



Neutral Citation Number: [2022] EWHC 10 (Pat)

Case No: HP-2019-000032

IN THE HIGH COURT OF JUSTICE
CHANCERY DIVISION
PATENTS COURT

Royal Courts of Justice
Strand, London, WC2A 2LL

Date: 6th January 2022

Before :

THE HON MR JUSTICE MELLOR

Between :

(1) **INTERDIGITAL TECHNOLOGY CORPORATION** **Claimants**
(2) **INTERDIGITAL PATENT HOLDINGS, INC.**
(3) **INTERDIGITAL, INC.**
(4) **INTERDIGITAL HOLDINGS, INC.**

- and -

(1) **LENOVO GROUP LIMITED** **Defendants**
(2) **LENOVO (UNITED STATES) INC.**
(3) **LENOVO TECHNOLOGY (UNITED KINGDOM) LIMITED**
(4) **MOTOROLA MOBILITY LLC**
(5) **MOTOROLA MOBILITY UK LIMITED**

Adrian Speck QC, Mark Chacksfield QC and Edmund Eustace (instructed by **Gowling
WLG**) for the **Claimants**

James Abrahams QC, William Duncan and Kyra Nezami (instructed by **Kirkland & Ellis
International LLP**) for the **Defendants**

Hearing dates: 22nd-25th, 29th-30th June 2021
Draft Judgment distributed 20th December 2021

Approved Judgment

I direct that pursuant to CPR PD 39A para 6.1 no official shorthand note shall be taken of this Judgment and that copies of this version as handed down may be treated as authentic.

COVID-19: This judgment was handed down remotely by circulation to the parties' representatives by email. It will also be released for publication on BAILII and other websites. The date and time for hand-down is deemed to be Thursday 6th January 2022 at 10am.

.....

THE HON MR JUSTICE MELLOR

Mr Justice Mellor:

TABLE OF CONTENTS

INTRODUCTION	3
THE EXPERT WITNESSES	4
THE SKILLED PERSON	5
COMMON GENERAL KNOWLEDGE	5
HSUPA	23
Multiplexing of data in HSUPA	26
Scheduled and non-scheduled data	28
CGK points in dispute	28
THE PATENT	28
CONSTRUCTION / CLAIM SCOPE	41
‘MAC-d flow data’	43
‘smaller than’	45
‘means for’	47
ESSENTIALITY/INFRINGEMENT	47
VALIDITY	47
TERMINOLOGY	48
JI-3 and X3	49
Resolution of the CGK points in dispute	50
The role of UE maximum transmit power in TFC and E-TFC restriction.	51
TFC selection on the legacy channels	52
The skilled person’s understanding of E-TFC selection and multiplexing in TS 25.309 V6.2.0	53
The relationship between common Guaranteed Bit Rate traffic flows and transport formats in E-DCH	55
THE DEVELOPMENT OF THE CASE ON FILIATRAULT	55
LENOVO’S CASE ON FILIATRAULT IN CLOSING	61
Alleged Anticipation by Filiatrault	62
What Filiatrault disclosed to the skilled person	69
The disclosure in Filiatrault	69
Points A and B.	74
Obviousness	75

Obviousness as presented in Dr Irvine's report(s)	75
Obviousness on IDC's reading of Filiatrault	77

DXX/14 – Alleged anticipation on IDC's reading of Filiatrault and construction of claim 1.	77
---	-----------

CONCLUSION	80
-------------------	-----------

Introduction

1. This is my judgment for Trial B concerning EP(UK) 3,355,537 B1 (EP537 or the Patent). The action is in familiar form by which the Claimants ('IDC') seek to make the Defendants ('Lenovo') take a FRAND licence to their portfolio of mobile telephony patents said to be standard essential patents (SEPs). EP537 has a priority date of 29 April 2005 which is not contested. It is said to be essential to the UMTS (3G) standard, release 6 onwards ('the Standard').
2. The Patent relates to a feature of Enhanced Uplink (EU) also referred to as High Speed Uplink Packet Access or HSUPA which was introduced in Release 6. Release 6 was under development at the priority date but not finalised and introduced into products until later. An Enhanced Downlink version (HSDPA) was introduced in Release 5, so *some* of the features were already familiar.
3. This case concerns the way in which in HSUPA it was proposed that data were assembled for transmission on the physical layer (PHY). Only data blocks of certain pre-determined sizes (called E-TFCs) are allowed to be transmitted. Furthermore, the amounts of data which could be sent on the Enhanced Uplink were controlled. The 'useful' data to be sent did not necessarily match an allowed E-TFC, so the system was one in which padding bits were added to ensure the E-TFC was filled. The Patent is concerned with minimising the amount of padding which is sent by adjusting or quantising the amount of data multiplexed into an E-TFC to more closely match the E-TFC size. There are different ways of doing this, none of which provide a perfect solution. The differences arise from (a) what limits are chosen, (b) how they are used and (c) the point in the sequence that various limits are applied.
4. Lenovo contend that the Standard does not require the patented method to be used. There is no dispute as to the requirements in the Standard and Lenovo's non-infringement arguments resolve to points of construction of claim 1 of the Patent.
5. The case on validity narrowed to anticipation or obviousness over a single piece of prior art called Filiatrault, which is a marked up version of 3GPP TS 25.309 v6.2.0 entitled '*FDD Enhanced Uplink; Overall Description; Stage 2 (Release 6)*' circulated by the editor, Mr Charles Filiatrault, on 27 April 2005 to the members of the RAN WG2 (Radio Access Network Working Group 2) by email to the WG2 email reflector, which makes their contents accessible to the public on the ETSI website.

6. At the heart of this case is a very substantial dispute over what Filiatrault disclosed. Lenovo contend that Filiatrault disclosed the invention, but if not, then they say the Patent was obvious over Filiatrault. IDC say it missed the invention completely, so that Patent is neither anticipated nor obvious.
7. The case based on Filiatrault evolved over time. In the first round of expert evidence, the parties' respective experts (Mr Townend for IDC and Dr Irvine for Lenovo) set out diametrically opposed views as to what Filiatrault disclosed. Dr Irvine only set out an explanation as to why the Patent was anticipated by Filiatrault. In their respective second reports, each expert responded to the view espoused by the other and Dr Irvine explained why he considered the Patent was obvious on Mr Townend's view of Filiatrault. That necessitated further short reports in reply from each expert, for which provision was made in my Order dated 27 May 2021 at the PTR. Although Lenovo dispute this, IDC say that Lenovo's case on Filiatrault changed during the trial. From my point of view, the case did change. Indeed, the various ways in which the case based on Filiatrault morphed made this Judgment more difficult to write because it was never clear which parts of the complicated evidence had either been abandoned or were no longer significant. Of course, an anticipation case which undergoes repeated change does not inspire confidence, but in this case the devil is in the detail and the detail must be examined. Finally, by way of introduction, I must apologise for the time it has taken to complete this judgment. Due to the subject matter, it was not easy to resume work after breaks to attend to other cases.
8. This brief introduction explains why the main sections of this judgment are:
 - (a) CGK;
 - (b) The Patent;
 - (c) Issues of Construction;
 - (d) Essentiality/Infringement;
 - (e) Filiatrault – disclosure, anticipation and obviousness.
9. Adrian Speck QC, Mark Chacksfield QC and Edmund Eustace argued the case for IDC and James Abrahams QC, William Duncan and Kyra Nezami argued Lenovo's case. The issues were argued with great tenacity and in considerable detail on each side.

The Expert Witnesses

10. As I have mentioned, Mr Jonathan Townend was the expert witness called by IDC and Dr James Irvine by Lenovo.
11. Both experts were well-qualified to give evidence in this case. Both were very knowledgeable, both were trying to assist the Court and I am grateful to both of them for their evidence. The differences between them were generated by the way in which each read Filiatrault and certain disputes on CGK which resulted. I will deal with the differences and disputes below. In addition, I must also

examine two important points concerning Dr Irvine's approach: the first is whether Dr Irvine approached his reading of Filiatrault with a mindset which reflects that of the skilled person at the priority date; the second is whether he took any steps to eliminate hindsight when he came to give his evidence on obviousness in his second report.

The Skilled Person

12. The parties were agreed that the skilled person/team would be a systems engineer working on HSUPA technologies and focussed on the MAC layer. Such a person would have a degree in electrical engineering or computer science and around 3-5 years of experience in the mobile communications industry. Although I recognise that it is very likely a team would have been involved, I will for convenience refer simply to the skilled person.
13. There was a minor dispute as to exactly what the skilled person would be working on. Mr Townend considered they would be working on HSUPA and developing products which incorporated HSUPA technology, as part of a wider team including other members with knowledge of the UMTS system. Dr Irvine was of the view that they would be working on developing the standards relating to HSUPA or implementing the requirements of those standards. I did not detect that this mini dispute made any difference to any of the issues, not least because it was common ground that, on either view, the skilled person would have a high-level understanding of how the various layers in UMTS worked, and a good understanding of the MAC functionality. He or she would be aware of the relevant 3GPP technical specifications and be able to refer to them for details as needed and these included the latest versions of the MAC specification (TS 25.321 v6.4.0) and the HSUPA stage 2 overall description (TS 25.309 v6.2.0).

Common General Knowledge

14. As is usual in these cases, there was very significant agreement as to the CGK, and much of it is familiar since 3G mobile telephone systems have been the subject of a number of judgments in the Patents Court. Much of what follows I have taken from the first expert report of Mr Townend, with a few additions from Dr Irvine's reports.
15. Mobile telecommunication systems are characterised by the support of user mobility, that is the ability for users to use the system while moving throughout the area covered by the system. In practical terms, user devices, such as mobile phones, are connected to the system using radio frequency communications. Other form factors (e.g. cellular data modems in computers, tablets, vehicles etc.) for present purposes operate in fundamentally the same way as a mobile phone.
16. The vast majority of commercial mobile telecommunication systems (from the 1980s to the present day) are so-called cellular radio networks, where the coverage area of the system is divided into a number of cells. Each cell is served by a base station transmitter and receiver (or transceiver) through an antenna system.

17. The base stations are connected to other pieces of inter-connected network equipment, such as controllers, switching centres, routers and subscriber databases, that work together to provide the telecommunication services to the end users of the system. The list of network elements, the functional split between them and how they communicate are all specific to particular network technologies, but can be split at a high level into the Radio Access Network (**RAN**) (which manages radio communications and the radio interface between base stations and mobile phones) and the Core Network (which manages services provided to the end users).
18. As mentioned above, the user devices communicate with the network equipment (and vice versa) using radio frequency communications. The communication from the user equipment is described as the **uplink**, and the communication to the user equipment is known as the **downlink**.
19. Uplink and downlink communications within a cell are usually distinguished because they are transmitted either:
 - (a) using different frequencies (i.e.: one set of frequencies are used in uplink, another in downlink, known as Frequency Division Duplex (**FDD**)); or
 - (b) at a different time, (known as Time Division Duplex (**TDD**)).
20. At the Priority Date communications from different users to a base station, and from different base stations, were distinguished using:
 - (a) different frequencies (i.e.: each user used a different uplink frequency, and different base stations used different downlink frequencies, Known as Frequency Division Multiple Access (**FDMA**));
 - (b) different timeslots (known as Time Division Multiple Access (**TDMA**); and/or
 - (c) different codes (known as Code Division Multiple Access (**CDMA**)).
21. Real life systems often use combinations of these technologies. For example, from before the Priority Date, and until the present, in the UK, 3G networks used codes (known as wideband CDMA, or **WCDMA**) to differentiate between communications from a base station to different users, and from different users to a base station, and different frequencies to distinguish uplink from downlink communications. In the UK, UMTS was therefore an **FDD WCDMA** system.

1G AND 2G

22. Mobile telecommunication system technology has evolved considerably from its introduction in the 1980s to the present day. While much of that evolution has been continuous, the industry has settled upon a classification of technologies into generations, with each incorporating a significant change in the radio interface technology.

23. The first generation of mobile telecommunication systems (**1G**) were introduced in the 1980s. These systems all used analogue technology to support voice telephony service, and were all incompatible with each other.
24. The second generation systems (**2G**) were introduced in the early 1990s. The most commonly adopted 2G technology worldwide was the GSM system, which was based upon industry standards published by ETSI. GSM used digital technology to support voice telephony, low speed circuit switched data and introduced the Short Message Service (**SMS**, now commonly referred to as text messaging).

UMTS - the Introduction of UMTS (Excluding HSDPA and HSUPA)

25. In the late 1990s, industry standards bodies turned their attention to 3G systems, which was designed to support the transmission of considerably higher data rates, leading to the support of a wide array of data applications, including multimedia.
26. Amongst the 3G systems developed was the UMTS system developed by the 3rd Generation Project Partnership (3GPP).
27. The first release of UMTS (Release 99) was "frozen" in 2000, and commercial networks launched in Europe in 2003. A release is a complete set of Technical Specifications which together specify a particular iteration of a mobile telecommunications system. Features are "frozen" to allow manufacturers and networks to implement the standard.
28. In an FDD WCDMA system, multiple adjacent cells use the same pair of frequencies (one for downlink and one for uplink), and use specially selected spreading and scrambling codes to distinguish transmissions from different sources (user devices in the uplink or cells in downlink). I do not need to explain in any detail how this is accomplished, however, there are a number of important consequences of this technology:
 - (a) Every other (i.e. unwanted) transmission appears as interference to the receiver. This means that it is critically important for all transmissions to use the bare minimum of transmit power required to be properly received and decoded. Therefore, one key feature of WCDMA is fast power control (discussed in more detail in paragraph 64 below).
 - (b) The capacity of the system is determined in part by the amount of interference being generated by other users, which is a direct result of their transmit power.
 - (c) When operating close to the boundary between two (or more) cells, a technique known as "soft handover" (or "macrodiversity") must be used, meaning that the user device is connected to two (or more) cells at once. This is different from previous systems, where devices are only connected to a single cell at a time.

29. When considering transmit power, it is important to understand that there is a relationship between data rate and required power. It is well known that for a given transmission technology (in particular, for a given set of modulation and coding techniques), higher data rates will require higher transmit power in order to be received and decoded with the same error rate.
30. The remainder of this section contains a description of the aspects of UMTS which are relevant to the dispute.

System Architecture

31. Figure 1 (adapted from Figure 5.2 of *WCDMA for UMTS* (3rd Edition) Holma and Toskala) below shows the overall system architecture of the UMTS system, and its connections to external networks.

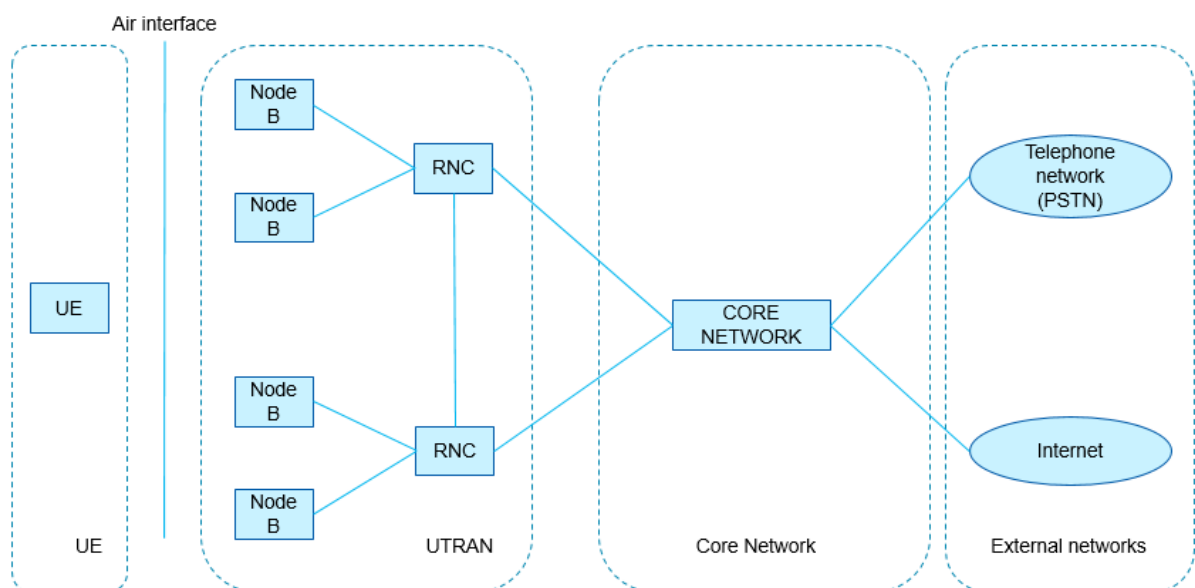


Figure 1: UMTS overall system architecture, functional network elements and logical entities

32. The UMTS system can be thought of as comprising the following functional network elements, visible in Figure 1 above:
- The User Equipment (**UE**), e.g. a mobile phone.
 - The UMTS Terrestrial Radio Access Network (**UTRAN**), which handles the radio-related functionality.
 - The Core Network, which provides switching and routing of calls and data connections to external networks (shown to the right of the Core Network).
33. Data sent may terminate beyond the Core Network (for example, a call made from a mobile phone to a landline will terminate in the Public Switched Telephone Network). In addition, data received may originate from beyond the Core Network (for example, a video streamed from the internet). Other than

this, only the UE and the UTRAN are relevant to the Patent, and so it is not necessary to consider further the functionality of the Core Network.

Logical Entities of the UTRAN

34. As can be seen in Figure 1 above the UTRAN comprises the following logical entities:
- (a) One or more Node Bs, each of which can provide the UTRAN with a connection to a UE over the air interface. Node Bs are the UMTS-specific version of the base stations described in paragraph 16 above.
 - (b) One or more Radio Network Controllers (**RNC**), which control the Node Bs. Each RNC is connected to one or more Node Bs, and connects those Node Bs to the Core Network (and, from there, to any external networks).
35. RNCs can serve several different roles simultaneously, and are named according to which role is relevant. Some of the documents in this case refer to an RNC being the Serving RNC (**SRNC**). This means it is the RNC which controls Radio Resource Control (RRC) signalling (the relevant functions being explained below) for a UE and maintains a connection to the Core Network for that UE. No other roles of the RNC are relevant.
36. Node Bs and RNCs are commonly referred to as *logical* entities because there is nothing in the standards which requires them to be two physically separate entities. However, for present purposes each logical entity can be considered a physically separate entity.

The Protocol Stack

37. Like many data networks, UMTS uses the Open Systems Interconnection (OSI) model (an architectural model for facilitating the interconnection of different types of data networks). This case is concerned with the bottom three layers described in the OSI model as Layer 3, Layer 2, and Layer 1. The higher layers are not relevant.
38. UMTS characterises the functions of the UE, the Node B and the SRNC by specifying the protocols in them. A protocol specifies a set of predefined rules that allow devices to communicate with each other.
39. The relevant UMTS protocols for this case are:
- (a) Layer 3: Radio Resource Control (RRC);
 - (b) Layer 2:
 - (i) Radio Link Control (RLC); and
 - (ii) Medium Access Control (MAC) (the functioning of MAC is of particular relevance to the Patent); and

- (c) Layer 1: Physical (PHY).
40. These protocols exist on both the UE side and the UTRAN side of the telecommunication system and, for example, the RRC functionality in the UE is frequently referred to as being in the UE's "RRC protocol entity". The functions of each of the protocols is explained further below.
 41. Each of the UE and the UTRAN will contain a "protocol stack" comprising all these protocols (i.e.: RRC, RLC, MAC and the physical layer). The relevant part of the UE's protocol stack comprises RRC, RLC, MAC and the physical layer. So far as the UTRAN in UMTS is concerned, a simplified view (which suffices for the Patent) is that the Node B's protocol stack comprises only the layer 1 (physical layer) functionality (i.e.: only a physical layer protocol entity), and the SRNC's protocol stack comprises only the layer 2 and 3 functionality (i.e.: RRC, RLC and MAC functionality or protocol entities). There is in fact some physical layer functionality in RNCs and some (very specific) MAC functionality in the Node B. Furthermore, in HSDPA and HSUPA some of the MAC layer functionality is moved to the Node B and I discuss this further below. In broad terms, the Node B and the SRNC combined provide the corresponding functionality (or protocol entities) for all the relevant protocols in the UE.
 42. Each protocol entity communicates with its "peer entity" on the other side. For example, the RRC protocol entity in the UE communicates with its peer RRC protocol entity in the SRNC. However, such communication between peer entities is indirect. The peer RRC protocol entities cannot communicate directly – information sent from the UE RRC protocol entity to the SRNC RRC protocol entity must make use of relevant lower-layer protocol entities in the protocol stack both at the UE and the SRNC, and also in any intervening logical entity (here, the Node B). The peer entities are not concerned with the operation of the lower-layer protocols, nor do the lower layer protocols have to know the contents of the higher-layer messages they are transmitting. Figure 2 below (Adapted from Figure 11 of TS 25.301 V6.2.0.) shows the protocols in the UE, the Node B and the SRNC (the physical layer functionality in the SNRC is to deal with the user device being connected to two or more cells at once during 'macrodiversity' as mentioned above). Protocol entities connected with horizontal lines communicate with each other, using lower protocol entities to carry the data.




Figure 2: Data flows through protocols and between logical entities

43. There are a few points to note about this scheme:
- (a) Different data will have different originating points and different terminating points. For example:
 - (i) An email will go from a UE to an endpoint outside the network;
 - (ii) RRC control messages configuring a service will need to go from the SRNC to the Node B; and
 - (iii) Layer 1 control messages such as power control commands will need to go from the Node B to the UE.
 - (b) Not all logical entities (UE, Node B, etc.) will need to be able to interpret all the data sent over the network (e.g.: a Node B does not need to understand the user data it is forwarding from a UE to an external network, or RRC control messages from the UE to the SRNC).
44. Flows of data between these protocols are described as belonging to different **radio bearers** or **channels** depending on which protocol entities they are passing between.
45. Data which is sent over the UMTS system will either be:
- (a) User data (e.g.: a webpage) – also referred to as user plane data; or
 - (b) Control data (e.g.: instructions to set up a connection to allow the transfer of user data) – also referred to as control plane data or control signalling.
46. Figure 3 below shows relevant protocols, radio bearers and channels. The left hand side of Figure 3 shows flows of control data through the protocol stack. The right hand side shows flows of user data through the protocol stack. The terms used in Figure 3 are explained in paragraphs 47 and 48 below.

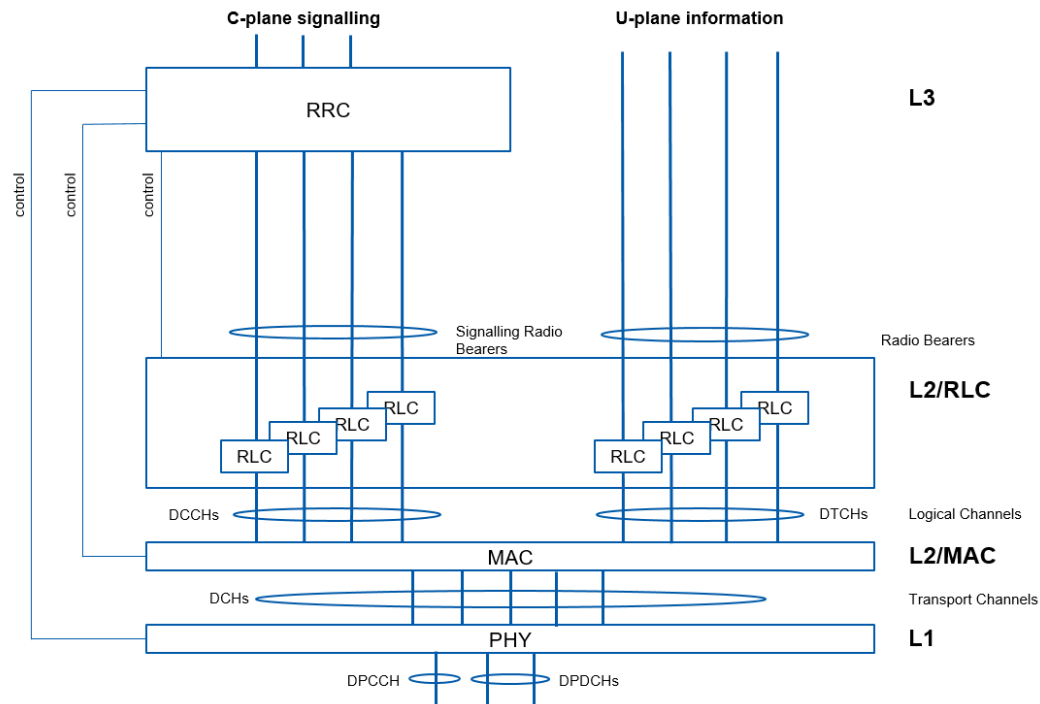


Figure 3: Layers, Protocols, Radio Bearers and Dedicated Channels of UMTS
(Adapted from Figure 2 of TS 25.301 V6.2.0.)

47. In particular, as can be seen in Figure 3 (reading from top to bottom, and left to right):
- (a) A "**Signalling Radio Bearer**" is a flow of control data (described as C-plane signalling) between the RRC and RLC protocol entities.
 - (b) A "**Radio Bearer**" is a flow of user data (described as U-plane information) into the RLC protocol entity.
 - (c) A "**logical channel**" is a flow of data (either user (referred to as "traffic") or control) between RLC and MAC protocol entities. In HSUPA, as I explain below, a conceptual grouping of multiple logical channels is referred to as a MAC-d flow.
 - (d) A "**transport channel**" is a flow of data (user, control or both (as a result of C/T multiplexing – see below) between MAC and the physical layer protocol entities.
 - (e) A "**physical channel**" is the signal sent over the physical layer between the physical layer protocol entities in the UE and the Node B (in either direction).
48. In the UE, for uplink communications, the protocols perform the following functions:
- (a) **RRC** is a control-only protocol. It does not carry user data (except for SMS), but is responsible for setting up, modifying, and releasing Layer 2 (RLC and MAC) and Layer 1 (physical layer) protocol entities required to carry user data. It is responsible for controlling the Quality

of Service (QoS – see further below) provided by lower layers in response to requests from higher layers in the protocol stack.

- (b) **RLC** provides services for user data (i.e.: Radio Bearers) and control data (i.e.: Signalling Radio Bearers) which it receives from higher layers in the protocol stack, and then passes that data on to MAC on logical channels. Its precise function is not relevant to the Patent, but it processes each flow of data it receives (from a Radio Bearer, usually by adding a header), and then maps it onto its own logical channel (i.e.: if it receives three Bearers, it will output three logical channels).
 - (c) **MAC** is very relevant for the subject matter of the Patent. It is responsible for the selection and transmission of data (discussed below), and for mapping logical channels onto transport channels. Several logical channels may be mapped onto a single transport channel if they can be treated in the same way by the physical layer (i.e.: they are carrying a similar type of user or control data, with similar QoS). If more than one logical channel is mapped onto a transport channel, the logical channels are described as being "multiplexed" onto the transport channel. There are a number of specific MAC protocol entities to deal with specific transport channels, but in UMTS only the MAC protocol entity which deals with dedicated transport channels (the MAC-d) is relevant. The MAC-d is responsible for mapping the **Dedicated Control Channel(s)** (DCCH) (i.e.: the logical control channel(s)) and the **Dedicated Traffic Channel(s)** (DTCH) (i.e.: the logical user data channel(s)) to the **Dedicated (transport) Channel(s)** (DCH). As Figure 3 above only shows dedicated logical channels being mapped onto dedicated transport channels, the MAC layer shown in it is effectively the MAC-d. As will be seen below, in HSUPA, further relevant MAC entities are introduced – the MAC-e and MAC-es.
 - (d) The **physical layer** processes data received from MAC on transport channels in accordance with instructions received via RRC signalling in order to transmit it over the air interface to the Node B.
49. In the UTRAN, for uplink data each protocol performs the reverse process of that performed in the UE. For RRC, this means that the SRNC communicates with the Node B in order to set up etc. the lower layer entities, and also provides various instructions to the UE. For downlink data, much of the functionality is swapped between the UE and the UTRAN, but it is not necessary to consider those aspects for the purposes of this case.
50. In UMTS, the flows of data discussed above are transmitted between layers in "packets". A packet of data, which a lower layer in the protocol stack receives from the layer above, is known as a Service Data Unit (**SDU**) for the receiving layer. Generally, each layer will at least add its own header to an SDU before passing it on to a lower layer. A packet of data (including header) which a higher layer passes to a lower layer is known as a Protocol Data Unit (**PDU**). The part of the new PDU which contains data from the layer above (i.e.: the lower layer's SDU) is also known as the PDU's "payload".

51. Figure 4 below illustrates the mapping of data into the data units of two layers in a given network element, namely the upper (N)-layer and the lower (N-1)-layer. The "payload" of (N-1)-PDU is (N-1)-SDU. The figure describes the header as being "Protocol control information" or "PCI".

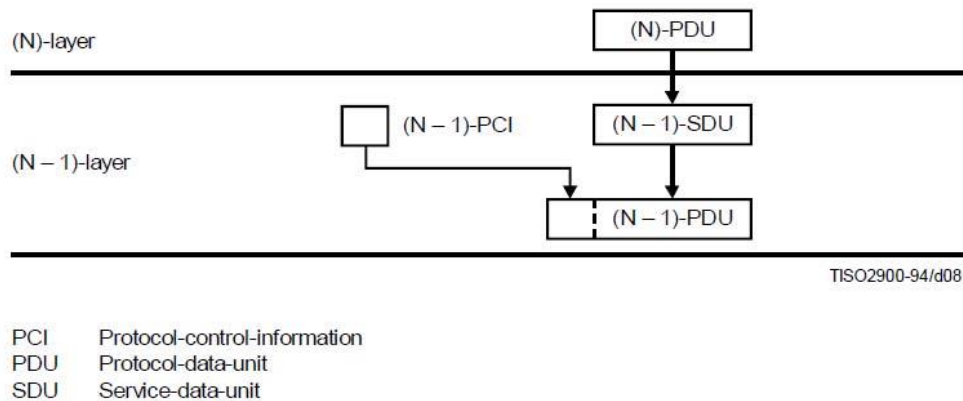


Figure 4: PDU construction (Taken from Figure 9 of the International Telecommunication Union – Telecommunications Standardization Sector (ITU-T) X.200 standard (07/94), commonly known as the OSI 7-layer model).

52. Whilst a lower layer will process a higher layer PDU (for example, by adding a header and passing it on to a lower layer), it will not interpret the contents of a higher layer PDU. The higher layer PDU is said to be "transparent" to the lower layer. This is because ultimately the critical communication is only between peer entities within the protocol stack as explained in paragraph 42 above.
53. Thus, the RLC receives RLC SDUs and outputs RLC PDUs. The MAC receives RLC PDUs (as MAC SDUs) and outputs MAC PDUs. However, the physical layer receives a MAC PDU as a **Transport Block (TB)** (effectively the physical layer SDU).
54. The relevant logical, transport and physical channels are all **dedicated channels**. That is to say they are dedicated to a single user (in other words, any data on that channel must be from or for that user, and no further addressing is required). The relevant UMTS channels are as follows:

	<u>User Data</u>	<u>Layer 3 Control Data</u>	<u>Physical Layer Control Data</u>
Bearer	Radio Bearer (RB)	Signalling Radio Bearer (SRB)	
Logical Channel	Dedicated Traffic Channel (DTCH)	Dedicated Control Channel (DCCH)	
Transport Channel	Dedicated Channel (DCH)		
Physical Channel	Dedicated Physical Data Channel (DPDCH)		

Table 1: Relevant UMTS channels

55. The overall scheme to bear in mind is that:
- (a) RLC maps higher layer control plane data (Signalling Radio Bearers - SRB) to one or more (dedicated) logical control channels (DCCH).
 - (b) RLC also maps higher layer user plane data (Radio Bearers - RB) to one or more (dedicated) logical (user) traffic channels (DTCH).
 - (c) MAC maps (dedicated) control logical channels (DCCH) and (dedicated) traffic logical channels (DTCH) to (dedicated) transport channels (DCH).
 - (d) The physical layer multiplexes (dedicated) transport channels (DCH) which are to be transmitted using the same type of (dedicated) physical channel (DPDCH) into a Coded Composite Transport Channel of a particular type, so that they can be transmitted using the same dedicated physical data channels (DPDCH). There can be more than one dedicated physical data channel (DPDCH) if required due to a sufficiently high data rate, or if there is more than one type of Coded Composite Transport Channel.
 - (e) Control information which originates in the physical layer will be carried on the dedicated physical control channel (DPCCH).

Quality of Service (QoS)

56. Quality of Service (QoS) is used within UMTS to describe the required service attributes which a communication (e.g.: a phone call, streaming a video, downloading a webpage, uploading an email) must have.
57. The standards set out various ways of defining QoS (including priority of the data, traffic classes or QoS classes, required bit rate and residual bit error rate). In particular, RRC can control the scheduling of uplink data by giving each logical channel a priority between 1 (highest priority) and 8 (lowest priority). Other than this, how priority is dealt with is not particularly relevant.

58. There are four QoS traffic classes in UMTS:
- (a) conversational class;
 - (b) streaming class;
 - (c) interactive class; and
 - (d) background class.
59. Delay is the main distinguishing factor between the four QoS traffic classes, although there are also differences between the error correction mechanisms which are used as a result, and consequently on the transmit power which must be used for a given bit error rate. Furthermore, conversational class and streaming class traffic (which are both real time classes) may have a guaranteed bit rate specified.
60. Although there is no strict relationship between a QoS traffic class and the type of data it can carry, there are typical types of data for which a QoS traffic class may be used.
61. Table 2 below sets out the fundamental characteristics of each QoS traffic class, and the typical use case (where RT indicates ‘real time’ services).

Traffic class	Conversational class conversational RT	Streaming class streaming RT	Interactive class Interactive best effort	Background Background best effort
Fundamental characteristics	<ul style="list-style-type: none"> - Preserve time relation (variation) between information entities of the stream Conversational pattern (stringent and low delay) 	<ul style="list-style-type: none"> - Preserve time relation (variation) between information entities of the stream 	<ul style="list-style-type: none"> - Request response pattern - Preserve payload content 	<ul style="list-style-type: none"> - Destination is not expecting the data within a certain time - Preserve payload content
Example of the application	<ul style="list-style-type: none"> - voice 	<ul style="list-style-type: none"> - streaming video 	<ul style="list-style-type: none"> - Web browsing 	<ul style="list-style-type: none"> - background download of emails

Table 2: QoS traffic classes and typical use cases (Table 1 of TS 23.107 V5.13.0)

62. Within each class priority levels may be specified, which the relevant network elements can use to prioritise data flows with a high priority value over data flows with a lower priority value. Alongside these priorities, QoS will also take into consideration required data rate and residual bit error rate.
63. RRC is responsible for configuring the radio interface between the UTRAN and the UE in accordance with the required QoS traffic class (and other attributes) for a connection. This includes configuring RLC, MAC (including logical channel priorities), transport channel(s) and the physical layer.

Power control and the DPCCH

64. As discussed in paragraph 28(a) above, in any system using CDMA technology, fast power control is critically important. Power control is especially important in the uplink where it ensures that the UE transmits at the minimum power required so that the Node B receives the UE's transmissions with the desired

block error rate (BLER) at the receiver. Using minimum power means that the transmission from one UE is less likely to block or unnecessarily interfere with transmissions from other UEs in the same cell.

65. When the UE is configured to transmit using dedicated channels, the SRNC estimates a signal to interference power ratio (SIR) target for the UE's transmissions required to ensure the desired block error rate at the Node B and sends this signal to interference power ratio target to the Node B. The Node B instructs the UE to increase or decrease transmit power to maintain the signal to interference power ratio as close as possible to the target.
66. The power control commands are sent to the UE, and so the UE can change transmit power level, every 0.667ms (or 15 times every 10ms).
67. While the power control commands manage the overall transmit power of the UE, the UE can transmit multiple physical channels at different power levels in order to support their individual QoS requirements.
68. This is accomplished by establishing a single physical channel as a baseline and then defining all other channel powers relative to that baseline. The power control commands then cause the power for every channel to increase or decrease together, while preserving the proportional relationship between their power levels.
69. In UMTS, the dedicated physical control channel (DPCCH), which carries physical layer control information, is used as the baseline power level relative to which all other dedicated channel powers in the UE are set.

Multiplexing

70. UMTS allows a user to use several "applications" at the same time. For example, a user might be simultaneously using three applications: talking on the phone, browsing the web and sending an email.
71. Each application in use will require logical traffic channels (DTCH) to send the user data, and logical control channels (DCCH) to manage the channels carrying the user data. So in the example above:
 - (a) a user would send three flows of user data (or Radio Bearers) being the user data for each of the phone call, web browsing and sending of an email to the RLC;
 - (b) RRC would need to:
 - (i) configure higher layer control channels (i.e.: Signalling Radio Bearers) to control the dedicated logical, transport and physical channels (i.e. DTCH(s), DCH(s) and DPDCH(s)) required for each flow of user data, to send the user data over the air interface; and
 - (ii) configure logical, transport and physical control channels to allow control signalling to reach the appropriate protocol entity

(e.g.: a peer RLC entity) in the appropriate logical entity (e.g.: in the SRNC).

72. Typically, RRC will configure three logical control channels (DCCH) and part of this configuration will be that they are multiplexed onto a single transport channel (DCH) (rather than configuring one transport channel per logical control channel). Indeed, it might also be possible to multiplex one of the logical traffic channels (DTCH) onto the same transport channel (DCH). The multiplexing of different logical channels onto the same transport channel in the MAC is known as "C/T multiplexing". When two (or more) logical channels are C/T multiplexed a MAC header comprising a C/T field (a number identifying the logical channel to which the MAC SDU belongs) is added to each PDU individually based on its logical channel, so that whilst the physical channel will not distinguish between packets from different logical channels, the peer MAC protocol entity (for example, in the SRNC) can de-multiplex the received transport channel, and send packets to the correct logical channel (whether control (DCCH) or traffic (DTCH)). Each MAC PDU is still dealt with as an individual Transport Block by the physical layer.
73. The physical layer can also multiplex transport channels into a single Coded Composite Transport Channel for further processing by the physical layer. For example, the transport channels carrying the user data for the three applications referred to above, being the phone call, the web browsing, and the email might be multiplexed in the physical layer into a single Coded Composite Transport Channel if they are going to be sent over a single dedicated physical channel (or multiple physical channels of the same type). In contrast to C/T multiplexing, physical layer multiplexing allows for different levels of error protection between the different transport channels being multiplexed because the physical layer can distinguish between packets based on the transport channels they arrive on, and process them differently.

Scheduling & Statistical Multiplexing

74. As noted above, scheduling is the process by which radio resources are allocated between different users in the system, a process which should ensure that users are able to maintain a certain QoS. Scheduling is very straightforward for constant bit rate services, but most services have a variable bit rate. The variability means that allocating resources based on the peak data rate would guarantee access but would be very wasteful. To avoid this, scheduling usually involves the process of statistical multiplexing, whereby users' resources are allocated based on a specific average data rate, plus a margin for safety. This increases the overall efficiency of the system with very little impact on individual users. Obviously, this 'multiplexing' of signals from various users is not the multiplexing of data in the MAC-e/-es entity described below.
75. Dr Irvine provided a simple illustration (Fig. 5 below), in which four users (blue, orange, grey and yellow) are to be scheduled, their maximum resource requirements being 2, 7, 3 and 4 resource units respectively. Their cumulative maximum resource requirement is 16 and in UMTS this would be an overall power level. The average cumulative resource use is 8.58. Adding half as much again would yield a resource requirement of 13 (as shown on the right-hand

side). Thus, using statistical multiplexing reduces the resource requirement from 16 to 13 with no loss of performance. There remains a significant amount of unallocated resource (shown in green), but this cannot be assigned to users in the normal way, otherwise the average load would go up and the margin of safety would be reduced, increasing the chance that the instantaneous load would exceed 13 and that when an assigned user tried to transmit they would not be able to. These unallocated resources can be allocated on a short term basis when the actual use of the allocated users is known.




Figure 5: Statistical Multiplexing

Sending Data

- 76. Data is transmitted within defined periods known as Transmission Time Intervals (**TTI**). Every TTI (for example, every 20ms) a UE will need to determine which data to send to the Node B, and send it.
- 77. In UMTS, MAC is responsible for determining which data to send from a UE to a Node B. However, how it does so is tightly controlled.
- 78. Each TTI a UE can send a single Transport Format Combination (**TFC**) to the Node B. Strictly a TFC is a combination of Transport Formats, each of which implicitly specifies the number of Transport Blocks which can be sent in a TTI from a given transport channel (by specifying the size of the Transport Block and total size of the Transport Format), and how those blocks are to be processed. This can be seen in Table 3 below, where each TFC (e.g.: TFC₅) specifies a TF for each channel (TF_{1,2} and TF_{2,2}). However, this level of detail is not relevant to the Patent. A TFC specifies the size of packet from each transport channel (Transport Block size (TB Size in Table 3 below)) and the total size of data sent for each transport channel (i.e.: total bits per TTI). It therefore implicitly specifies the number of packets of data (Transport Blocks) from each transport channel which can be sent. The UE can only select a TFC from the set of TFCs (the Transport Format Combination Set, **TFC Set or TFCS**) configured by RRC (partly because the physical layer is configured only to process data of certain predetermined sizes).
- 79. For example, if a user is browsing a web page, there might be two transport channels configured for uplink data:

- (a) Transport channel 1 (DCH1): Control signalling.
- (b) Transport channel 2 (DCH2): User data (i.e.: the user interacting with the webpage).

80. The TFC Set will specify how many PDUs for each transport channel can be sent. For example, it might look like this (in a UMTS system):

Transport Format Combination	Transport Channel 1 (DCH1)		Transport Channel 2 (DCH2)		Total Bits
	Transport Format	Transport Blocks/TB Size/ Total Bits per TTI	Transport Format	Transport Blocks/TB Size/ Total Bits per TTI	
TFC ₁	TF _{1,1}	0/148/0	TF _{2,1}	0/336/0	0
TFC ₂	TF _{1,1}	0/148/0	TF _{2,2}	1/336/336	336
TFC ₃	TF _{1,1}	0/148/0	TF _{2,3}	2/336/672	672
TFC ₄	TF _{1,2}	1/148/148	TF _{2,1}	0/336/0	148
TFC ₅	TF _{1,2}	1/148/148	TF _{2,2}	1/336/336	484
TFC ₆	TF _{1,3}	2/148/296	TF _{2,1}	0/336/0	296
TFC ₇	TF _{1,3}	2/148/296	TF _{2,2}	1/336/336	632
TFC ₈	TF _{1,4}	3/148/444	TF _{2,1}	0/336/0	444

Table 3: Example Transport Format Combination Set

81. Each TTI, the UE selects a TFC from the set of valid TFCs from the TFC Set which satisfies the three following criteria in the order in which they are listed:

- (a) No other TFC shall allow the transmission of more highest priority data than the chosen TFC.
- (b) No other TFC shall allow the transmission of more data from the next lower priority logical channels (a test which is applied recursively for the remaining priority levels).
- (c) No other TFC shall have a lower bit rate than the chosen TFC.

82. The technical specification published by 3GPP dealing with MAC is TS 25.321. The latest version at the Priority Date (TS 25.321 V6.4.0, s11.4) specifies exactly what is meant by a "valid" TFC in some detail. For present purposes, the only relevant criterion is the state of the TFC, which depends on whether use of a TFC is likely to exceed the allowed transmission power. In particular:

- (a) If a TFC is not likely to exceed the allowed power, it is in "Supported state".
- (b) If the estimated required UE transmit power for the TFC is greater than the maximum UE power for a given period of time, it is in "Excess-power state".
- (c) If a TFC is in "Excess-power state" for a given period of time, it is considered to be in "Blocked state".

- (d) I interpolate here that Mr Townend's evidence to the effect set out above was a routine recitation of what is set out in TS 25.321. However, in his second report, Dr Irvine was keen to add a qualification. In relation to (a) and (b) above, Dr Irvine added the qualification that a UE is not allowed to transmit above the maximum allowed transmission power (or the maximum output power of the UE), defined as the "maximum UE TX power" in TS 25.331 v6.5.0. He said this applies to all TFC states including the "excess-power state" TFCs: in the "excess-power state" case, a UE would send the TFC at less than the usual power so that it does not exceed the maximum UE TX power. This qualification (and similar points) generated most of the CGK disputes I discuss at paragraph 201 below. It is also the reason behind the experts' disagreement as to what Filiatrault disclosed.
83. A UE:
- (a) can use TFCs in "Supported state";
 - (b) cannot use TFCs in "Blocked state"; and
 - (c) can decide whether to use TFCs in "Excess-power state".
84. That is to say:
- (a) TFCs in "Supported state" are "valid" TFCs;
 - (b) the UE can decide whether TFCs in "Excess-power state" are "valid" TFCs; and
 - (c) TFCs in "Blocked state" are not "valid" TFCs.
85. The procedure of determining the state of a TFC is also referred to as "TFC restriction".
86. In summary, the overall scheme is that the UE starts with the set of TFCs in the TFC Set, those which require too much power are excluded, and then of the remainder the UE selects the TFC which maximises the transmission of highest priority data, and subject to that, lower priority data recursively, without selecting a TFC which is larger than necessary (for example, a TFC which allows for the transmission of more PDUs from a given logical channel than the UE has to transmit).
87. There are two consequences arising from this scheme:
- (a) It maximises the transmission of the highest priority data, even if this means the total data sent is less than might otherwise be sent.
 - (b) The TFC size is always the same size as the sum of Transport Blocks included in it.

HSDPA - Introduction

88. HSDPA was the key new UTRAN feature included in Release 5. It was introduced with the objectives of improving user experience by enhancing system capacity and efficiency. In particular, HSDPA aimed to provide higher data rates for downlink user data sent from a Node B to a UE to support particular services such as web browsing and streaming.
89. By the Priority Date, the standardisation of HSDPA was sufficiently complete that HSDPA was already part of commercially available products.
90. HSDPA introduced various techniques into UMTS, including:
- (a) A shorter (2ms) TTI (compared to the previous shortest TTI of 10ms);
 - (b) A more complicated, but more efficient, form of error correction (known as Type II HARQ) (discussed further below);
 - (c) Node B controlled scheduling (i.e.: the Node B was responsible for some of the scheduling decisions which were previously made in the SRNC). By locating the decision making in the Node B, physically close to the air interface, this meant that HSDPA was more responsive than the legacy channels and thus facilitated the use of the shorter TTI and Type II HARQ; and
 - (d) Higher order modulation.
91. In particular Release 5 introduced a new high speed packet access channel on the downlink, the High Speed Downlink Shared Channel (**HS-DSCH**). This channel was shared between users and, with control being devolved from the RNC to the Node B, the system was much more responsive in scheduling which resulted in higher efficiency in data transmission.
92. Fast scheduling at the Node B allowed the system to make use of the spare resources which cannot be allocated in the normal way due to the requirement to allow a margin of safety in resource allocation for statistical multiplexing on a TTI by TTI basis. Thus, the spare resources indicated in green on the right-hand side of Figure 5 above can be allocated on a short term basis.

Type II Hybrid ARQ

93. Type II Hybrid ARQ (**HARQ** for short) was introduced in HSDPA, and is a combination of methods for correcting errors in a transmission. HARQ combines:
- (a) The use of Cyclic Redundancy Check bits. This allows the receiver to make a quick initial determination of whether a received packet is error free or not;
 - (b) The use of an Automatic Repeat Request (ARQ) mechanism. ARQ uses Acknowledgements (ACK) to indicate that a packet has been correctly

received, and a Negative Acknowledgement (NACK) to indicate that an errored packet has been received, and a retransmission is required;

- (c) The use of Error Correction Coding (for example, Forward Error Correction coding), which allows the receiver to "correct" some errors in a transmission (the use of CRC bits, ARQ and ECC together is commonly referred to as Type 1 HARQ, which has been used since GSM); and
 - (d) The combination of different (re)transmissions, instead of the discarding of errored transmissions.
94. The particular version of HARQ specified for HSDPA is called "stop-and-wait", which involves the transmitter waiting to send a second block of data until it has received an ACK for the previous block (or until a maximum number of retransmissions has been reached, or, perhaps, a timer has expired). In order to prevent the obvious delays that this would cause, HSDPA calls for multiple (up to 8) HARQ processes. This allows the HARQ processes to operate in sequence and cyclically, such that in the first TTI, HARQ process 1 transmits, then in subsequent TTIs, HARQ processes 2 to 8 transmit whilst HARQ process 1 waits for its ACK or NACK, then in TTI 9, HARQ process 1 can either transmit the next packet (i.e.: the one after the packet sent on HARQ process 8) or retransmit the packet it transmitted last time (in case there has been a NACK (or no ACK)). HARQ in HSDPA is not required to operate in this way – each TTI the Node B can decide which HARQ process to transmit. HARQ in HSUPA is, however, required to transmit each HARQ process in turn.

HSUPA

95. At the Priority Date the skilled addressee would be aware that Release 6 of UMTS was under development and that the main UTRAN work item within it would be HSUPA. I set out a summary of the skilled addressee's common general knowledge regarding HSUPA below.
96. The aim of HSUPA was to provide similar enhancements to the uplink as HSDPA provided to the downlink, in particular in terms of user experience (throughput and delay) and/or capacity. The skilled addressee would understand that there was a desire to keep HSUPA as simple as possible, taking into account the level of system performance.
97. HSUPA aimed to do this by introducing:
- (a) new channels (comprising a new transport channel, mapped to a new physical channel) for sending user data on the "enhanced uplink";
 - (b) new protocol entities in the MAC to allow access to the new channels in the UE (and new peer MAC protocol entities in the Node B and the RNC to join the new channels back to the pre-existing system); and
 - (c) new enhanced uplink functionality, including Node B controlled scheduling (within limits set by the RNC), HARQ, the ability to choose

a shorter, 2ms, TTI (instead of the longer, required, 10ms TTI), and non-scheduled transmissions (configured by the SRNC). Non-scheduled transmissions allow the SRNC to specify separate Non-Scheduled Grants for specific services, which are not varied by Node B controlled scheduling.

98. Node B controlled scheduling, HARQ, and a 2ms TTI were all functionality which the skilled addressee would have known were included in HSDPA, although the skilled addressee would be aware that their implementation in the uplink would entail some significant differences to their implementation in the downlink, including that:
- (a) As discussed above, UE power control by the network is very important in the uplink, to avoid excessive interference. In addition, UE power control has an important impact on battery life. These issues are much less acute in the downlink direction.
 - (b) In HSDPA the scheduling function (i.e. the allocation of resources based on data to transmit to a given UE, and the use of those resources to transmit data to a given UE) is all located in the Node B, whereas in HSUPA it must be split between the Node B (responsible for resource allocation) and the UE (responsible for making use of the resource allocation to send data). In each case, the UE provides information for the scheduling decisions made in the Node B (downlink channel quality information in the case of HSDPA, and buffer occupancy in the case of HSUPA).
99. The skilled addressee would also be aware that in the UMTS standards (and in GSM) the behaviour of the UE is very tightly specified, while the behaviour of the network is, to the greatest extent possible, left to implementation. Thus the extent of prescription in the standards would have been expected to be very different between HSDPA and HSUPA.
100. The new channels introduced by HSUPA were the **Enhanced Dedicated Channel (E-DCH)**, one per user, a transport channel, mapped to a new physical channel (**E-DPDCH**, or **Enhanced Dedicated Physical Data Channel**), plus various other channels to support HARQ and indicate the scheduled grants which could be used.
101. Data which was to be sent over these channels would be forwarded from the MAC-d to a new protocol entity in the UE, the **MAC-e/es**.
102. The MAC-e/es in the UE would be responsible for determining what data (both user and higher-layer control data) would be sent over the new channels each TTI. As in uplink TFC selection by the MAC-d in UMTS, how the MAC-e/es selects data for transmission is tightly controlled. The MAC-e/es would multiplex the selected data in order of its logical channel priorities, as in UMTS. Data from each logical channel would be concatenated (i.e. inserted in order) into a separate **MAC-es PDU** and one or more MAC-es PDUs would be multiplexed into a single **MAC-e PDU**, which would be carried on the enhanced uplink dedicated physical channel (E-DPDCH) as a single transport block, the

Enhanced Dedicated Channel Transport Format Combination (the **E-TFC**, see further below). Precisely how the multiplexing process into the MAC-e PDU works is a matter in dispute. I return to this issue in the context of Filiautrault.

103. The single MAC-e PDU could contain data from a number of logical channels (for example, it could contain data for a video call, a webpage, and an email, as well as associated higher layer control information). Some of these logical channels might be grouped together within the MAC-e/es in the UE as **MAC-d flows** (a conceptual grouping of logical channels with similar QoS requirements). This is set out at a high level in Figure 6 (see below).
104. This would therefore require multiplexing of MAC-d PDUs from different logical channels into a single MAC-e PDU in the MAC.
105. The peer MAC-e protocol entity in the Node B determines whether a HARQ retransmission of the MAC-e PDU is required, and if not it de-multiplexes the MAC-e PDU into its constituent MAC-d flows. The peer MAC-es protocol entity in the SRNC combines data from different Node Bs to which the UE is sending data (due to soft handover), separates the MAC-d flows into their different logical channels (if the MAC-d flow is made up of more than one logical channel), and forwards them to the MAC-d which forwards them to the RLC.
106. The skilled addressee would know that access to the enhanced uplink channel (E-DCH) was to be controlled by grants which limited the amount of data that could be sent in a TTI. There would be scheduling grants from the Node B that could vary the amount of scheduled data that could be sent on a per-TTI basis. The UE receives various scheduling grant instructions over time and stores the cumulative outcome of those instructions as the Serving Grant variable. There would also be slower-to-change “non-scheduled” grants by the SRNC for specific services (for example, higher layer control signalling like Signalling Radio Bearers (SRB), and Guaranteed Bit Rate (GBR) services such as Voice over IP (**VoIP**)). A MAC-d flow can only contain scheduled or non-scheduled data (but not both) and this would be configured by RRC.
107. A UE stores data to send on the uplink in a buffer. The UE would also need to send Scheduling Information (**SI**) to the Node B to allow the Node B to make scheduling decisions. The UE would send Scheduling Information on the E-DPDCH. The Scheduling Information would include information on the UE's buffer status (in particular, the highest priority channel for which it had data to send, the amount of data in the buffer for that highest priority channel and the UE's total buffer status). The Scheduling Information would also include an estimate of the available power (relative to the DPCCH), in order to allow the Node B to make scheduling decisions. Periodic reporting of Scheduling Information was described, although whether there should be event-based triggering, and how the Scheduling Information would be included in the MAC-e PDU, were both "FFS" (i.e. For Further Study).
108. The UE maintains a Serving Grant (**SG**) parameter (which is dependent upon scheduling grant messages received from the Node B – known as Absolute Grants or Relative Grants), for controlling the transmission of scheduled data.

The Serving Grant is the maximum power that can be used by the Enhanced Dedicated Physical Data Channel (**E-DPDCH**), expressed as a ratio to the power at which the existing UMTS Dedicated Physical Control Channel (**DPCCH**) is transmitted.

- 109. In addition, as mentioned above, the UE would receive Non-Scheduled Grants in respect of specific MAC-d flows (for services like VoIP or higher layer control information) from the SRNC.

Multiplexing of data in HSUPA

- 110. The MAC layer multiplexing in HSUPA would therefore be more complicated than the legacy C/T multiplexing as MAC-d PDUs from different logical channels with potentially different data rates and QoS requirements would need to be incorporated into a single MAC-e PDU in accordance with the various grants applicable in that TTI.
- 111. Figure 6 below shows Figure 7.2.1-2 of the **Prior HSUPA Stage 2 Specification** (i.e. TS 25.309 V6.2.0, the latest version (V6.2.0) of the stage 2 Technical Specification of HSUPA (TS 25.309) published by 3GPP at the Priority Date). This figure is also present in Filiatrault and it explains how data which is to be transmitted using HSUPA will be multiplexed into a MAC-e PDU.

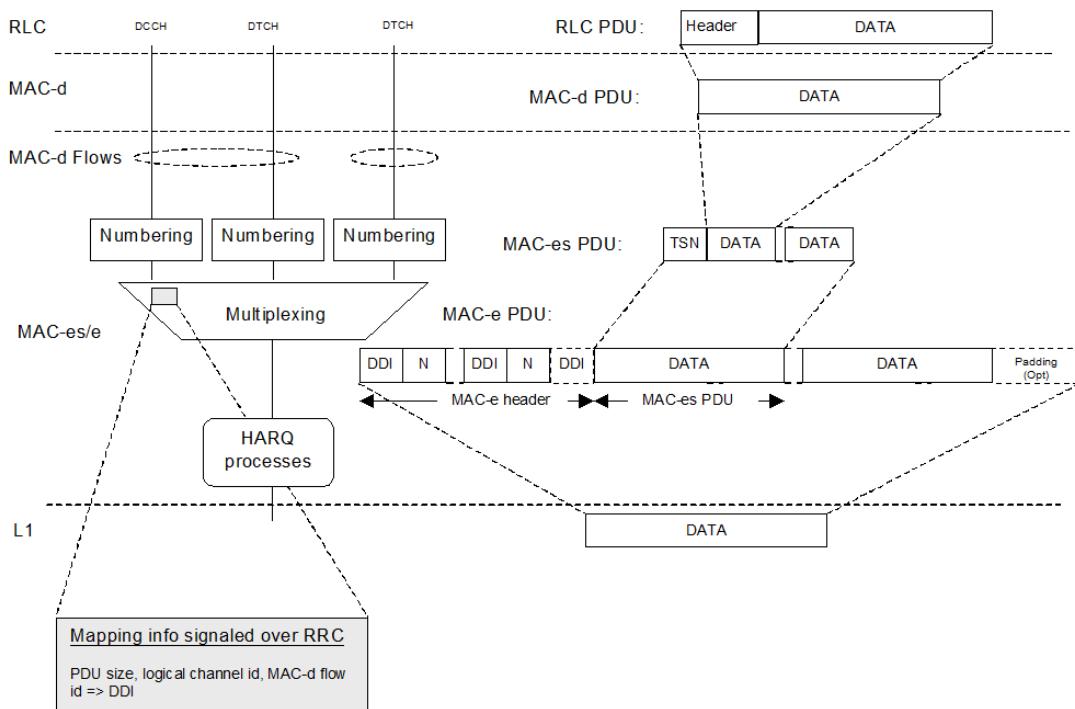


Figure 6: Overview of data multiplexing in HSUPA

- 112. The left-hand side of Figure 6 shows the functional view of the MAC protocol entities in the UE (together with (at the top) the RLC protocol entities feeding into the MAC-d, and (at the bottom) the MAC-es/e feeding into the physical layer (L1)).

113. The right-hand side of Figure 6 shows PDU construction at each layer, by the addition of header elements for a layer to the PDU from the layer above, and for the MAC-e layer, the multiplexing of different MAC-es PDUs into a MAC-e PDU.
114. The skilled addressee would understand that although there is correlation between the two sides of Figure 6, the right-hand side does not show PDU construction for all the data flows shown in the functional view on the left, as what is shown is sufficient for the skilled addressee to understand how PDU construction as a whole would work. In particular:
- (a) At the RLC layer, the functional view (on the left-hand side) shows three logical channels (a logical control channel (the DCCH), and two logical traffic channels (the DTCHs)), but the PDU construction (on the right-hand side) only shows one RLC PDU (comprising a header and data), rather than three.
 - (b) This RLC PDU is carried down into the MAC-d: the single RLC PDU is shown as a single MAC-d PDU at the MAC-d layer.
 - (c) On the left-hand side the functional view shows the three logical channels as two MAC-d flows indicated by dashed ovals around the relevant logical channels: the one on the left comprising a DCCH and a DTCH; the one on the right a single DTCH. Each logical channel goes to a separate numbering entity. On the right-hand side of the Figure, the PDU construction shows a single MAC-es PDU (rather than three) comprising a TSN (the Transmission Sequence Number of the MAC-es PDU) and (at least) two MAC-d PDUs.
 - (d) The multiplexing entity in the functional view shows the three logical channels being multiplexed into a single data stream. In contrast, the PDU construction only explicitly shows two MAC-es PDUs being combined into a MAC-e PDU.
115. Finally, the right-hand side of Figure 6 shows a single data block being passed to L1 (and the skilled addressee would understand that the functional view on the left-hand side shows the output of the multiplexing function (i.e.: the MAC-e PDU) being passed to the HARQ process entity for transmission).
116. Figure 6 shows that by the time MAC-d flow data arrives at the physical layer, it would have various headers added to it (specifically, a MAC-es header for each MAC-es PDU and a MAC-e header), and optional ("*opt*") padding may be included at the end of a MAC-e PDU in order that the MAC-e PDU would be the same size as the E-TFC selected for transmission.

DDI, TSN and N

117. The MAC-e header comprises, for each MAC-es PDU in the MAC-e PDU, a Data Description Indicator (DDI - a field of 6 bits) which identifies the logical channel, MAC-d flow and MAC-d PDU size for a MAC-es PDU, and N (another field of 6 bits) which identifies the number of MAC-d PDUs in each

MAC-es PDU. The MAC-es header for each MAC-es PDU comprises a TSN (another field of 6 bits) which provides the transmission sequence number on the E-DCH. Where present, the final DDI contains a special value indicating that no more data is contained in the remaining part of the MAC-e PDU (i.e. there is padding).

Scheduled and non-scheduled data


118. There is one final point which I found helpful to keep in mind. The terminology of *scheduled* and *non-scheduled* data is somewhat counterintuitive because non-scheduled transmissions are able to be made on a predictable basis, and scheduled transmissions not so. However, the naming is understandable when one bears in mind that the scheduled transmissions are scheduled by the Node B, whereas non-scheduled transmissions are specified by the SNRC and are not the subject of Node B controlled scheduling.

CGK points in dispute

119. Before trial, the parties identified a series of CGK points in dispute, together with cross-references to relevant paragraphs of the experts' reports. By the time of closing argument, one dispute (concerning rate matching) was no longer relevant and I can leave it on one side. Having reflected on the disputes which remain, all of them were closely related, in one way or another, to the critical issue in this case – how the skilled person would read Filiatrault – and some were proxies for that very dispute. For this reason (but also because I don't believe any of the points in dispute affect consideration of the Patent or the issues of construction), it is better to address the CGK points in dispute in conjunction with Filiatrault.

The Patent

120. The Patent is entitled 'MAC multiplexing and TFC selection procedure for enhanced uplink'. In view of the construction arguments and their significance in this case, it is necessary to set out a fair amount of the teaching in the Patent.
121. The Background section runs from [0002] to [0017] and the experts were agreed that [0002]-[0014] is based on Filiatrault or some other version of TS 25.309. Figs 1 and 2 set out very high level schematics of a 3G cellular system (which nonetheless takes the trouble to identify the RNC as separate to the Node B) and the UE (which the Patent refers to as a WTRU (Wireless Transmit/Receive Unit)) protocol architecture respectively.



122. By reference to Fig 2, [0003] explains that an Enhanced Uplink (EU) medium access control (MAC-e) is used to support EU operation between a dedicated channel MAC (MAC-d) and a physical layer. It continues:

The MAC-e 206 receives data for EU transmission from channels known as MAC-d flows. The MAC-e 206 is responsible for multiplexing data from MAC-d flows into MAC-e protocol data units (PDUs) for transmission, and for selecting proper EU transport format combinations (E-TFCs) for EU transmissions.

[0004] To allow for EU transmissions, physical resource grants are allocated to the WTRU 106 by the Node-B 102 and the RNC 104. WTRU UL data channels that require fast dynamic channel allocations are provided with fast "scheduled" grants provided by the Node-B 102, and channels that require continuous allocations are provided with "non-scheduled" grants by the RNC 106. The MAC-d flows provide data for UL transmission to the MAC-e 206. The MAC-d flows are either configured as scheduled or non-scheduled MAC-d flows.

[0005] A "serving grant" is the grant for scheduled data. A "non-scheduled grant" is the grant for non-scheduled data. The serving grant is the power ratio that is converted to a corresponding amount of scheduled data that can be multiplexed, thus resulting in the scheduled data grant.

[0006] The RNC 104 configures non-scheduled grants for each MAC-d flow using radio resource control (RRC) procedures. Multiple non-scheduled MAC-d flows can be configured simultaneously in the WTRU 106. This configuration is typically performed upon radio access bearer (RAB) establishment, but may be reconfigured when necessary. The non-scheduled grant for each MAC-d flow specifies the number of bits that can be multiplexed into a MAC-e PDU. The WTRU 106 is then allowed to transmit non-scheduled transmissions up to the sum of nonscheduled grants, if multiplexed in the same transmission time interval (TTI).

[0007] Based on scheduling information sent in rate requests from the WTRU 106, the Node-B 102 dynamically generates scheduling grants for scheduled MAC-d flows. Signaling between the WTRU 106 and the Node-B 102 is performed by fast MAC layer signaling. The scheduling grant generated by the Node-B 102 specifies the maximum allowed EU dedicated physical data channel (E-DPDCH) / dedicated physical control channel (DPCCH) power ratio. The WTRU 106 uses this power ratio and other configured parameters to determine the maximum number of bits that can be multiplexed from all scheduled MAC-d flows into a MAC-e PDU.

[0008] Scheduled grants are "on top of" and mutually exclusive of non-scheduled grants. Scheduled MAC-d flows can not transmit data using a non-scheduled grant, and non-scheduled MAC-d flows can not transmit data using a scheduled grant.

123. [0009] then explains that the WTRU knows the set of all possible E-TFCs. For each EU transmission, an E-TFC is selected from the set. However since other uplink transmissions take precedence over EU transmissions [0010] explains that the power available for EU transmissions on E-DPDCH is the remaining power available after the power required for DPDCH, HS-DPCCH and E-DPCCH is taken into account. Thus, based on the remaining transmit power available, E-TFCs in the set are either blocked or supported, states which are continuously determined by the WTRU.

[0011] Each E-TFC corresponds to a number of MAC layer data bits that can be transmitted in an EU transmission time interval (TTI). Since there is only one MAC-e PDU per E-TFC that is transmitted in each EU TTI, the largest E-TFC that is supported by the remaining power defines the maximum amount of data, (i.e., the number of bits), that can be transmitted within a MAC-e PDU.

[0012] Multiple scheduled and/or non-scheduled MAC-d flows may be multiplexed within each MAC-e PDU based on absolute priority. The amount of data multiplexed from each MAC-d flow is the minimum of the current scheduled or non-scheduled grant, the available MAC-e PDU payload from the largest supported TFC, and the data available for transmission on the MAC-d flow.

[0013] Within the supported E-TFCs, the WTRU 106 selects the smallest E-TFC that maximizes the transmission of data according to the scheduled and non-scheduled grants. When scheduled and non-scheduled grants are fully utilized, available MAC-e PDU payload is fully utilized, or the WTRU 106 has no more data available and allowed to be transmitted, MAC-e PDUs are padded to match the next largest E-TFC size. This multiplexed MAC-e PDU and corresponding TFC are passed to the physical layer for transmission.

[0014] The serving and non-serving grants specify the maximum amount of data that can be multiplexed from specific MAC-d flows into MAC-e PDUs each EU TTI. Since the scheduled grants are based on the E-DPDCH/ DPCCCH ratio, the number of data bits allowed to be multiplexed per MAC-e PDU can not be explicitly controlled only to allow certain sizes which match the limited number of data sizes of the supported E-TFCs within the E-TFCS.

124. Then in [0015]-[0017], the Patent explains some of the disadvantages of the prior art system so far described:

[0015] The remaining transmit power for EU data transmission determines the list of supported E-TFCs within the E-TFCS. Since the supported E-TFCs are determined from a limited number of E-TFCs in the TFCS, the granularity of allowed MAC-e PDU sizes will not allow for all possible MAC-d flow and MAC-e header combinations. Therefore, since the amount of MAC-d flow data allowed by the grants to be multiplexed into a MAC-e PDU will frequently not match the size of one of the supported E-TFCs, padding will be applied to the MAC-e PDU to match the smallest possible E-TFC size within the list of supported E-TFCs.

[0016] It is expected that when EU cells are operating at maximum capacity the MAC-e PDU multiplexing is frequently limited by the serving and non-serving grants, and not limited by the largest supported E-TFC or the WTRU EU data available for transmission. In this case, depending on the granularity of specified E-TFCs within the E-TFCS padding required to match the selected E-TFC may exceed the multiplexing block size of MAC-d flow data including associated MAC-e header information. In this case, the effective data rate is unnecessarily reduced from what is allowed by the selected E-TFC and the physical resources required for its transmission.

[0017] Figure 3 illustrates a MAC-e PDU 300. A MAC-e PDU header 302 and MAC-d flow data 304 allowed by scheduling and non-scheduling grants are multiplexed. Among a set of supported E-TFCs, the WTRU 106 selects the smallest E-TFC from a list of supported E-TFCs that is larger than MAC-e PDU header 302 and MAC-d flow data 304. Padding 306 is then applied to the MAC-e PDU to match the selected E-TFC size. However, the padding 306 may exceed the multiplexing block size of MAC-d flow data. In this case, physical resources used in the EU transmission are under-utilized and the effective WTRU data rate is unnecessarily reduced. Accordingly, it is desirable to have alternate approaches to multiplexing EU data.

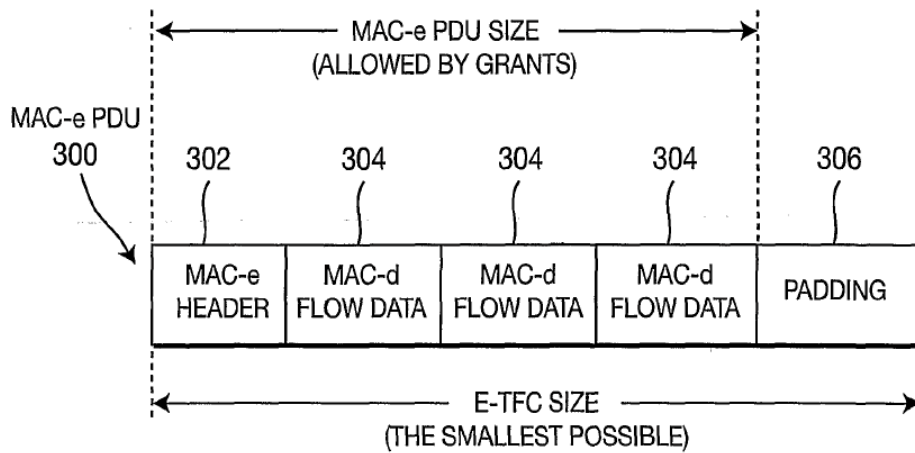


FIG. 3
PRIOR ART

125. Accordingly, under the heading ‘Summary’, the inventive concept is explained as follows:

[0018] The present invention is related to quantizing the amount of multiplexed data allowed by grants to closely match a selected E-TFC transport block size is disclosed. The amount of scheduled and/or non-scheduled data allowed to be transmitted is either increased or decreased relative to the grants so that the amount of data multiplexed into a MAC-e PDU more closely matches the selected E-TFC transport block size.

[0019] When the amount of scheduled data is adjusted to more closely match a selected E-TFC, the maximum amount of scheduled data to multiplex, the scheduled payload to transmit, is determined by the sum of the scheduled and non-scheduled data available to be transmitted and allowed by the grants quantized to the next larger or smaller E-TFC size, minus the amount of available to be transmitted non-scheduled data that is allowed by the non-scheduled grants.

[0020] This quantization is applied when multiplexing is grant limited, and not limited by the maximum E-TFC size resulting from E-TFC restriction or limited by E-DCH data available for transmission.

126. Thus, ‘quantizing’ is an important concept whereby the amount of data to be transmitted is either increased or decreased. The purpose of quantizing is to make the amount of data multiplexed into a MAC-e PDU more closely match an E-TFC transport block size. [0019] explains how to quantize (i.e. adjust). The amount of non-scheduled data to be transmitted is subtracted from the next larger or smaller E-TFC to give an adjusted amount of scheduled data that can be multiplexed. After the usual brief description of the drawings, the Detailed Description begins at [0022]. In this section [0024] repeats the inventive

concept as described in [0018]-[0020], but introduces the remaining paragraphs in the following way:

[0024] The following modifications to MAC-e PDU multiplexing logic are proposed for more efficient data multiplexing and improved radio resource utilization for the cases where MAC-e PDU multiplexing is limited by scheduled and/or non-scheduled grants, and not limited by the largest supported E-TFC or available EU data for transmission.

127. By reference to Figures 4 and 5, [0025] and [0026] explain the basic idea of adjusting (i.e. quantizing) the multiplexing limit set by the grants to match a selected E-TFC size, in two embodiments. Both comprise three steps:

- (a) First, the UE receives a "*scheduled data grant from a Node B and/or non-scheduled grants from an RNC*".
- (b) Second, "*an E-TFC transport block size is selected based on the amount of data allowed to be multiplexed according to the scheduled and non-scheduled grants*"
- (c) Having selected an E-TFC based on the grants, Figure 4 requires the amount of data allowed to be transmitted by the grants to be quantized so that the amount of data multiplexed more closely matches the selected E-TFC. In view of one of the construction points which I discuss later, it is to be noted that the quantizing step is described in [0026] '*so that the sum of the scheduled and non-scheduled data (including MAC header and control information) is multiplexed into each EU MAC-e PDU more closely matches the selected E-TFC transport block size*'.
- (d) Figure 5 requires the amount of buffered data be quantized, so that the sum of data (including headers and control information) multiplexed more closely matches the selected E-TFC size. This recognises that in the grant limited scenario, it is not just the grants which will be quantized, but also the amount of data which can consequently be sent. It also recognises that header and control information (e.g. Scheduling Information) will need to be included alongside the data in the quantized E-TFC size.
- (e) In either case, the outcome is that the UE has identified an E-TFC based on the amount of data allowed by the grants, and modified the amount of data that can be multiplexed into it (i.e. it has adjusted, or "quantized", the multiplexing limit set by the grants).

128. At this point it is worth noting the following:

- (a) The Skilled Addressee would understand that the background paragraphs were based on the Prior HSUPA Stage 2 Specification (i.e. either Filatrault or an earlier version of it).

- (b) Because the Serving Grant is a power ratio, it cannot explicitly control the number of bits which can be sent in a MAC-e PDU to match an E-TFC size. Thus, frequently, the data to be multiplexed into a MAC-e PDU will not match an E-TFC size, so padding will be required.
 - (c) If the amount of data which can be multiplexed is limited by the grants, the padding required may exceed the multiplexing block size (see [0016]). Therefore, "physical resources ... are under utilized and the effective [UE] data rate is unnecessarily reduced" (see [0017]). That is to say, the Patent discloses that this is inefficient because more data could be sent without creating more interference.
 - (d) Having identified these circumstances, the Patent proposes a solution: select an E-TFC size which closely matches the grants, and then quantize the amount of data allowed to be multiplexed by the grants to more closely match that E-TFC size (see [0018]).
 - (e) The Patent proposes a series of solutions which are "modifications" to the prior MAC-e PDU multiplexing and E-TFC selection method.
 - (f) The Patent introduces the idea (in [0025] and [0026]) of selecting an E-TFC based on the grants (rather than selecting an E-TFC based on the data multiplexed into a MAC-e PDU in accordance with the grants), and creating a MAC-e PDU in accordance with that selected E-TFC
129. [0028]-[0058] and [0060]-[0061] then explain the remaining Figures 6-14 in turn. Mr Townend provided a useful overview of these figures, as follows:
- (a) Figures 6 to 9 ([0028] to [0032]) describe how to use the **next-smaller** (Figures 6 and 7) or **next-larger** (Figures 8 and 9) E-TFC as the multiplexing limit for a MAC-e PDU, and illustrate a resulting MAC-e PDU.
 - (b) Figure 11 ([0043] to [0047]) and Figure 14 ([0056] to [0058] and [0060]-[0061]) explain how to adjust (i.e. quantize) the amount of scheduled data multiplexed using the multiplexing limits determined using the **next-smaller** (Figure 11) or next-larger (Figure 14) E-TFC.
 - (c) Figure 10 ([0033] to [0042]) explains how to integrate the multiplexing logic of Figure 11 or Figure 14 into a system to be used when a UE might not be grant limited.
 - (d) Figure 13 ([0050] to [0055]) describes a modified version of Figure 10 in which, instead of determining the E-TFC which will be the multiplexing limit in advance of multiplexing, the E-TFC to be used as the multiplexing limit is determined during the multiplexing procedure itself. Again, the multiplexing logic can be used either with the next-smaller or the next-larger E-TFC as the multiplexing limit.

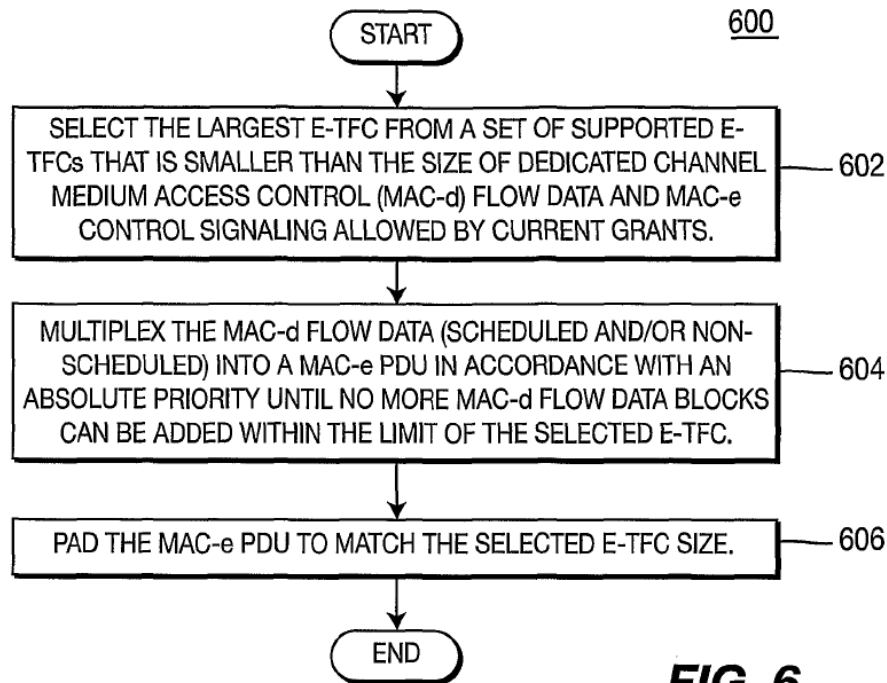
- (e) Figure 12 ([0048] and [0049]) illustrates the functional elements of the UE and the UTRAN which are required for HSUPA, and additional elements which are required to work the invention.

130. Thus, the three schemes described are:

- (a) the Next Smaller Scheme (Patent Figures 6 and 7);
- (b) the ‘Multiplex Then Select Scheme’ (Patent Figures 8A and 9); and
- (c) the ‘Next Larger Scheme’ (Patent Figures 8B and 9).

131. I will concentrate in the Next Smaller Scheme, since at least a version of this is what is claimed. Integer 1F of claim 1 (see below) uses very similar phraseology to the passage underlined in [0028] which describes Fig 6 as follows (emphasis added):

[0028] Figure 6 is a flow diagram of a process 600 for generating a MAC-e PDU in accordance with another embodiment. A largest E-TFC is selected from a set of supported E-TFCs that is smaller than the size of MAC-d flow data and MAC-e control signaling allowed by current grants 602. As a result, the selected E-TFC permits a decreased amount of data to be multiplexed onto the MAC-e PDU relative to the amount allowed by the grants, to more closely match the largest E-TFC size that is smaller than the amount required by scheduled and non-scheduled grants. The MAC-d flow data (scheduled and/or non-scheduled) is multiplexed into a MAC-e PDU in accordance with an absolute priority until no more MAC-d flow data blocks can be added within the limit of the selected E-TFC 604. The MAC-e PDU is padded to match the selected E-TFC size 606.



132. Then [0029] describes Fig 7

[0029] Figure 7 illustrates the decreased MAC-e PDU 700B size that more closely matches a selected E-TFC size in accordance with the embodiment of Figure 6. A MAC-e PDU header 702 and MAC-d flow data blocks 704a-704c are supported by the current scheduled and non-scheduled grants. Referring to Figures 6 and 7, the largest E-TFC that is smaller than the size of MAC-d flow data allowed by current grants is selected from the set of supported E-TFCs (step 602). MAC-d flow data blocks, (in this example, the two MAC-d flow data blocks, 704a, 704b), are multiplexed into the MAC-e PDU 700B in accordance with an absolute priority until no more MAC-d flow data blocks can be added within the limit of the selected E-TFC size (step 604). MAC-d flow data block 704c is not multiplexed since it will exceed the limit of the selected E-TFC. Preferably, only the amount of multiplexed scheduled data is adjusted to more closely match the selected E-TFC size. Padding 706 is then applied to the MAC-e PDU 700B to match the selected E-TFC size (step 606). One technique for the padding is accomplished implicitly by insertion of an end-of-data indicator in the MAC-e PDU header information.


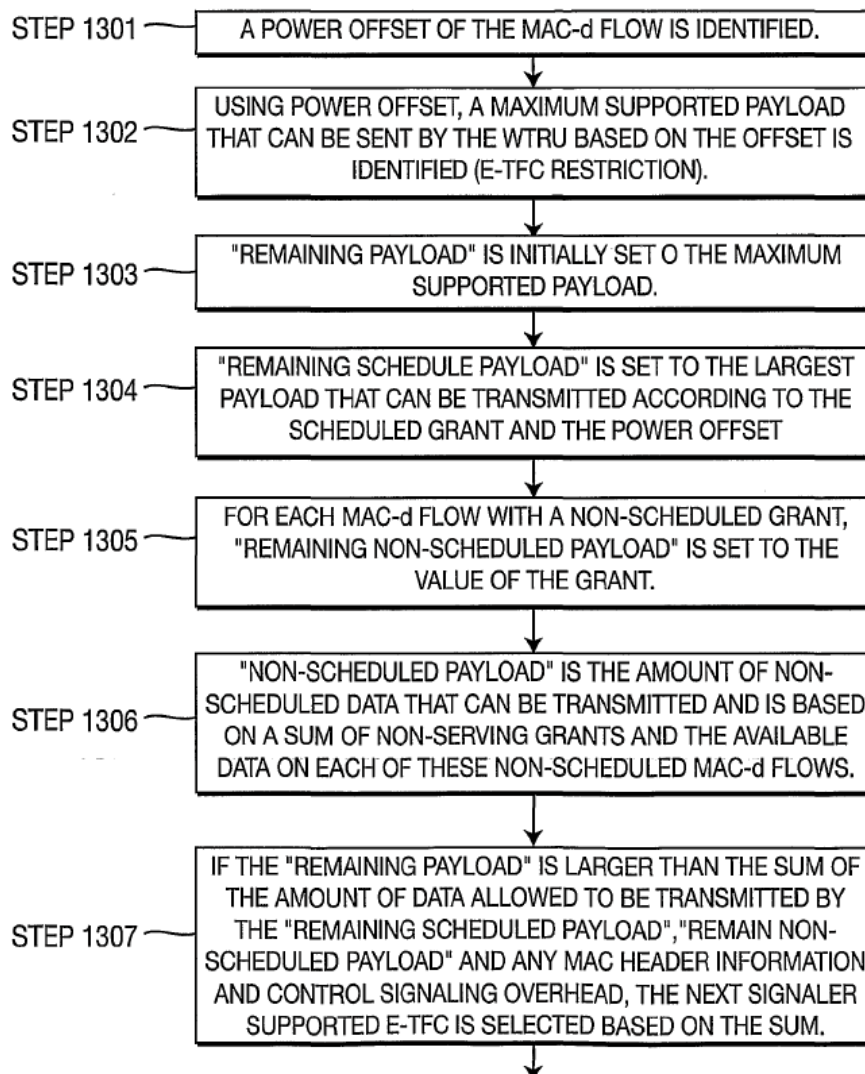


FIG. 7

133. More detail of how the Next-Smaller scheme works is provided in [0043]-[0047] by reference to Figures 11A and 11B.

FIG. 11A

134. E-TFC restriction involves two steps – 1301 and 1302. In 1301, the power offset of the ‘selected’ MAC-d flow is identified. Mr Townend said the Skilled Reader would understand this to be the highest priority MAC-d flow for transmission in that TTI. That power offset is used to determine the maximum supported payload that can be sent.
135. Then the E-TFC Selection operation determines whether the WTRU is grant limited (in steps 1303-1305 and 1307) by comparing the ‘Remaining Payload’ variable to the sum of two other variables – the ‘Remaining Scheduled Payload’ and the ‘Remaining Non-Scheduled Payload’ – referred to as ‘the sum’. The Remaining Scheduled Payload is set in 1304 to the number of bits of data allowed to be sent by the Serving Grant. The Remaining Non-Scheduled Payload is set in 1305 to the number of bits of data allowed to be transmitted by all the Non-Scheduled Grants.
136. As [0045] explains:

[0045] The next smaller supported E-TFC is the largest supported E-TFC that does not carry more data than the sum. In other words, the selected E-TFC is the next smaller E-TFC based on the serving grant, non-scheduled grants, the power offset, available data, including any MAC header information and control signaling overhead, such as scheduling information.

This (along with [0044]) indicates (a) that the E-TFC selected must be less than or equal to the sum of the grants, (b) that (as the Skilled Team would naturally expect) MAC headers etc are included, (c) that the word SIGNALER in the penultimate line of 1307 should read SMALLER (there are other more obvious typos as well) and (d) that 1307 uses the sum of the grants, whereas [0044] refers to the amount of 'available data allowed to be transmitted', which separates the serving grant from the non-scheduled grant(s).

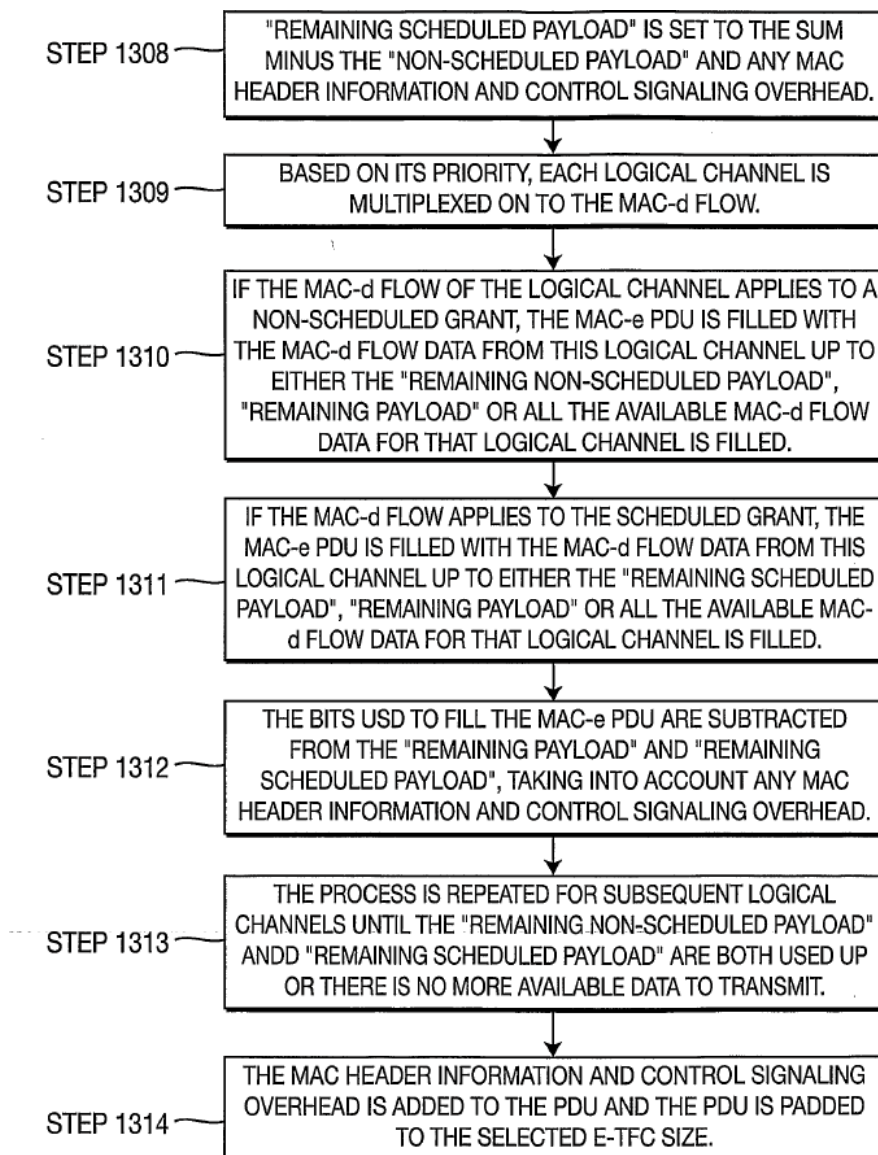


FIG. 11B

137. Then, in 1308 the E-TFC Selection quantizes the Remaining Scheduled Payload to the selected E-TFC size less the *‘Non-Scheduled Payload and any MAC header information and control signalling overhead.’*
138. Multiplexing accounts for the remaining steps. The logic comprises two steps: first, data multiplexing in 1309 to 1313; and second, the addition of MAC header information, control signalling and padding.
139. The data multiplexing starts with the logical channel with the highest priority. For each non-scheduled MAC-d flow of a given priority, data is multiplexed up to the minimum of the remaining Non-Scheduled Grant for that MAC-d flow (Remaining Non-Scheduled Payload), remaining space in the E-TFC (Remaining Payload), or available data for that MAC-d flow (1310). The variables "Remaining Non-Scheduled Payload" and "Remaining Payload" are then reduced accordingly. This is not shown in Fig 11B, but [0046] says *‘The bits used to fill the MAC-e/es PDU are subtracted from the "Remaining Payload" and the "Remaining Non-scheduled Payload", taking into account any MAC header information and control signaling overhead’*.
140. For scheduled MAC-d flows of a given priority, data is multiplexed up to the minimum of the remaining Serving Grant (Remaining Scheduled Payload), remaining space in the E-TFC (Remaining Payload), or available data for that MAC-d flow (1311). The variables Remaining Scheduled Payload and Remaining Payload are then reduced accordingly (1312).
141. In relation to both scheduled and non-scheduled MAC-d flows, the Patent notes at [0046] that the "bits used to fill the MAC-e/es PDU are subtracted ... taking into account any MAC header information and control signaling overhead".
142. The process is then repeated for subsequent MAC-d flows, in priority order, until the Remaining Payload is used up, or there is no more data to transmit (1313). Thus, although Fig 11B does not show a loop, the skilled addressee would understand a return to the start of the multiplexing process at 1309.
143. Finally, in 1314, header information and control signalling overhead are added and the PDU is padded to the selected E-TFC size. Although the header information is only added here, it is clear that the space required for it is taken into account throughout the multiplexing logic.
144. [0047] completes the description of Fig 11, by pointing out the procedure allows the WTRU operation to be *‘deterministic’* and that the Node B scheduler can *‘therefore accurately predict how resource grants will be used by the UE’* and *‘more efficiently allocate resources’*.
145. Similar additional detail is provided in the description of other figures for the Next-Larger and Multiplex then Select Schemes, but it is not necessary to set it out. However, the description is consistent in pointing out at various places that headers and control information need to be taken into account, even if the Figures tend to refer simply to the MAC-d flow or MAC-D flow data.

146. In this regard, I mention [0059]. Although this paragraph is sandwiched between paragraphs [0056]-[0058] and [0060]-[0061], all of which deal with Figure 14, it seems to me it is a paragraph of general application, along with [0062] and [0063]. It provides:

‘[0059] References to the amount of data multiplexed according to grants, and the amount of data that can be multiplexed according to a selected E-TFC takes into account MAC header information and other control signalling overhead required in the formatting of a MAC-e PDU.’

Construction / Claim Scope

147. This case has not involved any arguments as to equivalents, so my task is to undertake a ‘normal’ interpretation of the claims: see *Eli Lilly v Actavis UK Ltd* [2017] UKSC 48. This is a very familiar test, but I am reminded that it remains an exercise in purposive construction (*Icescape Ltd v Ice-World International BV* [2018] EWCA 2219 at [60] per Kitchin LJ (as he then was)). It is an objective exercise and the question is always what a skilled person would have understood the patentee to be using the words of the claim to mean.

148. In this case, the terms in issue are underlined in what follows. Claim 1 of the Patent, broken down into integers, reads:

1A *A wireless transmit/receive unit, WTRU, comprising:*

1B *means for receiving at least one serving grant and at least one non-scheduled grant,*

1C *wherein the at least one serving grant is a grant for scheduled data transmission*

1D *and the at least one non-scheduled grant is a grant for non-scheduled data transmission,*

characterized by:

1E *means for determining supported enhanced dedicated channel transport format combinations, E-TFCs, within a E-TFC set, E-TFCS, based on a remaining transmit power for enhanced uplink, EU, transmission; and*

1F *means for limiting medium access control for dedicated channel, MAC-d, flow data multiplexed into a medium access control for enhanced uplink, MAC-e, protocol data unit, PDU, to a largest E-TFC size that is smaller than a size of MAC-d flow data allowed by the received serving and non-scheduled grants and available for transmission.*

149. Lenovo emphasised two points:

- (a) First, that this claim is laboriously written out so that all the abbreviations are defined in the claim: WTRU, E-TFCs, E-TFCS, EU, MAC-d, PDU.

- (b) Second and consequently, it is possible to simplify claim 1 by removing unnecessary definitions and abbreviations:

1A *A WTRU comprising:*

1B *means for receiving at least one serving grant and at least one non-scheduled grant,*

1C *wherein the at least one serving grant is a grant for scheduled data transmission*

1D *and the at least one non-scheduled grant is a grant for non-scheduled data transmission,*

characterized by:

1E *means for determining supported E-TFCs, within a E-TFC set based on a remaining transmit power for enhanced uplink transmission; and*

1F *means for limiting MAC-d flow data multiplexed into a MAC-e PDU to a largest E-TFC size that is smaller than a size of MAC-d flow data allowed by the received serving and non-scheduled grants and available for transmission.*

150. In view of some of the arguments, it is necessary to be aware of claims 2 and 3:

- (a) Claim 2 adds this integer (my emphasis):

Means for selecting a smallest possible E-TFC that is required to support the amount of data allowed to be multiplexed by the at least one of the serving and non-scheduled grants.

- (b) Claim 3 reads:

The WTRU of any preceding claim wherein the amount of MAC-d flow data to be multiplexed into the MAC-e PDU is based on the largest E-TFC that is smaller than a size of MAC-d flow data allowed by the received serving and non-scheduled grants on a condition that there is no E-TFC equal to the size of MAC-d flow data allowed by the received serving and non-scheduled grants and available for transmission.

151. Before I turn to deal with the three short issues of construction which arise, there are a couple of comments to make about these claims.

152. First, as IDC pointed out, the E-TFC selection step is introduced in claim 2.

153. Second, the additional integer in claim 3 (dependent, directly or indirectly on claim 1) is, in effect, use the largest E-TFC which is smaller than the MAC-d flow data unless there is an E-TFC which is the same size as the MAC-d flow data. In the scheme of the Patent, this would appear to the Skilled Team as entirely logical. Furthermore, this also suggests that claim 1 must include this situation.

‘MAC-d flow data’

154. The issue is whether this expression includes the headers which are inevitably added in the process of the MAC-d flow data being multiplexed into the MAC-e PDU. It will be recalled (see Figure 6 above) that an instance of MAC-d data is first incorporated into a MAC-es PDU with the addition of a TSN. Then multiple MAC-es PDUs are multiplexed into a MAC-e PDU, each with its own MAC-e header comprising N plus DDI.
155. Lenovo contend that ‘MAC-d flow data’ must be read literally, such that the expression (and thus integer 1F) is only concerned with the bits in MAC-d PDUs and nothing else. Lenovo say that the reference just to ‘MAC-d flow data’ is deliberate.
156. In support of their argument, Lenovo stress that the patent is a technical document in which technical terms are used with precision, relying in particular upon the laborious way in which all defined terms are spelled out in the claim
157. In my view, the Skilled Person would regard Lenovo’s argument as unreal. As IDC submitted, Lenovo’s argument is untenable technically and unsupported linguistically.
158. First, the Skilled Person would be very familiar with the way in which (on the UE side) data is passed down from layer to layer, with a PDU passed from an upper layer becoming the SDU of the lower layer, with that SDU being incorporated into the PDU of the lower layer along with a header or more formally, PCI. The Skilled Person would be aware that this basic process continues down to the physical layer, even if at that point the nomenclature changes.
159. It would make no technical sense to the Skilled Person to have to consider data from an RLC PDU in, say, a MAC-es PDU without the headers added at the MAC-d and MAC-es protocol layers. So too, it makes no technical sense for the Skilled Person to have to consider MAC-d data alone at the MAC-e layer without the headers which *necessarily* accompany it by the time it reaches the MAC-e layer.
160. Second, Mr Townend explained the major technical problem with Lenovo’s construction in the following paragraph, which was not challenged in cross-examination:

100. If, as Dr Irvine appears to suggest, claim 1 referred only to the MAC-d PDUs then it would lead to an absurd situation. A UE operating in accordance with claim 1 would select the E-TFC lower than the grants, and would multiplex data until no more data could be multiplexed in. Having done so, it would then add in the necessary header information, meaning that the total data to be sent would then increase, potentially by a substantial number of bits. This would or could likely push the total multiplexed data over the level of the selected E-TFC. The UE would therefore then have to select a higher E-TFC, which by

definition would exceed the sum of the grants. That makes a nonsense of the idea in claim 1 of the Patent, and indeed the overarching inventive concept described in the Patent. Further, since that E-TFC is likely to be filled only partially with data (since it will only contain such of the header as is unable to be fitted within the originally selected E-TFC), then a substantial amount of padding may be needed (which may indeed exceed the amount of padding which would be required in Filiastrault on my understanding of it). This would be completely contrary to the inventive concept of the Patent, and the idea of the Next Smaller Scheme. The skilled addressee would think that technically that was not what the patentee could have meant, and is clearly not what is disclosed by Patent Figure 7.

161. In other words, Lenovo's construction would increase the amount of padding and reduce the efficiency of the system, as compared not only to what was disclosed in the Patent itself, but it would even be worse than the acknowledged prior art. That would, as Mr Townend observed, be an absurd position.
162. Lenovo's argument appears to be purely linguistic: that claim 1 refers to 'MAC-d flow data' and therefore cannot include any of the associated header overhead generated when MAC-d flow data is multiplexed into the MAC-e PDU. However, no Skilled Person could think this term was to be understood literally. As I pointed out above, the idea in claim 1 is explained in [0028] and [0029] by reference to Figures 6 and 7. [0028] is explicit that the E-TFC which is chosen is one which is smaller than the size of the MAC-d flow data and the MAC-e control signalling allowed by the current grants. This is also explicit in both Figs 6 and 7.
163. There are numerous other points which indicate that Lenovo's argument is hopeless, including:
 - (a) On Lenovo's construction, claim 1 would claim a system which is nowhere described in the Patent.
 - (b) Furthermore, it would be a system which is not directed to the problem with which the Patent is concerned i.e. *increasing* padding efficiency.
 - (c) On Lenovo's construction, claim 1 does not include header overhead, but claim 2 does. Yet claim 2 is clearly refers to the same data allowed to be multiplexed as claim 1.
 - (d) On Lenovo's construction, the promise in [0047] of a 'deterministic' system would not be fulfilled
164. By contrast, with the teaching in the specification in mind, the Skilled Addressee would readily understand the expression 'MAC-d flow data' to be a shorthand which includes the necessary headers. The Skilled Addressee would also understand why this shorthand was used: it is because the sequence in which data is multiplexed into a MAC-e PDU is determined by identification of individual MAC-d data flows (see Figure 6 above) but, as I have indicated, it is

second nature to the Skilled Addressee that as each MAC-d PDU is incorporated first, into a MAC-es PDU and then into a MAC-e PDU, header overhead is added. The headers *must* accompany the data so that the E-TFC transport block can be unpacked at the Node B.

165. Paragraph [0015] is an early and clear example of the shorthand being used. The second sentence indicates clearly that it is talking about the combination of MAC-d data and headers – see the reference to ‘all possible MAC-d flow and MAC-e header combinations’. The third sentence continues on the same theme, starting with the word ‘Therefore’, but then simply referring to ‘MAC-d flow data’ being multiplexed into a MAC-e PDU. [0016] then reverts to specifying ‘MAC-d flow data including associated MAC-e header information’. There are many other examples of the shorthand being used in the Patent.
166. I have well in mind Lenovo’s point about the laborious way in which the definition of each term was included in claim 1 and the apparent conflict with the expression ‘MAC-d flow data’ being used as a shorthand but this is another essentially linguistic argument which has no weight against the technical considerations I have set out.
167. During closing argument, I asked Mr Abrahams QC if he could identify any technical reason why the patentee would have restricted himself to just talking about the MAC-d flow data in the context of a MAC-e PDU (i.e. excluding the header). In essence, he was unable to do so: simply saying that the patentee had a choice and one can have a rule which operates at any layer.
168. Finally, [0059] seems to me to put the matter beyond any doubt.
169. For all these reasons, I reject Lenovo’s contention.

‘smaller than’

170. Here the issue is whether the words would be read literally or as ‘not larger than’. In other words, if all the data (and headers) multiplexed into the MAC-e PDU fitted exactly in a supported E-TFC with no padding (what was termed the ‘Matching Scenario’), is that within the claim or not?
171. Lenovo say that the term ‘smaller than’ is a mathematical term with a single, fixed meaning and its use in the claim is deliberate. Lenovo characterise the Figs 6 & 7 embodiment as different from the embodiment described by reference to Fig 11 because, in Figs 6 & 7, the E-TFC used is ‘smaller than’ the grants, whereas in the Fig 11 embodiment, the ‘next smaller’ E-TFC is used.
172. Lenovo argue that the Patent discloses several different embodiments but it clearly does not claim all of them. Lenovo also point out that the Patent is, on its face, part of a ‘large-ish’ family of divisionals. Accordingly, so the argument goes, the skilled addressee is taken to know that there are or may be other aspects of what is described in the Patent claimed in other patents divided out from the original application, by reference to *Virgin Atlantic Airways Ltd v Premium Aircraft Interiors UK Ltd* [2009] EWCA Civ 1062 at [10]-[11] per

Jacob LJ. Lenovo therefore argue that the Skilled Addressee would assume that the 'next smaller' embodiment is claimed in another divisional.

173. Whether Figs 6-7, on the one hand and Fig 11 on the other are different embodiments or the same embodiment described in differing levels of detail rather begs the question, but it is a question which must be determined on what the Skilled Person would understand the language of the claim to mean, in the light of the technical disclosure of the Patent.
174. In my judgment, with the teaching in the specification in mind, the Skilled Addressee would understand this part of integer 1F to mean the largest E-TFC that can accommodate but does not exceed the MAC-d flow data allowed by the serving and non-scheduled grants and available for transmission i.e. 'smaller than or equal to', for the following reasons.
175. First, as I pointed out above, in [0045], the Patent explicitly teaches that the "*next smaller E-TFC is the largest supported E-TFC that does not carry more data than the sum*". It therefore expressly teaches that in the Matching Scenario, the matching E-TFC should be used, and that this is included within the concept of the "*next smaller E-TFC*".
176. Second, any other interpretation would be technically absurd. The skilled addressee would recognise from the teaching in the Patent that in the Matching Scenario it would be completely illogical and inefficient to use an E-TFC other than that matching E-TFC in terms of the benefits which the Patent identifies in relation to the Next-Smaller Scheme.
177. It is true that the invention in the Patent is not really addressed to the Matching Scenario, since the problems addressed by the Patent do not arise in that scenario. The Patent teaches that the Next-Smaller Scheme (and, indeed, the other schemes as well) can be used in the Matching Scenario. In the Matching Scenario, since there is no need to pad to start with, it is not possible to adjust the amount of multiplexed data "*so that the amount of data multiplexed into a MAC-e PDU more closely matches the selected E-TFC transport size*" ([0018]), since they already match absolutely.
178. It follows that if the next-smaller E-TFC did not include the matching E-TFC in the Matching Scenario, applying the process would reduce data throughput, and would likely increase the amount of padding (from zero). This would be the reverse of the aims of the system. Once again, it would not provide a 'deterministic' system ([0047]).
179. All these reasons, which are based on the Skilled Addressee's understanding of the teaching in the Patent, also explain why it would make no sense for the 'next smaller' scheme to be claimed in a separate divisional. In these circumstances, even if (which I very much doubt) the Skilled Addressee even started to wonder about divisionals, he or she would dismiss immediately the possibility that the Patentee wished to claim the 'next smaller' scheme separately in another divisional.

180. Accordingly, Lenovo's literal interpretation makes no technical sense. Finally, my conclusion on the previous point of construction tends to point away from an overly literal interpretation, which is what Lenovo's argument involves. The argument has no merit and I reject it.

'means for'

181. Lenovo argue that the means in integer 1F form part of a scheme such that the means must operate all the time. For its part, IDC makes the conventional argument that provided the apparatus has the requisite means, it infringes even if the means are employed only some of the time.

182. This argument arises because in the Standard as implemented, the integer 1F means are employed when scheduled data is to be transmitted but not for the transmission of non-scheduled data alone. Since there will be occasions where only non-scheduled data are sent, Lenovo argue that there is no infringement.

183. Lenovo's argument is wrong. It is not necessary in this apparatus claim that the means in integer 1F must operate all the time. The accused products have means as specified in claim 1, including in integer 1F. Even if I assume this argument works in claim 1, Lenovo would infringe and would be responsible for multiple infringements of claim 5, the method equivalent to claim 1.

184. My finding has consequences however because this integer 1F is expressed in very broad terms: it is satisfied by any means which bring about the stated result.

Essentiality/Infringement

185. The Claimants' case on essentiality and infringement is that EP537 is infringed by certain of the Defendants' devices that implement sections of the following standards:

- (a) TS 25.321 v6.18.0;
- (b) TS 25.309 v6.6.0; and
- (c) TS 25.319 v7.0.0.

186. It was common ground that Lenovo's points on essentiality and infringement depended on (and only on) the three points of construction which I dealt with above. Since I have decided all three points against Lenovo, it follows that there is no remaining dispute. If the Patent is valid, then Lenovo's identified devices infringe and the Patent is essential. I must now consider whether the Patent is valid.

Validity

187. The main attack was anticipation by Filiatrault. Hence there was a great deal of discussion which involved direct comparisons between what was said to be disclosed in Filiatrault and what was claimed in the Patent. Various shorthands were used by each of the experts and the following discussion of the terminology will also highlight the key arguments on anticipation.

Terminology

188. This case turns on the difference(s) if any between what is disclosed by Filiatrault and claim 1 of the Patent. The arguments concern the sequence in which various steps are taken to process MAC-d data into a MAC-e PDU, which is also the Transport Block which is transmitted by the UE on the physical layer.

189. In order to aid understanding of the difference(s) which Mr Townend was trying to explain in his first report, he attempted to encapsulate what he saw as the key difference between what was disclosed in Filiatrault and the Patent in paragraph 305 of his first report. He also coined the nomenclature (which I have adopted) to describe the various embodiments in the Patent, viz: 'Next Smaller', 'Next Larger' and 'Multiplex then Select'. For his part, Dr Irvine also coined largely equivalent expressions when discussing the Patent, viz: 'round down' and 'round up' and 'round nearest'.

190. Mr Townend's paragraphs 305 and 306 read as follows:

'305. The invention claimed in the Patent is therefore different from the disclosure in Filiatrault because:

(a) Filiatrault uses the grants as a multiplexing limit in a MAC-e PDU and then selects the smallest E-TFC which is larger than the multiplexed data, and adds any necessary padding. This allows the selected E-TFC to be larger than the sum of the grants, and/or the padding to be larger than a MAC-d PDU which could otherwise be sent.

(b) The claimed invention in the Patent uses the grants as a limit on the largest E-TFC which can be used and then uses the size of that largest E-TFC as a multiplexing limit. This guarantees that the selected E-TFC (i.e.: the multiplexed data, the headers, and any padding that comprise the MAC-e PDU) will not be larger than the sum of the grants, nor will the size of padding be greater than the size of a MAC-d PDU which could otherwise be sent.

306. It is therefore my understanding that Filiatrault discloses a fundamentally different scheme for MAC-e PDU creation and E-TFC selection compared to the invention claimed in claims 1 and 5 of the Patent.'

191. Dr Irvine encapsulated his view that Filiatrault anticipated in this sentence:

'It is clear that Filiatrault is describing a system whereby the grants indicate a maximum amount of data which can be sent, and limits the scheduled and non-scheduled data to fit within that limit. This is exactly the same as the claim....[he added a qualification which concerns the issue of construction on 'smaller than', but it is not relevant to anticipation].

192. In his second report, Dr Irvine applied his nomenclature to paragraph 305 of Townend 1 in his exhibit JI-3. In its opening skeleton argument, Lenovo adopted JI-3 and the ‘round up’ and ‘round down’ terminology with some enthusiasm, and JI-3 was the focus of significant attention at trial. By the time it came to closing arguments, Lenovo had lost this enthusiasm, suggesting that the terms ‘do not add clarity’ and it was much better ‘to focus on the purpose of the grants’, by which it meant ‘a limit on uplink resources’.

JI-3 and X3

193. In his second report, Dr Irvine explained his JI-3 in this way. He said that one way to describe the difference between Mr Townend’s 305(a) and (b) was in the sequence of steps. He explained that in 305(a), multiplexing was the first step, limited by the grants, and then the E-TFC selection takes place after data has been multiplexed into the MAC-e PDU, resulting in the round up approach. In 305(b) he described E-TFC selection as the first step, based on the grants, then data is multiplexed up to the size of the chosen E-TFC, resulting in the round down approach. In contrast to Mr Townend, Dr Irvine was of the view that Filiatrault disclosed a round down approach, as in the Patent.

194. As I indicated above, JI-3 was useful up to a point. In cross-examination, Dr Irvine accepted that his JI-3 omitted, on the right-hand side, the actual E-TFC selection step. Mr Speck put to him a corrected version of JI-3 which became X3. With the true E-TFC selection step included (as step 3A on the right-hand side) X3 is shown below.

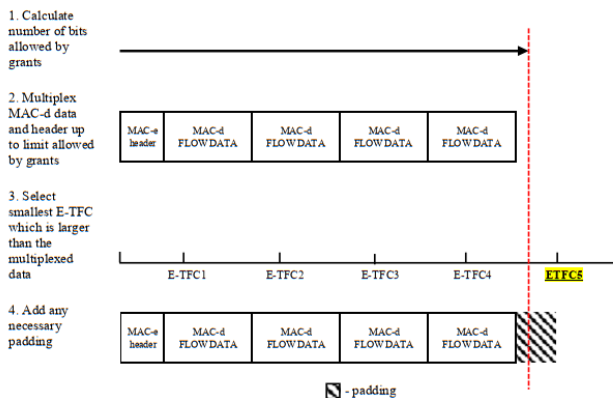
195. There are one or two additional points to be made about X3:

- (a) First, on each side E-TFC restriction is assumed to have taken place already and is not illustrated.

Exhibit JI-03

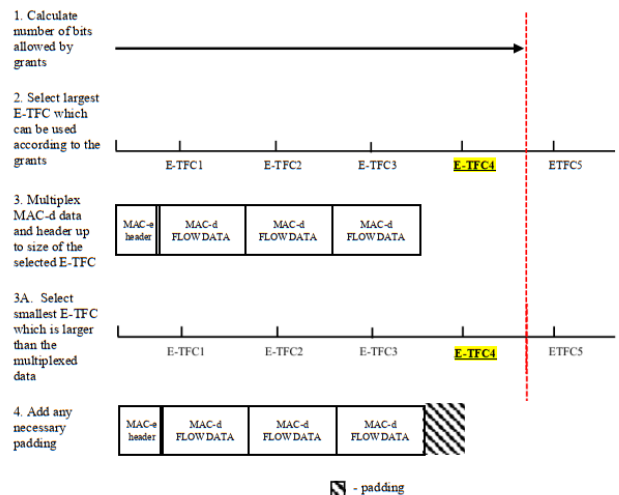
Townend 1 paragraph 305

(a) “... uses the grants as a multiplexing limit in a MAC-e PDU and then selects the smallest E-TFC which is larger than the multiplexed data, and adds any necessary padding”



This is the round up scheme

(b) “... uses the grants as a limit on the largest E-TFC which can be used and then uses the size of that largest E-TFC as a multiplexing limit.”



This is the round down scheme

- (b) Second, the right-hand side of X3 highlights the additional and earlier selection step, Step 2, which the Patent's scheme requires. Note this is not E-TFC restriction but a selection from the E-TFCs permitted following E-TFC restriction.
196. Another way in which Dr Irvine described the difference between Mr Townend's paragraph 305(a) and (b) in his second report was as follows:
- ‘..in the former, the round up approach, the grant is a limit on the amount of MAC-d flow data and associated header which may be sent, and in the latter, round down approach, the grant is a limit on the total amount of resources that may be used (i.e. including MAC-d flow data, header and padding).’
197. Finally, I should mention that on the left-hand side of JI-3 or X3, the representation of Filiatrault assumes that it disclosed that the number of bits allowed by the grants should be added together. This assumption reflected paragraph 170 of Mr Townend's first report where he was initially of the view that, although the expression in Filiatrault ‘*Scheduled grants will be considered on top of non-scheduled transmissions*’ was ‘not particularly precise language’, it suggests that the Scheduled grant (converted to bits) and any Non-Scheduled grants (already in bits) are to be added together to work out the total number of bits and hence the total transmit power which the UE may use.
198. In his second report, Mr Townend was of the view that he had slightly misstated matters in paragraph 170 of his first report and that it would be better to say that the two grants ‘*would have to be considered together*’ to work out a total number of bits and hence total transmit power. I accept Mr Townend's qualification, but I don't believe it makes any material difference.
199. The analysis so far shows that, in order to decide which approach is correct, one can address these questions:
- (a) First, do the grants indicate a maximum amount of data which can be sent?
- (b) Second and more precisely, is ‘the grant’ a limit on the total amount of resources that may be used (including MAC-d flow data, header and padding)?
- (c) Third, does Filiatrault disclose the additional selection step i.e. step 2 on the right-hand side of X3?

Resolution of the CGK points in dispute

200. I can now return to deal with the CGK points which were in dispute and I do so under a series of headings. As far as I can see, all of them were generated by the opposing views on Filiatrault.

The role of UE maximum transmit power in TFC and E-TFC restriction.

201. As I mentioned at paragraph 82(d) above, Dr Irvine was very keen to make the point that a UE is not allowed to transmit above the maximum allowed transmission power, defined as the ‘maximum UE TX power’ in TS 25.331. He said this applied to TFC restriction and to E-TFC restriction.
202. It was tolerably clear why Dr Irvine was so keen on making this point. Evidently, he saw it as a critical stepping stone to characterising Filiatrault as a ‘round down’ system, as claimed in the Patent.
203. In his third report, Mr Townend explained why he did not agree. His explanation involved the following stages:
 - (a) First, he agreed that the maximum UE transmit power is a limit on all of the uplink transmissions (i.e. not just transmissions on the Enhanced Uplink).
 - (b) Second, he repeated his point from his first report that the maximum UE transmit power is used to transition TFCs to one of three states: supported, excess power or blocked, with only ‘blocked’ TFCs being prohibited from use by the maximum UE transmit power.
 - (c) Third, he said the use of supported and excess power TFCs can result in the UE selecting a TFC which would result in it exceeding the maximum UE transmit power value for any of three reasons:
 - (i) First, the elimination criterion looks back over the previous 30 measurement periods (each being a 0.667ms slot) to determine whether the TFC would have required more power than was available for at least 15 of those measurement periods and only if so, transitions the TFC to excess-power state. Thus, for at least 15 slots before that transition, it will have been estimated by the UE to require more than the maximum UE transmit power, despite being in ‘supported’ state.
 - (ii) Second, for a 10ms TTI, a TFC will be in excess-power state for at least 40ms (4 TTI) before the blocking criterion is met.
 - (iii) Third, UE power control changes the UE transmit power every 0.667ms, or 15 times per 10ms TTI. This means that even if a UE has sufficient power to transmit a given TFC at the required power at the start of the TTI, there is no guarantee that it will have sufficient power at the end.
 - (d) Fourth, he pointed out that TS 25.214 explains how the UE is to handle a situation where it has selected a TFC that requires it to transmit with more than the maximum power. The UE must apply additional scaling to the total transmit power so that it is equal to the maximum allowed power. In other words, the UE must scale back all physical channels and scaling back can occur every 0.67ms.

- (e) Fifth, whilst the maximum UE transmit power is treated as an absolute limit, the UE is not obligated to select the ‘next smaller’ TFC relative to it. Instead, so Mr Townend explained, this is a factor in TFC restriction (not selection). The state transition criteria are used to reduce the chance that the UE will be required to scale back the transmit power below the required level for a TFC. Only then does the UE apply the three TFC selection criteria (set out in paragraph 81 above).
 - (f) Mr Townend concluded by saying that it is clear that the scaling of the transmit power required for a TFC has nothing to do with TFC selection or the amount of data transmitted, because the transmit power can be scaled back after the TFC has been selected and up to 15 times during a 10ms TTI.
 - (g) In the course of this explanation, Mr Townend pointed out an error in Dr Irvine’s first report where he had stated that UE power control operates once per TTI for the legacy channels, whereas UE power control operates 1500 times a second (i.e. once every 0.667ms).
204. Mr Townend applied the same reasoning in response to Dr Irvine making the same point as regards E-TFCs. He said that UE maximum transmit power is a limit not a grant and a limit applied in a similar manner in Filiatrault in E-TFC restriction (not selection) and to control the power at which each slot is transmitted.
205. I found Mr Townend’s reasoning persuasive. Furthermore, there was precious little or no independent material presented to support Dr Irvine’s view. I address below the materials relied upon by Lenovo to support the ‘basic concept’ which it was said the skilled person would bring to his or her reading of Filiatrault.

TFC selection on the legacy channels

206. In his second report, Dr Irvine referred to the maximum transmit power in TFC selection as being akin to a ‘round down’ system.
207. One important point to note is that there were no grants involved for the legacy channels. Instead, Dr Irvine said that the skilled person would be aware that the UE had a maximum amount of power that it was allowed to use and that this could not be exceeded. Dr Irvine therefore stated that this performed the equivalent of a ‘round down’ approach.
208. Mr Townend did not agree, for the reasons I have already set out in paragraph 203 above. It follows that even in the context of the DCH, the maximum UE transmit power is not a ‘round down’ system.
209. There was no dispute that TFC Selection on the legacy channels was governed by three criteria:

“1. No other TFC shall allow the transmission of more highest priority data than the chosen TFC.

2. No other TFC shall allow the transmission of more data from the next lower priority logical channels. Apply this criterion recursively for the remaining priority levels.

3 No other TFC shall have a lower bit rate than the chosen TFC" (see TS 25.321 v6.4.0, s.11.4, p.60-1).

210. The dispute arose over the third criterion. Dr Irvine said ‘...*all other things being equal, the MAC will always choose the TFC transmitting the smallest number of data bits, resulting in a lower bit rate.... This check ensures that the UE is not transmitting more bits than necessary. The MAC is not allowed to add padding for the Release 99 and Release 4 TFC, although the size of the transmitted data can be adjusted using rate matching.*’
211. Mr Townend responded by saying that Dr Irvine was conflating the size of the TFC with the data rate of the TFC. He explained that the third criterion relates to minimising the amount of padding added into RLC PDUs at the RLC layer, which will then affect the amount of data which is handed onto the MAC and which needs to be sent. The MAC in DCH does not add any padding itself, nor can it require the RLC to produce padding PDUs. The point is that the Transport Blocks to be used by the MAC and the sizes of the RLC PDUs handed down by the RLC are under the control of the SRNC and are coordinated so that each RLC PDU fits exactly into the MAC Transport Block, once the MAC header has been added.
212. Once again, I found Mr Townend’s explanation convincing and I do not believe it was challenged.

The skilled person’s understanding of E-TFC selection and multiplexing in TS 25.309 V6.2.0

213. It will be recalled that this is a reference to what Mr Townend called the Prior HSUPA Stage 2 Specification. For ease of reference I will simply call it ‘the Prior Specification’ or v6.2.0. Although presented as a dispute on CGK, this issue was very close indeed to the critical disagreement over what Filiatrault disclosed to the skilled person. As is clear from Filiatrault itself, it has a great deal in common with the Prior Specification. Furthermore, as Mr Townend explained, some of the changes made in Filiatrault over the Prior Specification simply made explicit what the skilled person would have understood from the Prior Specification anyway.
214. This issue began with Mr Townend’s description (by reference to the Prior Specification) of one or more MAC-es PDUs being multiplexed in accordance with pre-defined limits into a single MAC-e PDU (see paragraph 111 above). In his second report, Dr Irvine agreed with this, but stated (as was the case) that it appeared from the rest of Mr Townend’s first report that they disagreed about what the limit should be, in the situation where the UE is grant-limited, as opposed to being power or data limited. Of course, Mr Townend also detected this same fundamental disagreement, and addressed it in his second report.

215. It is not necessary to describe all the slight differences between the two specifications. It suffices to say that I accept Mr Townend's points that:
- (a) In Filiatrault the E-TFC restriction and selection processes are more clearly separated from one another than in the Prior Specification.
 - (b) The E-TFC restriction process is more defined in Filiatrault, but in respects which the skilled person would have understood from the Prior Specification.
216. So far as E-TFC selection is concerned, the key sentence in section 11.2 of the Prior Specification read as follows:
- The UE adds the resulting power offset for the MAC-e PDU to the previously calculated reference power offsets for different E-TFCs. It then selects the E-TFC, taking into account the obtained power offsets, the required E-TFC-dependent power back-off, the UE's remaining power and the amount of data to be transmitted.
217. The first sentence is concerned with E-TFC restriction (and in Filiatrault, considerations of this nature are moved up in section 11.2 into the points dealing with E-TFC restriction). The amount of data to be transmitted is the data allowed by the Serving Grant and the Non-Scheduled Grant(s). As Mr Townend pointed out, although there is no express reference in section 11.2 of the Prior Specification to the data being sent being limited by the grants, that this is the case is plain from the document as a whole and in particular from section 10, where it is stated that "*Logical channels will be served in the order of their priorities until the non-scheduled granted rate and scheduled grants are full, or the maximum transmit power is reached*". For each logical channel, data is multiplexed until the relevant grants are used up, or the maximum transit power is reached. Applying the E-TFC selection process set out in section 11.2, an E-TFC is then selected to fit the size of this multiplexed data.
218. So, Mr Townend said the skilled person would understand from the Prior Specification that:
- (a) The UE works out which E-TFCs it can use (E-TFC restriction);
 - (b) It multiplexes data up to the limits specified by each of the Serving Grant and the Non-Scheduled Grant(s);
 - (c) It chooses an E-TFC from the list of E-TFCs it can use which can carry the multiplexed data – in other words, it selects an E-TFC from the E-TFC Set which is larger than the multiplexed data; and
 - (d) It would be implicit that the smallest such E-TFC which was large enough to carry the multiplexed data would be selected.

219. In anticipation of Filiatrault, it is relevant to note that what I have called the key sentence in section 11.2 in the Prior Specification is deleted in Filiatrault and effectively replaced by the ‘according to the... grants’ sentence.
220. This CGK dispute was a proxy for the dispute as to what Filiatrault disclosed, and I decide it in the same way (see below), even though Lenovo developed a different argument to justify the supposed multiplexing limit.

The relationship between common Guaranteed Bit Rate traffic flows and transport formats in E-DCH

221. In his second report in paragraph 3.4, Dr Irvine stated that only a finite number of transport formats would be used in HSUPA in order to simplify system design. He said the transport formats could be specified with common CBR (constant bit rate) traffic flows in mind, so there would be a transport format that a Voice over IP flow or signalling bearer could exactly fill. He went on to acknowledge that would not be possible for the variable services, so they would have to be mapped to one of the transport formats. He went on to describe what he said were the three options – effectively, round down, nearest and next larger or round up. He said all three options are routine resource allocation strategies which the skilled person would have known about at the priority date.
222. It may be noted that this could be said to provide a foundation for an obviousness attack based on CGK alone. An attack based on CGK alone was neither pleaded nor pursued. However, Dr Irvine’s paragraph 3.4 was a key part of his argument that claim 1 was obvious.
223. Mr Townend responded by noting first of all that Dr Irvine appeared to use CBR (constant bit rate) as a synonym for GBR (guaranteed bit rate) services. Whilst he acknowledged that being able to specify transport formats so they could be exactly filled was ideal, the complexity of trying to specify a system which could address all of the common GBR services and combinations of them being sent at the same time would very quickly become extremely complicated. The net result was that the skilled person would expect that for services configured to use Non-Scheduled Grants there would often be a mismatch between the Non-Scheduled Grants and the E-TFCs that might be used to transmit the non-scheduled data so padding would need to be used.
224. I was not persuaded that these considerations set out by Dr Irvine in his second report in paragraph 3.4 were CGK. I accept that the skilled person might well have known of various resource allocation strategies but the missing link, in my view, was any notion that any of these resource allocation strategies would occur to the skilled person when reading Filiatrault. I note that none of the material in his paragraph 3.4 was mentioned in Dr Irvine’s first report. I am of the view that Dr Irvine wrote his paragraph 3.4 to provide himself with a basis for the obviousness argument he presented later in his second report.

The Development of the case on Filiatrault

225. Filiatrault was published 2 days before the earliest priority date of the Patent. It comprises a marked-up version of the Prior Specification, the draft of the Stage

- 2 document for Release 6 published on 24 March 2005. It defines its scope as a technical specification of the overall support of FDD Enhanced Uplink in UTRA.
226. As I have already mentioned, the case based on Filiatrault changed over time. However, I do not believe it is either practical or sensible for me to deal with Lenovo's case on Filiatrault as it was finally presented without having in mind how both experts got to that point, largely because their views on earlier iterations of the case undoubtedly influenced their views later.
227. I have already set out the basic positions taken by each of the experts on Filiatrault so I can start here with the way in which Lenovo set out their case on Filiatrault in their Opening Skeleton argument. As I understand matters, this reflected the views which Dr Irvine had expressed in his three expert's reports.
228. In Section 4 of Filiatrault, headed 'Background and Introduction', it indicates the technical objective is to improve the performance of the uplink dedicated transport channels i.e. to increase throughput and reduce delay. By reference to the Feasibility Study, it states that the following techniques are part of the work item:
- (a) Node B controlled scheduling: possibility for the Node B to control, within the limits set by the RNC, the set of TFCs from which the UE may choose a suitable TFC,
 - (b) Hybrid ARQ: rapid retransmissions of erroneously received data packets between UE and Node B,
 - (c) Shorter TTI: possibility of introducing a 2 ms TTI.
229. The main sections are: 6 Overall Architecture of enhanced uplink DCH; 7 MAC architecture; 8 HARQ protocol; 9 Node B controlled scheduling; 10 Non-scheduled transmissions; 11 QoS control; and 12 Signalling parameters.
230. I have already covered some of the content of Filiatrault in the CGK section above. For example, Figure 6 above is Figure 7.2.1-2 in Filiatrault.
231. Lenovo's Opening Skeleton argument drew attention to the following particular parts of Filiatrault, in which I have retained Lenovo's underlining. For ease of reference below, I identify the main points which Lenovo combined in their argument as A-G. Paragraphs taken directly from Lenovo's Opening Skeleton are in quotation marks, with my interpolations in square brackets.
232. "Details of the MAC-e in the UE are given in s7.2.5:
- Multiplexing:

The multiplexing entity is responsible for concatenating multiple MAC-d PDUs into MAC-es PDUs, and to multiplex one or multiple MAC-es PDUs into a single MAC-e PDU, to be transmitted at the next TTI, and as instructed by the E-TFC selection function. It is also responsible for managing and setting the TSN per logical channel for each MAC-es PDU.

- E-TFC selection:

This entity is responsible for E-TFC selection according to the scheduling information (Relative Grants and Absolute Grants) received from UTRAN via L1, and for arbitration among the different flows mapped on the EDCH. The detailed configuration of the E-TFC entity is provided by RRC over the MAC-Control SAP. The ETFC selection function controls the multiplexing function.”

233. “This is illustrated in figure 7.2.5-1 (reproduced below with Dr Irvine’s mark-up to show data flows):




Figure 7.2.5-1: UE side MAC architecture / MAC-es/e details

234. “This shows two MAC-d data flows coming from the top into the MAC-e/es multiplexing entity (the solid lines). The multiplexing entity is controlled by the E-TFC Selection entity (as stated in s7.2.5) [I interpolate: this is **Point A**]. The scheduling grants are shown coming into the E-TFC Selection entity on the dotted lines: “E-AGCH” and “E-RGCH(s)” are the channels on which the scheduling grants are sent. After E-TFC selection and then multiplexing, the data is sent downwards (solid line) to HARQ processing and then down to lower layers for transmission. (HARQ is a system for retransmission of erroneously received data packets.)”

235. Much of that was not contentious. The exception was Lenovo’s contention that the figure (and the preceding text) indicated that E-TFC selection occurred first, followed by multiplexing [This is **Point B**].

236. “Section 9 deals with scheduled data and scheduling grants. In s9.1 (General Principle) it provides:

In the downlink, a resource indication (Scheduling Grant) is required to indicate to the UE the maximum amount of uplink resources it may use...

The Scheduling Grants have the following characteristics:

- Scheduling Grants are only to be used for the E-DCH TFC selection algorithm (i.e. they do not influence the TFC selection for the DCHs);
- Scheduling Grants control the maximum allowed E-DPDCH/DPCCH power ratio of the active processes.

...

- There are two types of grants:
 - The Absolute Grants provide an absolute limitation of the maximum amount of UL resources the UE may use;
 - The Relative Grants increase or decrease the resource limitation compared to the previously used value;”

237. “And under s.9.2.1 (Grants from the Serving RLS):

- The UE maintains a “Serving Grant” (SG);
- The SG is used in the E-TFC selection algorithm as the maximum allowed E-DPDCH/DPCCH power ratio for the transmission of the active HARQ processes;”

238. So, as Lenovo pointed out, section 9 specifies three times that the serving grant (**SG**) indicates to the UE the maximum UL resources it can use, once in those very words and twice more in terms of the maximum allowed E-DPDCH/DPCCH power ratio of the active processes. This is **Point C**. Although Filiatrault doesn’t explain this in terms, the experts were agreed that while the serving grant is provided to the UE in the form of a power ratio, it would have to be converted from a power ratio to a number of bits.

239. “Section 10 deals with non-scheduled data and non-scheduled grants. It provides:

‘When non-scheduled transmission is configured by the SRNC, the UE is allowed to send E-DCH data at any time, up to a configured number of bits, without receiving any scheduling command from the Node B...

- The resource for non-scheduled transmission is given by the SRNC in terms of maximum number of bits that can be included in a MAC-e PDU, and is called non-scheduled grant;
- ...
- Scheduled grants will be considered on top of non-scheduled transmissions;”

240. **Point D** is that the Non-scheduled Grant (**NSG**) is given by the SRNC in terms of a maximum number of bits that can be included in a MAC-e PDU.
241. **Point E** is that SGs are considered ‘on top of’ NSGs. Lenovo submitted that the experts agreed that the reference to the scheduled grant being “*on top of*” non-scheduled transmissions mean that the number of bits allowed by each grant would be added together. This submission relied on the passage in Mr Townend’s first report (at [170], which I have mentioned above) yet ignored the qualification he made to that paragraph in his second report that the grants ‘would have to be considered together’ in order to work out the total number of bits, and hence the total transmit power, a UE may use.
242. “Section 11 is headed “TFC and E-TFC Selection”, the Skilled Team would note that it concerns two operations: E-TFC restriction and E-TFC selection.”
243. “**E-TFC restriction [Point F]** is the process of restricting the E-TFCs which the UE can use, based on whether transmitting the E-TFC would require more or less than the remaining transmit power available. This involves an assessment of:
- (a) The UE remaining power (the amount of power the UE has available to send the enhanced channel, after transmitting the legacy channels).
 - (b) The power offset for the logical flow channel.
 - (c) The power offset for the particular E-TFC.
 - (d) The power back-off needed for a particular E-TFC.”
244. “E-TFCs for which the UE has sufficient remaining power are “*supported*” (i.e. they may be selected for use), whereas E-TFCs that require more than the remaining transmit power available are “*blocked*”.”
245. “In relation to **E-TFC selection**, s11.2 says [and this is **Point G**]:
- Among the supported E-TFCs, the UE selects the smallest E-TFC that maximises the transmission of data according to the non-scheduled grant(s) and the serving grant.”
246. From the above, it will be noted that multiplexing is touched on, almost in passing, in Points A and B. Furthermore, Lenovo’s argument as to when multiplexing occurs is heavily dependent on the model described by reference to figure 7.2.5-1.

247. Lenovo's case on anticipation of claim 1 was originally pleaded relying solely on references in Filiatrault. In particular, as originally pleaded, Lenovo relied upon the combination of (a) the sentence from 7.2.5 "The E-FTC selection function controls the multiplexing function" [i.e. Point A] together with (b) the sentence quoted above from section 11.2 [i.e. Point G] as disclosing (in the context of Filiatrault as a whole) integer 1F of claim 1. Following service of the first round of expert's reports, Lenovo then amended their statement of case on validity to rely in addition in relation to integer 1F on:

'(i) The description of the Multiplexing and E-TFC selection entities in Section 7.2.5 and as illustrated in Figure 7.2.5-1 [adding Point B to Point A];

(ii) The description of the Scheduling Grants in Section 9.1 [Point C]; and

(iii) The description of the non-scheduled grant in Section 10.[Points D & E]'

248. This was essentially the case set out in Lenovo's Opening Skeleton, as can be seen from the extracts I set out above. However, Lenovo's Opening Skeleton also contained a hint that the anticipation case had changed. At [50], Lenovo indicated that the central dispute was as to how 'the skilled person would have understood the grants in HSUPA and how they would be used.'

249. In the course of cross-examination, it became clear that Dr Irvine accepted that the previously key sentence in section 11.2 [Point G] described the conventional final step of E-TFC selection, in the sense of identifying the E-TFC that was actually to be used, being the one which is big enough to accommodate the data which has been multiplexed, with padding as necessary.

250. The emphasis in Lenovo's argument now lay on a specific part of that sentence, namely 'according to [the grants]'. It emerged that Lenovo did not rely on any specific wording in Filiatrault itself as explicitly disclosing integer 1F of claim 1. Instead, Lenovo now relied on a specific part of the CGK which, it is said, the skilled person would bring to his or her reading of Filiatrault.

251. The change of heart was prompted by Dr Irvine's review of the cross-examination documents presented to him in the CXX bundle. He was perfectly frank about this in his cross-examination.

Q. What you are explaining here is a different way of looking at
18 it. It is not in your report and is in fact something that we
19 heard Mr. Abrahams advanced with Mr. Townsend yesterday, is it
20 not?

21 A. My reading of this sentence came about after I saw the
22 cross-examination documents, because I was thinking to myself
23 "How do people have these different world views?" When you
24 look again at the text, you see this.

25 Q. So the answer, I think, to my question is yes, it is something
2 that you have thought of in the last few days?

3 A. It is a realisation that I have come to. I have been troubled

4 from square one by the fact that an expert in the area, and
5 I am very happy to consider Mr. Townsend to be an expert in
6 this area, has read a document and come to a view on that
7 document that is diametrically opposed to my view on that
8 document.

9 Q. Sure. I understand.

10 A. I was very, very puzzled by that. However, you can read this
11 text both ways. When I read Filiatrault and I read the
12 "according to", and even when I read that statement within the
13 patent, I am thinking down here because I am thinking in
14 terms of a radio system. We have discussed the fact that, you
15 know, you could get relative grants, because you are worried
16 about the rise over thermal and you have to turn down what is
17 going on because you are worried about how much interference
18 you are generating. So, I am thinking in terms of this grant
19 is really important. I do not want to go over a particular
20 level. This is where I am thinking about.

252. In the extract (T3/pp406-407) 'this sentence' is the sentence in Filiatrault at section 11.2 which says '*Among the supported E-TFCs, the UE selects the smallest E-TFC that maximises the transmission of data according to the non-scheduled grant(s) and the serving grant.*' The reference to 'down here' is at the level of the physical layer in Dr Irvine's drawing X1. Almost as soon as his cross-examination started, Dr Irvine was very keen to draw X1 in order to explain his view of what the Patent said in [0011] when characterising the prior art position. X1 remained on the clipboard, so he was able to and did make various gestures to it later in his cross-examination to indicate at which of the three levels (MAC-d payload, MAC-e (i.e. MAC-d data plus header) or at the physical layer (where the E-TFC contained headers, MAC-d payload and padding)), he was considering particular teaching.

Lenovo's case on Filiatrault in Closing

253. In its closing arguments, Lenovo contended that claim 1 was invalid over Filiatrault on four different bases, which I will consider in turn:
- (a) First, that Filiatrault clearly and unambiguously taught that the grants (both types) are a limit on uplink resources and therefore Filiatrault anticipates.
 - (b) Second, that it was obvious to use the grants (both types) as a limit on uplink resources, therefore claim 1 involved no inventive step.
 - (c) Third, it contends that the previous point is good even if the skilled person understood Filiatrault in the way that Mr Townsend read it.
 - (d) Fourth, that if I find in favour of IDC on the 'means for' construction issue, that Filiatrault anticipates claim 1 even on Mr Townsend's reading. This is the case based on DXX/14.

Alleged Anticipation by Filiatrault

254. Lenovo's case on Filiatrault in its final iteration can be summarised quite shortly. It embraces various parts of Filiatrault, but the critical element is not disclosed in Filiatrault at all. Instead, the critical element is a piece of CGK which the Skilled Person is said to bring to his or her reading of Filiatrault. This was explained in Lenovo's closing as follows (in the quote I have stripped out the numerous references in the original):

“...the original concept for Node B controlled scheduling in HSUPA was to use the spare uplink resources in the cell, i.e. the available headroom taking into account the resources allocated for transmission of the legacy channels. Fundamentally, the resource being allocated is the contribution to interference at the Node B. There would need to be tight control over this because uplink power control is critical in a CDMA system and Node B controlled scheduling would have to operate within the cell's margin of safety. The way such tight control would work is that the Node B would grant UEs a share of that resource, in the form of a maximum amount of the uplink resource which they could use. That grant would bound/limit the E-TFC which the UE could use. The effect would be to limit the maximum uplink resources which the UE could use, because it would provide a limit on the size of the MAC-e PDU, and therefore on the size of the Transport Block, and therefore on the amount of physical resources that could be used. All of this was CGK and would be in the skilled person's mind when reading Filiatrault.”

255. So, what Lenovo called 'the original concept' was that the grants allocated to a UE were a maximum amount of the uplink resource which that UE could use. There is some sleight of hand in this formulation of the original or basic concept. As will be seen in due course, the original 'grants' were those making up *the single* grant for Scheduled data, namely the Absolute grant and the Relative grant, which were naturally taken together to make up the 'Serving grant' or the Scheduled grant. In the context of Filiatrault, however, the grants (plural) were different. In Filiatrault, the 'grants' meant the Serving Grant and the Non-Scheduled Grant(s). The critical issue is whether Filiatrault disclosed that those two grants should be added together and treated as the maximum amount of uplink resource which could be used in a particular TTI. Ultimately, it was common ground (but in any event I find) that Filiatrault does not disclose that – hence the need for the basic concept to be brought to the Skilled Person's reading of Filiatrault.

256. One of the reasons why the disclosure of Filiatrault was open to the arguments deployed in this case was because the document represents a snapshot at a particular point in time of the development of the enhanced uplink functionality. It is plainly a document which was a work in progress. Sometimes it indicates this expressly by saying a particular topic was FFS – for further study. Other parts were not marked FFS but were clearly still in development. There are also parts which were introduced at an earlier stage and one has to question whether they needed to be changed to reflect later developments. Furthermore, the

document was a work in progress towards a finalised Stage 2 specification, the details of which would be worked on in a Stage 3 document. These points have three important consequences. First, Filiatrault has passages which clearly derive from earlier drafts of the specification when certain features had not been incorporated. Second, it is unsafe to treat Filiatrault as fully worked out or even necessarily internally consistent. Third, parts of Filiatrault are open to interpretation as to precisely how the concepts described would actually be implemented.

257. With those introductory remarks, I can turn to the way in which Lenovo presented its argument for reading Filiatrault with the so-called original concept fully in mind. This was developed for the first time in the cross-examination of Mr Townend. When referring to the particular documents on which Lenovo relied, I do not restrict myself to the parts to which Lenovo drew attention, since it is important to understand how the ideas and specifications were developing over time.
258. The argument starts with an early proposal (from Lucent Technologies, presented at a RAN1 meeting in November 2002) for fast Node B controlled scheduling in the proposed enhanced uplink. The document explains that, conceptually, one of the main advantages of Node B controlled scheduling is better resource management, and that, on the uplink, this would imply better noise rise management. Figure 1 is presented as illustrating all the quantities which contribute to the noise rise at the Node B.

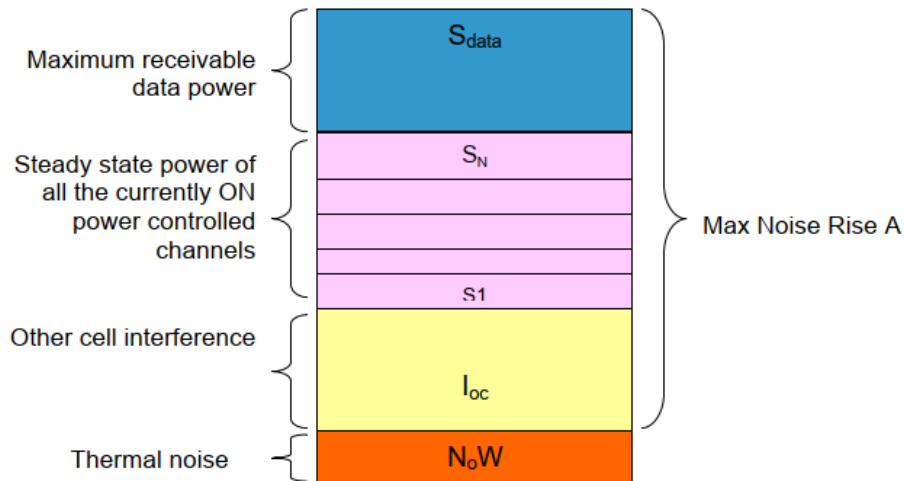


Figure 1. Noise rise management in Node B controlled fast scheduling.

259. Mr Townend characterised this as a ‘conceptual model’ of the basic operation of uplink fast scheduling, but was not reflective of a real system which could be built.
260. Mr Abrahams QC was keen to invoke the blue rectangle at the top, at various points in his submissions, as if it signified the maximum resource available for an individual UE. However, the document clearly explains that S_{data} denotes the maximum available received power that can be tolerated from *all* data

transmissions and is described as the resource which can be shared between multiple data transmissions.

261. Next, Lenovo’s argument moved to the Feasibility Study (3GPP TR 25.896 v6.0.0) and the proposal in section 7 for Node B controlled scheduling, published in March 2004. The basic principle was ‘to allow Node B [to] set and control a new restriction to the TFC selection mechanism of the UE by fast L1 signalling.’ An example was illustrated in Figure 7.1.1, where the TFCs in a TFCS are shown in descending order (with regard to the power required to transmit them). The text explains that:

In Figure 7.1.1, the TFCs in a TFCS are shown ordered in descending order (with respect to the power required) starting from zero. Both TFC pointers are initialised to both the Node B and to the UE by the RNC in the beginning of the connection. After initialisation the Node B can command the UE pointer up/down with the restriction that UE pointer may not exceed Node B pointer. The TFC selection algorithm in the UE may select any TFC up to the TFC indicated by the UE pointer. The purpose here is to control the UE's power usage by restricting its TFC (i.e. data rate) selection.

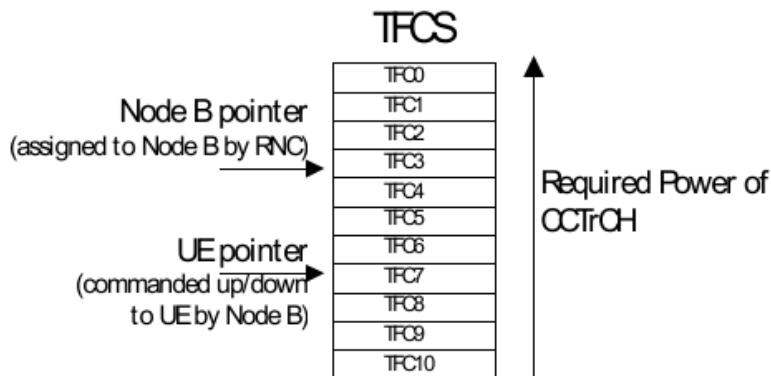


Figure 7.1.1: Depiction of the TFC pointers

262. The Feasibility Study also includes (at Figure 7.1.2) a version of Figure 1 from the Lucent proposal.

263. Lenovo also relied upon a Motorola proposal from June 2004 to establish (in the XX of Mr Townend) that persons working on HSUPA were able to assume that the resource grant would come either in the form of an explicit maximum TFC or in the form of a power ratio which would effectively tell the UE the maximum power which it could use on the physical data channel.

264. By September 2004, the proposals had started to be reflected in a very skeletal draft technical specification TS 25.309 v6.0.0 which was published that month. Section 9 ‘Node B controlled scheduling’ provided ‘*In the downlink, a resource allocation (scheduling grant) is required to indicate to the UE the maximum amount of uplink resources it may use*’ and ‘*It is FFS whether the scheduling*

grant controls the maximum allowed in terms of E-DPDCH/DPCCH power ratio, E-DCH TF index, E-DPDCH+DPDCH/DPCCH power ratio, other....'.

265. It also explained there would be two types of grant – absolute and relative – and that ‘*The combination of absolute and relative grants to get the total grant is FFS.*’
266. This version also included a short section 10 entitled ‘TFC selection’. After some details concerning TFC selection for the DCHs, it turns to E-TFC selection, providing:
- ‘- Every E-DCH TTI, the UE shall estimate the remaining power;
 - Then it performs the TF selection for the E-DCH, with the estimated remaining power, based on logical channel priorities like in the R99;
 - In addition, the UE may need not to go below a minimum rate for the E-DCH. In some case, this means that the UE may have to power scale down all physical channels present;
 - In order to be backward compatible, some E-DCH minimum set support is needed. Details are FFS.
267. Version 6.1.0 of that specification was published in December 2004. Considerable detail had been added. Of particular relevance for present purposes is that this included the diagram of the simplified architecture showing MAC interworking in the UE i.e. Figure 6 above, and section 7.2.5 with Figure 7.2.5-1, showing the UE side architecture and MAC-es/-e details. In section 9, section 9.1 ‘General Principle’ appears to have been almost fully worked out, in the sense that there are no points FFS. In that section, Lenovo drew particular attention to the statement ‘*Scheduling Grants control the maximum allowed E-DPDCH/DPCCH power ratio*’.
268. Section 9.1 also provided:
- (a) ‘The Absolute Grants provide an absolute limitation of the maximum amount of UL resources the UE may use;
 - (b) ‘The Relative Grants increase or decrease the resource limitation compared to the previously used value;
 - (c) ‘Relative Grants (updates) are sent by the Serving and Non-Serving Node Bs as a complement to Absolute Grants’, with further details provided. However, the range of possible grants were UP, HOLD or DOWN, depending on particular conditions.
269. In this version, section 10 has been renamed as ‘QoS control’ in which section 10.1 General Principle explains, inter alia:
- ‘The Node B controls the resources allocated to a UE versus other UEs by means of scheduling as specified in clause 9. The

UE controls the QoS of all its logical channels mapped on E-DCH by means of E-TFC selection as specified in subclause 10.2 and by HARQ operation, as specified in clause 8.'

270. It goes on to introduce non-scheduled transmission for the first time:

'In addition to these mechanisms, guaranteed bit rate services for MAC-d flows/ logical channels (FFS) are also supported through non-scheduled transmission. A flow using non-scheduled transmission is defined by the SRNC and provided in the UE and in the Node B. The UE can transmit data belonging to such flow without first receiving any scheduling grant.'

271. In section 10.2 the TFC selection for the E-DCH, with the estimated remaining power, is explained by reference to an expanded set of rules. I need not set out all those rules, but I draw attention to the last four bullet points:

- For each transmission, the MAC-e entity gives the selected power offset of E-DPDCH(s) relative to DPCCH to the L1 in addition to the E-TFC;

- What should be done when this exceeds the L1 maximum transmission power is FFS.

- In addition, the UE may need not to go below a minimum rate for the E-DCH. In some case, this means that the UE may have to power scale down all physical channels present;

- An E-DCH minimum set because of power limitation is needed. Details are FFS.

272. It is important to note that prior to v6.1.0, all the previous proposals were only concerned with one type of data – scheduled data. In v6.1.0 although non-scheduled data was introduced, it is clear that none of the details had been worked out or included.

273. From these foundations, Lenovo submitted that the Skilled Person would understand and have well in mind the following features of Node B controlled scheduling:

- (a) Node B controlled scheduling is about the allocation of spare resources (i.e. the blue rectangle).

- (b) The resource being allocated is contribution to uplink interference.

- (c) The allocation of these resources would be tightly controlled and the skilled person would think it essential to have tight and accurate control over the resource.

- (d) The grant would be a bound or a limit on E-TFC size. In particular Lenovo submitted that by the date of TS 25.309 v6.0.0 '*it was clear that*

there would be a ‘resource indication (scheduling grant) to indicate to the UE the maximum amount of uplink resources it may use’.

274. The next version (TS 25.309 v6.2.0) was published on 24 March 2005 (i.e. what I have been referring to as the Prior Specification). The content of this version can be discerned from Filiatrault, which is a markup of the additions and deletions to v6.2.0.
275. One important point from v6.2.0 concerns the earlier introduction of non-scheduled data. The Prior Specification included more detail on how non-scheduled data, and in particular, non-scheduled grant(s) would be handled. As Mr Townend explained, the introduction of multiple grants made the system more complicated. The ‘model’ which the skilled person would have in mind was now very different from the ‘single grant’ model in the Feasibility Study: *‘...fundamentally the difference between them is you would have a single grant versus now you have multiple grants for different things.’*
276. Subject to that important point, I don’t believe anything turns on other differences between v6.2.0 and Filiatrault. So I can pass over v6.2.0 and move to Filiatrault.
277. Lenovo had a prior point (before we reach the key parts of Filiatrault) concerning the situation where the UE is power limited. The E-TFCs are restricted by the E-TFC restriction procedure and the largest E-TFC available is smaller than the sum of the grants. In this scenario, Lenovo submitted that Mr Townend accepted that the largest supported E-TFC is used as a multiplexing limit. Lenovo further submitted that this is significant because it puts the skilled person in mind of the idea of using the largest permissible E-TFC as a limit on the total amount of data to be multiplexed and that it would be a simple system to also use the grant limit as an absolute limit in the same way as the power limit. In my view, this submission illustrates the significant degree of hindsight in Lenovo’s whole approach.
278. As explained in its written closing, Lenovo’s case on Filiatrault rested on only a few statements and comprised five elements, which I summarise below. I need to set out each of the statements relied upon. Various submissions were made in relation to each element and, for the most part, I need not set them out even though I have had them fully in mind. On the fourth and fifth elements, it will be seen that Lenovo’s argument depends only indirectly on the words extracted from Filiatrault.
279. 1. ‘Text of Filiatrault; description of the scheduling grant’
62. In relation to the **scheduled grant**, Filiatrault says (at section 9):
- (a) *“a resource indication (Scheduling Grant) is required to indicate to the UE the maximum amount of uplink resources it may use”;*

(b) “Scheduling Grants are only to be used for the E-DCH TFC selection algorithm...”;

(c) “The SG is used in the E-TFC selection algorithm as the maximum allowed E-DPDCH/DPCCH power ratio for the transmission of the active HARQ processes”;

(d) “Scheduling Grants control the maximum allowed E-DPDCH/DPCCH power ratio of the active processes”; and

(e) “The Absolute Grants provide an absolute limitation of the maximum amount of UL resources the UE may use”.

280. 2. ‘Text of Filiatrault: description of the non-scheduled grant’

‘68. In relation to the **non-scheduled grant**, it must be remembered that this was a recent addition to the enhanced uplink and the details were less worked out. Filiatrault says (in section 10):

(a) “When non-scheduled transmission is configured by the SRNC, the UE is allowed to send E-DCH data at any time, up to a configured number of bits”

(b) “The resource for non-scheduled transmission is given by the SRNC in terms of maximum number of bits that can be included in a MAC-e PDU”

281. 3. ‘Text of Filiatrault: relationship between the grants’

‘71. Section 10 of Filiatrault says: “Scheduled grants will be considered on top of non-scheduled transmissions”.’

282. 4. ‘Text of Filiatrault: description of E-TFC selection and multiplexing’

‘81. Filiatrault does not spell out in detail how E-TFC selection and multiplexing is to be done. In particular the role of the grants is not spelled out in section 11.2 itself – all that is stated is that the “smallest E-TFC that maximises the transmission of data according to the non-scheduled grant(s) and the serving grant [sc is selected]”. Therefore, the role of the grants has to be collected from the description of the grants in other sections of Filiatrault.

82. As we explain above, that description of the grants is (clearly and unambiguously) in terms of being a maximum amount of uplink resources, and accordingly, the skilled person would understand that the grants operated (individually or together depending on which grants were available) to limit the total amount of data in a MAC-e PDU.’

283. 5. ‘Text of Filiatrault: “according to [the grants]”’

‘85. This leaves the wording which Lenovo originally pointed to in its pleadings. In relation to E-TFC selection, s11.2 says:

“Among the supported E-TFCs, the UE selects the smallest E-TFC that maximises the transmission of data according to the non-scheduled grant(s) and the serving grant.”

86. As we pointed out in opening, in order to understand what “according to the non-scheduled grant(s) and the serving grant” means, the court needs to first determine what the grants are, and therefore, what maximising the transmission of data according to the grants entails.

87. If the grants are understood to be a limit on the size of the MAC-e PDU (as we discuss above), then selecting “the smallest E-TFC that maximises the transmission of data according to the non-scheduled grant(s) and the serving grant” would mean selecting an E-TFC which was not bigger than the grants (Townend XX T2/190/7-18).’

What Filiatrault disclosed to the skilled person

284. In order to assess whether Lenovo’s analysis is valid, I set out images of extracts from section 9.1 General Principle (under the heading to section 9 Node B controlled scheduling) and the whole of sections 10 Non-scheduled transmissions, and 11.2 TFC and E-TFC selection. The images help to show the amendments made in Filiatrault to the Prior Specification. These fuller extracts assist in showing that Lenovo’s approach to Filiatrault is highly selective.

The disclosure in Filiatrault

9 Node B controlled scheduling

9.1 General Principle

The Node B controlled scheduling is based on uplink and downlink control together with a set of rules on how the UE shall behave with respect to this signaling.

In the downlink, a resource indication (Scheduling Grant) is required to indicate to the UE the maximum amount of uplink resources it may use. When issuing Scheduling Grants, the Node B may use QoS-related information provided by the SRNC (see subclause 11.1.1) and from the UE in Scheduling Requests (see subclause 9.3.1)

The Scheduling Grants have the following characteristics:

- Scheduling Grants are only to be used for the E-DCH TFC selection algorithm (i.e. they do not influence the TFC selection for the DCHs);
- Scheduling Grants control the maximum allowed E-DPDCH/DPCCH power ratio of the active processes. For the inactive processes, the power ratio is 0 and the UE is not allowed to transmit;
- All grants are deterministic;

- Scheduling Grants can be sent once per TTI or slower;
- There are two types of grants:
 - The Absolute Grants provide an absolute limitation of the maximum amount of UL resources the UE may use;
 - The Relative Grants increase or decrease the resource limitation compared to the previously used value;

285. I make a general point about the statements in section 9.1 on which Lenovo rely (as set out above). A number of these statements were *almost* literally true when only scheduled transmissions were in contemplation but could not be taken as literally true by the Skilled Person when Filiatrault was published e.g. “*a resource indication (Scheduling Grant) is required to indicate to the UE the maximum amount of uplink resources it may use*”, not least because non-scheduled transmissions had been added. (I say ‘almost literally true’ in the former case because even then scheduling information could and would be sent which was not included in the Scheduling Grant). The same point applies to “*The SG is used in the E-TFC selection algorithm as the maximum allowed E-DPDCH/DPCCH power ratio for the transmission of the active HARQ processes*” (noting that in fact Lenovo’s quote is slightly inaccurate).

10 Non-scheduled transmissions

When non-scheduled transmission is configured by the SRNC, the UE is allowed to send E-DCH data at any time, up to a configured [number of bits/bitrate](#), without receiving any scheduling command from the Node B. Thus, signalling overhead and scheduling delay are minimized.

Typical examples of data that may use non-scheduled transmission are the SRBs and GBR services.

Non-scheduled transmissions have the following characteristics:

- Non-scheduled transmissions are defined per MAC-d flow;
- The resource for non-scheduled transmission is given [by the SRNC in terms of maximum number of bits that can be included in a MAC-e PDU, and is called non-scheduled grant in terms of bit rate by the SRNC, and is called non-scheduled-granted-rate](#);
 - Scheduled logical channels cannot use the non-scheduled [granted-rate](#);
- UTRAN can restrict a non-scheduled MAC-d flow to use a limited number of HARQ processes in case of 2ms TTI. The same restriction for the 10ms TTI is FFS;
- UTRAN can reserve some HARQ processes for non-scheduled transmission (i.e. scheduled data cannot be sent using these processes, [they are considered disabled](#)) in case of 2ms TTI. The same reservation scheme for the 10ms TTI is FFS;
- Multiple non-scheduled MAC-d flows may be configured in parallel by the SRNC;
 - The UE is then allowed to transmit non-scheduled transmissions up to the sum of the non-scheduled [granted-rates](#) if multiplexed in the same TTI;
- Scheduled grants will be considered on top of non-scheduled transmissions;
- Logical channels mapped on a non-scheduled MAC-d flow cannot transmit data using a Scheduling Grant;
- Logical channels mapped on a non-scheduled MAC-d flow can only transmit up to the non-scheduled [grant granted-rate](#) configured for that MAC-d flow;
- The multiplexing list restricting the set of HARQ profiles that can be used by a given logical channel will apply both for scheduled and non-scheduled logical channels;
- Logical channels will be served in the order of their priorities until the non-scheduled [granted-rate](#) and scheduled grants are [exhausted/full](#), or the maximum transmit power is reached;
 - When multiple logical channels are assigned the highest priority, the selection of the HARQ profile for these logical channels is not specified.

286. The key phrase “*Scheduled grants will be considered on top of non-scheduled transmissions*” – in the eighth bullet, has to be read in context. Treating the bulleted points as numbered, I draw particular attention to the third, sixth, seventh, ninth, tenth and twelfth bullet points. These convey (and emphasise) a key point to the skilled reader – that the scheduled grant (and transmissions) must be considered separately from the non-scheduled grants and transmissions (recognising that there may well be multiple non-scheduled grants, each of which has to be considered separately). It then follows that neither the scheduled grant nor the non-scheduled grants (nor their sum) include any padding since any padding which is necessary can only be determined and added once the scheduled data and non-scheduled data have been multiplexed *and* the E-TFC has been selected.
287. It also follows that Dr Irvine was wrong to view the grants at the level of the physical layer (i.e. as including padding). The correct level at which the grants must be considered is at the MAC-e level, so including the MAC-d data payload plus the MAC-e header control information.

11.2 TFC and E-TFC selection

Logical channels mapped on the DCHs are always prioritised over those mapped on E-DCHs.

The principle of the TFC selection across E-DCH and DCH is the following:

- The UE performs TFC restriction for the CCTrCH of DCH type;
- The UE performs the TFC selection for the DCHs;
- E-TFC restriction is performed with the following characteristics:
 - The E-TFC restriction mechanism is independent of the existing TFC restriction;
 - The E-TFC states defined per MAC-d flow are managed independently of the TFC states;
 - The UE uses the power offsets for the reference E-TFC(s), the signalled power offset attributes for its MAC-d flows, the required E-TFC dependent backoff, and the UE remaining power to determine the E-TFC states;
 - The HS-DPCCH, DPCCH, DPDCH and E-DPCCH powers are taken into account when calculating the remaining power;
 - The result of E-TFC restriction is a state (blocked or supported) per E-TFC and MAC-d flow;
 - The minimum set of E-TFCs is defined as the number of bits that can be transmitted in a TTI independent of the power situation in the UE, provided there is nothing sent on the DCH, and is configurable from the RNC as one E-TFC per UE. When there is nothing sent on DCH, the E-TFCs belonging to the minimum set are in supported state;
 - ~~The HS-DPCCH and DPCCH are taken into account;~~
- Every E-DCH TTI, the UE shall estimate the remaining power versus DPCCH (taking into account HS-DPCCH);
- Then the UE# performs the E-TFC selection for the E-DCH, with the estimated remaining power, taking into account DCH power and the following rules:
 - The E-TFC selection is based on logical channel priorities like in the Release '99, i.e. the UE shall maximise the transmission of higher priority data;
 - The UE shall respect the allowed combinations of MAC-d flows in the same MAC-e PDU;
 - The UE shall use the multiplexing list of the different MAC-d flows to see if a certain MAC-d flow can use the power offset of the highest priority MAC-d flow to be transmitted;
 - The power offset of E-DPDCH(s) relative to DPCCH associated-~~The supported/blocked E-TFCs for to a~~ MAC-e PDU including MAC-d PDUs coming from one or several MAC-d flows are obtained#~~is set~~ as follows.;

- The UE uses the E-TFC restriction result (i.e. blocked/supported E-TFCs) selects the power offset of the HARQ profile associated to the MAC-d flows with the highest priority logical channel in the MAC-e PDU;
 - Among the supported E-TFCs, the UE selects the smallest E-TFC that maximises the transmission of data according to the non-scheduled grant(s) and the serving grant.
- The UE adds the resulting power offset for the MAC-e PDU to the previously calculated reference power offsets for different E-TFCs. It then selects the E-TFC, taking into account the obtained power offsets, the required E-TFC-dependent power back-off, the UE's remaining power and the amount of data to be transmitted.
- For each transmission, the MAC-e entity gives the selected power offset of E-DPDCH(s) relative to DPCCH to the L1 in addition to the E-TFC;
 - In case the maximum UE transmit power is exceeded, the UE shall scale down the E-DPDCH only on TTI level for retransmissions. At slot level all uplink physical channels are equally compressed. The scaling down of only E-DPDCH at slot level remains FFS.
 - In case the selected E-TFC leads to a case where the maximum UE transmit power is exceeded because the selected E-TFC is part of the minimum set in addition, the UE may need not to go below a minimum rate for the E-DCH. In some case, this means that the UE shall may have to power scale down all physical channels present;
 - An E-DCH minimum set because of power limitation is needed. Details are FFS.

288. As noted above, E-TFC restriction is separated from E-TFC selection. As for the latter stage, in his second report, Dr Irvine described the key 'according to the [grants]' sentence in section 11.2 as 'imprecise', to which Mr Townend responded by saying that in context, the language would be understood by the skilled person as fairly clear. He said it is telling the reader that the UE should choose the smallest E-TFC which sends as much (payload) data as possible until the grants become limiting. In other words, if it was possible to send more payload data in a larger E-TFC, then the UE should do so, until it reaches the limits of the grants. Mr Townend observed that on Dr Irvine's approach, the UE does not do so, but chooses an E-TFC which does not maximise the sending of payload data and stops before all of the grants are exhausted.
289. Accordingly, in my view, the key sentence 'according to the [grants]' cannot sensibly be read by the skilled reader as directing him or her to add the scheduled and non-scheduled grants together for the purposes of an initial E-TFC selection step, as Lenovo's case requires. In any event, I note that Lenovo eventually did not contend that this phrase in section 11.2 had that effect.
290. Instead, (see the extract from paragraph 87 of Lenovo's closing quoted above), Lenovo rely on a piece of alleged CGK to achieve that effect. The suggestion here is that the skilled reader would have in mind that the grants were an absolute limit on the uplink resources that the UE could use.
291. I have no hesitation in rejecting this argument for a number of independent reasons, any one of which requires its rejection:
- (a) First, it is necessary to be aware of the sleight of hand involved here. When only scheduled data was involved, the grants (plural because the single grant was made up from two - the Absolute and Relative grants) concerned the same type of data and consequently it was appropriate to take them together. Not so when one is considering a scheduled grant and non-scheduled grant(s). As section 10 emphasises, they must be considered separately.

- (b) So, this supposed piece of CGK regarding scheduled data only is simply not CGK for the situation under consideration in Filitrault. It is not a good basis for further action at all because, as Mr Townend indicated, the model had changed. Furthermore, section 10 directs the skilled reader away from using the added grants as a limit.
 - (c) Lenovo's case is also based on a misunderstanding of the role of pieces of knowledge which are (or might be) CGK. CGK is not a reservoir from which pieces of information can be plucked to supplement an alleged anticipation where the document in question is lacking. It is only if that piece of knowledge would come to the skilled person's mind directly during or on their reading of the document, that it can be taken into account. As already indicated, in the circumstances of Filitrault, what Lenovo requires would not.
 - (d) In this case, it would be far too much of a stretch to find that the skilled person reading Filitrault would read into that document this additional role for the grants.
292. The second point which arises on section 11.2 is that it illustrates that there is no absolute limit on UE transmission power in a TTI (see the last bullet point in particular). IDC also drew what seemed to me to be a valid distinction between cell-wide control of noise/interference (which is what the blue rectangle is concerned with) and control of individual UEs (which is what is in issue).
293. There are some other points to clear up. Mr Abrahams relied upon 'the ordinary meaning' of certain key parts of Filitrault, namely, in sections 9, 10 and the 'according to' sentence. He secured some answers from Mr Townend in cross-examination which were said to support these 'ordinary meaning' submissions. However, the answers relating to particular parts of sections 9 and 10 were extracted by ignoring the wider context of Filitrault. His question concerning the 'according to' sentence was expressly put on the assumption that the skilled reader would understand that the grants *are* a limit on the size of the MAC-e PDU. Since the premises for the argument were not valid, this way of arguing or supporting anticipation also fails.
294. It is not necessary for me to deal with all the slightly different ways in which Mr Abrahams sought to establish that the disclosure of Filitrault was either clear and unambiguous or that it anticipated. Not only did I conclude that the disclosure did not anticipate, but, even if the skilled person brought to his reading of Filitrault a conviction that the grants operate as an absolute limit on the uplink resources which can be used by the UE, I am not persuaded that the disclosure would have been clear and unambiguous. Filitrault does not disclose or even hint at the additional selection step.
295. Overall, I formed the distinct impression that Lenovo's case on Filitrault was driven by the knowledge of what their case required to get within the claim, rather than any objective assessment of what the text actually said, when read by the notional skilled person at the priority date.

Points A and B.

296. In his second report, Dr Irvine relied on what I have characterised as Points A and B as pointing towards the ‘round down’ approach of Mr Townend’s paragraph 305(b). In particular, he said that the phrase in Point A namely that ‘*[t]he E-TFC Selection function controls the multiplexing function*’ in section 7.2.5 suggested that E-TFC selection should occur before multiplexing.
297. As far as I could detect, by the time of closing submissions, what I identified as Points A and B above did not form part of Lenovo’s arguments. This was not surprising. Section 7.2.5 and Figure 7.2.5-1 are expressed at far too high a degree of generality to support the argument that Filiatrault anticipated. In any event, as Mr Townend explained in his third report, the skilled person would understand from the explanation in section 7.2.5 that the E-TFC Selection entity and the multiplexing entity would operate in co-operation rather than in strict hierarchy. So the E-TFC Selection function would have to tell the multiplexing function how many PDUs from which logical channel will need to be multiplexed together and which E-TFC to use. However, for example, the E-TFC Selection function would need to know how much data was in the relevant buffers and Figure 7.2.5-1 shows the connection between the MAC-e/-es and the MAC-d is through the multiplexing entity. Mr Townend gave other examples illustrating his co-operation point:
- (a) If the UE is data limited, the multiplexing entity would have to communicate with the E-TFC Selection entity to inform it there was no more data available;
 - (b) If the UE is power limited, that would also have to be fed into the E-TFC Selection entity, but that is not shown in the figure either.
 - (c) The figure also does not show any input for the non-scheduled grant(s), but Dr Irvine and Lenovo clearly understood they affected E-TFC selection.
298. Further the text in section 7.2.5 against the ‘E-TFC Selection’ bullet point indicates that it had not been updated to take into account the introduction of non-scheduled grant(s).
299. For these reasons, I agree with Mr Townend’s view that Dr Irvine had read far too much into section 7.2.5 and figure 7.2.5-1.
300. Finally, in case the answers are not already clear, I return to consider the questions I mentioned above:
- (a) First, do the grants indicate a maximum amount of data which can be sent? No. Furthermore, I am satisfied that Dr Irvine’s mindset (to the effect that the grants do indicate a maximum amount of data which can be sent) was not one shared by the skilled person or brought by him or her to their reading of Filiatrault at the Priority Date.

- (b) Second and more precisely, is ‘the grant’ a limit on the total amount of resources that may be used (including MAC-d flow data, header and padding)? No, see paragraph 287 above.
- (c) Third, does Filiatrault disclose the additional selection step i.e. step 2 on the right-hand side of X3? Not only does Filiatrault not disclose this additional selection step, but the argument that it did depended on the skilled person sharing Dr Irvine’s mindset, a point I have firmly rejected.

301. For all these reasons, I reject all of the various ways in which Lenovo argued that Filiatrault anticipated.

Obviousness

302. It is not necessary to set out any of the well-known propositions of law which apply to an obviousness attack. Suffice to say I have them fully in mind. I will deal with the two ways in which the attack was presented.

Obviousness as presented in Dr Irvine’s report(s)

303. In his second report at 4.17, Dr Irvine explained why he thought the patent was obvious over Filiatrault, even if Filiatrault did not anticipate. His summary was:

‘4.17 Therefore the step involved to make a Filiatrault system operate in the round down fashion that Mr Townend contends the Patent operates in can be described in three ways:

4.17.1 Make the grants a limit on the total number of bits sent; or

4.17.2 Perform the E-TFC selection before multiplexing; or

4.17.3 Use the grants as part of the E-TFC selection process rather than in the multiplexing process.’

304. Dr Irvine put forward three reasons in support. The first started with his skilled person having in mind two pieces of knowledge:

- (a) The first was that the ‘round down’ approach was a known resource allocation strategy (referring to his paragraph 3.4);
- (b) The second was that the skilled person would be aware of how TFC selection was performed for the legacy channels, a process which Dr Irvine said the skilled person would recognise as a ‘round down’ approach.

305. I have rejected both of these as being CGK, for the reasons set out above. Even if they had been CGK, I agree with Mr Townend that the skilled person would not think that the role of maximum transmit power in TFC restriction had any particular parallel with the use of the grants in E-TFC Selection. Indeed, to characterise each as a ‘round down’ approach struck me as redolent of hindsight. They were said by Dr Irvine to make it obvious to perform E-TFC Selection

according to a round down approach. I reject this reason as being contrived and the result of hindsight.

306. I recognise that in his second report Dr Irvine had reminded himself of the instruction he had set out in his first report when approaching obviousness and that it was important to avoid hindsight. Notwithstanding that, it is not clear to me that Dr Irvine took any steps to attempt to avoid hindsight when addressing obviousness in his second report. Although he identified the basis on which he was considering obviousness i.e. Mr Townend's interpretation of Filiatrault, it was not clear what the starting point was for the skilled person or what the skilled person's motivation was for altering Filiatrault. Of course, by that point, it was extremely difficult for Dr Irvine to avoid hindsight due to the intense focus in the first reports on Filiatrault and the Patent.
307. Dr Irvine's second reason was that each of his three ways '*are clearly contemplated and suggested by Filiatrault*' in certain of the key passages. I need mention only three to deal with this reason: the first was the sentence in section 7.2.5 [i.e. Point A], the second was a sentence from section 9.1 that '*Scheduling Grants are only to be used for the E-DCH TFC Selection Algorithm*', the third was Point C. Dr Irvine's reasoning was that each points towards the round down approach of paragraph 305(b). I have to say I consider this was wishful thinking on his part, knowing where he needed to get to and how he had to get there.
308. Dr Irvine's third reason was that even in a round up system, it will necessarily operate in a round down style if it is power limited. In other words, if the UE is constrained by the maximum power it can send at, it would therefore have to choose an E-TFC smaller than the limit of the grants i.e. in a round down fashion. Therefore, so Dr Irvine reasoned, it would then operate by first choosing the E-TFC and then multiplexing the data to suit.
309. In response to this third reason, Mr Townend thought it similar to Dr Irvine's first point but that it also confused two separate concepts – on the one hand, the maximum UE transmit power which is relevant to E-TFC restriction and, on the other, E-TFC selection. As Mr Townend pointed out, the maximum UE transmit power does not cause the grants to be 'rounded down': instead, it affects which E-TFCs are considered supported or blocked. Once again, I agree with Mr Townend. I also find Dr Irvine's third reason to be hindsight reasoning.
310. Returning to the three ways identified by Dr Irvine in his 4.17, although the step can be described in simple terms, his whole analysis appeared to me to be shot through with hindsight, and I reject it. Nowhere was it suggested that the skilled person would perceive any issue over padding or a need to increase the efficiency with which padding was used. In any event, as IDC submitted, this was not the obviousness case put to Mr Townend in cross-examination. This was not a surprise, since the obviousness case put forward by Dr Irvine was dismantled in Mr Townend's third report, as I have found in the paragraphs above.

Obviousness on IDC's reading of Filiatrault

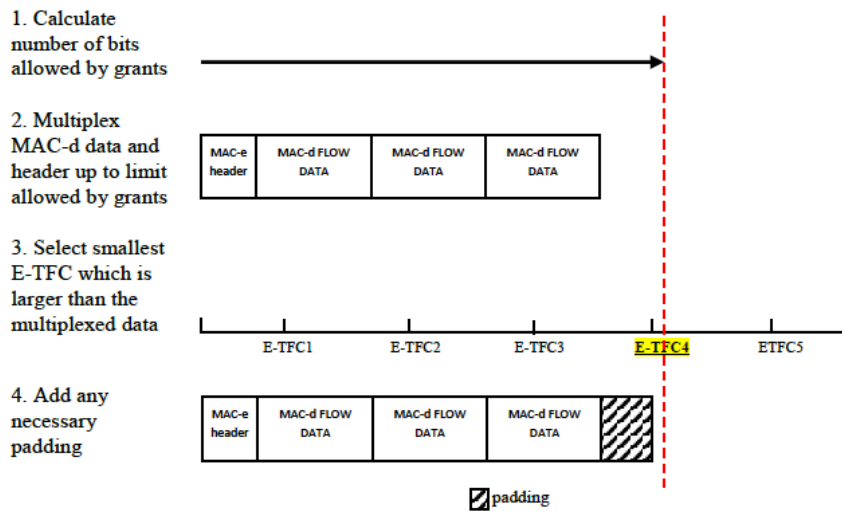
311. Lenovo's obviousness argument, as presented in its closing submissions, proceeded as follows:
- (a) If the skilled person read Filiatrault as Mr Townend did, he or she would recognise the change from the original concept presented in the Feasibility Study.
 - (b) The skilled person would recognise the problem whereby the UE could use more uplink resources than granted and would be concerned that something could go 'horribly wrong' if one went over this limit (per Dr Irvine). Furthermore, limiting resources in this way was the standard way of specifying resources for a radio engineer.
 - (c) There are clear pointers in the text of Filiatrault to a system where the grants specify the maximum uplink resources that the UE can use – relying on the 'ordinary meaning' argument again.
 - (d) It requires no invention to prefer a system in which the grants operate as a limit on uplink resources used by the UE; and, if more than one grant is applicable, for the maximum to be the sum of the grants.
312. I was not persuaded that the skilled person would think back to the original concept presented in the Feasibility Study. The proposal for the Enhanced Uplink had progressed a long way since that point. The original concept made sense when the system was dealing with only a single grant – for scheduled data (even if conveyed by means of the Absolute and Relative Grants). However, the introduction of non-scheduled grant(s) and the clear teaching that the two types of grant had to be considered and used up separately gave the skilled person no reason to think back to a single grant situation.
313. I was also not persuaded that the skilled person would share Dr Irvine's degree of concern about the consequences of the UE using more uplink resources than the grants allowed. The skilled person would know of (a) power control and (b) the margin of safety built in by statistical multiplexing.
314. Overall, I formed the view that this obviousness argument was also driven by hindsight and I reject it.

DXX/14 – Alleged anticipation on IDC's reading of Filiatrault and construction of claim 1.

315. The final point Lenovo raised on Filiatrault is very short. Lenovo submit that on Mr Townend's reading of Filiatrault (which I have agreed with), the UE will on occasion limit MAC-d flow data multiplexed into a MAC-e PDU to a largest E-TFC size which is smaller than a size of MAC-d flow data allowed by the serving and non-scheduled grants and available for transmission. In its cross-examination bundle for Mr Townend, Lenovo presented this illustration (DXX/14):

Townend 1 paragraph 305(a)

(a) “Filiatrault uses the grants as a multiplexing limit in a MAC-e PDU and then selects the smallest E-TFC which is larger than the multiplexed data, and adds any necessary padding. This allows the selected E-TFC to be larger than the sum of the grants, and/or the padding to be larger than a MAC-d PDU which could otherwise be sent.”



316. Shortly before this illustration was put to Mr Townend, he clarified that his paragraph 305(a) was a bit of a shorthand, in that *‘the grants are used, sort of one by one to collectively create that multiplexing limit, as opposed to sort of being used together in the way we were discussing earlier...’* He also made it clear that his paragraph 305 was considering the scenario where the UE is not power-limited, but grant-limited. So his text from paragraph 305(a), quoted in DXX/14, contemplates that the E-TFC restriction step has already taken place. In the illustration therefore, all the illustrated E-TFCs are supported.

317. The cross-examination on DXX/14 went as follows:

- Q. So in this case, the grants, and again subject to the clarification you made, amount to slightly fewer bits than they did in JI-03. So at Stage 2, adding a fourth MAC PDU would cause the multiplexed data to exceed the grants. You can see that just by eyeballing the page; yes?
- A. Yes, and so I think this is what you were referring to when you were saying "subject to the same thing". So effectively that adding an additional MAC-d PDU would cause one of the grants to be exceeded. It does not much matter which, I suspect.
- Q. So in this case, Filiatrault, as you understand it, would produce a MAC PDU as shown here at Stage 4; yes?
- A. Yes.
- Q. So in this scenario, the UE would limit the amount of MAC-d flow data and associated MAC-e header, to be within the largest E-TFC size that is smaller than the sum of the grants?

- A. In this particular case, by virtue of where the MAC-d PDU sizes and the particular grants that were being assumed, yes.
- Q. Okay. Thank you. That was just checking that we had understood your evidence correctly, and we obviously have.
- A. Yes.

318. The force of this point is that Mr Townend accepted that, with certain particular MAC-d PDU sizes and particular grants as illustrated in DXX/14, Filiatrault had means which produced the result required by integer 1F of the Patent.
319. IDC's only response to this evidence was to brush it aside, saying in their written closing:
- 'All this shows is that a phone that does not have the patented means and therefore never applies the patented limit on multiplexing, depending on circumstances in which it is operating (particularly the E-TFC sizes, the multiplexing block size and where the grants fall) may by use of completely different means, applying different multiplexing limits, end up selecting an E-TFC for use that might have been arrived at by a phone using the patented means.'
320. It is true (as I have found above) that Filiatrault does not disclose or teach the solution of the 'next smaller' scheme as described in the Patent. However, as I pointed out above, the means required by integer 1F are means which produce a particular result, and the scenario presented in DXX/14 achieves that result.
321. There was no attempt by IDC to establish that the occurrence of this type of scenario was *de minimis* and IDC had no other answer to the point. Lenovo's argument struck me as akin to the argument rejected by Graham J. in *Hickman v Andrews* [1983] RPC 147, where the prior art bookbinder's press was alleged to anticipate a patent for a workbench, and where the Judge (upheld by the Court of Appeal) said at p168:
- 'I think one must be realistic about these things when construed in a patent specification and must avoid, if one can, falling into the trap of being astute after the event by ex post facto synthesis to build up an anticipation out of a prior document or prior user in order to make it fit the claim.'
322. Anticipation was avoided in that case through the construction of the term 'workbench' in the claim – the bookbinder's press was not a 'workbench'. In the present case, I have already construed 'means for' in integer 1F and it provides no basis for excluding the result shown in DXX/14 from the claim (cf also *Merrell Dow v HN Norton* [1996] RPC 76 at p.82 per Lord Hoffmann).
323. Accordingly, I find that, on the basis of this DXX/14 argument, claim 1 of the Patent is anticipated by Filiatrault.

324. It also follows that if I had construed ‘means for’ as contended for by Lenovo, the Patent would be valid but neither infringed by Lenovo nor essential to the Standard.

Conclusion

325. For all the above reasons, I find EP 3 355 537 B1 to be invalid and therefore not capable of being infringed by Lenovo’s identified devices or capable of being essential to the Standard.