Backhauling X2

Cambridge Broadband Networks Limited

Summary

The X2 is a new type of interface introduced by the LTE Radio Access Network. It connects neighbouring eNodeBs in a peer to peer fashion to assist handover and provide a means for rapid co-ordination of radio resources.

This paper considers in detail the functions that are performed over the X2 in order to derive performance requirements for the transport networks which will need to carry it. X2 does not require a significant bandwidth. RAN Vendors expect that traffic levels will be only a few percent of the main 'S1' backhaul interface. Latency requirements, on the other hand, are more challenging. X2 delay for user traffic is additional to that on the S1, so should be as small as possible. Furthermore, interference co-ordination features proposed for future releases of LTE-Advanced are highly sensitive to small delays.

The paper goes on to consider the impact of different last mile backhaul topologies on X2 connectivity and latency. The hub-and-spoke topology achieved with point-to-multipoint (PMP) microwave provides a lower and more consistent latency than achieved with the tree-topologies resulting from point-to-point (PTP) networks.

Finally, it shows that VectaStar point-to-multipoint microwave backhaul from Cambridge Broadband Networks is well suited to the transport of LTE X2 traffic.
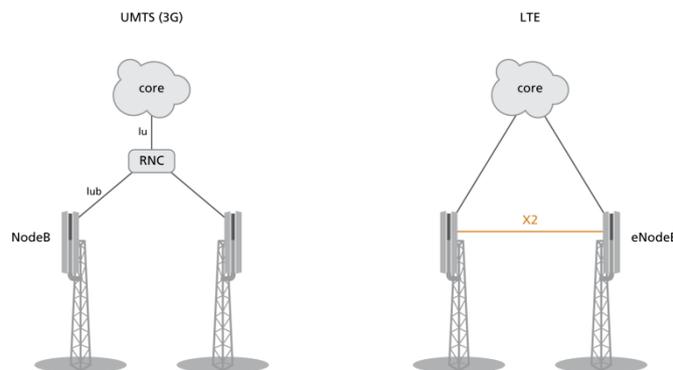
Introduction

The rapid adoption of data services has forced the wireless industry to rethink the way that mobile services are delivered. Compared to voice, data requires significantly more traffic and must be delivered at a much lower cost per bit in order to be profitable for operators. The industry has been working for many years on standards for the LTE RAN (Long Term Evolution Radio Access Network) and EPC (Evolved Packet Core), which address these evolving needs. Higher performance is achieved through a new OFDM Air Interface, with flexible bandwidth, native MIMO support and optimisations for delivering packet-based data services.

Lower cost per bit is achieved with a simpler, 'flatter' architecture based on the widely used IP protocols. The use of IP allows operators to use Ethernet packet-based transport networks which, thanks to the success of the Internet, are now lower cost than legacy ATM-based backhaul networks. While Ethernet may be lower cost, it is 'packet' rather than 'circuit' switched, which provides 'best-effort' rather than 'guaranteed' delivery of data. LTE/EPC standards therefore include end-to-end Quality of Service features to ensure essential traffic is prioritised during times of congestion.

The flatter architecture reduces cost by reducing the number of core network nodes that operators need to buy and maintain. In particular, the RNC (Radio Network Controller) which was used to co-ordinate NodeBs in the UMTS RAN (Universal Mobile Telecommunications System Radio Access Network) has been removed. In LTE, most of the RNC functionality has been moved to the eNodeBs (LTE Base Stations), which now co-ordinate in a peer-to-peer fashion over a new interface called the 'X2'. These architectures are illustrated in Figure 1 below.

Figure 1 The LTE Radio Access Network moves RNC functionality to the eNodeB and introduces a new X2 Interface

In this paper we take a look inside the X2 interface, analysing the tasks it needs to perform in today's 'release-8' LTE networks as well as future LTE-Advanced networks. For these tasks we consider the implications for bandwidth and latency and develop requirements for backhauling of X2. Many operators were initially concerned that X2 would require physical connections between all eNodeBs. However, X2 is a 'logical' interface between only neighbouring eNodeBs and can be switched over the existing transport network. We consider different options for backhauling X2 and the implications on performance.

X2 functionality

Overview

The X2 interface supports exchange of information between eNodeBs to perform the following functions:

- **Handover**: mobility of UEs between eNodeBs

- **Load Management**: sharing of information to help spread loads more evenly

- **CoMP** (Co-ordinated Multi-Point transmission or reception): Neighbouring eNodeBs co-ordinate over X2 to reduce interference levels

- Network Optimisation

- eNB configuration update, cell activation, including neighbour list updates

- **Mobility Optimisation:** co-ordination of handover parameters

- **General Management:** initialising and resetting the X2

Many of the key functions are described by the 'X2 Application Protocol', specified in [2]. These 'control plane' signalling procedures have been standardised in order to ensure eNodeBs from different vendors are interoperable. Some of these functions additionally require the sending of 'user plane' data over the X2. This can be carried in bearers using the same GTP[1] protocol used for the other core network interfaces, as specified in 36.424 [12].

Of the above X2 functions, it is anticipated that handover and CoMP will be the most demanding in terms of bandwidth and latency, and these are analysed in more detail in the following sections.
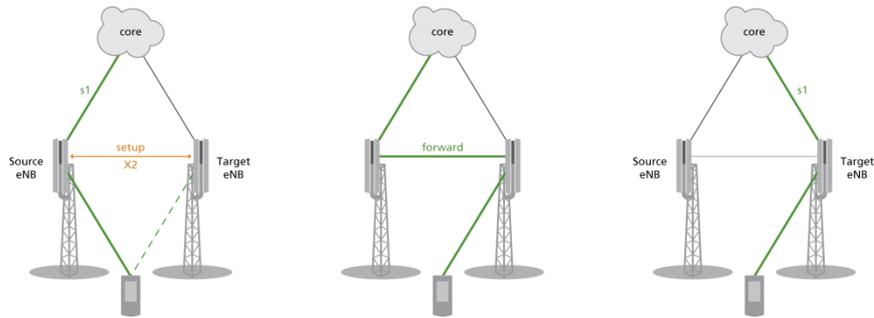
X2 handover

Description

Handover is required when a UE moves between coverage areas of different eNodeBs. Note that in 3GPP standards terminology, "mobility" refers to handover, and not just the general movement of a user. 3G systems used a make-before-break 'soft handover' to ensure connectivity was maintained as users moved between cells. LTE however uses a fast break-before-make 'hard handover' which makes more efficient use of both spectral and backhaul resources. It does however create a short interruption in the connection. Packet-based services are by nature short parcels of information with gaps in between, so providing the handover interruption time is short, it will not impact the quality of the service.

3G handover is co-ordinated by the RNC shown in Figure 1. This is not present in LTE, and so eNodeBs themselves co-ordinate the handover. The X2 interface is not mandatory, and in its absence eNodeBs arrange the handover over the S1 interface via the core network. However, X2 handovers have the benefit of reduced preparation times and lower core processing load.

Figure 2 illustrates the role of the interface in an X2 assisted handover:

---

[1]GTP = GPRS Tunnelling Protocol. GPRS = Generic Packet Radio Service

Figure 2 X2 assisted handover

1. Preparation

   a. A UE approaching the edge of its cell will start detecting signals from neighbouring cells.

   b. Measurements of signal strength and quality will be reported back to the serving or *source* eNodeB.

   c. At some point the source eNodeB may decide that the UE would be better served by one of its neighbours, and so it sends a handover request to the *Target*, over the X2.

   d. The target will perform admission control and if it has sufficient resources, it will confirm the handover, and provide various IDs and security information back to thesource.

2. Execution

   a. The source eNodeB sends a handover command to the UE which includes the information needed to connect to the target.

   b. The UE then retunes and connects to the target.

   c. At this point the core network is still sending the users' data to the source eNodeB, which is forwarded over the X2 to the target and then to the UE. This data forwarding is the most significant component of X2 traffic.

3. Completion

   a. Once the UE is connected, the target eNodeB sends a 'path switch' command to the core.

   b. The core updates the S1 connectivity for the UE to route directly to the target eNodeB.

The X2 handover reduces the handover setup time and hence reduces the chances of handover failure when the network cannot keep up with the rapidly changing radio conditions of fast moving UEs. X2 forwarding also helps minimise packet delay variation, as forwarded packets arrive at the target eNodeB sooner. X2 handover reduces core processing and backhaul requirements, as the signalling can be kept local to the source and target eNodeBs. The interested reader can find a detailed description of the handover procedure in 3GPP specification 36.300 [3]

Handover-related X2 functions

The X2 Application protocol also includes procedures to optimise handover performance:

Mobility parameter management: Neighbouring eNodeBs can recommend updates of handover thresholds

Handover history reports: Records of historical handovers can help in optimisation, e.g. to reduce ping-ponging

Radio link failure: The target eNodeB can quickly inform the source, if the UE loses contact during a handover. The source may then re-establish the link for another attempt.

Backhaul requirements for X2 handover

As mentioned earlier, the forwarding of the user plane data dominates bandwidth requirements for the X2 interface, and it may be in the order of Mbps. Control plane signalling is likely to be of the order of kbps, and can therefore be considered negligible.

Although LTE might be capable of delivering data rates of the order of hundreds of Mbps, these only occur when the user has an excellent radio link with an eNodeB and no interference from others. This is not the case for a UE at the cell edge in handover. By definition, a UE in handover receives similar signal levels from the source and target eNodeBs. This results in low signal quality and lower data rates.

eNodeB vendors have provided us with the following estimates of X2 bandwidth:

- "...the X2 interface needs little bandwidth–a maximum of 3% of the amount required by the S1 interface", Huawei [4]

- "X2 is < 2% of S1 traffic, more likely 1.6%", Ericsson [6]

- "typically see an average of 5% and that's consistent with customers' views" Alcatel-Lucent [5]

- "3 to 5% of S1 is a generous allowance for X2", FT/Orange, Transport Networks for Mobile Operators Conference, 2010

All vendors express X2 throughput as a percentage of that over S1, which in turn depends on the configuration of the eNodeB supported. Wider channel bandwidths (up to 20MHz) and higher order MIMO (up to 4x4 for early LTE releases) result in more traffic over S1 and therefore X2. Further information on LTE backhaul provisioning can be found in a CBNL webinar [7].

X2 traffic is not constant, and depends on the degree of mobility between cells. For example, an X2 link between eNodeBs covering a train line may generate more traffic than an X2 between eNodeBs covering nomadic users in an office. One paper addresses the subject of 'batch handover', where many users - for example on a train - handover at once. In such cases, X2 bandwidth may increase to over 10% of S1 [5].

Latency

X2 handover is no more sensitive to latency than S1 signalling or user traffic, so the same requirements apply. NGMN requirements are 10ms end-to-end round trip delay, and 5ms is recommended [8]. When used for data forwarding, the X2 delay must be added onto the S1 delay, and so should be as small as possible. Sub 1ms delays are therefore desirable.

Interference co-ordination and CoMP

Description

In a cellular network, some UEs will be able to hear (or be heard by) more than one eNodeB. In the first release of LTE, the serving cell is the one with the strongest signal, and the other signal is considered an interferer. If these two signals are at similar levels, they will interfere with each other and result in poor performance for the user. It can be shown that capacity can be improved if the two cells co-ordinate to use different frequency or time slots [9]. In LTE-Advanced, rather than just *avoiding* interference, cells can instead *harness* the energy from the multiple signal paths in order to achieve further capacity gains.
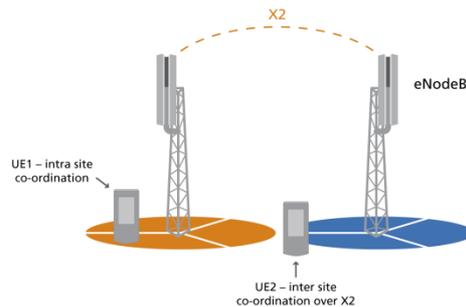

Figure 3 UEs requiring intra- and inter-site co-ordination

Most cellular networks today use tri-cellular Base Stations shown in Figure 3. This diagram illustrates two different scenarios for co-ordination:

- UE1 sits on the boundary of two cells which are controlled by the same eNodeB. In this scenario, any co-ordination is internal to the eNodeB and so can benefit from almost zero latency and wide bandwidth signalling.

- UE2 sits on the cell boundary between two different eNodeBs. Any co-ordination information must therefore pass over the X2 interface, and may be limited by its bandwidth and latency capabilities.

LTE and LTE-Advanced include several schemes for "*Co-ordinated Multi Point transmission and reception*" [9,13], each with varying degrees of information sharing as follows:

- Inter-Cell Interference Co-ordination (ICIC, Rel-8 LTE):

  - For the downlink eNodeBs share information with their neighbours about frequency and time resources. The neighbours can avoid using these same resources for their cell edge users.

  - For the Uplink, eNodeBs share information on the interference they are receiving from their neighbours. The serving cell then learns which UEs are causing the interference, and can co-ordinate their transmissions.

  - LTE-rel8 provides semi-static co-ordination of resources of the order of a few seconds.

- Co-ordinated Scheduling/Beamforming (CS/CB) (LTE-Advanced):

  - LTE Advanced enhances the rel-8 ICIC scheme to dynamic co-ordination where resources can be co-ordinated as rapidly as on a per 1ms sub-frame basis.

  - Frequency, time and beamforming information can be co-ordinated.

  - This requires lower latency and higher bandwidth signalling over X2.

- Joint Processing (LTE-Advanced):
    - Joint Processing (JP) differs in that the user data is present at multiple eNodeBs.

    - Transmission or reception can be co-ordinated to rapidly select the best cell, or even combine signals from multiple cells coherently.

    - JP requires significantly more backhaul bandwidth than the other schemes.

    - eNodeBs must be tightly synchronised, requiring a low latency X2.

Backhaul requirements for CoMP

The first release of LTE (rel-8) supports static and semi-static interference co-ordination, with information exchange in the order of seconds. Bandwidth requirements are expected to be negligible compared to X2 handover.

CoMP schemes in LTE-Advanced may require considerably more bandwidth. An analysis of co-ordinated scheduling suggests as much as 770 kbps of signalling would be needed for an X2 interface between tri-cell eNodeBs co-ordinating every 1 millisecond[10].

Joint Processing schemes may additionally require sharing of user data over the X2, which for the cell edge UEs requiring co-ordination could be of the order of several Mbps. Note that it would not be necessary to forward user data over X2 for handover if this was already being done to support a Joint Processing CoMP scheme.

Latency

In release 8 LTE, the semi-static schemes co-ordinate over timescales of several seconds, and are not sensitive to latencies of order 10ms required for X2 data forwarding.

LTE-Advanced (release 10) introduces dynamic co-ordination, which can reconfigure resources as often as every 1ms sub-frame. Such high rates of adaptation would only be needed for rapidly changing radio conditions, and in practice less frequent changes may be sufficient. Rapid adaptation requires that information arriving over the X2 from neighbour cell is up to date. Latency on the X2 connection will therefore reduce the benefit of co-ordination schemes as illustrated in Figure 4. X2 Latency as low as 5ms reduces cell edge and median user throughputs by over 20% in this example of a Joint Transmission CoMP scheme. Simpler schemes may not be impacted as much, but this shows that the sophisticated CoMP schemes require very low X2 latencies to achieve their full potential.
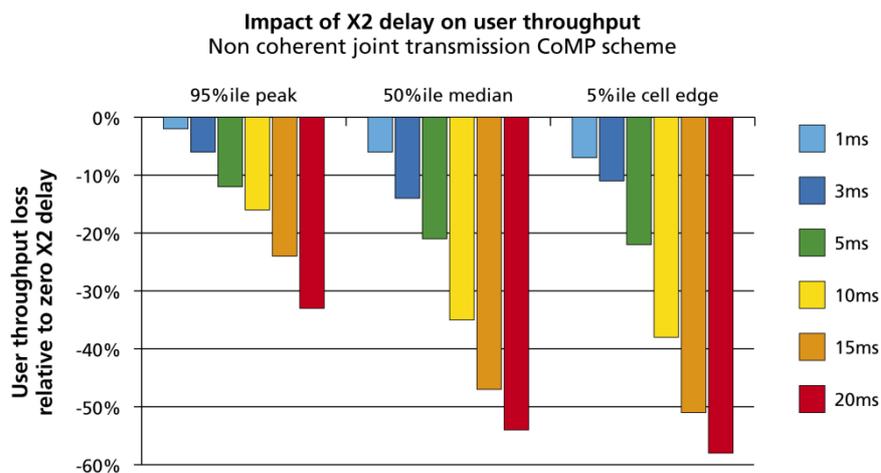
Impact of X2 delay on user throughput
Non coherent joint transmission CoMP scheme

Figure 4 Impact of X2 delay on user throughput with CoMP scheme. 3km/h users assumed. Source Qualcomm [11]

X2 connectivity

- X2 does not require a dedicated physical connection between eNodeBs

- It is a logical interface that can pass over the existing IP transport network.

- X2 does not require L3 Routing

- Where possible, switching can be used and is preferable for the higher performance achieved.

- X2 interfaces are not needed between ALL eNodeBs in a network

- X2 interfaces are only needed between *neighbouring* eNodeBs (i.e. those that control cells with overlapping coverage areas). It is only between neighbours that handovers will occur or interference co-ordination will be needed.
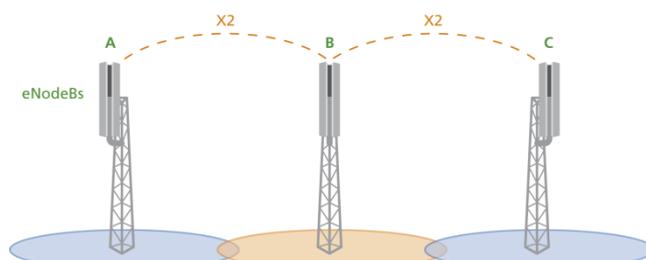

Figure 5 X2 interfaces are only needed between neighbouring eNodeBs

Figure 5 provides a simplified illustration, where A&B are neighbours as well as B&C. However, A&C would not be neighbours as their coverage does not overlap. Real world propagation is not as clear cut, and there are many cases where coverage overlap occurs between non-adjacent cells, especially in dense networks covering city centres, for example. The following vendor comment and NGMN requirement suggest that anything up to 32 X2 interfaces will be needed per eNodeB.

- "With mixture of macro/micro eNBs, one eNB may need connectivity up to ~20..30 neighbouring cells", Ericsson [6]

- "Typically up to 6 operators and 32 X2 interfaces MAY be envisioned per e-NB". NGMN Requirements [8]

Point-to-multipoint backhaul and X2

Network topologies for mobile data

Mobile networks themselves are based on a point-to-multipoint topology, where each base station connects to many user devices. This allows base stations to share their resources between many users. The PMP topology is well suited to mobile broadband where end user traffic is bursty, and it makes sense for many users to 'share a fat pipe' rather than providing a fixed 'circuit' to each user whether or not they use it. This is the principle of packet switching on which LTE is based.

If we now look at the backhaul network, we can see that exactly the same principles apply. PMP microwave backhaul is an effective way of sharing the backhaul resource between several Base Stations, each generating bursty traffic. Figure 6 illustrates how the PMP backhaul network overlays onto an LTE RAN, and gives the terminology for the components of the network.
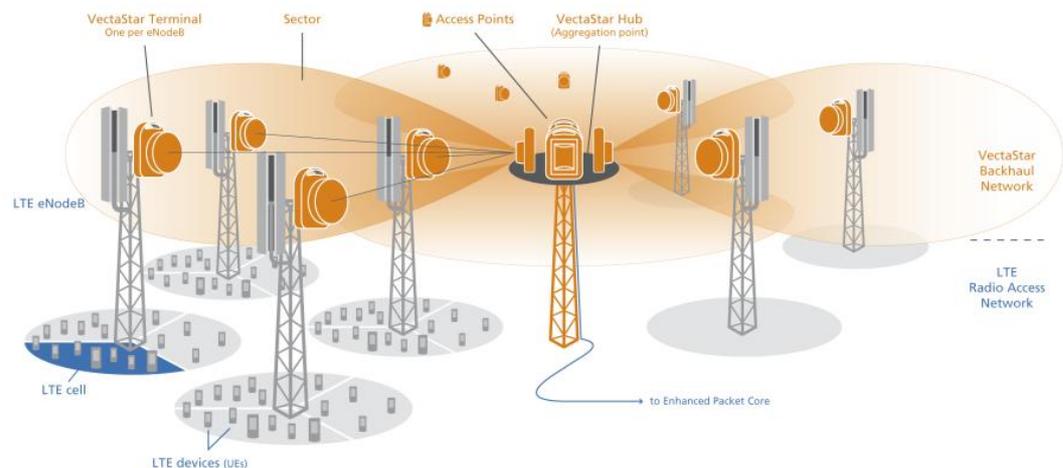

**Figure 6 Topologies for LTE and PMP backhaul overlay**

In the same way that many LTE devices connect to one cell of an eNodeB, many Vectastar Terminals can connect to one *Sector* of a VectaStar Hub. The sector of coverage is provided by an *Access Point*, and each *Hub* can control several Access Points.

Just like end user data traffic, backhaul traffic from eNodeBs is uncorrelated and bursty in nature [7]. It therefore makes more sense for several eNodeBs to 'share a fat backhaul pipe' rather than provisioning a fixed bandwidth 'circuit' to each one. This represents a shift away from the point-to-point (PTP) backhaul topologies used in legacy 'circuit switched' voice networks. The move to data requires an evolution to packet switching both in the LTE Radio Access Network, and in the backhaul network.

Connectivity and latency

The X2 requires a *logical* connection between neighbouring eNodeBs. For lowest latency and minimum loading of the transport network, the path of the *physical* X2 connection should be as short as possible. When using PMP backhaul, many neighbouring eNodeBs are covered by the same hub site, as shown in Figure 6. X2 interfaces between eNodeBs within the same hub site can connect over only two hops, and so will have low latency and will not need any other transport resources.

Figure 7 shows the network topology for PMP and PTP backhaul networks. In theory, a PTP network could be arranged to provide the same hub-and-spoke topology as the PMP network. In practice however, this would require tens of PTP antennas on one mast at the hub site. For this comparison we consider a practical limit of four antennas at any one site. This results in a 'tree' topology for backhaul based on PTP microwave. Analysis of these topologies reveals that the *average* hop count between adjacent is only slightly higher in a tree than in the hub-and-spoke. However, the *variation* in hop count is much wider in the tree, with some neighbours having up to five hops between them. While extra links can be added into the tree to reduce hop count and add redundancy, the same is true for a PMP. In summary, the hub-and-spoke topology resulting from a PMP deployment provides a more consistent connectivity for X2 links between neighbouring eNodeBs.
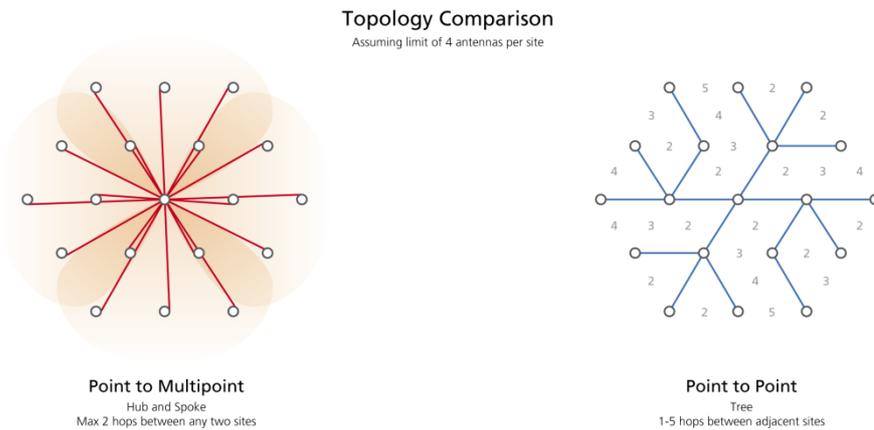
Figure 7 Comparison of X2 hop count over different backhaul topologies

Figure 8 shows latency performance of the VectaStar PMP product measured with over 17million packet delays. For 'normal'S1 backhaul between eNodeBs and the core network, one-way link latencies are applicable. Median values are 320µs for DL and 570µs for UL. The uplink is slightly longer because unlike the downlink, the VectaStar Terminal must request resources from the Access Point before transmitting. This mechanism is also observed in the latency performance of the LTE RAN in [14].
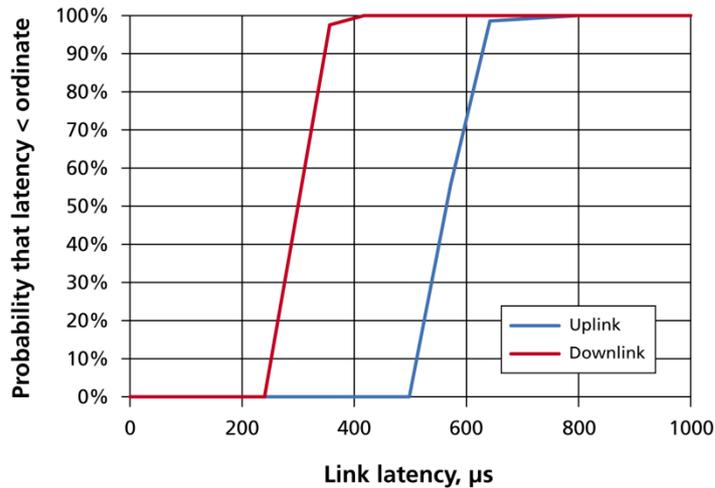
Figure 8 Link latency for VectaStar PMP product

For two-hop X2 connections, latency is the combination of the uplink and downlink figures. Two-hop latencies of 930µs can be achieved with 99% availability. This confirms that VectaStar is well within X2 latency requirements for both forwarding of user data during handover, and future interference co-ordination schemes in LTE-Advanced.

Summary

The new X2 interface provides a means for sharing information between neighbouring eNodeBs to improve handover and to reduce mutual interference. The X2 carries both control plane signalling and forwarded user plane data, and it is the latter which dictates the bandwidth requirement. Vendor views vary, but most consider that 5% of the S1 bandwidth is a generous allowance for the X2. The S1 bandwidth itself depends on the operating bandwidth and MIMO configuration of the eNodeB it serves. Other studies have shown a typical 10MHz 2x2 downlink requires a total of around 30-40 Mbps of backhaul per tri-cell eNodeB during busy times. This figure includes S1, X2 and transport protocol overheads.

When used to forward user data, the X2 adds delay to the overall budget, and so low latencies are highly desirable. In the future, LTE-Advanced includes a capacity-enhancing technique called Co-ordinated Multipoint Transmission and Reception. It has been shown that even relatively small latencies of 5ms can reduce the gains of such schemes significantly.

PMP microwave backhaul is well suited to meeting the low latency requirements of the X2 interface. Neighbouring eNodeBs are nearly always covered by the same multipoint sector, which enables rapid turnaround of the X2 traffic. PMP networks form a hub-and-spoke topology, which is much better suited to X2's peer-peer connectivity than the tree topologies formed by PTP networks. With VectaStar PMP backhaul from Cambridge Broadband Networks, sub 1ms latency can be reliably achieved for X2.

PMP is the best microwave backhaul choice for urban/suburban LTE networks.

References

[1] "*X2 general aspects and principles*", 3GPP TS 36.420 V9.0.0 (2009-12)

[2] "*X2 application protocol (X2AP)*", 3GPP TS 36.423 V9.4.0 (2010-09)

[3] "Evolved Universal Terrestrial Radio Access Network, (E-UTRAN); Overall description" 3GPP TS 36.300, (handover described in section 10.1.2)

[4] "*LTE requirements for bearer networks*", Huawei Publications, June 2009, http://www.huawei.com/publications/view.do?id=5904&cid=10864&pid=61

[5] "*Sizing X2 Bandwidth For Inter-Connected eNodeBs*", I. Widjaja and H. La Roche, Bell Labs, Alcatel-Lucent, IEEE VTC 2009-Fall, Sept 2009

[6] "*Right Sizing RAN Transport Requirements*", P. Unell; Ericsson, Transport Networks for Mobile Operators, 2010

[7] "LTE Traffic Webinar", European Comms, 6 October 2010, http://www.cbnl.com/news/articles/European_comms_LTE_Webinar_v2.pdf

[8] "*NGMN Optimised Backhaul Requirements*", Next Generation Mobile Networks Alliance, August 2008, http://www.ngmn.org/uploads/media/NGMN_Optimised_Backhaul_Requirements.pdf

[9] "*LTE: The UMTS Long Term Evolution: From Theory to Practice*", section 12.5 p294, Stefania Sesia, Issam Toufik and Matthew Baker, Wiley 2009

[10] "*Backhaul Requirements for Centralized and Distributed Cooperation Techniques*", S, Brueck, Qualcomm, Presented at ITG Heildleberg, 8 July 2010 http://www.ikr.uni-stuttgart.de/Content/itg/fg524/Meetings/2010-07-08-Heidelberg/index.html

[11] "*Centralized Scheduling for Joint-Transmission Coordinated Multi-Point in LTE-Advanced*", S. Brueck, L. Zhao, J. Giese, M. A. Awais, proc. ITG/IEE Workshop on Smart Antennas, Bremen (WSA'10), Germany, 23. – 24. February 2010

[12] "*X2 data transport*", 3GPP TS 36.424, v8.5.0

[13] "*Further advancements for E-UTRA physical layer aspects*", 3GPP TR 36.814 http://ftp.3gpp.org/specs/html-info/36814.htm

[14] "*Latest results from the LTE/SAE Trial Initiative*", February 2009 http://www.lstiforum.org/file/news/Latest_LSTI_Results_Feb09_v1.pdf