------------------|--|----------------------|
| 48 | 43 | 0.021 | 0.018 |
| 80 | 72 | 0.033 | 0.030 |
| 96 | 87 | 0.046 | 0.043 |
| 144 | 87 | 0.075 | 0.072 |
| 192 | 87 | 0.083 | 0.080 |

Table 6. CBR queuing delay in RNC Site Switch.

The queuing delays in the Hub Site overbooking cases are analyzed next. In the example scenarios the uplink from the Hub Site is either STM-1/VC-4 or STM-1/VC-12. In both cases the interface to the Node Bs is E-1. Each Node B has 6 x E-1 interfaces and the bandwidth is shared between CBR and UBR circuits. Figure 17 shows the Hub Site Switch with STM-1/VC-12 uplink. There are 24 E-1 circuits sharing the 12 x VC-12 uplink so the overbooking factor is 2. Table 7 shows CBR and UBR queuing delays with different ratios of CBR and UBR traffic. CBR delay is fairly constant independent of the load on the link while there is a greater variation of the UBR delay when the load approaches 100%.

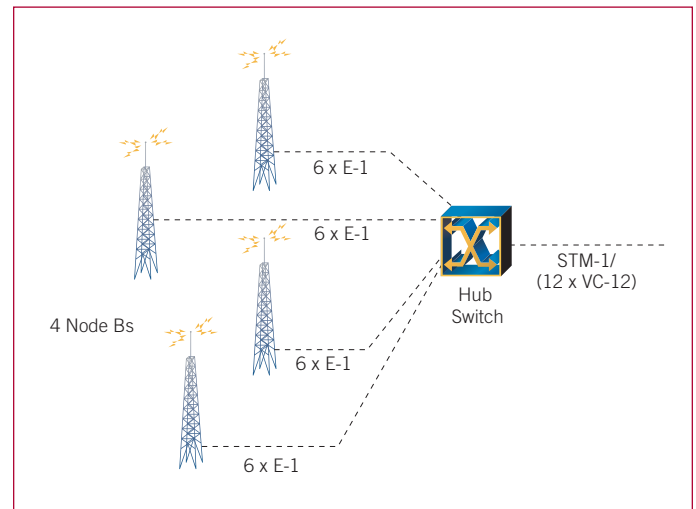


Figure 17. Hub site with STM-1/VC-12 uplink.

| CBR/UBR ratio (E-1 quantity) | Ave. uplink utilization % | CBR delay (ms) 10 ⁻⁵ quantile | UBR delay (ms) 10 ⁻⁵ quantile |
|------------------------------|---------------------------|--|--|
| 16/8 | 90 | 0.377 | 1.93 |
| 12/12 | 90 | 0.339 | 1.39 |
| 8/16 | 90 | 0.316 | 1.53 |
| 16/8 | 95 | 0.395 | 11.2 |
| 12/12 | 95 | 0.343 | 7.9 |
| 8/16 | 95 | 0.319 | 6.3 |

Table 7. CBR and UBR queuing delay in Hub Site with STM-1/VC-12 uplink.

The next scenario assumes 20 Node Bs. There are 120 E-1 circuits accessing the Hub Switch with an STM-1/VC-4 uplink so the overbooking ratio is 1.54 (120 * 1.92 Mbps/ 149.76 Mbps). Table 8 shows Constant Bit Rate (CBR) and Undefined Bit Rate (UBR) queuing delays when the average STM-1 uplink utilization is 90% or 95%. Even though the utilization level of the uplink is similar to Table 7, UBR queuing delay stays at the one ms range. When the average uplink utilization is 95% or more, the delay of bursty packets increases rapidly. Alternatively, packet drops will increase if queues are kept shorter. To keep the UBR queuing delay, packet loss and queue lengths at reasonable levels, the average uplink utilization should max. 90–95%.

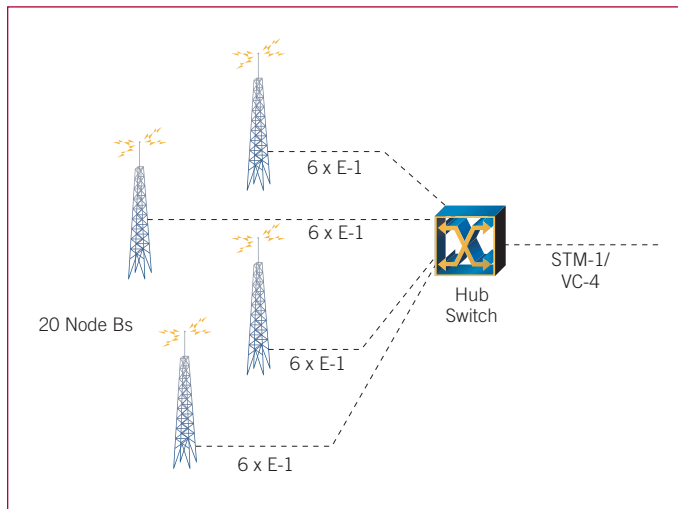


Figure 18. Hub site with STM-1/VC-4 uplink.

| CBR/UBR ratio (E-1 quantity) | Ave. uplink utilization % | CBR delay (ms) 10 ⁻⁵ quantile | UBR delay (ms) 10 ⁻⁵ quantile |
|------------------------------|---------------------------|--|--|
| 80/40 | 90 | 0.033 | 0.27 |
| 60/60 | 90 | 0.025 | 0.18 |
| 40/80 | 90 | 0.019 | 0.15 |
| 80/40 | 95 | 0.036 | 0.88 |
| 60/60 | 95 | 0.026 | 1.22 |
| 40/80 | 95 | 0.020 | 0.93 |

Table 8. CBR and UBR queuing delay in Hub Site Switch with STM-1/VC-4 uplink.

The final example doubles the amount of E-1 links per Node B. The scenario assumes the following:

- 20 Node Bs per Hub Switch, 12 x E-1 per Node B
- CBR and UBR traffic in a 12 x E-1 IMA bundle
- CBR capacity 2 x E-1 per Node B
- UBR capacity 10 x E-1 per Node B with 5:1 overbooking (2 x E-1 effective bandwidth)
- UBR packet size 8 ATM cells (488 bytes), CBR packet size 1 ATM cell (68 bytes)

This scenario illustrates the case where the majority of the traffic is data in UBR service category. The larger packet size for UBR eliminates the overhead caused by MPLS encapsulation. Data services tolerate more delay which enables higher overbooking factors. The CBR traffic uses 66% of the STM-1 uplink bandwidth and the rest is available for UBR. The serialization delay for CBR and UBR packets is 3.6 and 26.1 us respectively.

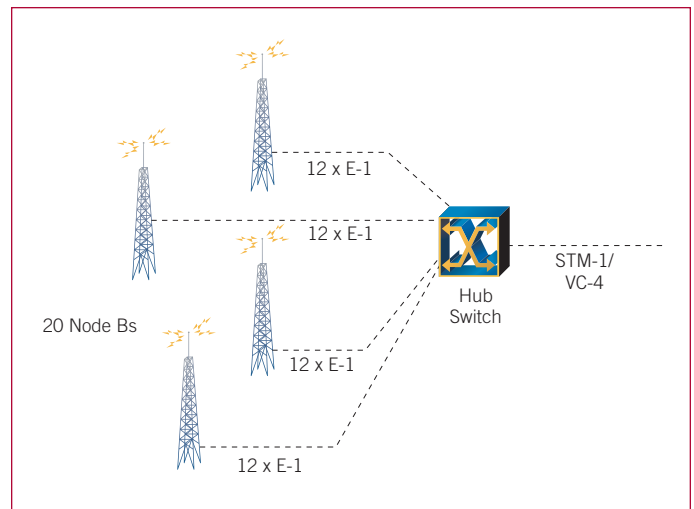


Figure 19. Hub site overbooking scenario.

| CBR share of uplink capacity | UBR share of uplink capacity | Ave. uplink utilization | CBR delay (ms) 10 ⁻⁵ quantile | UBR delay (ms) 10 ⁻⁵ quantile |
|------------------------------|------------------------------|-------------------------|--|--|
| 66% | 14% | 80% | 0.065 | 0.71 |
| 66% | 19% | 85% | 0.065 | 0.99 |
| 66% | 24% | 90% | 0.067 | 1.77 |
| 66% | 29% | 95% | 0.068 | 2.65 |
| 66% | 34% | 100% | 0.069 | — |

Table 9. CBR UBR queuing delays with high overbooking.

UBR queuing delay is largely dependent on uplink utilization. Queuing delay increases steeply as utilization approaches 100%. The overbooking ratio for UBR is 5:1 in terms of physical interfaces but the average UBR bandwidth should be such that the average uplink utilization stays below 100%. CBR delay stays almost the same independent of the uplink utilization which is important for delay critical services.

Conclusions

MPLS switching with ATM pseudowires provides an effective alternative to native ATM switching. The overhead caused by ATM pseudowire encapsulation can be compensated with ATM cell concatenation. The increased serialization delay caused by the larger packet size can be compensated by bundling links with ML-PPP or by using high-speed links like STM-1 or GigE. The CBR service category carrying real-time traffic always has minimal delay, and the traffic load on UBR service category does not affect the delay in the CBR service category. Therefore, the UBR (or VBR) traffic can be overbooked and that enables savings in the transport costs and network equipment. Overbooking and statistical multiplexing is essential for data services to ensure cost-effective networks.