



Neutral Citation Number: [2023] EWHC 172 (Pat)

Case No: HP-2019-000032

IN THE HIGH COURT OF JUSTICE
BUSINESS AND PROPERTY COURTS OF ENGLAND AND WALES
INTELLECTUAL PROPERTY (ChD)
PATENTS COURT

Rolls Building,
Fetter Lane,
London, EC4A 1 NL

Date: 31st January 2023

Before :

THE HON MR JUSTICE MELLOR

Between :

- | | |
|---|--------------------------|
| (1) INTERDIGITAL TECHNOLOGY CORPORATION | <u>Claimants</u> |
| (2) INTERDIGITAL PATENT HOLDINGS, INC. | |
| (3) INTERDIGITAL, INC. | |
| (4) INTERDIGITAL HOLDINGS, INC. | |
| - and - | |
| (1) LENOVO GROUP LIMITED | <u>Defendants</u> |
| (2) LENOVO (UNITED STATES) INC. | |
| (3) LENOVO TECHNOLOGY (UNITED KINGDOM)
LIMITED | |
| (4) MOTOROLA MOBILITY LLC | |
| (5) MOTOROLA MOBILITY UK LIMITED | |

Douglas Campbell KC and Joe Delaney (instructed by Bird & Bird LLP) for the Claimants
James Abrahams KC, Ben Longstaff and Kyra Nezami (instructed by Kirkland & Ellis
International LLP) for the Defendants

Hearing dates: 10th-13th, 18th May 2022

APPROVED JUDGMENT

Remote hand-down: This judgment will be handed down remotely by circulation to the parties or their representatives by email and release to The National Archives. A copy of the judgment in final form as handed down should be available on The National Archives website shortly thereafter but can otherwise be obtained on request by email to the Judicial Office (press.enquiries@judiciary.uk). The deemed time and date of hand down is 10.30am on Tuesday 31st January 2023.

Mr Justice Mellor:

TABLE OF CONTENTS	Page
INTRODUCTION	4
APPLICABLE LEGAL PRINCIPLES	7
THE EXPERT WITNESSES	7
THE SKILLED PERSON	8
COMMON GENERAL KNOWLEDGE	9
CGK points in dispute	9
Simulations	10
THE PATENT	13
CONSTRUCTION / CLAIM SCOPE	18
‘in response to’	23
‘smaller than needed’	25
ESSENTIALITY/INFRINGEMENT	26
VALIDITY	28
THE PRIOR MAC SPECIFICATION – PMS	28
Alleged Anticipation by the PMS – Case 2.	29
Obviousness over the PMS - Case 1	30
How the evidence on the PMS developed	32
Lenovo’s ‘mapping’ argument	37
Lenovo’s simulations argument	41
The significance of the transmission-blocking problem.	44
Simulations to investigate the various values for T_SIG	46
Simulating T_SIG = ‘no report’	47
Where is the evidence that others identified the problem (and its solution) from their simulations?	47
Secondary Evidence	48
Discussion and Conclusion	51
KIM	53
Disclosure	53
Anticipation - discussion	62
Obviousness	64

Claims 2 & 4	65
OVERALL CONCLUSION	66
CGK ANNEX	67
Mobile Telecommunications Systems	67
UMTS System Architecture	71
Logical entities of the UTRAN	72
UMTS protocol stack	73
Radio bearers and channels	75
PDUs and SDUs	79
Quality of Service (QoS)	80
Power control and the DPCCH	81
Multiplexing	82
Sending data	83
HSDPA	84
HSUPA	84
1 The Introduction of HSUPA into the 3GPP Standards	85
2 High-level overview of HSUPA	86
3 HSUPA protocol stack and relevant channels	90
4 RRC role in HSUPA	92
5 Relevant MAC functions in HSUPA	93
6 MAC-e/es functions in Node B scheduling	94
7 First Function: resource allocation and Scheduling Information	95
8 Terminology	95
9 Resource allocation	97
10 The Absolute Grant	97
11 Relative Grants	98
12 The Serving Grant Update procedure	99
13 Scheduling Information	99
14 Contents of Scheduling Information	100
15 Scheduling Information reporting	101
16 Grants for the transmission of triggered Scheduling Information	103
17 Second Function: E-TFC restriction	103
18 Third Function: MAC-e PDU generation and E-TFC Selection	104
19 Fourth function: Configuring control information sent on the E-DPCCH	107
20 Fifth function: HARQ retransmissions	108
21 MAC-e and MAC-es header information	108

Introduction

1. This is my judgment for Trial C concerning EP(UK) 2 421 318 B1 (**EP318** or **the Patent**). EP318 has a priority date of 21st August 2006 which is not contested. It is said to be essential to the UMTS (3G) standard, Release 7 onwards (**the Standard**). It is owned by the Claimants (hereinafter '**IDC**').
2. The Patent relates to the sending of Scheduling Information (or **SI**) in the context of the Enhanced Uplink (**EU**) also referred to as High Speed Uplink Packet Access or **HSUPA** which was first introduced in Release 6, that release having largely been completed around September 2005. By the Priority Date, work had moved onto Release 7 which included specifying further HSUPA functionality. Early Release 7 specifications started to be released shortly before the Priority Date, although the first commercial Release 7 network did not launch until at least February 2007.
3. Release 6 and TS 25.321 v7.1.0 specified certain trigger conditions which would cause a UE to send Scheduling Information to the base station (or Node B). The Patent is concerned with what it calls the transmission blocking problem and the invention is said to be a new trigger condition for sending Scheduling Information which overcomes that problem.
4. IDC says that the invention in the Patent was first incorporated into TS 25.321 in v7.6.0, released on 24th October 2007.
5. On validity, Lenovo allege the Patent was invalid based on two pieces of prior art:
 - (a) The first is a 3GPP Technical Specification TS 25.321 v7.1.0 entitled 'Universal Mobile Telecommunications System (UMTS); Medium Access Control (MAC) protocol specification', published in June 2006. This was referred to at trial as the **Prior MAC Specification** or **PMS**. It was one of the early Release 7 specifications.
 - (b) The second is US Patent Application US 2006/0146761 A1 (**Kim**).
6. At the PTR, it appeared to be the case that there was no issue on infringement/essentiality. In those circumstances, Lenovo said they should begin the trial, since the burden of proving invalidity lay on them and I so ruled.
7. As sometimes happens in these cases, at a late stage the case underwent something of a metamorphosis, in two respects: first, as to Lenovo's invalidity attacks and second, on a key construction issue.
8. At the time when expert evidence in chief was exchanged and only claim 1 was alleged to be independently valid, Lenovo's invalidity case alleged anticipation and/or obviousness over the Figure 8 embodiment in Kim, and obviousness over the PMS. However, in Dr Irvine's first report for Lenovo a different set of attacks emerged. In Kim, now the Figure 9 embodiment was relied upon and the PMS was suggested also to anticipate, depending on the construction of '*in*

response to'. The evolution in Lenovo's validity attacks caused IDC to rely also on claims 2 and 4 as independently valid.

9. The exchange of the opening trial skeleton arguments also prompted a flurry of activity around the construction of claim 1 and Lenovo's previous acceptance that there was no issue on essentiality on the basis of the construction adopted in IDC's statement of case on essentiality. It has been clear throughout that the principal issue was the true construction of the words '*in response to*' in claim 1. On reading what IDC had said on construction, Lenovo served a fourth expert report and a Supplemental Skeleton, in which Lenovo took the position that if the construction put forward in IDC's Opening Skeleton was correct, then the Patent was not essential or infringed. IDC responded with a further report from its expert.
10. Fortunately, the parties co-operated to deal with these developments, sensibly agreeing appropriate amendments to pleadings, the service of further expert evidence and that IDC should open, since there was now an issue on infringement/essentiality.
11. Ultimately, Lenovo advanced the three invalidity attacks set out below. I say 'ultimately' because these attacks underwent some change and development during the exchange of experts' reports which I will have to consider later:
 - (a) Case 1: that claim 1 was obvious on the basis of an obvious modification to the PMS. The modification was also said to render claims 2 & 4 obvious.
 - (b) Case 2: that claim 1 is anticipated by the PMS. If that allegation succeeds, then it is said that claims 2 & 4 involve a single, obvious modification to the PMS.
 - (c) Case 3: that claim 1 is anticipated by Kim in his figure 9 embodiment, alternatively that it was obvious, with claims 2 and 4 obvious as well.
12. Lenovo's arguments, particularly when it came to obviousness, were developed and presented with great skill but comprised a lot of detail. When preparing this judgment I found it necessary to pay very close attention to the written and oral closing arguments and the references therein (more so than usual).
13. By way of introduction, I should also say that the alleged invention in this case is one of those rare cases where the invention (if there be one) really lies in identifying the problem. Once the problem was identified, Lenovo's case was that the solution to it is obvious. Lenovo accepted that it is possible for a patent to be inventive on the basis of the identification of the problem at the priority date but submitted that can only apply in a rare and unusual case. Lenovo referred to the explanation of the point by Henry Carr J. in *TQ Delta v Zyxel* [2019] EWHC 562 (Pat):

243. ...Counsel were unable to point to any case where a patent whose solution was obvious was nonetheless held to be inventive on the basis of perception of the problem.

244. Nonetheless, I recognise that such a case is possible. The EPO has recognised that the identification of a technical problem could give rise to an invention, although it has emphasised that such identification will rarely amount to an inventive step. The principles are summarised in the following passages from *Case Law of the Boards of Appeal*, at I.D.9.10.

The posing of a new problem did not represent a contribution to the inventive merits of the solution if it could have been posed by the average person skilled in the art (T 109/82, OJ 1984, 473). It also had to be taken into consideration that it was the normal task of the skilled person to be constantly occupied with the elimination of deficiencies, the overcoming of drawbacks and the achievement of improvements of known devices and/or products (see T 15/81, OJ 1982, 2; T 195/84, OJ 1986, 121). In T 532/88 the board confirmed the established principle that to address a problem simply by looking for ways of overcoming difficulties arising in the course of routine work did not constitute inventiveness.

[...]

In T 971/92 the board emphasised that the appreciation of conventional technical problems which formed the basis of the normal activities of the notional person skilled in the art, such as the removal of shortcomings, the optimisation of parameters or the saving of energy or time, could not involve an inventive step. The appreciation of a technical problem could thus only contribute to the inventive step in very exceptional circumstances.

245. I would not use the expression “very exceptional circumstances”. Each case depends on its own facts. However, I agree with the logic of the Board of Appeal in the case law referred to above. In a field where the person skilled in the art regularly confronts technical problems and is used to solving them, if a real problem exists, she or he is likely to be aware of it.

14. I have kept this point very much in mind and I return to it below.
15. This brief introduction explains why the main sections of this judgment are concerned with:
 - (a) CGK;
 - (b) The Patent;
 - (c) Issues of Construction;
 - (d) Essentiality/Infringement;
 - (e) Validity over the PMS and Kim.

Applicable Legal Principles

16. No issues arose as to how to identify the Skilled Person or the CGK and I do not lengthen this judgment by setting out the well-known principles. The bulk of this case was about obviousness but it is not necessary to set out the well-known passages from *Pozzoli* at [23] or *Actavis v ICOS* at [58]-[73], which I have kept in mind. In *Actavis*, I have paid particular attention to the following passages:
- (a) [63], where Lord Hodge not only (re-)endorsed the fact-specific approach set out by Kitchin J in *Generics v Lundbeck* at [72] but also said that another factor which needed to be considered in that case was the routineness of the research, a factor of relevance in this case. I discuss this point somewhat further below.
 - (b) [70] regarding motive, when the relevant motive in Lenovo's Case 1 is a very general one of seeking to implement a commercially realistic HSUPA system (as opposed to one where every problem has been solved).
 - (c) [72] regarding the need to avoid hindsight.
17. One particular issue concerning construction arose and I discuss that below. Otherwise, the vast majority of this case concerns the facts and my evaluation of them.

The Expert Witnesses

18. The parties called the same expert witnesses as in Trial B: Mr Jonathan Townend was the expert witness called by IDC and Dr James Irvine by Lenovo. Having served three reports in Trial B, Mr Townend's reports in this Trial C were his fourth, fifth, sixth and seventh reports. For his part, Dr Irvine styled his reports for this Trial as 1C, 2C, 3C and 4C. For the most part and for simplicity, I simply refer to their first, second etc reports as filed for this trial. Errata sheets were served for each expert, which saved time at trial.
19. IDC levelled a number of criticisms at the evidence given by Dr Irvine: first, that his evidence was a mixture of hindsight and speculation; second, concerning the way he was instructed; third, that he had the wrong idea as to what was CGK; fourth, that he failed to distinguish between his own perspective and that of the skilled addressee; and fifth, he was not sufficiently careful in his evidence. These criticisms were developed in IDC's written closing, to which Lenovo responded in a supplemental document, which I have taken into account. In summary:
- (a) On the first point, it is true that Counsel managed to get Dr Irvine to accept he was speculating on a few points. On some of these, Dr Irvine accepted the point rather too readily since part of his job as a suitably qualified expert witness is to predict what the skilled person would have done in circumstances where there is no evidence available as to what

real people actually did at the time. I have assessed his evidence accordingly.

- (b) Hindsight is the other part of the first point and really the whole of the second. I have to consider this below.
 - (c) The third and fourth points are essentially the same. Dr Irvine is an enthusiastic teacher of the technology and this can, on occasion, give the impression that he has strayed beyond the viewpoint of the man of ordinary skill. Whether in fact he did so is again something I have to keep in mind below.
 - (d) The fifth point was valid, but only on a particular part of his evidence. I discuss this below.
20. Overall, however, although both experts are very knowledgeable in this technology, in this trial I am inclined to place slightly more reliance on Mr Townend. This is partly because Dr Irvine had to make certain corrections to his evidence, the significance of which I consider below, but also because Mr Townend seemed to be able to speak with more authority than Dr Irvine. Indeed, on a few points, Dr Irvine deferred to Mr Townend.
21. Having made that point, both experts were trying to assist the Court, and I am very grateful to both of them for their assistance. Dr Irvine is a very engaging witness who is anxious to get things right and who was clearly very concerned to avoid hindsight. Whether he succeeded is something I have to consider below.

The Skilled Person

22. 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.
23. 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. Mr Townend accepted that the products included both UEs and equipment in the network (Node Bs). 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 v6.9.0) and the HSUPA stage 2 overall description (TS 25.309 v6.2.0).

Common General Knowledge

24. As is usual in these cases, there was very significant agreement as to the CGK, and many of the concepts are familiar since 3G mobile telephone systems have been the subject of a number of judgments in the Patents Court, including in particular my own judgment between these parties in Trial B: [2022] EWHC 10 (Pat). I was tempted to incorporate substantial passages from the CGK section in that judgment and to supplement those passages as necessary to bring the CGK up to the position at the priority date of this Patent. However, on reflection, there were some subtle changes required even to rather basic passages and some detail which was not relevant for this trial, which would have made the presentation somewhat messy. Accordingly, I decided it was better to set out a self-contained CGK section.
25. In accordance with an agreed direction in the PTR Order, the parties addressed themselves to indicating which paragraphs in the CGK sections of the experts' reports were agreed to be CGK and to preparing a list of points in dispute. I observe this particular direction is not as useful as the preparation of a statement of agreed CGK, since the product still requires examination of two sets of apparently agreed CGK to determine whether there are any differences which matter. The preparation of a statement of agreed CGK should be relatively easy, precisely because it is not necessary for the lawyers to get involved in arguing over detailed points of wording.
26. Having made those points, the process resulted in a number of paragraphs from Mr Townend's reports being agreed following corrections and clarifications provided by Dr Irvine. Where there were disagreements, the process also brought out the reasons for them. This was very useful work. As IDC submitted, the major point which needed to be clarified concerned E-TFC Selection.
27. Accordingly, much of the CGK Annex is taken from Townend 4, with the corrections and clarifications provided by Dr Irvine and a few additions from Dr Irvine's reports and Townend 5. Aside from the points discussed in the next section, it was convenient to address other relatively minor disagreements in context, which I have done in the CGK Annex. The remainder of this Judgment assumes the content of the CGK Annex has been read and understood.

CGK points in dispute

28. There was a mini-dispute between the experts as to the status of Release 6 at the Priority Date. Dr Irvine disagreed with Mr Townend's characterisation of Release 6 'largely being finalised' at the Priority Date. It appears Release 6 was unusual in that many Release 6 specifications were 'frozen' at the end of 2004 but in practice many 3GPP working groups continued development. Dr Irvine accepted that there were a number of fairly significant updates to the specifications through the course of 2005 to add 'implementation details', but he thought for almost all the specifications this had been completed by

September 2005, with only corrections and clarifications being added after that date. Mr Townend indicated that the enhanced uplink functionality had only just started to be added when Release 6 was frozen.

29. The significance of this dispute concerned when the Skilled Team would start simulations of their HSUPA system. I address this below, but the dispute seemed to me to make no real difference. The content of TS 25.321 at the Priority Date was not in dispute, nor that all the relevant material was in the specification by September 2005.
30. A significant dispute over CGK concerned the issue of E-TFC Selection. In his second report, Mr Townend pointed out various respects in which he said Dr Irvine's description of E-TFC Selection (which included the quantisation of the Serving Grant) was incorrect.
31. Although Dr Irvine presumably saw Mr Townend's comments shortly after the second round of expert's reports were exchanged around 30th March 2022, and although Dr Irvine served two further short reports on 22nd April and 7th May, it was not until 9th May 2022 that an "Errata Sheet" for Dr Irvine was served. Although the amendments he made to paragraphs 6.75 and 6.77 of his first report appear relatively minor in nature, they were significant.
32. In cross-examination, Dr Irvine accepted that the mistake in his first report was failing properly to describe the E-TFC Selection process. He said he had summarised it poorly because he was 'writing quickly'. IDC pointed out that this was a particularly unfortunate error for Dr Irvine to have made given that a similar issue arose in Trial B. I do not propose to re-visit the point which arose in Trial B, although I am aware of it. It is sufficient for me to note that Dr Irvine accepted that this part of his evidence was not prepared as carefully as the Court was entitled to expect.
33. Ultimately, with Dr Irvine's corrections to his first report and his explanations in cross-examination, no dispute remained as to the process of E-TFC Selection. However, Dr Irvine's 'Errata Sheet' contained further changes on other points in his evidence. I will pick those up at the appropriate point.
34. By far the most significant point identified by way of CGK disputes was the scope of the simulations which the Skilled Person would have undertaken and his or her focus on and understanding of the results of such simulations. This point is fundamental to Lenovo's Case 1 (obviousness over the PMS). It is relevant to consider how the evidence on simulations developed.

Simulations

35. It was not in dispute that a network operator which was seeking to implement HSUPA would undertake simulations of the system operation under different configuration options in order to try to optimise system efficiency. However, Lenovo's invalidity attacks based on the PMS effectively required the carrying out of simulations which would have revealed the transmission blocking problem. This case raised important questions as to (a) what simulations would have been run and (b) what the results would have been.

36. The evidence relating to simulations underwent significant development. In his first report for this trial, as IDC submitted, Dr Irvine mentioned simulations in his CGK section almost in passing, in this paragraph:

‘The fact that the standard allows flexibility in system operation means that the signalling is designed to support a wide range of different operational choices. This results in a very large range of configurable parameters, and system operators and equipment vendors are faced with the challenge of having sensible options for these parameters. For example, different MAC-d flows can have different priorities assigned to them, but operators tend to use a subset of the possibilities, or even configure all traffic to have the same priority. As the system is deployed, operators and vendors would then simulate system operation under different configuration options in order to try to optimise system efficiency. This would be at a greater level of detail than the conformance tests. One of the challenges of this approach is that different operators’ networks will have different characteristics (mixes of traffic, cell layouts, etc.), which may mean that the optimal values for different parameters will change between networks and even over time within the same network. In a number of cases, the anticipated flexibility allowed for in the standard has never in fact been implemented in live deployments, as the perceived benefit was not felt to be worth the effort in optimisation.’

37. Later, after being asked to consider what the Skilled Person would do with the PMS in August 2006, Dr Irvine stated the Skilled Person would have been interested in how the standard could be improved, in particular improvements suggested by the experience of implementing the new features of Release 6 through simulations and initial deployments.

38. By August 2006, Dr Irvine stated, Release 6 had been frozen and the possible values of E-TFC sizes, Serving Grant and T_SIG and T_SING had been decided, and testing and simulations of the operation of the system would have been conducted as a matter of routine. He continued (in his paragraph 7.28):

“One objective would have been to identify good values of T_SIG and T_SING from a system efficiency viewpoint in order to identify possible improvements for Release 7. This would involve running simulations of a system operating according to the standard as specified in TS 25.321 and assessing parameters such as throughput and latency. Latency is the delay of transmission of information through the network. This simulation would need to be done in order to trade off the signalling load against scheduling responsiveness. Such optimisation may not have been undertaken as part of the standardisation of Release 6 but would be done as the networks started to be deployed, which occurred in the first half of 2007. This would have indicated to the Skilled Person where there was any latency in the system and therefore where the system could be optimised.”

39. In his second report, Dr Irvine attempted to bolster his evidence in relation to simulations with these passages (bearing in mind that Mr Townend had not yet responded to any part of Dr Irvine’s first report).

“In any event, the Skilled Person working on developing the Release 7 standards, including those with an interest in developing HSUPA equipment, would have started work on implementations of Release 6 once it was complete. As part of this work, they would have carried out simulations of the type I described in Irvine 1-C (see paragraphs 7.27 and 7.28). These simulations would be undertaken to help inform Release 7 improvements, focussed on the key quality metrics of throughput and latency.

Wider simulations would also be performed, focussed either on the performance of specific parts of the HSUPA system, or on the more general performance of the system as a whole from the user’s point of view. Examples of both of these types of simulations can be found in Holma and Toskala (Eds), HSDPA/HSUPA for UMTS, John Wiley, published 4 months prior to the Priority Date. Chapter 8 titled “*HSUPA bit rates, capacity and coverage*”, shows the results of simulations of physical layer performance, HARQ, and of example Node B scheduling algorithms. Chapter 9, “*Application and end-to-end performance*” looks at end to end performance. The importance of latency to the latter can be seen in the discussion of web browsing satisfaction with delay (Figures 9.20 and Figures 9.21), along with the discussion specifically on TCP (Section 9.3.2).”

40. However, cross-examination established the simulation in Holma & Toskala Chapter 8 is an example of doing a simulation prior to finalising a standard. By reference to the underlying model, Dr Irvine was constrained to accept that this simulation would not have revealed the existence of any transmission blocking.
41. The simulation described in Holma & Toskala Chapter 9 was not a simulator for assessing a HSUPA system. Again, Dr Irvine accepted it was not sophisticated and nothing like precise enough to identify the transmission blocking problem.
42. Cross-examination also revealed that whilst Dr Irvine had experience of simulations of earlier and later systems, he had no personal experience of the optimisation of an HSUPA network. In relation to HSUPA, Dr Irvine accepted there were ‘probably thousands’ of configurable parameters in the 3GPP standards to choose from when setting out to simulate system performance and that there must be ‘at least thousands, perhaps millions of possible simulations depending on what parameters you chose, how you configure them and what combinations’. Cross-examination also established there was no ‘routine’ way to carry out optimisation simulations – different people would have carried out different simulations:

- 4 Q. The reality is that you just do not actually know, other than
5 what you can read in documents, what the detail of simulations
6 would have been, do you? You are having to speculate.
7 A. I mean, not at this period but, a few years earlier, we were

- 8 running simulations for companies to look and see what
9 different parameters they would use for these types of
10 networks. I am familiar in the general case with the sort of
11 simulations that were being undertaken to optimise different
12 parts. I have not been involved and do not have personal
13 experience of the optimisation of an HSUPA network.
- 14 Q. But even the example you just gave, that is not one you refer
15 to in 100 pages of evidence, is it?
- 16 A. No, it is not.
- 17 Q. It is not as if there is an "official list" of generally
18 agreed simulations that should be performed on any given
19 release, is there?
- 20 A. No, well, apart from the conformance tests.
- 21 Q. Yes, but that is talking about something else, is it not?
- 22 A. Yes.
- 23 Q. There was not even so much as an unofficial list, either, was
24 there?
- 25 A. I think there are certain things you would have to do in order
2 to make the system work.
- 3 Q. But you would agree that not all people working on the
4 development of HSUPA (whether it is Release 6, Release 7 or
5 any release) would necessarily have decided to perform the
6 same types of simulation?
- 7 A. That is right.
- 8 Q. What simulations, if any, to perform was a matter of choice to
9 a very large extent?
- 10 A. They would have to be directed at optimising the system
11 operation, that is the whole point of doing it. So that would
12 direct you down a fairly narrow path.
- 13 Q. What are you calling "a fairly narrow path"? We discussed
14 earlier thousands of variables and millions of combinations
15 and now you are saying a "narrow path". It is not that
16 narrow, is it?
- 17 A. For different scenarios, and there are agreed reference
18 scenarios, you would perform simulations across a range of
19 different loads for the different parameters that you had to
20 optimise.
- 21 Q. Is there any particular document you can take me to which
22 might show me what you are talking about?
- 23 A. In terms of the reference?
- 24 Q. No, in the case papers, I meant. A list of how you do
25 simulations.
- 2 A. No.

I revert to the topic of simulations in the context of Lenovo's Case 1 (obviousness over the PMS).

The Patent

43. The Patent is entitled '*Method and apparatus for transmitting scheduling information in a wireless communication system*'. TS 25.321 v6.9.0 is acknowledged prior art on the front page. Lenovo established to my satisfaction that there is no material difference between v6.9.0 and the PMS, so effectively the PMS is acknowledged prior art in the Patent.

44. [0002] explains that the invention is related to HSUPA systems and, more particularly, concerns a method and apparatus for preventing transmission blocking in an HSUPA system.
45. The Background section extends from [0004] to [0013] and sets out some relevant CGK. [0004] places the invention firmly in the context of the Third Generation Partnership Project (i.e. 3GPP) Release 6 and refers to fast Node B scheduling in HSUPA. [0004] continues:
- ‘This faster control results in better control of the uplink (UL) noise rise, which allows operation at a higher average UL load without exceeding the threshold, thereby increasing system capacity. In HSUPA, control and feedback occurs through different physical control channels and information elements (IEs).’
46. [0005]-[0008] set the scene for [0009] to explain the transmission blocking problem. The Patent uses the phrase wireless transmit/receive unit (WTRU) to refer to the mobile device and later explains that WTRU includes user equipment - UE – which is the abbreviation used in the Standard(s). To summarise [0005]-[0008]:
- (a) The Patent explains principles of Node B controlled scheduling. These would have been well-known to the Skilled Person (see paragraphs 65.3 and 71-87 in the CGK Annex). The explanation covers a Node B’s ability to send power control commands on the Absolute Grant and Relative Grant channels and a UE’s ability to send SI and a happy bit. [0005] explains the procedure for a UE to send a happy bit, RSN and E-TFCI on the E-PDCCCH.
 - (b) The Patent explains that data received from higher layers may be buffered at the RLC layer, and may be segmented at the RLC layer, so that the RLC PDU passed to the MAC is of a size configured by RRC signalling. It says that there is currently no further segmentation at the MAC layer.
 - (c) The Patent goes on to explain that an integer number of RLC PDUs (which can include zero PDUs) must be sent in each transmission. A fraction of an RLC PDU cannot be sent. As with the preceding paragraphs, this information would be within the common general knowledge of the skilled addressee. The Patent concludes that, in order to transmit an RLC PDU, a certain minimum bit rate or, equivalently, power ratio, is required in each TTI.
47. Against that background, we reach [0009] which starts as follows:
- “[0009] During scheduled operation, WTRU transmissions from a given MAC-d flow can be completely interrupted, or "blocked," if the granted power ratio falls under the minimum required to transmit the RLC PDU at the head of buffer. Such a

situation may occur out of the control of the serving radio link set, (i.e., Node-B) for a number of reasons.”

48. Three examples are then given:

- (a) a non-serving relative grant requesting a decrease of power is received by the UE from another Node-B (i.e. a DOWN non-serving Relative Grant);
- (b) the UE erroneously decodes a relative or absolute grant command from the serving Node-B;
- (c) a MAC-d flow may have several different RLC PDU sizes configured, and a larger PDU is up for transmission.

49. The consequences are described in the following two paragraphs:

[0010] When such a situation occurs, the WTRU cannot transmit until the time it is scheduled to transmit an SI. Until then, and unless the previous SI has been transmitted recently enough for the Node-B to infer that the WTRU buffer is not empty based on its subsequent transmissions, the Node-B has no ability to determine whether transmission stopped because the power ratio fell under the minimum, or simply because the WTRU has nothing to transmit. Accordingly, transmission from the WTRU is delayed until the SI can be transmitted.

[0011] This issue imposes a configuration of a small periodicity of SI transmission (T_SIG) for delay-sensitive applications, thereby increasing overhead. Furthermore, even if the Node-B was aware that transmission stopped because the power ratio is too low, when multiple RLC PDU sizes are configured, the Node-B does not know what power ratio to apply to correct the situation. Thus, the Node-B has to find out by trial and error what the correct power ratio is. This results in inefficient resource allocation and/or excessive scheduling delays.

50. [0012] is important in at least two respects:

- (a) First, because it not only addresses the acknowledged prior art directly, it acknowledges the relevant part of the PMS by describing the conditions under which SI is triggered in TS 25.321 at the Priority Date.
- (b) Second, because when identifying possible but unsatisfactory solutions to the problem, an aim of the invention is indicated, in my view: ‘a solution to prevent blocking that would be compatible with the mechanisms defined in the current state of the art.’

[0012] In the current state of the art, transmission of scheduling information (SI) is only allowed under certain conditions such as those described in 3GPP TS 25.321, such as if the user has a

grant (power ratio) of zero or has all its processes de-activated and has data to transmit, upon a change of E-DCH serving RLS (base station), or periodically, with a configurable period depending on whether the user has a grant or not. Accordingly, a solution to prevent blocking that would be compatible with the mechanisms defined in the current state of the art may include configuring periodic reporting of the SI with a very low period, such that the SI is transmitted along with almost every transmission of new data. However, overhead may be significantly increased since each SI takes up 18 bits. For instance, assuming a MAC service data unit (SDU) size of 280 bits and a MAC-e header size of 18 bits, this would represent an additional overhead of approximately 6%.

51. Having identified the problem, in a familiar way [0013] says it would be beneficial to provide a method and apparatus for transmission blocking in an HSUPA wireless communication system that is not subject to the limitations of the current state of the art. [0014] then says the invention is related to the method as defined in claim 1 and to a WTRU as defined in claim 7.
52. It was common ground that the Patent describes four embodiments. In brief:
 - (a) The first embodiment is described in [0027] and [0028] with reference to Patent Figure 2. It relates to a method for preventing transmission blocking by sending Scheduling Information in response to a ‘trigger condition’, for example, when it is not possible to transmit a single PDU of a given MAC-d flow.
 - (b) The second embodiment is described in [0029] and [0030] with reference to Patent Figure 3. It relates to a method for preventing transmission blocking by, in response to a ‘trigger condition’, transmitting nothing on the E-DPDCH and setting all 10 bits of the E-DPCCH to zero, and transmitting the E-DPCCH alone.
 - (c) The third embodiment is described in [0031] to [0040] with reference to Patent Figure 4. It relates to a method for preventing transmission blocking by sending a new type of control information, Minimum Grant Information (**MGI**), from the UE to the Node B.
 - (d) The fourth embodiment is described in [0041] to [0045]. It relates to a hybrid scheduled/non-scheduled system.
53. It was common ground that the claims only relate to the first embodiment, so I concentrate on that. Figure 2 is extremely simple:

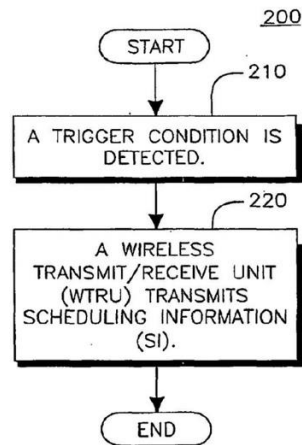


FIG. 2

54. The relevant text is as follows (with my emphasis):

[0027] Figure 2 is a flow diagram of a method 200 for preventing transmission blocking in an HSUPA wireless communication system in accordance with the present invention. In the present embodiment of the present invention, new conditions for the transmission of the SI are created. In step 210, a trigger condition for transmitting an SI is detected. For example, the transmission of the SI alone may occur when the transmission of any, or [in] a specifically defined, MAC-d flow is stopped because the current non-zero grant is smaller than the minimum required to transmit the next MAC SDU, or RLC PDU, of the particular MAC-d flow. The trigger condition, in this case, may occur when it is not possible to transmit a single PDU of a given MAC-d flow. Preferably, a MAC-d flow is a group of logical channels that may be identified, or specified, with an index.

[0028] Once the trigger condition is determined, a particular WTRU 110 transmits the SI (step 220). This transmission may occur once when the triggering condition is met and periodically thereafter, (e.g. over a configurable period), or the transmission may occur at any time the triggering condition occurs. Additionally, the list of MAC-d flows subject to triggering the transmission of SI due to blocking may be signaled by higher layers, as well as the configured periodicity of transmission once the condition is met.

55. Mr Townend indicated that the word ‘in’ (halfway through [0027]) was an error. Certainly the text makes more sense with that word omitted.

56. It is apparent that the reference to the SI being sent ‘periodically thereafter’ is the subject of claim 2, and the configuration of the period by the network is the subject of claim 4.

Construction / Claim Scope

57. 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.

58. The particular issue which arises in this case invokes a canon of construction which can be described in different ways. It is convenient to refer to the following passage from the judgment of Floyd LJ in *Stretchline Intellectual Properties Ltd v H&M Hennes & Mauritz Ltd* [2017] EWCA Civ 199:

'8.....In *Beloit v Valmet (No 2)* [1995] RPC 255 at 270, Jacob J, as he then was, explained that it was not normally legitimate to construe claims using the prior art, because there was normally no reason to suppose that the patentee, when he set the limits of his claims, knew of any individual item of prior art. However, he continued:

"Of course the position is different if the prior art is specifically acknowledged in the patent. The purposive construction would lead to a construction of a claim which did not cover that acknowledged prior art: it can hardly have been the inventor's purpose to cover that which he expressly recognises was old."

9. Jacob J's conclusion that a patentee would not readily be taken to be intending to claim something which he expressly recognised as old is, as Henry Carr J pointed out, no more than an application of the well-known principle referred to by Lord Russell in *Electric and Musical Industries v Lissen* (1936) 56 RPC 23 at 39:

"... if possible, a specification should be construed so as not to lead to a foolish result or one which the Patentee could not have contemplated."

10. Mr Burkill QC, who appeared on behalf of H&M with Mr Geoffrey Pritchard, relied on a passage in *Terrell on the Law of Patents* 18th Edition at 9-254 which suggests that the same approach may be appropriate where the consequence of a construction would be that the patent would be obvious in the light of common general knowledge. However, in *Adaptive Spectrum v British Telecommunications Ltd* [2014] EWCA Civ 1462 I had some reservations about that:

" ... as soon as one departs from documents specifically acknowledged in the specification, the skilled reader has no basis for assuming that the patentee was aware of the document in question. Still further, where the objection is one of obviousness rather than lack of novelty, a value judgment is involved on which widely differing views are possible. It is true that if the document is said to form part of the common general knowledge, it might be said to be more likely that the patentee is aware of it. But a patentee may have been isolated from the common general knowledge, or may, despite the later finding of obviousness, have genuinely believed that he had made an invention over it."

11. I remain of the view that the fact that a particular construction may lead to a conclusion that the invention is obvious in the light of the common general knowledge is likely to be a much less potent factor weighing against that construction than a specific acknowledgement of prior art onto which the claim will read directly.

12. Finally, it is of course well settled that one cannot disregard obviously intentional elements in a claim. Deliberate limitations must have a meaning: see *STEP v Emson* [1993] RPC at 522 summarised in *Virgin Atlantic* at paragraph 5(vii).'

59. The point was clearly flagged by IDC in its Opening Skeleton argument and Lenovo responded in some detail in its written Closing, making three submissions on the law:

- (a) First, that generally a claim should not be construed with an eye on prior material, in order to avoid its effect, relying on *Terrell* at 9-72.
- (b) Second, a mere reference to a prior patent does not require the skilled addressee to find it and study it in detail before arriving at an understanding of the meaning of the language used in the claims, relying on *Terrell* at 9-76, citing *Jushi Group v OCV* [2019] RPC 1 CA.
- (c) Third, a validating construction is not to be preferred, relying on *Terrell* at 9-85, citing *Warner-Lambert v Generics* [2018] UKSC 56, per Lord Briggs at [95] and [98].

60. Lenovo's second point refers to *Jushi* at [39] where Floyd LJ summarised and approved this observation from Lewison J in *Ultraframe*:

'...much depends on the way in which the prior art is acknowledged. A mere reference to a prior patent does not necessarily require the addressee of the patent to dig it out and study it in detail. On the other hand if the specification identifies some particular feature of the prior patent as disclosing a problem which the inventor claims to have overcome, it may be of considerable relevance in interpreting the width of the claim. It is not that the prior patent is irrelevant in interpreting the patent; it is a question of what to do with it.'

61. As for Lenovo’s third point, I have considered the whole section in the Judgment of Lord Briggs from [91]-[98] where he discussed the principles of construction in issue in that case in the context of a second medical use claim. The two particular passages picked out in *Terrell* read as follows:

‘[95] Nonetheless, in my opinion, validating construction will not usually have a significant place in modern patent law. ...

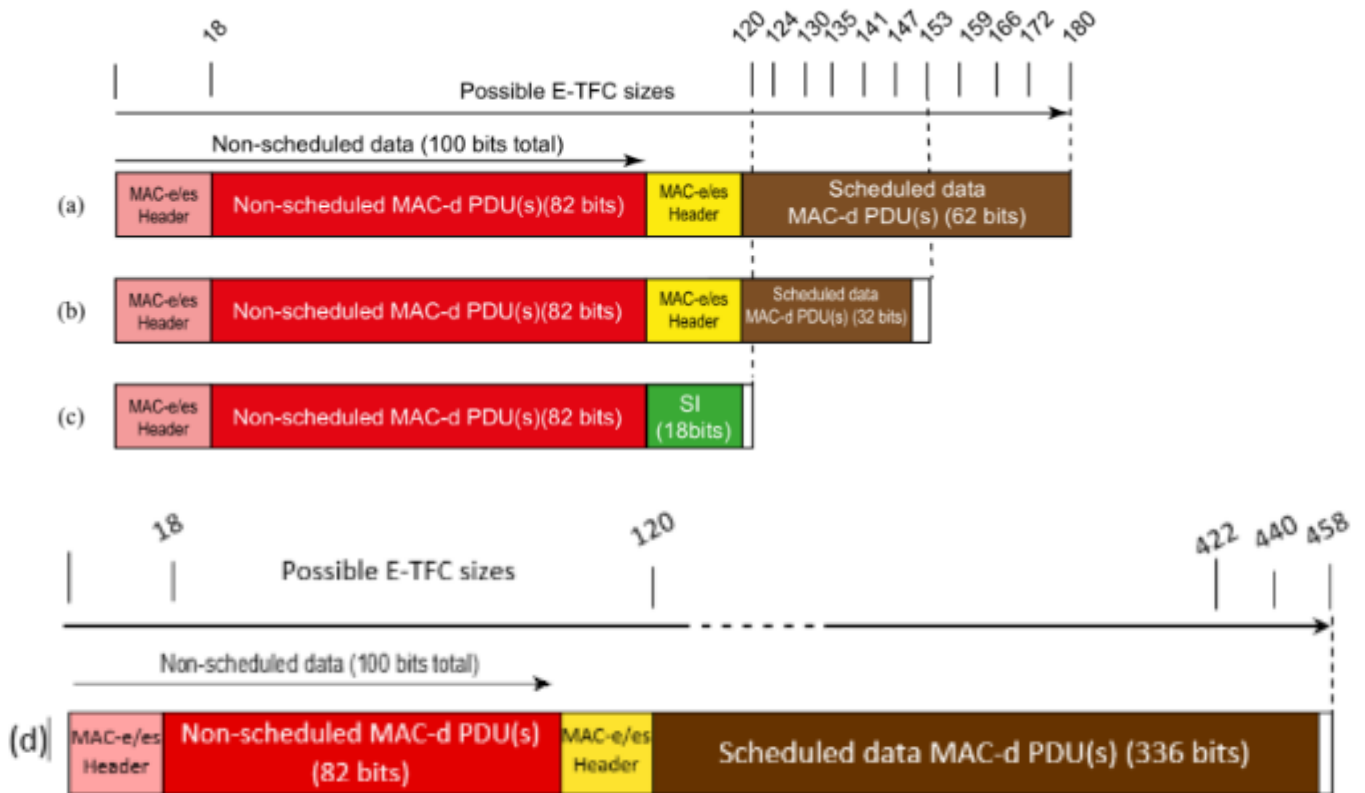
[98]....There are therefore sound reasons of policy for requiring clarity in the claims of patents of this kind. None of this means that claims are to be construed with a predisposition to find fault, or the description read with a mind that is not willing to learn. But it does require that an issue as to the construction of a claim should be addressed, as far as possible, by deciding what it really does mean, rather than by too easily accepting there is an ambiguity, and resolving it by inventing a meaning which saves the claim from invalidity.’

62. I propose to follow this approach.
63. I have also considered the two sections in *Terrell* under the headings ‘Relevance to construction of the legal consequences which follow’ (9-71 to 9-79) and ‘No rule of benevolent construction’ (9-80 to 9-85). These sections usefully gather together all the passages from the relevant authorities. I agree that the correct general approach is that stated by Jacob J. in *Beloit v Valmet (No.2)*, quoted above. I also agree that much depends on the way the prior art is acknowledged, as per Lewison J in *Ultraframe*, as endorsed by Floyd LJ in *Jushi*. In the circumstances of this case, much also depends on the prior art itself i.e. that the prior art is itself CGK, and its content.
64. In this case, the terms in issue are underlined in what follows. IDC contends that claims 1, 2 and 4 are independently valid. The relevant wording of those claims is as follows (with reference numbers removed):

1	A method for use in wireless communication, the method comprising: sending scheduling information, SI, by a wireless transmit/receive unit, WTRU, <u>in response to having a non-zero grant smaller than needed to transmit a protocol data unit, PDU, of a scheduled medium access control, MAC-d, flow.</u>
2	The method of claim 1, wherein the SI is sent periodically.
4	The method of claim 2, further comprising: receiving an indication by the WTRU of a period for periodically triggering transmission of the SI.

65. It was common ground that in EP318 there was no intention to do away with the existing triggers for the sending of SI already specified in the PMS, this being consistent with the use of the word ‘comprising’ in claim 1.

66. Lenovo argued that claim 1 provides for SI to be sent when a specified condition is satisfied and that the specified condition is that the UE has a non-zero grant smaller than needed to transmit a scheduled MAC-d PDU. Lenovo were keen to stress the word ‘having’, arguing that claim 1 refers to ‘having’ a non-zero grant which is too small to transmit a MAC-d PDU and therefore that claim 1 covers a method in which SI is sent in consequence of the UE having a grant which matches that description.
67. Lenovo made this argument in order to pave the way for their anticipation case. Lenovo did not articulate what they said ‘in response to’ actually meant, other than to say it had a broad meaning. However, the true breadth of the meaning for which Lenovo was contending is illustrated by their anticipation case, as explained in Dr Irvine’s evidence.
68. Dr Irvine acknowledged that the PMS ‘does not specify that SI should be sent in precisely’ the situation spelt out in claim 1. He continued (emphasis added):
- ‘However, TS 25.321 deals with a situation where if a UE has a grant and it becomes zero while there is still data to send, scheduling information is triggered. Although TS 25.321 envisages this happening because a grant to transmit data has been set to zero, the Skilled Person would realise that having a very small grant is effectively the same as not having a grant, i.e., in both cases, transmission is not possible.
- In fact, when implementing TS 25.321, Scheduling Information will sometimes be sent as a result of having a Serving Grant which is insufficient to send any MAC-d PDU. Following the requirements in section 11.8.1.4, where there are non-scheduled grants which allow non-scheduled data to be sent, but a non-zero Serving Grant which is not sufficient to send any MAC-d PDU, this will sometimes result in Scheduling Information being sent which would not have been included in the MAC-e PDU if the Serving Grant were sufficient to send a MAC-d PDU (because there would be no space left in the selected E-TFC). Whether or not this satisfies Claim 1 depends on the meaning of “*in response to*” in that claim which I understand, as a matter of construction, is for the Court.’
69. In his second report, Dr Irvine provided an illustrative example of how this could occur, in instances (a) to (c), set out in his Figure 1 - Transmission of SI along with non-scheduled and scheduled data. Instance (d) was added in the DXX bundle (unfortunately not set out to the same scale):



70. Dr Irvine’s example assumed (1) the UE has a non-scheduled grant of 100 bits per TTI, 10ms TTIs and E-TFC Table 0 are in operation, (2) the UE is not power limited and (3) there is a Serving Grant at a level which, when converted to a number of bits using the signalled reference E-TFC levels and HARQ offset, allows 80 bits to be sent. This would allow the UE to transmit an E-TFC of up to 180 bits (plus SI if triggered). Of the possible E-TFCs in Annex B.3 of the PMS, E-TFC 11 is 180 bits, so this could be used if scheduled and non-scheduled data is transmitted.
71. Fig 1(a) shows the situation where 80 bits is sufficient to send a MAC-es PDU, that size corresponding to a MAC-d PDU of 62 bits or less. No SI is sent because there is no spare space for it. The same is true in Fig 1(b) where the MAC-d PDU is 32 bits.
72. However, if the 62 bits was not sufficient to send a MAC-d PDU, then no scheduled data could be sent. As a result, the UE would use the smallest E-TFC large enough to carry the non-scheduled data. This is E-TFC 1 which has a size of 120 bits. As Fig 1(c) shows, there is enough space for SI (18 bits) so this would be included. Although Dr Irvine acknowledges his specific example focussed on the 10ms TTI and Table 0, he said that similar situations would have occurred for all other possibilities (Table 1 with 10ms TTIs and for 2ms TTIs with either Table).
73. Mr Townsend responded in his second report. He made valid points that (1) the scenario described by Dr Irvine is not explicitly disclosed in the PMS; (2)

whether SI is sent (or not sent) in either of the hypothetical cases he describes, it is not sent because the reporting of SI has been triggered, but because there is space remaining in the E-TFC selected by the UE. His opinion was that the sending of SI is not as a result of the Serving Grant. This was because the Serving Grant is but one of a large number of variables that can affect the amount of space remaining in an E-TFC. Thus, although the Serving Grant may contribute indirectly to SI being sent, it is not the determining factor. He pointed out that Dr Irvine’s illustration depended on the condition for sending SI in section 9.2.4.2 and not in section 11.8.1.4 (as Dr Irvine had indicated, which is concerned with E-TFC Selection).

74. Mr Townend reinforced his reasoning by presenting four worked examples (in his JRT-10). Example 1 shows a scenario in which the UE will send SI in circumstances where it has a Serving Grant not sufficient to send a MAC-d PDU of scheduled data, when it would not have sent SI given a Serving Grant sufficient to send the MAC-d PDU. Together his examples 1 to 4, in the two scenarios where (1) the Serving Grant was sufficient to send a MAC-D PDU and (2) the Serving Grant was too small to do so, gave all the possible permutations of SI sent/not sent, as illustrated:

	Serving Grant sufficient to send one PDU (Case A)	Serving Grant too small to send one PDU (Case B)
Example 1	Scheduling Information not sent	Scheduling Information sent
Example 2	Scheduling Information sent	Scheduling Information not sent
Example 3	Scheduling Information not sent	Scheduling Information not sent
Example 4	Scheduling Information sent	Scheduling Information sent

75. Thus, as Mr Townend said, Dr Irvine’s use of ‘sometimes’ essentially indicates that a broad range of further pre-requisites and variables (other than the Serving Grant) will actually determine whether or not SI will be sent.
76. Thus, this is one of those cases where it is helpful to consider the validity attack in order to understand the rival contentions on the key issue of construction. With that background in mind, I can turn to the principal issue.

‘in response to’

77. Although the issue arises on these particular words in claim 1, they must necessarily be understood in the context of claim 1 and in the wider context of the Patent as a whole, read with the Skilled Person’s CGK in mind.

78. By the end of the trial the respective positions of the parties were well-defined. In short, IDC contended for a precise and relatively confined meaning of the term so that it was limited to the trigger described in [0027] and [0028], whereas Lenovo contended for a meaning which was considerably broader. Mr Abrahams KC for Lenovo accepted that [0027] disclosed a new trigger for SI and that claim 1 *covers* this new trigger. His point was that claim 1 is much broader. He invited me to ‘just read the words’, which is usually an invitation to construe the claim acontextually.
79. I agree that the expression in claim 1 *can* be read each of the ways contended for by the parties, but Lenovo’s reading effectively requires the context to be ignored.
80. It seems to me that there are a number of related problems with Lenovo’s construction. First, it means the claim does not conform to the embodiment of the invention described in [0027] and [0028]. Although there are four different embodiments described in the Patent, it is clear that the claims do not concern the other three and only concern this one.
81. There is, of course, a well-known technique for patentees to phrase the principal claim broadly and then introduce further limitations in subsidiary claims and sometimes one finds the preferred embodiment only described in a subsidiary claim. However, the Skilled Person (with his/her knowledge of patent claim drafting techniques (cf *Virgin*)) would note that, with Lenovo’s suggested construction in mind, none of the subsidiary claims limit the invention to the new trigger described in [0027] and [0028]. In my view, that would strike the Skilled Addressee as unusual.
82. The second point engages the principles of law I discussed above. Lenovo’s construction means that the claim covers what is described in the acknowledged prior art. This case is perhaps a paradigm case for the application of the principle described in *Beloit v Valmet (No.2)*, in that the prior art is not only expressly acknowledged on the face of the Patent, it is CGK, the problem addressed by the Patent is described by reference to it and, on Lenovo’s construction, it anticipates claim 1. This would be a foolish result which the patentee could not have possibly contemplated.
83. On the facts, Lenovo submitted that, whilst the Skilled Person would be familiar with TS 25.321, the Patent only mentions it in general terms and does not mention or seek to distinguish the part relied upon for anticipation – this being a reference to section 9.2.4.2 which describes so-called ‘opportunistic SI’.
84. It is not correct to characterise the reference to TS 25.321 in the Patent as ‘in general terms’. At [0012], not only is TS 25.321 specifically identified, it is in the context of the patentee explaining that in the current state of the art, SI is only allowed under certain conditions ‘such as those described in 3GPP TS 25.321 such as...’ and it proceeds to mention examples. It is evident from the wording that the patentee is not purporting to give chapter and verse on all the circumstances in which SI is sent, as specified in TS 25.321.

85. However, in my judgment, the discussion of SI in the Background section of the Patent, and particularly in [0012], would put the Skilled Addressee in mind of the sections of TS 25.321 which relate to SI. He or she would either recall them or refer to them – i.e. sections 9.1.5, 9.2.4.2, 9.2.5.3.2, 11.8.1.4 and 11.8.1.6.
86. It is true, as Lenovo submitted, that the circumstances in which so-called ‘opportunistic SI’ are sent, as set out in section 9.2.4.2 of TS 25.321, are not specifically mentioned in the Patent. However, it must be remembered that the whole of that specification is agreed to be CGK. It is unreal to think, which I consider is the consequence of Lenovo’s construction, that the Skilled Addressee would somehow forget about section 9.2.4.2 or think the patentee had deliberately not referred to that section because it wanted to include those circumstances in claim 1. After all, there are cross-references in the sections I identified above. In particular section 11.8.1.6 has an explicit cross-reference to section 9.2.4.2. The consequence of this second point is, as I have already indicated, that on Lenovo’s construction claim 1 covers something which was CGK in the acknowledged prior art.
87. The third problem with Lenovo’s construction is that it flies in the face of these words in [0027]: *‘In the present embodiment of the present invention, new conditions for the transmission of the SI are created.’* The Skilled Addressee would expect the patentee to claim what is new and not to claim what was old. Yet Lenovo’s construction means that claim 1 claims both.
88. Finally, I should deal with one of Lenovo’s objections to IDC’s construction which was that it was based on an attempt to write the word ‘trigger’ into claim 1. Naturally, Mr Abrahams made the point that the patentee could have used the word ‘trigger’ in claim 1 (not least because that word is used in [0027] and [0028]) but chose not to. When I asked him what the purpose could have been of claim 1 being broader than the trigger described in [0027], his response was: to make the claim as broad as possible. In the context of this particular patent, I did not find this explanation convincing.
89. Overall, for the reasons set out above, I find that the process of deciding what this claim actually does mean is relatively straightforward. In the circumstances of this case, a purposive construction requires a construction which does not claim what was old. In context, the expression ‘in response to’ means the new trigger condition described in [0027] and [0028]. I reject the broader construction for which Lenovo contended.
90. It is also relevant to note the origins of Lenovo’s construction argument. It was not identified until Dr Irvine’s first report was being finalised. It was at that point that the possible anticipation argument over the PMS was first identified. Thus, it would appear that Lenovo’s construction was driven by hindsight, the aim being to read claim 1 so that it read onto the CGK content of the PMS.

‘smaller than needed’

91. Neither party identified this as an issue of construction during the trial but I consider it does arise from some of the submissions made. Although I have identified this point using the shorthand expression, naturally I approach it with

the whole context of claim 1 in mind, and indeed the context of the Patent in the light of CGK.

92. The issue arose as part of the debate over Essentiality, which I deal with in the next section. As I explain in that section, Mr Townend said the expression related to the state of the buffer in the UE. If it was empty, the UE has no data to send in a MAC-d PDU and therefore does not *need* any transmission resources. By contrast, Lenovo submitted the ‘need’ did not refer to the UE or its requirements, but to the grant. This submission made no sense and I reject it for the more detailed reasons set out below.

Essentiality/Infringement

93. The relevant trigger condition for sending SI is at section 11.8.1.6.2 of TS 25.321 v7.19.0, and is as follows:

11.8.1.6.2 Report Triggering when SG <> “Zero_Grant” and at least one process is activated

If SG becomes too small to allow transmission of a single PDU from any scheduled MAC-d flow or if the SG is too small to allow transmission of a single PDU from any scheduled MAC-d flow and TEBS becomes larger than zero, the transmission of Scheduling Information should be triggered.

94. IDC submitted that essentiality stands or falls with construction, and that Lenovo’s original concession that EP318 was essential to the standard was rightly made.
95. Lenovo’s position on essentiality was complicated by the manoeuvring just before the trial started. To recap, Lenovo reacted to the way that IDC had explained, in [69]-[71] of IDC’s Opening Skeleton, their construction of ‘*in response to*’ to the effect that the specified condition *necessarily* causes SI to be sent. Dr Irvine’s fourth report was then prepared in response, then answered by Mr Townend. Note that, having seen the reaction, IDC eschewed this way of putting their case. IDC opened the trial on the construction which I have accepted above.
96. In his fourth report, Dr Irvine gave three reasons why, in his opinion, if IDC’s ‘necessarily’ construction was correct, claim 1 would not be essential to the Standard. These reasons were repeated in Lenovo’s written closing. However, in his oral closing, Mr Abrahams accepted that if claim 1 was properly construed as a trigger, then it was essential.
97. Nonetheless, I found it helpful to consider Dr Irvine’s three reasons, in order to confirm (and refine) my findings on construction.
98. In essence, Dr Irvine identified scenarios in which, when the UE has a non-zero grant smaller than needed to transmit a MAC-d PDU, this will not *necessarily* cause the sending of SI. I will identify his three points and then discuss each in turn:

- (a) The first was if TEBS = 0, i.e., the buffer in the UE is empty.
 - (b) The second was when a UE was power limited such that all E-TFCs are blocked.
 - (c) The third point refers to the requirements of section 11.8.1.6.2 of TS 25.321 v7.19.0. The experts were agreed that the point was that SI was triggered on the initial occurrence of the specified condition, but SI would not be sent in subsequent TTIs even if the condition persists.
99. On the first point, Dr Irvine acknowledged that Mr Townend had said, when discussing the PMS in his first report, that there was a general prohibition on sending SI when TEBS is zero, even if SI is otherwise triggered. V7.19.0 contains the same prohibition in explicit terms.
100. Mr Townend said that it was CGK that the buffer in the UE might be empty from time to time. He pointed out that in that situation, the UE has no MAC-d PDUs to send and does not need any transmission resources. Therefore, he said, it makes no technical sense to ask whether or not the UEs Serving Grant is 'smaller' than 'needed to transmit a MAC-d PDU'. Equally, a larger Serving Grant would not impact on the UE's ability to transmit a MAC-d PDU.
101. Mr Townend also made the point that this was consistent with his understanding of the claimed method, which is not performed by a UE whose buffer is empty since there is nothing to send. In short, this first situation does not give rise to the claimed trigger condition.
102. Lenovo's answer to Mr Townend's evidence focussed on his reference to the UE not needing any transmission resources. Lenovo submitted that the word 'need' in claim 1 is not used in that sense, but in the context of describing the grant, not the state of the UE or its requirements. I disagree. In my view if the UE has no data to send, it has no need in the relevant sense, whether one describes that as a need for transmission resources for a MAC-d PDU or the grant needed to send such a (non-existent) MAC-d PDU.
103. As to the second point, that a UE could be power limited was also part of the CGK. In this second situation, the UE is prevented from sending any data (assuming it has some to send), but the reason it is prevented from sending data has nothing to do with its Serving Grant. All E-TFCs could be blocked because all of the UE's power is being used for transmissions on other channels. This is E-TFC restriction.
104. Mr Townend pointed out that in this situation once again the triggering condition is not satisfied. It is the power limit which means there is no 'need' for transmission resources. The size of the grant is irrelevant. The only thing the UE needs is more power.
105. Once again, Mr Townend was of the view that this was consistent with his view of the claimed method, which is not relevant to a UE that is blocked from sending data for reasons other than having too small a grant.

106. In response, Lenovo made the same contention about the meaning of ‘need’ which I rejected above, but added that in the power limited case, the UE still needs a grant sufficient to send a MAC-d PDU, but *also* needs more power. However, in this case, the limiting condition is not the size of the Serving Grant. The UE might have an ample Serving Grant but still be prevented from transmitting by the power limitation.
107. On the third point, on the initial occurrence of the condition, SI is triggered and sent. However, Mr Townend said it would make no technical sense to the Skilled Person if claim 1 required the UE to assess the relevant condition every TTI and (if the condition persisted) to repeat the sending of SI every TTI until the condition ended, for the following reasons:
- (a) First, such a requirement would be contrary to the teaching in [0028], where it explains SI can be sent once when the condition occurs and periodically thereafter, with the period being configurable (see claims 2 and 4).
 - (b) Second, such a requirement would be contrary to the teaching in [0012] concerning increased overhead, which the claimed method avoids.
 - (c) Third, such a requirement would also be contrary to the manner in which the other triggers for the reporting of SI operate. As Mr Townend put it, SI is sent once upon reporting having been triggered, albeit subject to the configuration of any periodic triggering for the reporting of SI.
108. On this third point, Lenovo agreed that the Skilled Person would expect the trigger disclosed in the Patent to work in the same way as the triggers in HSUPA at the priority date. Lenovo submitted that is one of the reasons why IDC’s construction is untenable. This submission was not further explained and I confess I do not understand it.
109. For all these reasons, I agree with Mr Townend that the UE sending SI ‘*in response to having a non-zero grant smaller than needed to transmit*’ a MAC-d PDU is satisfied if the UE sends SI on the initial occurrence of that condition even if it did not then continue to send SI every TTI whilst the condition existed. On the basis of the construction of ‘*in response to*’ which I have set out above, claim 1 is essential to section 11.8.1.6.2 of the Standard. Lenovo accepted that the additional requirements of claims 2 and 4 were essential.

Validity

The Prior MAC Specification – PMS

110. Lenovo initially pleaded that it would be obvious to amend the PMS to provide a trigger within claim 1; and that having done so, the requirements of claims 2 & 4 would simply follow from what was already in the standard. This is ‘**Case 1**’.
111. In the course of preparing Irvine 1, Dr Irvine explained that he spotted that the PMS already provided that, in some circumstances, SI should be sent ‘in

response to' having a non-zero grant which was too small to send a scheduled MAC-d PDU. On this basis, Lenovo amended its statement of case to plead that anticipation case, plus obviousness of claims 2 & 4 on the basis that they only involve a single, minor, obvious modification to that scheme. This is 'Case 2'.

112. In their submissions, IDC placed some emphasis on the way in which Case 2 emerged and the way in which the evidence on Case 1 developed through the expert reports. For those reasons, I will need to summarise these developments.
113. The PMS is a substantial document. Including Annexes, it comprises more than 90 pages of dense material. Lenovo submitted that their cases of anticipation and obviousness over the PMS involve no selection from the document, since both cases assume that the Skilled Person is implementing HSUPA based on that document.
114. IDC submitted that Lenovo's obviousness cases amounted to cases of obviousness over the CGK and relied upon the well-known warnings which attach to such cases – Birss J in *Accord v Medac* [2016] EWHC 24 (Pat), citing Floyd J. in *Napp v ratiopharm* [2008] EWHC 3070 (Pat). However, I consider that the 'inconvenient details' warning has no real application here because Lenovo rely on the whole PMS. To the extent that there are any 'inconvenient details' they are in the PMS itself, and in the very parts on which Lenovo's case depends, or in the developments from the starting point.
115. IDC's second overarching point was that Lenovo's cases were infused with hindsight. I will deal with that in context.

Alleged Anticipation by the PMS – Case 2.

116. Although Lenovo preferred to put their obviousness case first (i.e. Case 1), I will deal with alleged anticipation first i.e. Case 2. Lenovo accepted that Case 2 turns on the construction of claim 1, namely that, if IDC's construction was correct, then Case 2 fails. Since I have found that IDC's construction was the correct one, it follows that Case 2 fails.
117. Nonetheless, it is relevant to the obviousness arguments to consider how the anticipation case arose which led to Dr Irvine's explanation of it in paragraph 9.8 of his first report. This emerged in cross-examination of Dr Irvine:

6 Q. Okay. The point I am talking about now, your para 9.8 point
7 or the opportunistic SI, the non-triggered SI, did you
8 personally identify that point or was it suggested to you?

9 A. We were talking about non-scheduled and scheduled data and
10 going through some examples. (After a pause) We were knocking
11 around some numbers to see what would happen when scheduling
12 information was being sent and that was in a discussion with
13 the lawyers. I honestly cannot remember what was the trigger
14 for that.

15 Q. Okay. Just as you honestly cannot remember the trigger, you
16 cannot honestly remember whose idea it was?

17 A. Oh, that was what I meant.

- 18 Q. Okay.
19 A. It was, we were having a discussion on Zoom and talking about
20 whether or not when Scheduling Information would be sent when
21 there was a mix of scheduling and non-scheduling data, and
22 came up with that situation.
23 Q. When did this idea surface?
24 A. This idea surfaced when we were in the final stages of review
25 of the first report.

118. On the basis that claim 1 was anticipated, Lenovo argued that claims 2 and 4 were obvious. In view of my rejection of the anticipation case, it is not necessary to analyse the arguments in detail. However, in order to get to claim 2 on this hypothesis, it was necessary for Lenovo to argue that it was obvious to change the rules in the PMS relating to the sending of periodic SI, such that the timer was also re-set after sending ‘non-triggered SI’ (also referred to as ‘opportunistic SI’). Dr Irvine suggested it was, but Mr Townend said there were technical reasons why the Skilled Person would be deterred from changing the operation of the periodic timers in this way, not least because of the safeguard that triggered SI can be re-transmitted if delivery of the SI to the serving Node B fails. By contrast, ‘non-triggered SI’ is simply a bonus for which no similar safeguard is required. Furthermore, if the periodic timers were re-set for both types of SI, a situation could arise where there was an extended period where the serving Node B was not receiving any SI at all, because each time the timer was re-set the UE would send another ‘non-triggered SI’ that may or may not be received correctly. Dr Irvine responded with yet a further change, namely, that the Skilled Person would apply the safeguard to both triggered and non-triggered SI. This ‘double modification’ featured in a Huawei Tdoc, but the minutes from the relevant meeting showed the proposal ‘died a death’. Dr Irvine agreed and also accepted that this was entirely consistent with Mr Townend’s evidence that it was not a good idea.
119. For these brief reasons, if claim 1 had been anticipated by the PMS, I would not have found claims 2 and 4 obvious.

Obviousness over the PMS - Case 1

120. In my recitation of Case 1 which follows, there are two aspects to which I return later. The first is the way in which the time or period in which certain key work would have been carried out tended to vary. The second is the way in which the evidence relating to simulations continued to develop through expert reports and indeed through cross-examination.
121. In their Amended Statement of Case on Validity, Lenovo identified the parts of the PMS on which they relied. These were extracts from 9.1.5 MAC PDU (E-DCH), 9.2.4.2, MAC-e header parameters, 9.2.5.3.2 Scheduling Information and 11.8.1.4 E-TFC Selection. It may be noted that the salient parts of each of these sections is covered in the CGK Annex.
122. By the time of closing arguments, Lenovo was contending that its case comprised two steps:

- (a) It would be obvious to identify the transmission blocking problem.
 - (b) If the Skilled Person wished to address that problem, an obvious solution was to define a SI Trigger.
123. Lenovo submitted this case required a straight application of *Pozzoli*, but they stressed two additional points. First, Lenovo referred to the extract from *TQ Delta v Zyxel* which I quoted above, submitting that principle only applied in rare and unusual cases. Second, Lenovo placed particular emphasis on the decision of the UK Supreme Court in the *Tadalafil* case – *Actavis v ICOS* [2019] UKSC 15. Whilst that case is frequently referred to as containing the latest set of principles relating to obviousness (see in the speech of Lord Hodge [58]-[73], as mentioned above), Lenovo rely on it in addition for the proposition that uninventive steps which the skilled team would take after the priority date in carrying out routine work to implement the prior art are not excluded from consideration in assessing the obviousness of the alleged invention at the priority date (see [59]). This proposition was the answer to Lilly’s case which (as Lord Hodge observed at [52]) would require the court to disregard the work which a skilled person would carry out after the priority date to implement the teaching of the prior art – in that case the Daugan patent.
124. I also detected hints in Lenovo’s case that they were seeking to draw or at least interest the Court in a factual analogy with *Tadalafil* but any such invitation must be declined since it is trite that each case depends on its own facts.
125. In *Tadalafil* the Daugan patent was published on 6 February 1997 and the patent in issue claimed priority from 30 April 1999. Thus, apart from Lilly’s argument, there was no temporal element which stood in the way of the work involved in the obviousness argument, although it is clear it would have taken possibly months to carry out. Precisely because it was routine work which led to the identification of a dosage regime of 5mg daily dose of tadalafil, the case, it seems to me, embraces several legal fictions. I say ‘several’ because the Court is required to consider the notional skilled team, possessed of the common general knowledge at the priority date of the patent, being presented with and reading the prior art with interest at that date. The key legal fiction for present purposes is that the skilled person is effectively assumed to carry out all the routine work on or as of the priority date so as to render the result that the patent was obvious at that date.
126. In *Tadalafil* there was no evidence or suggestion that anyone had actually carried out the ‘routine work’ around the Priority Date (or even afterwards). The present case is very different in that simulations must have been carried out by all those manufacturers who were building equipment for HSUPA networks and all those operators seeking to incorporate HSUPA capability into their networks. The key issue is whether those simulations would have given rise to the identification of the so-called transmission blocking problem and the solution to it in the Patent.
127. Lenovo, largely but not exclusively through Dr Irvine’s evidence, developed two ways in which this obviousness argument was put. Each argument reached a stage where it was said, the problem was *staring the Skilled Person in the face*.

These submissions emphasise the especial importance of two issues which often feature in obviousness cases: first, the relevance of the question: ‘if it was obvious, why was it not done before?’; and second, the influence of hindsight.

128. On the second point, I have to consider whether the arguments that the problem was ‘staring the Skilled Person in the face’ stem from hindsight. On the first point, the question requires some modification in the particular circumstances of this case. I need to consider not just the period before the Priority Date but also some months after.
129. However, before I address the two ways in which Lenovo put their obviousness arguments over the PMS on case 1, I must relate how the evidence on these arguments developed, on which IDC placed some weight.

How the evidence on the PMS developed

130. Dr Irvine’s evidence as to what the Skilled Person would do with the PMS, reading it at the Priority Date, comprised a number of stages which I summarise as follows.
131. First, I note that Dr Irvine was provided with the prior art of the PMS and Kim and the relevant date of 27 August 2006 before he saw the Patent. It appears from his first report that he initially discussed the PMS generally. Although he provided an overview of the content of the PMS, Dr Irvine was asked to explain what the Skilled Person was taught about scheduling in the document. He said the Skilled Person already knew quite a lot about scheduling and referred to his CGK section. As he related, the PMS does not specify how scheduling is undertaken but it does specify information which has to be provided to the Node B which will allow scheduling to take place, the information being the SI and the happy bit.
132. In closing, IDC submitted that not only had Dr Irvine been instructed to focus on the issue of scheduling when he read the PMS, he had even been instructed to consider the transmission blocking problem itself. IDC relied on this passage of cross-examination (with IDC’s emphasis):

- 2 Q. Did the first discussion involved discussion of HSUPA or not
3 HSUPA. I thought you said a moment ago it was?
4 A. I would have to go and check my notes. It certainly started
5 as a general discussion of scheduling. Whether it become
6 specific on scheduling in HSUPA I would have to check my
7 notes.
8 Q. I am not asking you to check your notes, but it is obviously a
9 possibility because you would not have mentioned it otherwise.
10 A. Yes.
11 **Q. That was the first meeting and then you were given the TS**
12 **25.321. Was that the point when the solicitors asked you to**
13 **explain what the skilled person was taught about scheduling**
14 **from TS 25.321?**
15 **A. Yes.**
16 **Q. That must have directed you to focus on the scheduling**

- 17 sections of TS 25.321 in particular.
18 A. That is correct, yes.
19 Q. I imagine you read them quite carefully.
20 A. Yes.
21 Q. How long did this phase take?
22 A. As I say, I think it was about two weeks between when I was
23 given 25.321 and when I was given the patent.
24 Q. Over those two weeks you were thinking about the scheduling
25 aspects of TS 25.321?
2 A. Yes.
3 Q. Were you ever instructed to consider what happened to HSUPA
4 when a UE had a smaller grant than that required to transmit a
5 MAC-d flow?
6 A. Not in quite those terms, but yes.
7 Q. Was that before you saw the patent?
8 A. Yes, that was before.

133. As a result, IDC submitted that Dr Irvine's evidence was unavoidably tainted on the issues of both novelty and obviousness by - literally - asking him to have in mind the problem identified by the Patent (which was not CGK) when considering the prior art. I return to consider this point later, in conjunction with the passage from his cross-examination I quoted at paragraph 117 above. For the moment I continue with a description of how the evidence developed.
134. In the section of his report on the PMS, Dr Irvine explained his views (drawing particular attention to the periodic timers T_SIG and T_SING), as well as more general points about SI and the purpose of triggering SI i.e. to inform the Node B about things it will not or may not know. The Skilled Person would have considered SI as useful to allow the Node B to make scheduling decisions and sent by the UE to alert the Node B that resources were required. At the end of that section, he referred to three pre-priority documents, two of which were supplied to him by the solicitors and a Qualcomm document which he thought he had found.
135. Dr Irvine then considered potential developments of the PMS. He considered the Skilled Person would have been interested in improvements suggested by experience of implementing the new features of Release 6 through simulations and initial deployments. Given that by this time Release 6 had been frozen, he said (his paragraph 7.28) that testing and simulations would have been conducted as a matter of course, one objective of which would have been to identify good values of T_SIG and T_SING. He said this would involve running simulations of an HSUPA system to assess parameters such as throughput and latency in order to trade off the signalling load against scheduling responsiveness. He added that such optimisation may not have been undertaken as part of the standardisation of Release 6, but would be done as the networks started to be deployed which occurred in the first half of 2007.
136. In a later section, he identified the point that gave rise to the anticipation argument (his paragraph 9.8). I have already quoted the passage from his cross-examination (at paragraph 117 above) which elicited that this point was identified in discussions when his first report was in the course of being

finalised. He continued with a few paragraphs which set out the essence of the obviousness arguments over the PMS, on which various points arise. It is convenient to quote this section in full here (note that Dr Irvine used the expression ‘stall scenario’ to refer to the transmission blocking problem):

‘9.9 If the Skilled Person did not foresee the possibility of a stall scenario in which a non- zero Serving Grant was insufficient to send a MAC-d PDU from looking at TS 25.321 itself, I would expect this to come to light during routine simulation of the system in order (which I explained above at paragraph 7.28). For example, I would expect simulations to find effective values of T_SIG and T_SING to show increased latency where the Serving Grant is low due to the stall condition. Latency is a parameter that would be an area of interest to the Skilled Person at the Priority Date.

9.10 TS 25.321 describes a system which sends scheduling information whenever space is available with E-DCH data traffic, as well as on a periodic basis, supplemented by triggered Scheduling Information. The low values that the Serving Grant is allowed to take (see TS 25.321, Table 9.2.5.2.1.1), combined with the relatively large value of the smallest E-TFC other than the SI-only one, means that it is possible for the Serving Grant to be turned down by Serving and Non-Serving Node Bs to a value that does not allow the smallest E-TFC to be sent (excluding the E-TFC for SI only). I believe that the simulations referred to above, which would have taken place once Release 6 was frozen and fed into the Skilled Person’s work on Release 7, would indicate the issue with a stall should the Serving Grant be low enough.

9.11 As I stated above (paragraph 8.5), I do not think the Skilled Person would consider the stall scenario as being a significant problem which required fixing. For example, according to TS 25.321 Section 11.8.1.6.3, triggered Scheduling Information which is transmitted on its own is not retransmitted if it is received incorrectly. This supports my view that the Skilled Person would have been willing to accept the delay in waiting for the next periodic transmission of Scheduling Information.

9.12 However, if the Skilled Person thought it was worth addressing the stall scenario, for example, because they were concerned that periodic sending of Scheduling Information would not be responsive enough, an obvious solution would be to configure a trigger to send Scheduling Information in that scenario. This is for the following reasons.

9.13 First, the Skilled Person would be well aware that there were two ways for the sending of scheduling information to be initiated: periodic sending or sending based on triggers. Therefore, as periodic reporting would already be configured, the Skilled Person would consider supplementing that by defining a trigger to send Scheduling Information in this scenario.

9.14 Second, as I explain above at paragraph 7.25, the Skilled Person knows that the purpose of triggering Scheduling Information in TS 25.321 is to inform the Node B about things it will not or may not know. Although the Node B could in theory allocate such a grant deliberately, from a system point of view this would be inefficient, as resources would have been allocated which could never be used. A stall scenario is therefore one of those cases where the Node B (particularly the serving Node B) does not or may not know that the Serving Grant is insufficient to send a MAC-d PDU. The solution in such cases in TS 25.321 is to trigger the sending of Scheduling Information.

9.15 Third, adding a trigger to send Scheduling Information in this situation would be a very simple modification, as the fact that the standard already includes triggered Scheduling Information means that the UE will already have the bulk of the functionality to implement this new feature already be in place. The new trigger will represent only a few lines of code (simply a check at the end of the E-TFC Selection routine that if the number of scheduled MAC-es PDUs included in the MAC-e PDU is zero while the Serving Grant is not zero, trigger Scheduling Information).’

137. IDC submitted that the evidence in Dr Irvine’s first report about simulations was at a relatively high general level. I agree. It appears that Lenovo thought something similar. In his first report, Mr Townend addressed the disclosure of the PMS in general terms and identified eight areas of possible interest to the Skilled Person, including sections 9.1.5, 9.2.4, 9.2.5 and 11.8 (cf. Lenovo’s pleading). He rejected the notion that the key feature of claim 1 was disclosed in the PMS and then considered whether it would have been obvious to adapt the PMS to include that key feature. He indicated the Skilled Person would have to undertake a two-stage process: first, to recognise the transmission blocking problem and second, to decide to implement a new trigger of sending SI in response to the condition mentioned as a means of resolving the problem in place of some alternative method. He said the PMS does not assist on the first stage. At that stage the simulations argument(s) had not been explained, so Mr Townend did not consider simulations in his first report. It follows that in his second report, Dr Irvine had nothing regarding simulations to respond to. Nonetheless, he expanded on the topic – there is more on simulations in his second report than his first.
138. Dr Irvine repeated his point that the Skilled Person would have started work on implementations of Release 6 once it was complete and, as part of this work, they would have carried out simulations of the type he had described in 7.27 and 7.28 of his first report. He went on to say that wider simulations would be performed, focussed either on the performance of specific parts of the HSUPA system or on the more general performance of the system as a whole from the user’s point of view. He relied on and exhibited extracts from Chapters 8 and 9 of Holma & Toskala (published 4 months before the priority date) as including examples of both types of simulation.

139. He then responded to Mr Townend's evidence (that the Skilled Person would not have identified the problem) by stating, in his view, the Skilled Person would have done so through at least two types of routine simulations which he said would have been conducted before the priority date.
140. The first type was the mapping analysis, which I address below. I note that in reaching his conclusion that the Skilled Person's work would reveal the potential for the 'stall scenario' to occur, Dr Irvine relied on Mr Townend having noted in his first report that the Skilled Person would appreciate the concept of the stall situation. However, Mr Townend had made that point in the context of the Patent and it in no way supported Dr Irvine's conclusion.
141. As to the second type, Dr Irvine expanded on his simulations to find effective values of T_SIG and T_SING and stressed the importance of latency in internet services based on TCP. He then said it was not a surprise that reports of these types of simulation were not publicly available. He explained that these simulations would have been undertaken by equipment vendors and network operators to gain a competitive advantage and the results would have been sensitive proprietary information. Finally, he stressed the importance of web browsing for HSDPA and HSUPA, which would have manifested itself in simulations of downloading data over TCP/IP with a high bandwidth connection on the downlink using HSDPA and a low bandwidth uplink using HSUPA. He was of the view that the TCP acknowledgements (of 320 bits) sent on the low bandwidth uplink would identify any issues of latency.
142. In his second report, Mr Townend addressed the evidence on simulations given in Dr Irvine's first report. He disagreed that a 'stall' would necessarily be identified by simulations, giving the following reasons:
- (a) The first was that the Skilled Person would carry out modelling or simulations with a particular line of inquiry or objective already in mind e.g. to identify effective values of T_SIG and T_SING. He said there were two relevant effects of this approach.
 - (b) First, any data generated is dependent on the parameters set by the modelling. Addressing simulations generally, Mr Townend was unclear what question or objective the simulation would have been designed to answer or achieve, nor what parameters and assumptions that would have been used nor the particular parameters that would have returned results in which the transmission blocking problem would have been revealed. He was asked by the solicitors whether the Serving Grant would have been considered a relevant variable for simulations intended to find effective values of T_SIG and/or T_SING. He was of the view that varying the Serving Grant would have been an unusual choice, given it is not a variable that can be controlled in real world applications.
 - (c) The second effect is that the Skilled Person, when reviewing the generated data, would be focussed on determining the answer to his or her particular line of inquiry. On the T_SIG/T_SING example, Mr Townend said the Skilled Person would not have been inclined to

investigate other issues unrelated to that inquiry, of which there may have been many.

- (d) On the particular point of whether the stall would become evident to the Skilled Person from a given simulation, Mr Townend pointed out that simulations for optimisation at the priority date would commonly have employed methods that generated large sample sizes of data. He indicated that such methods tend to be poorly suited for identifying scenarios that do not commonly occur, as they can be obscured by the averaging of results.
 - (e) Finally, he said that even if the Skilled Person were to observe a stall or the issue of latency, this is distinct from identifying the underlying cause. He pointed out that the causes of a stall are not only the ones described in the Patent. So identifying the cause would require the Skilled Person to decide to investigate and then determine the cause of the stall or latency.
143. Dr Irvine's third report was limited to a response on claims 2 and 4 (since these had not been dealt with in the first round of reports). So the stage was set for Mr Townend's points to be challenged in the course of cross-examination. However, general assertions aside, it was not apparent that the cross-examiner had been armed with or had specific responses to the points made by Mr Townend in his written evidence, aside from the specific point put in cross-examination about investigating the effects of setting T_SIG to 'no report'. This point was not covered in Dr Irvine's written reports – the closest which Dr Irvine got to T_SIG = 'no report' was when he was discussing claim 4 in paragraph 7.14 of his second report. Even there, he referred simply to the default value. It was Mr Townend (in reply, in Townend 6) who mentioned that the default value of T_SIG was 'no report'.

Lenovo's 'mapping' argument

144. The first way in which Lenovo developed their obviousness argument – which I will call the 'mapping' argument – starts with Table 9.2.5.2.1.1 in the PMS – the Scheduling Grant Table or SG Table. This table sets out 38 different power levels which the Serving Grant may take. The next step involves the conversion of the Serving Grant power level into a number of bits of scheduled data, a necessary step so that the UE can perform E-TFC Selection and construct a MAC-e PDU. Although Mr Townend said that the PMS does not specify how the conversion should be done, Dr Irvine pointed out that the PMS contains a cross-reference to TS 25.214, section 5.1.2.5B.2.3 which gives the Skilled Person the necessary information to perform the conversion.
145. So, mapping the Serving Grant values to the number of bits being transmitted is simply the implementation of the conversion process and has to take place in order to carry out any simulation of the HSUPA transmission chain. Indeed, Lenovo pointed out that the Skilled Person would perform the calculations on paper in advance of actually constructing any simulation model.

146. Lenovo's submission was that it would be immediately apparent from the result of that conversion process that the Serving Grant could be non-zero but too small to send one MAC-d PDU. Lenovo pointed out that one of the main applications for HSUPA was envisaged to be the transmission of packet switched data which would entail frequent TCP acknowledgements. Common RLC PDU sizes were 336 bits and 656 bits, based on 320 or 640 bit RLC SDUs with a 16 bit RLC header in Acknowledged Mode (although Mr Townend made this point when discussing the Patent, I accept that it was part of the CGK).

147. When cross-examining Mr Townend, Counsel took him to a post-priority Tdoc from ZTE in which it appears ZTE spotted there was a problem but proposed a different solution. In the Tdoc, ZTE presented results from a 'simulation'. It does not matter whether the results were from calculations or an actual simulation, but I tend to agree with IDC that it looks like ZTE had just carried out some calculations. In support of this argument, Lenovo relied on two passages from Mr Townend's cross-examination:

Q. Let us assume they are doing as you said, which is figuring out how many bits of scheduled data could be sent according to each Serving Grant.

A. Okay.

Q. Once they have done that, they would see, immediately, that at the lower reaches of the SG table, the Serving Grants are too small for many common applications of the enhanced uplink?

A. Assuming that they came out with the numbers that we are looking at here, then yes, they would be able to see that those grants at the lower end were too small.

148. I interpolate here that, aside from the cross-examination on the ZTE document, Lenovo also submitted that it follows from the fact that transmission blocking is a real possibility that the Skilled Person would come out with numbers in which some grants were too small for many common applications.

149. The second passage came in a slightly later answer where Mr Townend said this:

... You will remember I said earlier, at least my understanding as to why the Serving Node B sending a grant that deliberately or accidentally caused this situation would not be something the skilled addressee would understand to be within the realm of a sensible Serving Node B implementation

What I am saying, just because the table can theoretically support a grant that is too small, the skilled addressee may well assume that, "That is fine, that means I need to make sure that my Serving Node B as part of its algorithm puts a lower floor on grants that it gives, so that it does not get itself into this mess."

150. Lenovo's response was to agree, but to add a further consideration and this was identified by Lenovo as the key step in the obviousness analysis. It was said

that the Skilled Person knows from their CGK that there are two reasons why a UE might end up below that floor, however good the Serving Node B scheduling algorithm is:

- (a) First, a UE may be sent a Non-Serving Relative Grant DOWN command (see section 9.2.5.2.1 of the PMS) which reduces the Serving Grant for two reasons (i) because the UE is causing too much interference in the non-serving cell or (ii) there is too much interference in the non-serving cell generally.
- (b) Second, there could be errors in the receipt of Relative Grant commands. On this point, ZTE considered the probability of such errors was 5% based on Holma and Toskala. (Reference was also made to conformance tests specified in TS34.121-1 which specify 5% and 10% error rates for Relative Grant UP/DOWN and HOLD commands respectively.) Mr Townend did not agree with the 5% figure and IDC pointed out these were not targets – for example, if the error rate was 4.5% it would be an indication the system was not working correctly.

151. Consequently, so Lenovo submitted, the possibility of transmission blocking, for either of those two reasons, ‘is staring the skilled person in the face’. The Skilled Person would be consciously considering the need to keep the UE above the floor in the SG Table to avoid getting into this ‘mess’. Hence, the Skilled Person could not fail to appreciate that that is not within their control: this ‘mess’ cannot be prevented.

152. However, if one reads the passage quoted at paragraph 148 above in context, Mr Townend’s evidence does not support Lenovo’s submission that the Skilled Person would be consciously considering the need to keep the UE above the floor. Immediately following that passage came this exchange (in which it will be noted that the ‘question’ is long and comprised of several parts):

- 14 Q. We agree, you and I agree that the skilled person would think
15 the way you just said. The skilled person would say to
16 themselves, "I must make sure that the Serving Node B does not
17 send a grant below that floor." But in terms of the causes,
18 two of those causes are part of the skilled person's common
19 general knowledge. They know that neighbouring cells can send
20 DOWN commands and they know that Relative Grants and indeed
21 Absolute Grants can be misinterpreted. So they know it is not
22 going to be good enough that their Node B scheduler, that the
23 Serving Node B does not set a grant below the floor. They
24 know there is a risk that things will happen so that the UE
25 goes below the floor, in any event?
- 2 A. Putting all of those pieces together is clearly what the
3 inventor did. I have not seen anything, as I have been
4 looking at this, that would suggest that the skilled addressee
5 would put all of those pieces together to come to that
6 conclusion. As an example, when you read the MAC
7 specification, and I am not suggesting for a moment that the
8 skilled addressee would not be aware or would not know you

9 could only put whole PDUs in a MAC-e PDU. Obviously, they
10 would know that, but the way that 25.321 is written, I do not
11 think it actually really makes reference to that. Possibly in
12 one line in the pseudo code, but in certainly all of the
13 normative text, it just talks about bringing in data according
14 to the grant or whatever. So the skilled addressee reading
15 the specifications is not really well set up to put together
16 all of the different pieces that we are talking about, even
17 though each of those pieces is within their mental toolbox.
18 That was why I had come to the conclusion that I had that it
19 was not apparent to me that the skilled addressee would put
20 all of those pieces together to identify the problem.

153. Counsel tried again:

7 Q. Let us imagine the skilled person has got to the point which
8 you have described, which is they have appreciated that there
9 is a floor within the SG table and the scheduler must not send
10 the UE below that floor.

11 A. Yes.

12 Q. Okay. To get from there to identifying the
13 transmission-blocking problem, the only step for the skilled
14 person to remember is that there are two reasons why the UE
15 might go down that table, below that floor, despite the best
16 efforts of the Serving Node B scheduler.

17 A. Yes. If the skilled person had got to that point, then, yes,
18 it would be one additional step to identify that.

154. Of course, I am not bound to accept Mr Townend's evidence. I have to place myself in the shoes of the Skilled Person and effectively ask whether he or she would have put all these pieces together. Before taking that step, it is relevant to consider the question: if this was obvious, why was it not spotted by several other participants in the standard setting/implementation process?

155. On this question, it is relevant that the SG Tables were first published in June 2005, although Dr Irvine said the tables on their own probably would not give the skilled person enough information and they would wait for the specification to be finalised, as it was in September 2005. Although the PMS itself was not published until June 2006 I was satisfied (as I have said) that the relevant content was the same as was published in September 2005. Thus, on Lenovo's mapping case, a period of about 11 months elapsed before the Priority Date during which period this problem was supposedly staring real-life teams in the face. If it had been, one would have expected many proposals to have emerged in pre-priority Tdocs identifying the transmission blocking problem and solving it, but there are none.

156. It is also relevant to note that Dr Irvine did not make the mapping argument in his first report, notwithstanding the point he made in his paragraph 9.8 which I regard as related in the sense that both arguments involve a Serving Grant which is insufficient to send any MAC-d PDU. Even in the circumstances of the discussion during the finalisation of his first report which led to his paragraph

9.8, where they were ‘knocking around some numbers’, the mapping argument did not surface. This consideration reinforces the particular force of hindsight – once you know what the problem is, it can ‘stare you in the face’ from various sources.

157. The two ways in which Lenovo put their obviousness arguments are related in the sense that the work relevant to the first way was a necessary pre-cursor and would have quickly given way to the extensive work involved in simulating an HSUPA set-up to develop or optimise it. For this reason, I discuss the post-priority Tdocs in the next section.

Lenovo’s simulations argument

158. The second way in which Lenovo put their obviousness case over the PMS was the simulations argument. In essence, the argument was that the transmission blocking problem would come to light during routine simulation of an HSUPA system in development.
159. In closing, Lenovo relied on two types of simulation: first, simulating the scheduler for the purposes of testing/optimisation and second, simulations to find effective values for the periodic timers T_SIG and T_SING. Mr Townend agreed in cross examination that both types of simulation were necessary and routine.
160. On the first point, Lenovo developed their argument as follows:
161. First, the experts agreed that the transmission blocking problem would not be a commonly occurring problem, but it would be a real problem with real effects. The UE is unable to transmit when it has data to transmit. Uplink resources are under-utilised because they are allocated to a UE which cannot use them. The effect would be on throughput, packet delay and resource utilisation. Mr Townend agreed it has a material effect on data throughput and latency.
162. Second, the fact that 3GPP chose to incorporate additional functionality into Release 7 to address the problem confirms it was considered a real problem with real effects. Dr Irvine considered the Skilled Person might be prepared to live with the problem in the context of frozen Release 6 but considered it was worth exploring in the context of developing Release 7.
163. Third, against that backdrop, Lenovo pointed out that the Skilled Person implementing HSUPA would have to write a scheduling algorithm and would have to simulate it as a matter of routine. Mr Townend agreed that the Skilled Person would have to devote particular attention to optimising the scheduler.
164. Fourth, the transmission problem would inevitably manifest as unexpectedly high latency at low Serving Grant values.
165. Lenovo submitted that Mr Townend did not actually dispute that at some point during routine testing and implementation, the problem would manifest itself as an issue of latency or reduced throughput, relying on this passage of cross-examination:

9 Q. You told us earlier that the transmission-blocking problem was
10 a real problem that would have a real effect. It follows from
11 that, does it not, that at some point during routine testing
12 of the scheduling algorithm optimisation and implementation of
13 the system, this problem would manifest itself as an issue of
14 latency or reduced throughput?

15 A. Again, it really comes down to the simulation assumptions that
16 are being used and, indeed, the capabilities of the simulator
17 that is being used which then comes down to really exactly
18 what it is you are trying to do with it as well.

166. This answer did not get Lenovo home. I was not persuaded that merely detecting latency in simulations of the scheduling algorithm would have led the Skilled Person to identify that transmission blocking was a or the cause of the problem.

167. The second simulation is more involved. As Lenovo submitted, the PMS provides for timers but leaves it to the implementer to decide (i) whether to implement them and (ii) what values to use for the period. Then the argument was developed as follows (in which I have omitted most of the evidence references relied upon, but have examined each of them):

(a) The skilled person would as a matter of routine carry out simulations to model the HSUPA system to determine whether to use a timer, in what circumstances, and if so, the optimal values for the periodic timers, from a system efficiency viewpoint.

(b) Mr Townend agreed that the skilled person would do simulations to determine the value of T_SIG and it would be worth looking at whether to configure a T_SIG period at all or use the default value of “no report”.

(c) In the context of T_SIG simulations, when T_SIG = “no report”, the transmission blocking problem would manifest itself as persistent blocking, as explained in the Qualcomm paper. So it clearly would be noticed.

(d) Lenovo repeated the point made above that the transmission problem would inevitably manifest as unexpectedly high latency at low Serving Grant values, again relying on the passage in the cross-examination of Mr Townend quoted above at paragraph 165 above.

(e) Simulations are intended to test for a range of scenarios and edge cases. Simulations also need to be realistic, otherwise the exercise would not provide a useful indication of an optimal value to use in practice (Irvine XX T4/473/7-14 explaining that the causes of transmission-blocking in the patent, such as transmission errors, would have to be modelled for the simulation to be accurate enough; Townend XX T2/126/3-11).

(f) The transmission blocking problem would therefore inevitably manifest – i.e. be identified by the skilled person. Increased latency indicates reduced transmission rate, and the skilled person would be motivated to

identify, and would be capable of identifying, when it was occurring, i.e., when the Serving Grant level was too small to send any scheduled data.

- (g) Ultimately, if the transmission blocking problem was a material problem, then it would be identified by the skilled person, for the reasons given by Henry Carr J in *TQ Delta v Zxyel* and in the EPO case-law.

168. In the final part of the argument, Lenovo submitted that '*Dr Irvine was absolutely clear on this and he has enormous experience in running simulations, and his CV is full of references to system performance assessment and the like*'.

169. On the first point, Lenovo relied on this answer given by Dr Irvine in cross-examination, which was certainly the high point of his evidence on this:

'I think whatever range of parameters you that used, if they were realistic to the operation of the system, then the transmission-blocking problem would show itself. The reason for that is because the transmission-blocking problem will occur whenever you are sending a fairly low rate uplink service. TCP acknowledgements are a classic example of a low rate uplink service. A lot of the time you will be using that service, I cannot envisage anybody simulating the system without simulating that service. As soon as you simulate that service, you will come across the transmission-blocking problem.'

170. As for Dr Irvine's experience in running simulations, his answers in cross-examination generated the following general picture. Although he was not being cross-examined on the content of his CV when he gave these answers, I read this section of the transcript with his CV open and picked up the references:

- (a) He described his 'first job', on the RACE II ATDMA project [1994-5], was on simulation for 3G and said everybody used simulations at all stages in the process.
- (b) Next, MOSTRAIN [1995-7], the design of mobile service provision for high speed trains.
- (c) Then the long-running [1997-2012] MobileVCE project, although I infer from his CV and other answers that the relevant simulation work occurred in the Core 1 period [1997-2000].
- (d) There are some papers in his CV in which simulations or modelling are indicated in the titles – e.g. in 2002 and 2003. No doubt other papers also involved simulations.
- (e) Dr Irvine accepted in cross-examination that he did not have experience in simulations of UMTS systems, nor of optimisation of an HSUPA network (although he said a few years earlier 'we were running simulations for companies to look and see what different parameters they would use for these types of networks' and that 'I am familiar in the general case with the sort of simulations that were being undertaken to

optimise different parts.’), nor in deploying a 3G network, but he said he had experience simulating 4G and 5G systems.

- (f) More generally he said ‘simulations is in the air that we breathe in terms of developing communications systems’.

171. Part of the reason why Dr Irvine’s experience of simulations came out only in cross-examination was because, although he realised the importance of simulations when preparing his second report, he did not update his experience in that regard. The lack of focus on simulations until later in the process also explains why Mr Townend did not focus on his experience of simulations in either his first or second reports. In addition, there was no challenge to Mr Townend in cross-examination as regards him being qualified to speak on simulations, and a review of his experience in industry before the Priority Date indicates it is highly likely he had experience of simulations in that work. Having seen Mr Townend in the witness box, he seemed to me to be able to speak with no less authority than Dr Irvine on simulations. In this regard, I refer not just to his demeanour (which it is difficult to ignore) but more to the content of his answers.

172. So although Lenovo’s submission which I related above in paragraph 168 above was largely justified (albeit a tad exaggerated), I do not proceed on the basis that Dr Irvine’s evidence on simulations trumped that given by Mr Townend.

173. Reverting to the substance of Lenovo’s argument, there are four critical points in particular which I must examine:

- (a) First, the significance of the problem and the attitude of the Skilled Team to it.
- (b) Second, Lenovo’s argument that Mr Townend agreed that the Skilled Person would do simulations to determine the value of T_SIG and it would be worth looking at whether to configure a T_SIG period at all or use the default value of ‘no report’.
- (c) Third, that in the context of T_SIG simulations, when T_SIG = ‘no report’, Lenovo submitted that the transmission blocking problem would manifest itself as persistent blocking, as explained in the Qualcomm paper, and so it clearly would be noticed.
- (d) Fourth, on the assumption that the transmission-blocking problem would be identified from routine simulations, why was the problem and/or its solution not identified by several participants in the standards-setting process?

174. I address these points under the relevant headings below.

The significance of the transmission-blocking problem.

175. This is a ‘background’ point, so I will deal with it first. Lenovo’s position (and Dr Irvine’s evidence) on the significance of the transmission-blocking problem

underwent some change. In his first report, in the context of discussing the Patent, Dr Irvine said this:

‘8.5 Although EP 318 suggests that the stall situation is a problem which results in inefficient resource allocation and/or excessive scheduling delays ([0011]), it should be noted that the Skilled Person would not necessarily see the stall situation as a problem as only a very small amount of resource is involved and goes unused. In addition, scheduling grants are designed for best effort services in any case. If the service involved was delay sensitive, it would probably be sent using non-scheduled grants as this guarantees a certain number of bits in the E-TFC bit rates and the stall situation will not occur. The Skilled Person would be aware that the period for sending Scheduling Information in the standard is configured to a value ranging from every TTI to every second (see paragraph 7.21), so at most the delay would be one second and on average it would be half that, even if the longest period was configured.’

176. This paragraph was consistent with Dr Irvine’s later paragraph 9.11 (quoted above).

177. In his second report, Dr Irvine addressed three post-priority TDocs from ZTE discussed at the WG2#58 meeting held 7-11 May 2007:

‘The meeting minutes record that the ZTE TDocs were “noted”, further discussion either offline or at the next meeting was suggested, but only for Release 7 or “further releases”, and explicitly not for Release 6 because it was not an essential correction. This also aligns with my view in Irvine 1-C paragraph 8.5 that the underlying stall condition was not an issue that was considered worth fixing.’

178. In his ‘Errata’, Dr Irvine clarified that when he wrote that last sentence, he had Release 6 in mind, accepting that he should have made that clear, and added that, for Release 7, the stall condition would have been worth exploring. By May 2007, Release 6 had been frozen for some time and it is understandable that only essential corrections would be permitted. However, Dr Irvine’s clarification does not make a lot of sense, in my view. The Skilled Person would either have considered the problem worth fixing or not. That conclusion would not have depended on whether the solution would be incorporated into Release 6 or Release 7.

179. As can be seen from the way Lenovo put this obviousness argument in closing, by the time it came to the cross-examination of Mr Townend, their point seemed to me to have firmed up somewhat. Perhaps realising that Dr Irvine’s equivocal evidence made the argument more difficult, the point was put to Mr Townend as being ‘a real problem with real effects’. Mr Townend accepted this, but it still begs the question as to whether the Skilled Person would identify the transmission blocking problem from his or her simulation(s) from amongst the ‘noise’ which the results would throw up.

180. On that point, Mr Townend agreed that simulations would be used to test different conditions (e.g. in city and in rural environments) and also agreed that minority cases are important to the skilled person (cf the edge cases in paragraph 167(e) above) but added the qualification that ‘they are not that easy to tease out with simulations’.

Simulations to investigate the various values for T_SIG

181. Mr Townend agreed that when the Skilled Person was considering what value to choose for T_SIG, there were two competing factors: first, the more frequently the UE sends SI, the more up-to-date the information the scheduler will be able to use to make scheduling decisions, and that factor suggested a low value for T_SIG; second, sending SI more frequently requires more signalling overhead which the Skilled Person is always looking to minimise, and that factor would suggest a high value for T_SIG or even not enabling it at all. Mr Townend also agreed that the Skilled Person would have to strike the right balance between those two factors. The next stage of the argument was the suggestion that the Skilled Person would do simulations and trials to see what values of T_SIG work best for their system. At that stage, Mr Townend did not agree that simulation would be required to make a decision on whether to adopt ‘no report’ for T_SIG:

7 Q. Obviously, again, one of the values they would try out would
8 be no report?

9 A. Not necessarily, just because no report is one that you could
10 actually make a decision on, frankly, without simulation,
11 because there is no question as to what the impact of that is
12 going to be, frankly, in terms of either of the two factors
13 that you listed. It clearly minimises the overhead, so it is
14 a clear win there. On the other side, it, obviously, you do
15 not get those periodic updates, so at that point you have some
16 very clear, that, some error condition might persist or your
17 algorithm cannot rely on, for example, long-term averaging of
18 data that it gets from the Scheduling Information. It is with
19 the no, the no reporting, it is a pretty much a matter of,
20 "Are you throwing away some functionality that you need for
21 your algorithm?" Because no amount of simulation is necessary
22 to tell you that.

23 Q. Would the skilled person not wonder whether opportunistic SI
24 would occur frequently enough that they did not need periodic
25 SI on top?

2 A. That would depend, I think, in part on, for example, whether
3 they were looking at doing averaging of measurements.

182. The cross-examination continued but in my view all that it demonstrated was a similarly circular point that if a simulation was conducted at a sufficiently detailed and specific level so as to be apt to reveal the transmission blocking problem, then it would do so. This begged the question as to whether that type of simulation would be carried out as part of ‘routine work’.

Simulating T_SIG = 'no report'

183. On the third point, Mr Townend was asked to assume a skilled person simulating various values of T_SIG but in particular the default value where T-SIG is set to 'no report'. He indicated that his earlier answers in cross-examination that morning did not support the assumption. However, the passage of cross-examination of Mr Townend relied upon as supporting this second point did not do so, in my judgment. This was the high point which produced an entirely circular answer:

22 Q. If the skilled person simulated T_SIG equals no report,
23 sometimes UEs in the simulation would suffer persistent
24 transmission-blocking?

25 A. That would depend on the simulation.

2 Q. If the simulation reflected the real world enough so that
3 transmission-blocking occurred.

4 A. Yes. If the simulation was sufficiently realistic that
5 transmission-blocking could occur, then if you were to do it
6 in this situation, then you would expect that
7 transmission-blocking would occur somewhere in your simulation
8 runs, yes.

9 Q. When it did occur, it would just be persistent?

10 A. Yes, because there was no -- yes.

Where is the evidence that others identified the problem (and its solution) from their simulations?

184. I have set out above the passages from Dr Irvine's evidence where he indicated that the Skilled Person might not consider the transmission-blocking problem as one worth fixing. Those pieces of evidence provide one possible answer as to why other participants in the standard-setting process did not draw attention to the problem or propose the solution in the Patent.

185. Lenovo evidently considered that there might be problems with Dr Irvine's evidence in that regard. This manifested itself in two ways: first, Lenovo's firming up as to the perceived significance of the problem; and second, in the explanation provided in Dr Irvine's second report to the effect that the those who conducted simulations would have kept the results secret.

186. Dr Irvine said it was no surprise that reports of simulations to find effective values of T_SIG and T_SING were (and are) not publicly available. He said that equipment vendors and network operators would have undertaken these simulations to try to gain competitive advantage with their Release 6 implementations, although he accepted that relevant results would also feed into Release 7 standardisation efforts. Thus, he said, although it would have been routine for these simulations to be conducted, the results would have been sensitive proprietary information which would not have been shared with others in the market.

187. Whilst I entirely accept that the *details* of simulations would have been kept confidential including resulting values of T_SIG and T_SING, this reasoning

does not explain why companies would not have identified improvements arising out of the simulations which could have been (a) patented and (b) presented for inclusion into the standard. In my judgment, if a company had identified the transmission blocking problem in the course of their simulations and identified the solution, it is highly likely, in the competitive field of UMTS development that they would have done both those things, not least to avoid another company patenting the solution they had identified. Furthermore, if the problem had been spotted and the solution in the Patent identified, those responsible would have realised that their additional trigger for SI required a change to the Standard. That, in my view, would have resulted in a proposal being put to RAN WG2 at one of its meetings. In other words, the RAN WG2 records provide, in my view, a reliable record as to whether anyone did identify the problem and identify the solution in the Patent at least up until around the time at which it was clear the Patentee's solution was being adopted into the Standard (see paragraph 194 below).

188. I must also keep in mind the period of time over which it was said these simulations would have been run. The following timeline is relevant:
- (a) In cross-examination, Dr Irvine indicated that simulations for T_SIG and T_SING would have been looked at from September 2005 onwards. This answer was entirely consistent with his written evidence that the Skilled Person working on developing the Release 7 standards would have started work on implementations of Release 6 once it was complete.
 - (b) Chipsets implementing HSUPA were available by the Priority Date. Dr Irvine's evidence was that the first HSUPA chipset for handsets was announced by Qualcomm on 4 May 2006 and was used in a demonstration of an HSUPA call using Nortel network equipment in October 2006.
 - (c) The first commercial HSUPA network launched in February 2007.
189. This timeline reinforces the evidence that numerous substantial entities in the field (both manufacturers of network equipment and operators implementing new HSUPA networks) would have been carrying out a wide variety of simulations over this period to develop and optimise their HSUPA offerings. No pre-priority Tdocs hinting at the transmission-blocking problem were identified. Therefore, it is necessary to examine various post-priority Tdocs which were referred to in order to see what picture is presented.

Secondary Evidence

190. The secondary evidence, such as it was, was primarily directed to the case over the PMS and so I will consider it at this point. In his second report, Dr Irvine indicated that after he had formed his opinions and substantive drafting of the earlier sections was complete, Lenovo's solicitors presented him with 8 post-priority documents from 3GPP RAN 2. Four of these were considered at WG2#58 (7-11 May 2007), two from WG2#58*bis* (25-29 June 2007) and two from WG2#59 (20-24 August 2007). Six related to a proposal from ZTE and two from Qualcomm.

191. In summary, these Tdocs show that:
- (a) ZTE identified a different type of blocking problem after the priority date, and proposed a solution to it that was different to the one in the Patent. Specifically, they had identified a problem to do with how the Relative Grant was updated, as Dr Irvine accepted. There was just a passing reference to the blocking problem, but that was not the one they were trying to solve, another point which Dr Irvine accepted.
 - (b) His point was that the problem which ZTE identified was *consequential* upon the transmission-blocking problem. ZTE identified that the UE cannot transmit with a non-zero Serving Grant but focussed on what needed to happen to the next Relative Grant command in order to get the Serving Grant up to index value 4 so that the next transmission would be sent.
 - (c) RAN2 rejected the ZTE solution, but during the WG2 #58bis meeting it is recorded that there was “*some interest in an alternative solution based on triggering the UE to send an SP*”. There is no attribution as to who proposed that ‘alternative solution’ or who expressed interest in it. There were numerous attendees at WG2 #58bis including the named inventor of the Patent on behalf of IDC and four representatives from Qualcomm.
 - (d) At the next meeting, WG2 #59, Qualcomm submitted a proposal to incorporate a new trigger for Scheduling Information in accordance with the Patent’s solution. It is apparent from Qualcomm’s proposal (R2-073430) that it did not originate from any simulation work but simply from the discussion at RAN WG2 #58bis.
192. Dr Irvine’s written evidence was that these Tdocs were consistent with his views that (1) the blocking problem would have been identified, (2) would not have been considered worth fixing and (3) the introduction of a new trigger for SI would have been an obvious solution. However, a different picture emerged in cross-examination in which Dr Irvine agreed that:
- (a) There is no suggestion that simulations played any role in the thought processes of ZTE. Indeed, he fastened on this point to repeat his view that you ‘*did not need simulations to get you to the fact that there is a gap in the grant table and the lower grant values will not be enough to send anything*’.
 - (b) ZTE was not proposing any change to the sending of SI.
 - (c) There is nothing in the ZTE Tdocs that makes it obvious to suggest any changes to the SI.
 - (d) Attendees at the WG2 #58bis meeting included the patentee and also 4 people from Qualcomm, and the idea discussed at that meeting could have come from the patentee.

- (e) There is no suggestion that Qualcomm had done any simulations and they could have got the idea from being in the same meeting as the patentee.
 - (f) The ZTE proposal ultimately went nowhere – it was another idea that “died a death”.
193. As IDC submitted, the Patents Court has commented on the limited assistance that can be gleaned from this type of ‘Tdoc’ analysis on numerous occasions. See for example the judgment in Trial A *Interdigital Technology Corporation & Anor v Lenovo Group Ltd & Ors* [2021] EWHC 2152 (Pat) at [120]-[121]), and Arnold J in *Koninklijke Philips NV v Asustek Computer Inc* [2018] EWHC 1224 (Pat) at [201]. Leaving aside the point that attendees at 3GPP meetings may not be representative of the relevant Skilled Person in a particular case, the fact that a claimed invention might be published by competitors of the patentee shortly after the priority date may or may not be significant depending on whether it is likely that the competitors had become aware of the work done by the patentee.
194. The solution in the Patent was incorporated into the standard in version 7.6.0, released on 24 October 2007. I can see that once it was incorporated, there would have been little point in any other participant mentioning either the transmission-blocking problem or the solution in the Patent in any Tdoc or WG2 meeting. I also consider it likely that participants would have been aware that that was coming in advance. However, this leaves an extended period in which the only relevant independent suggestion was that from ZTE in May 2007 which did not derive from simulations.
195. As at the date of WG2 #58bis (late June 2007), the solution in the Patent had not yet been adopted into the standard. In addition to the discussion of the ‘alternative solution’ at WG2 #58bis, it is likely that this solution had been the subject of prior discussion at a meeting shortly after the priority date. At WG2 #58bis it was in the inventor’s/IDC’s interests to promote their solution. Accordingly, it is more likely than not that it was the inventor who drew attention to this ‘alternative solution’ at that meeting and I so find. It is to be noted that the subsequent Qualcomm proposal was based simply on the discussion at the previous meeting – there is no indication whatsoever that any simulation work contributed to it.
196. IDC’s fall-back submission was that this whole ‘Tdoc analysis’ exercise was doomed to fail in any event, because, as Dr Irvine explained, the documents were produced by Lenovo’s solicitors, and he had no knowledge as to the search methods used and had taken no steps to ensure they were an objective selection of documents as opposed to ones that simply helped Lenovo. It is not necessary for me to determine this point.
197. After the evidence had been heard, Lenovo submitted that secondary evidence was not of assistance either way in this case but submitted that their case was entirely consistent with such secondary evidence as exists.

198. IDC took a measured view as to the value which could be extracted from these Tdocs, submitting that they were either entirely neutral to the issue of obviousness, or if anything they support the case of inventiveness. In my view, to the extent that weight should be given to this secondary evidence, it is slightly double-edged. The evidence suggests that ZTE spotted the problem without simulations (albeit some months after the Priority Date), but their proposed solution missed the invention. If ZTE had spotted the problem without simulations, two questions arise: first, why ZTE appear to have only spotted the problem several months after the Priority Date, when the relevant material in the PMS had been around for a considerable time – the SG Table in particular had been published in June 2005 but the specification was essentially complete in September 2005; second, why are there no indications of any others spotting the problem and proposing the solution claimed in the Patent. Overall, in my view, this evidence tends to damage Lenovo’s and Dr Irvine’s theories of obviousness over the PMS.

Discussion and Conclusion

199. It is now time to draw the threads together from the various topics I have discussed above.
200. There is no doubt that a proposed implementation of HSUPA would require a manufacturer or operator of network equipment to run simulations to optimise a significant number of variables. Various algorithms would have to be designed – in particular, a scheduling algorithm – which, once coded into software, would have to be optimised before live deployment. This would all be routine work of the skilled person. However, there was no official set of simulations and different simulations would give different results.
201. Although Dr Irvine presented a powerful analysis and the issue overall was finely balanced, ultimately I was not persuaded that routine work would have identified the transmission blocking problem to the Skilled Person leading him or her to the solution embodied in claim 1 of the Patent.
202. In reaching this conclusion, I paid particular attention to all the evidence references relied upon by Lenovo. I also re-read the whole relevant sections in the cross-examinations of Mr Townend and Dr Irvine. It is clear, as I have endeavoured to explain above, that Lenovo did not secure any clear acceptance from Mr Townend of either part of their obviousness analysis. If you know the end point you are trying to reach, it is easy to direct the questioning with a laser-like focus to lead to that end point. Mr Townend’s answers were a reminder that the uninventive Skilled Person proceeding without hindsight would not have such a focus.
203. Taking a step back from the detail, I observe that this transmission-blocking problem is one of those points that one cannot forget once it has been pointed out. And once alive to the point, one can then identify various scenarios where the problem is ‘staring you in the face’, the striking expression used by Dr Irvine in relation to both limbs of the obviousness arguments in Case 1.

204. In this regard, I regard it as of some significance that the problem was suggested to Dr Irvine. From that I infer that either he did not identify the problem simply by considering the values in the SG Table or through considering generally what sort of simulations would have to be carried out to implement an HSUPA system or (and this is possibly more likely) that Lenovo's solicitors were not confident that he would do so. I have been careful not to overstate the significance of this point because I recognise that, even with the progressive unveiling *Medimmune* approach, it could be very difficult (if not virtually impossible) to simulate what real-life teams would have done at the Priority Date to find out whether in fact the transmission blocking problem would have been identified.
205. However, since Dr Irvine was effectively led to the problem, the consequence ought to have been that those instructing him should have taken extra special steps to try to compensate for the hindsight that had thereby been injected into Lenovo's case. I realise this would have been difficult to achieve and counterintuitive for the team, but it is part of Lenovo's burden to prove their case.
206. Of all the factors I have discussed above, the ones which I consider have significant weight (in addition to the influence of hindsight on both arguments) are as follows:
- (a) On the mapping argument (notwithstanding the 'staring in the face' evidence):
 - (i) that Dr Irvine did not identify the mapping argument in his first report;
 - (ii) that there is no evidence that anyone in the industry spotted the transmission-blocking problem from studying the SG Table or implementing it before the Priority Date. In this regard, I have not forgotten the post-priority ZTE proposal (stemming from, as I have found, some calculations) but that did not lead to the solution in the Patent
 - (b) On the simulations argument:
 - (i) that despite the succinct statement of the case in Dr Irvine's first report, the case was forced to continue to develop. Ultimately (at least in one version), it depended on simulations investigating the consequence of T_SIG = 'no report'. Dr Irvine had not addressed that specific point and Mr Townsend did not accept either (i) that such a simulation would be carried out nor (ii) that the transmission blocking problem would be identified in such a simulation;
 - (ii) that the consequence of Lenovo's argument was that others must have identified the problem and, having done so, would have been motivated to find a solution, file for patent protection and file a Tdoc to 3GPP but there is no evidence that anyone developing or optimising HSUPA systems did spot the problem or independently propose the solution in the Patent.
207. For these combinations of reasons, I reject Lenovo's Case 1.

Kim

208. Kim is a US Patent Application filed by Samsung on 30 December 2005 entitled ‘*Method and Apparatus for Signalling Control Information of Uplink Packet Data Service in Mobile Telecommunications System*’. Priority is claimed from a Korean document dated 30 December 2004 (which I have not seen). Kim was published on 6th July 2006, some 6 weeks before the priority date of the Patent. The parties were agreed on the following points: first, that whether the document is read at the date of its publication or the priority date of the Patent makes no difference. Therefore, for convenience, I will simply refer to reading Kim at the priority date; second, that Kim provides a proposal for HSUPA; and third, between the date when Kim was written and its publication, the development of HSUPA had moved on.

Disclosure

209. The experts and the parties agreed on the central idea of Kim. It is that where a UE is sending a MAC-e PDU whose payload comprises only SI, the header of the MAC-e PDU is omitted. This can occur because the size of the E-TFCI implicitly indicates that only SI is being sent. Of course, the advantage of omitting the header is that the size of the E-TFCI is reduced, saving radio resources and improving efficiency on the enhanced uplink.

210. [0024] describes Kim’s invention as follows:

[0024] In accordance with another aspect of an exemplary embodiment of the present invention, there is provided a method for receiving scheduling information for an uplink service in a mobile telecommunication system. The method includes an MAC-e PDU received over an E-DCH and a transport format received corresponding to the MAC-e PDU over a control channel related to the E-DCH. The transport format is checked for a predetermined specific value. The MAC-e PDU is divided into a header and a payload, and the payload is disassembled according to the header, if the transport format does not have the predetermined specific value. The MAC-e PDU is determined as including no header, and the scheduling information is acquired from the MAC-e PDU, if the transport format has the predetermined specific value.

211. The experts were agreed that the Skilled Person would recognise that this invention had been incorporated into TS 25.321 by the priority date.
212. After setting out much of the familiar technical context of UMTS, Kim describes the operation of the “*enhanced uplink*” (i.e. HSUPA) from [0058] onwards. The skilled person would appreciate that what Kim refers to as “*MAC-e control information*” being inserted into a MAC-e PDU at [0058] is SI. At [0060] Kim also refers to an “*initial rate request packet*”, which is an SI that is sent on its own (i.e. without any data) when the uplink connection is set up.
213. I set out [0058] here because it featured in the arguments over the meaning to be given to Fig.9:

“[0058] If the Node B 437 grants uplink transmission resources to the UE 402 at a specific point of time, the UE 402 brings as much data as can be transmitted through the transmission resources from the RLC entities 405, 407 and configures and transmits a MAC-e PDU. In the multiplexing and TSN setting block 430, a MAC-e header is inserted into RLC PDUs transferred from the RLC entities 405,407 to get a MAC-e PDU. At this time, if a sufficient space remains in an E_DCH transmitter block or MAC-e control information to be transmitted exists, the MAC-e control information is also inserted into the MAC-e PDU and is transmitted together.”

214. The general structure of a MAC-e PDU containing RLC data is then described at [0062] to [0068], with the special DDI value being discussed in [0069]-[0071], all illustrated in Figure 6. It is not necessary to set this out because Fig.6 would be familiar to the Skilled Person from the standard (see section 9.1.5, along with the special DDI value of [111111] in section 9.2.4.2). Kim continues:

[0072] As having been already describe [sic], the size of the MAC-e PDU is notified from the UE to the node B by using an E-TF. The E-TF is a logical identifier which may have a size of 7 or 8 bits. One example of the relation between the E-TF and the MAC-e PDU is given below in Table 2:

215. Table 2 sets out an example set of the MAC-e PDU sizes corresponding to E-TF values from 0 to 3 and then from 125-127.
216. [0073] and [0074] explain certain points which are in any event obvious from Table 2: that E-TF 0 is the smallest E-TF and corresponds to a minimum size of MAC-e PDU, having a size of 36 bits, whereas the second largest E-TF value of 126 has a size of 18,879 bits.
217. [0075] explains what the MAC-e payload may comprise. If no RLC PDU is to be transmitted, but only MAC-e control information, then the MAC-e PDU has the minimum size because the MAC-e control information is significantly smaller than that of the RLC PDUs.
218. Kim notes at [0076] that when SI is sent alone (i.e. it is the MAC-e control SDU) it will be no more than approximately 30 bits. The skilled person would recognise that by the priority date the size had in fact been fixed at 18 bits. Kim suggests at [0078] that because the size of SI alone is considerably smaller than a MAC-e PDU containing RLC PDUs, which was typically 168 or 336 bits, an SI-only PDU could be indicated by the use of a “special” E-TFCI of 0, i.e. the smallest E-TFCI.
219. The possible structure of such a MAC-e PDU containing only SI is shown in Figures 7A and 7B. The latter of those figures, described at [0083], shows the header-less SI embodiment of the invention.


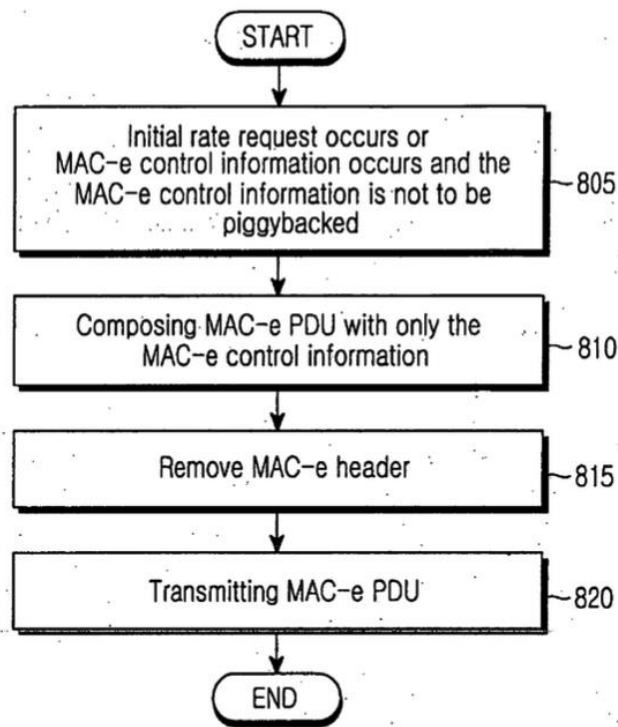


FIG.7A

FIG.7B

220. Kim then discloses the two embodiments on which Lenovo’s anticipation and obviousness case has focused. These are the embodiments illustrated in Figures 8 and 9. I consider it is necessary to have regard to both embodiments.
221. [0084] identifies Fig.8 as a flowchart illustrating an operation of a UE in accordance with an exemplary embodiment of the present invention:



222. This is explained further in [0085]-[0090]. In summary, at step 805, one of two events is recognised by the UE:
- (a) Event 1 is where an initial rate request occurs i.e.
- “...where a UE having no granted rate initially requests transmission resources. The initial rate request packet accommodates only a MAC-e control SDU, and the MAC-e control SDU includes scheduling information such as buffer status information and transmission power information of the UE.”
- (b) Event 2 is where MAC-e control information cannot be transmitted together with other MAC-es PDUs:

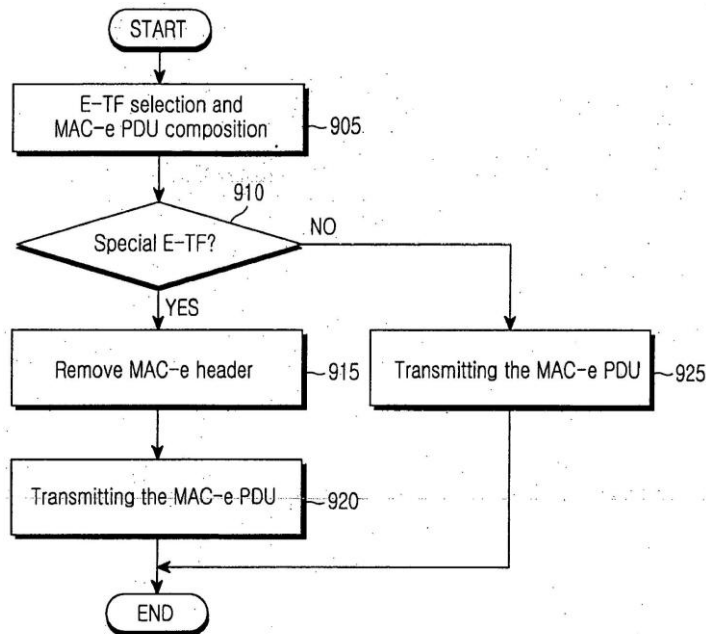
“This is a case where there is no MAC-es PDU to be transmitted when the MAC-e control information occurs or uplink transmission resources is not available for transmitting the MAC-es PDU. In this case, a MAC-e PDU is configured only with a MAC-e control SDU.”

223. The remaining steps are straightforward. The MAC-e PDU is configured by the UE with the MAC-e control SDU. The DDI header is then removed in step 815 and the headerless MAC-e PDU is transmitted. The final part of [0090] explains:

“in order to represent the MAC-e PDU from which the header has been removed, a special E-TF value is transmitted over the E-PDCCH.”

224. The specification moves straight onto Fig.9. [0091] explains:

“FIG. 9 is a flow chart illustrating an operation of a UE in accordance with an exemplary embodiment of the present invention. Similarly to the operation in FIG.8, the operation in FIG.9 also aims at transmitting a MAC-e PDU as configured as shown in FIG.7B.”



225. The description of the flow through the left-hand side of Figure 9 is at [0092] to [0094]:

“**[0092]** Referring to FIG. 9, in step 905, a UE selects an E-TF for the E-DCH at the present point of time, and configures a MAC-e PDU to be transmitted according to the E-TF. Considering transmission resources granted from the Node B, the UE determines the size of a MAC-e PDU to be transmitted in the next transmission, and selects an E-TF representing the determined size. The MAC-e PDU is then configured according to the selected E-TF. Accordingly, the UE receives data, suitable to the size denoted by the E-TF, from the E-DCH control block or an RLC buffer, configures a MAC-es PDU or a MAC-e control SDU with the received data, and then configures a MAC-e header for the MAC-es PDU and/or the MAC-e control SDU and attaches the MAC-e header to the MAC-es PDU and/or the MAC-e control SDU.

[0093] In step 910, the UE determines whether the selected E-TF has a predetermined special E-TF value. The UE goes to step 915 if the selected E-TF has the special E-TF value, and goes to step 925 if the selected E-TF does not have the special E-TF value. If the selected E-TF has the special E-TF value, the MAC-e PDU includes only a MAC-e control SDU. In contrast, if the selected E-TF does not have the special E-TF value, the MAC-e PDU includes MAC-es PDUs and a MAC-e control SDU together, or only MAC-es PDUs.

[0094] In step 915, the UE removes the header from the configured MAC-e PDU. As a result, only the MAC-e SDU remains in the MAC-e PDU as shown in FIG. 7B. In step 920, the UE transmits the MAC-e PDU from which the header has been removed. In this context, a MAC-e PDU including a header and a MAC-e control SDU is configured, and then the header is removed to generate a MAC-e PDU to be transmitted. However, when the E-TF has the special E-TF value, the UE may configure a MAC-e PDU that includes no header and consists of only a MAC-e control SDU, and then transmit the MAC-e PDU directly according to UE manufacturers' designs. Also, in order to represent the MAC-e PDU from which the header has been removed, the special E-TF value is transmitted over the E-DPCCH.”

226. The central issue on Kim is how the Skilled Person would read these paragraphs. In order to carry out that assessment, it is first necessary to note the following differences between Kim and the CGK of HSUPA at the Priority Date which were identified in the cross-examination of Mr Townend and acknowledged by him. Lenovo put forward this list:
- (a) Kim uses the term “*scheduling information*” in a different sense to the way it was used in HSUPA. (Townend 4 §273 and XX T2/190/18 – 191/3).
 - (b) The format of SI in Kim is different to the format of SI in HSUPA (Townend 4 footnote 140 and XX T2/190/4-7).
 - (c) Kim uses C/T multiplexing while that was not used in HSUPA (Townend 4 §287.3 and XX T2/192/8-10)

- (d) DDI described in Kim is used in a different way to DDI in HSUPA (Townend 4 §296 and XX T2/192/11-14).
 - (e) Kim uses a Transport Block Size Table which is inconsistent with the equivalent tables in HSUPA (Townend 4 §298 and XX T2/192/14 – 193/2).
 - (f) Kim talks about sending triggered SI alone in a case where there is no MAC-d flow data to send, but that is not permitted in HSUPA (Townend 4 §315 and XX T2/193/3-8).
 - (g) Kim implies the construction of a MAC-es PDU even where there is no transmission resource to send it – that does not happen in HSUPA (Townend 4 §§316-317 and XX T2/193/9 – 194/3, noting that Kim is inconsistent on this point).
 - (h) Kim [0092] describes a different method of E-TFC Selection and MAC-e PDU construction to that used in HSUPA.
227. Dr Irvine suggested that the Skilled Person would take the description of the various steps in the Fig 9 embodiment literally, whereas Mr Townend was of the view that the Skilled Person would read these paragraphs with his or her knowledge of HSUPA in mind and interpret them accordingly.
228. So Dr Irvine explained the scheme as follows:
- (1) the UE begins by selecting an E-TFC, based on the resources granted by the Node B;
 - (2) if the E-TFC selected is the “*special*” one of smallest size, then only the SI can be sent, and so the UE removes the MAC-e header and sends an SI-only PDU;
 - (3) otherwise, i.e. if the E-TFC is not the special one, a normal MAC-e PDU can be sent (step 925, on the right-hand side of Figure 9, which is described at [0095]).
229. For his part, Mr Townend did not agree with Dr Irvine’s summary of the Fig 9 embodiment. Mr Townend considered that it appeared to rely on an assumption that the UE’s selection of the special E-TFCI of Kim occurs when the UE does not have the resources needed to send any MAC-es PDUs but Mr Townend was of the view that this is not stated in Kim and is inconsistent with the Skilled Person’s CGK.
230. On Dr Irvine’s first step, the disagreement focussed on the first two sentences of [0092]. Dr Irvine focusses on the first sentence of [0092] – the selection of an E-TF. Yet that first sentence is merely a summary of step 905. The remainder of [0092] describes the process in greater detail. Thus, as Mr Townend pointed out, the selection of the E-TF is not the first thing that happens. Instead, as explained in the second sentence onwards, first the UE determines the size of a MAC-e PDU to be transmitted in the next transmission ‘considering

transmission resources granted from the Node B'. Then an E-TF is selected representing the determined size [of MAC-e PDU]. Mr Townend said the Skilled Person would recognise this two-step process for E-TFC Selection as consistent with their CGK and the PMS.

231. As to the second stage of Dr Irvine's approach, Mr Townend observed that the idea that selection of the special E-TFCI '*would occur when the UE does not have the resources needed to send any MAC-es PDUs*' did not appear to have any basis in Kim or in the CGK. Indeed, Mr Townend pointed out that if the correct E-TFC Selection process was followed, the result would be as follows:
- (a) A UE having a Serving Grant too small to send one PDU of scheduled data will determine it has no space for data in a MAC-e PDU.
 - (b) The UE will determine that it is unable to send a MAC-e PDU.
 - (c) Consequently, the UE will neither select the special E-TFCI nor send it.
232. Mr Townend acknowledged that the exception to this would be where SI has been triggered in accordance with section 11.8.1.6 of the PMS i.e. where SI has been triggered, it will always be sent. Thus, Mr Townend was of the view that the Skilled Reader of Kim's second embodiment describes sending SI in two different configurations:
- (a) First, where the UE selects the special E-TFCI and uses it to send SI alone only where (i) sending SI has been triggered in accordance with section 11.8.1.6 and (ii) the UE has no data that is allowed to be sent.
 - (b) Alternatively, where sending SI has already been triggered by one of the triggers in section 11.8.1.6 but where the UE has a grant that allowed data to be sent, SI will still be sent together with data.
233. This process would have reflected the PMS and would have been consistent with the Skilled Reader's CGK. Accordingly, Mr Townend was of the view that in Kim's second embodiment, there was no disclosure of SI being sent *in response to* having a non-zero grant smaller than needed to transmit a MAC-e PDU.
234. In cross-examination, Mr Townend acknowledged there are certain aspects of Kim which are inconsistent with HSUPA as it was specified at the priority date – I have set out the list above. However, he pointed to [0058] as supporting his reading of [0092] and suggested that the Skilled Reader would naturally use his or her CGK (of HSUPA) to fill in the gaps in Kim's teaching. When focussing on the process of E-TFC Selection, Mr Townend pointed out that the alternative reading (where the largest E-TFC that is supportable by the grant is selected) leads to an absurd result. He illustrated this by assuming an RLC PDU of 336 bits and a grant of 153 bits. On the alternative reading, E-TF 2 would be selected (of size 153 bits), SI would be inserted and the remainder would be filled by padding. As he put it, sending these 'padding-filled MAC-e PDUs' would not make any sense.

235. In Dr Irvine's cross-examination, it became clear that he had been asked to consider Kim on a particular basis (emphasis added):

12 Q. It is implicit, therefore, to someone reading Kim that if it
13 want to know about, for example, the E-DCH channel it mentions
14 in paragraph 0006 or the HARQ protocol they refer to in
15 paragraph 0007, they will need to refer to the 3GPP standards
16 to find out what these are?

17 A. If they want to find out what the 3GPP versions are, yes.

18 Q. Should they be looking for some other versions as well?

19 A. **I was not reading Kim to only relate to how HSUPA was**
20 **specifically implemented at the time.**

21 Q. What other systems were you reading that you are referring to?

22 A. **I was reading it more generally as looking for an inventive**
23 **step within it.**

24 Q. Were you looking for an inventive step within Kim?

25 A. Well, what is it telling me to do?

2 Q. Yes.

3 A. **I was given Kim and I was asked, "What does Kim tell you**
4 **about**
5 **scheduling?"**

And a little later, to the same effect:

20 Dr. Irvine, coming back to what you actually said, we are on
21 page 53 of your report and you are now talking about Kim
22 potential developments. When you wrote this, you were told to
23 focus on scheduling information; correct?

24 A. That is correct.

236. Whilst he maintained his reading of [0092] of Kim, Dr Irvine accepted (a) that the only paragraph in Kim which he considered might describe a new trigger for SI was [0058], but he acknowledged that could be read consistently with the CGK; (b) that on Mr Townend's view of the Fig 9 embodiment, it would not be understood as adding a new trigger for SI which was not in the CGK and (c) that Fig 9 could be read in that way.

237. Dr Irvine's additional point was that if Fig 9 of Kim was read in this way, the Skilled Reader would see there was no difference between the Fig 9 and Fig 10 embodiments of Kim. However, this Fig 10 point had not been discussed in any of the expert reports and had not been put to Mr Townend. I have no idea where any such discussion might have led, so I will not consider Fig 10 any further.

238. Counsel put to Dr Irvine that on his approach, the Skilled Reader reads Kim with his or her CGK of HSUPA in mind up to Fig 9, but then abandons CGK at that point. Although Dr Irvine denied he had suggested that, his answers were revealing:

5 A. My boss has given me this document and told me, "Go away
6 and

7 mine this for anything interesting". ...

- 15 Q. The way you have got there, you have been told to think about
16 Scheduling Information and you have asked yourself what else
17 you can get from Kim?
18 A. Yes, that is right.

239. In closing, IDC submitted that once one removes what they characterised as these misconceptions, the whole basis of Lenovo's argument, namely that the Skilled Reader would see something special in the Fig 9 embodiment which allows HSUPA to be stripped away, falls away, as Dr Irvine acknowledged:

- 21 Q. If you are just reading Figure 9, absent knowledge that you
22 had about focusing on scheduling information, in particular,
23 you would think this is just another illustration of the
24 central idea of Kim?
25 A. I think all the diagrams are just illustrations of the central
2 idea of Kim, but yes.

240. IDC's final point on the disclosure of Kim was that there is nothing in the Fig 9 embodiment which mentions or hints at transmission blocking. IDC submitted that the only reason why Dr Irvine went down this line of thinking was because of hindsight, taken in conjunction with his instructions:

- 12 Q. There is nothing in Figure 9 or anywhere in Kim that alerts
13 the skilled addressee to the existence of the
14 transmission-blocking problem, is there?
15 A. To the transmission-blocking problem? It does not
16 specifically flag it up to them. It simply says, "If you do
17 not have enough resources to send a MAC-e PDU, but you do
have
18 enough resources to send scheduling information, send
19 scheduling information". So that is what Kim says.

241. IDC submitted that there is nothing in the Fig 9 embodiment which 'says' anything of the sort.

242. In their closing, Lenovo argued that their anticipation case based on the scheme of Figure 9 of Kim was straightforward. The reasoning was explained as follows by reference to [0092]:

- (a) "*Considering transmission resources granted from the Node B*" – i.e. based on the grant – the UE works out "the size of a MAC-e PDU to be transmitted" in the next transmission. This step must of course entail the possibility that the UE determines that its grant is not zero but is insufficient to include any MAC-es PDU in a MAC-e PDU.
- (b) The UE selects the E-TFC accordingly. Again it is clear (and certainly so once the skilled person reads on to [0093]) that the UE may have to select the smallest possible E-TFC, i.e. the "special" E-TFC.

- (c) The MAC-e PDU is then configured according to the selected E-TFC: “Accordingly, the UE receives data, suitable to the size denoted by the E-TF, from the E-DCH control block or an RLC buffer...”. It inevitably follows that the amount of data “suitable to the size denoted by the E-TF” may in fact be only control information, i.e. no payload data at all. (Kim has already expressly disclosed at [0078] that the special E-TFC “may not be used for a MAC-e PDU accommodating RLC PDUs”.)
- (d) Lastly in step 905 the UE then “configures a MAC-es PDU or a MAC-e control SDU with the received data, and then configures a MAC-e header for the MAC-es PDU and/or the MAC-e control SDU and attaches the MAC-e header to the MAC-es PDU and/or the MAC-e control SDU”. In the case where there is no room for any MAC-es PDU payload, this step can only entail configuring the MAC-e control SDU alone (i.e. SI), and then adding a MAC-e header.
243. Steps 910, 915 and 920 at [0093] and [0094] are then simple to follow: if the E-TFC value associated with the MAC-e PDU configured according to [0092] is the special E-TFCI, the header is removed and the MAC-e control SDU (the SI) is sent alone (in the form illustrated in Figure 7B).
244. On this basis, Lenovo submitted that Kim anticipates claim 1 i.e. where the UE has a non-zero grant that is too small to transmit a MAC-es PDU, then it will send SI (provided that the grant is large enough to permit that).
245. In the alternative, Lenovo submitted that to the extent that such a system is not clearly and unambiguously disclosed in [0092], then an obvious way to implement [0092] would be as Dr Irvine suggests and as set out above.

Anticipation - discussion

246. When assessing this central dispute, I bear in mind that, in terms of its terminology, the experts agreed that the Skilled Person would recognise that Kim uses the term ‘scheduling information’ and ‘MAC-e control information’ to refer to what is SI in TS 25.321, and that Kim uses ‘transport format’ or ‘E-TF’ to refer to E-TFC/E-TFCI.
247. In my view, the answer to the central issue is provided by considering the context within which the teaching in Kim would be read and understood by the Skilled Person.
248. As far as I could detect, only two possible specific contexts were put forward: either the Skilled Person would read and consider Kim’s disclosure in the context of implementing HSUPA or in the context of a WiMAX system.
249. Dr Irvine did point out that HSUPA was an unusual communications system because there is so much variability in what it is that you can send. He referred to time division multiplex systems which use fixed transmission intervals into which ‘I can put whatever it is I am able to send’.

250. So far as WiMAX is concerned, it was not discussed in any of the written expert reports and it follows that I received no information about it, whether as part of CGK or otherwise. I do not know whether it was a TDM system or not. It happens to have been discussed in one of Dr Irvine's exhibits – an extract from the December 2006 edition (i.e. post-priority) of the IEEE Vehicular Technology magazine which was edited by Dr Irvine in the context of a proposal for 4G. The first mention of WiMAX in direct evidence came from Dr Irvine in an answer on Day 4.

251. Dr Irvine made it clear in his answers that he thought Kim was pointed/directed at HSUPA, but Counsel picked up that he was suggesting Kim could have been aimed at other systems and pursued that suggestion in this exchange:

- 7 Q. What else do you think they might be trying to implement,
8 other than HSUPA?
9 A. Any system where you have, you are sending Scheduling
10 Information and you want to improve the efficiency of sending
11 that Scheduling Information by not identifying it with a
12 header.
13 Q. Right. Would you at least accept that the most natural way to
14 read Kim is together with TS 25.321, explaining how the MAC
15 layer works in UMTS?
16 A. I would certainly accept it is a natural way of doing it.
17 Q. Why is it not the most natural way? Can you give another way
18 that is more natural than that?
19 A. If I was reading Kim and I was looking, for example, at WiMAX
20 and I had that in my mind, I might be reading Kim and saying,
21 "Is there anything in this that I could take from that which
22 might then apply to my particular situation?"

252. The only other mention of WiMAX came a few pages later:

- ...you drew attention to the use of the term "E-TF" in
2 Figure 9.
3 A. Yes.
4 Q. That is a term that relates to HSUPA rather than something
5 such as WiMAX; correct?
6 A. Yes, but because the patent had explained what it meant by
7 E-TF, then I would be able to take that concept and apply it
8 to other technologies.

253. Whilst I entirely accept that the Skilled Person would have been able to take the central idea of Kim and apply it in other communications systems, such as WiMAX, I do not consider the reference to WiMAX or other systems takes the analysis any further than that.

254. I accept Mr Townend's view that the Skilled Person would read Kim with his or her HSUPA CGK in mind. On that basis, I accept Mr Townend's view as to what the Fig 9 embodiment of Kim would disclose to the Skilled Person, as summarised in paragraphs 229-234 above. I also accept that Dr Irvine's alternative reading was a hindsight view prompted by the particular way in

which he was invited to consider Kim. For these reasons, I reject Lenovo's case that Kim anticipates.

Obviousness

255. 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.
256. I can deal with Lenovo's case of obviousness over Kim relatively briefly.
257. In Dr Irvine's first report, when assessing potential developments from Kim, he considered that the figure 9 embodiment would be "*potentially useful and would be of most interest to the Skilled Person*".
258. In their Opening Skeleton, Lenovo submitted that a real person working at the priority date would not use Kim as a starting point for commercial reasons, namely that they would not be keen on producing an alternative system to 3GGP HSUPA. (The submission was that such considerations are irrelevant to the Skilled Person). Dr Irvine was asked whether he agreed:

Q. Do you agree that a real person working at the priority date would not use Kim as a starting point, full stop?

A. I have personally been asked to go and look at old patents to see if they can be mined again for useful information. Would somebody take the system that Kim had described and implement it? There is HSUPA out there. Would they go back to a document like Kim and look for any interesting nuggets? Maybe.

Q. The highest you can go is "maybe", assuming somebody is looking for nuggets?

A. Yes.

...

Q. The skilled person interested in developing HSUPA Release 7 at the priority date is not going to go back in time and start developing a system they know to be obsolete, are they?

A. This is where you come back to your nuggets. You might look at older documents which had been following a particular theme, and knowing what you know now, go back and look at those to see if there are any interesting ideas that had not been followed through that you could now then follow through.

Q. You are assuming someone who has some motivation to go looking through these documents for the nuggets and they get to Kim and they do not read Kim as just talking about HSUPA; they read Kim as talking about something more general?

A. Yes.

Q. As to suggesting which had never in fact been done in HSUPA?

A. Yes.

Q. But have you to do all that to get the nugget?

A. Yes.

- Q. Can I suggest that these amount to good technical reasons as to why the skilled addressee would not go through that process?
A. They do amount to good technical reasons that the skilled addressee would not go through those; yes.

259. So far as considering Kim in the context of HSUPA is concerned, the Skilled Person would note (as I have said above) that the central idea of Kim had already been incorporated into HSUPA in Release 6. This was the reason why Dr Irvine accepted that Kim was ‘really old news’ by the priority date.
260. I do not think the Skilled Person would read [0092] to [0094] as a suggestion to re-write the process of E-TFC restriction and selection which was by then in the standard. It was not suggested that the Skilled Person had any motivation to reconsider, let alone re-write the processes of E-TFC restriction and selection which were in the standard (which is what Dr Irvine’s approach would have required). Those paragraphs would be read as just another way of explaining how the central idea in Kim could be implemented. On reading those paragraphs, the Skilled Person would, in my view, simply note that this was yet another respect in which HSUPA had moved on since Kim was written. In other words, having considered Kim and his Fig 9 embodiment with interest, the Skilled Person would have no motivation to do anything except put it on one side.
261. Overall, I formed the view that this obviousness argument was also driven by hindsight and I reject it.

Claims 2 & 4

262. In case I am wrong about the validity of claim 1 over Kim, I will consider the position of claim 2. I agree with IDC that claim 2 remains inventive. I need only cite the following two passages of cross-examination. In the first, Mr Townsend was effectively asked to assume Dr Irvine’s understanding of [0092]-[0095] of Kim, but his answer led back to the earlier discussion of Lenovo’s Case 2 (i.e. whether claims 2 & 4 were obvious based on a simple modification to the PMS):

- 2 Q. I understand. Let me rephrase my question. If they are
3 building a system in accordance with their understanding of
4 paragraphs 91-95 of Kim and following those instructions as
5 best they could, it would be an obvious thing for them to do
6 to include within their system a timer like T_SIG and T_SING
7 which provided for SI to be sent, triggered periodically.
8 A. Unfortunately, I think it does matter whether, how to read,
9 how we are reading the second embodiment, because under what I
10 will call my reading, the reading consistent with the CGK, it
11 would certainly be natural to add the triggers if they were
12 not already there, sorry, to add the periodic sending, because
13 my reading relies on there being triggers in the first place.
14 Under what I can call your reading, what it amounts to
15 is opportunistic SI, because what you are saying, I think, is
16 that you have picked an E-TFC and if there is space in it, you

17 will put in Scheduling Information. So we are then, I think,
18 back into that whole thing of does it make sense to add
19 periodic scheduling on top of opportunistic SI? I am not
20 convinced that that is an obvious thing to do, for all the
21 reasons we talked about this morning.

263. IDC also pointed to the extent of the ingenuity in Dr Irvine's answer of how to get to claim 2 from his understanding of Kim, and that in this answer even he was equivocal:

22 A. Claim 2 takes a little bit further. The technical reason for
23 doing what Kim suggests is if you do not have enough resources
24 to send a MAC-d PDU, but you do have enough resources to send
25 scheduling information, because the resources have already
2 been assigned to the mobile, you may as well use those
3 resources to send scheduling information. There is, however,
4 a small cost to the system, because you are transmitting and
5 that transmission will introduce interference. You gain,
6 because you get the scheduling information, you lose because
7 of that small interference rise. **Whether you would do that**
8 **continuously, well, the gain falls off very, very quickly.**
9 The loss is still there all the time. **Whether then that is**
10 **obvious, I think a good engineer would spot that, but then I**
11 **suppose that is a matter of opinion.** Immediately you then
12 turn round and say, "Well, I should not be doing it every time
13 that situation occurs, I should only be doing it a proportion
14 of the time that situation occurs, because I do not need to
15 send fresh Scheduling Information, I get a lot of value the
16 first time. Hey, I am stuck, I do not get that value the next
17 time, all I get is the cost of the interference." If I
18 recognise that, configuring the period for that, claim 4 falls
19 straightaway.

264. For completeness, I find that claim 4 was not inventive over claim 2. On that basis, it was obvious to allow the network to configure the period of the timer. Indeed, in IDC's Opening, this feature was acknowledged to be CGK.

Overall Conclusion

265. For all the above reasons, I find EP(UK) 2 421 318 B1 to be valid and essential.

266. I will hear Counsel as to the form of Order if it cannot be agreed. I direct that time for seeking permission to appeal shall not run until after the hearing on the form of Order (or the making of such Order if it is agreed). I draw attention to paragraph 19.1 of the Patents Court Guide, which says that a hearing on the form of Order should take place within 28 days of hand down. In the present case, 28 days from hand down will be 28 February 2023.

CGK Annex

Mobile Telecommunications Systems

1. Mobile telecommunication systems are characterised, as they were at the Priority Date, by the support of user mobility, which 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.
2. 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. This is illustrated in Figure 1 below. Each cell (shown as rough circular shapes in different shades) is served by a base station transmitter and receiver (or transceiver) (shown as small triangles) through an antenna system.




Figure 1: The concept of cellular coverage

3. In some mobile telecommunication systems (such as that shown in Figure 1 above) each base station would serve a single cell. In other mobile telecommunication systems, the base stations would often serve more than one cell each.
4. A single base station may have a number of directional antennas, each of which serves a cell. The cells of each antenna overlap and base stations might be situated so that their cells cover adjacent and overlapping areas. This is shown in Figure 2 below: three base stations are shown (1, 2 and 3) each of which has three directional antennas (A, B and C) which in turn each serve a cell (indicated by the blue, red and yellow cardioids).

5. If a mobile phone is in an overlapping area, it might have a connection to multiple cells, in which case it is said to be in **soft handover** between those cells. In the special case where the mobile phone has a connection to multiple cells of the same base station (and it receives a single power command through all of those base stations), it is said to be in **softer handover** between those cells, and the base station to which those cells belong is able to combine those signals. The set of cells with which a mobile phone is in soft (or softer) handover is known as the **active set**.

6. Figure 2 shows a representation of a mobile phone moving through the area covered by a mobile telecommunications system employing soft and softer handover, at five different instances, with five different active sets. It is not necessary to go through each instance, but note that at instance 5, the mobile phone's active set consists of cell 1B, cell 1C, and cell 2A. Because cell 1B and cell 1C are served by the same base station (base station 1), at instance 5 the mobile phone is in softer handover between those two cells. Additionally, at instance 5 the mobile phone is in soft handover between (i) cells 1B and 1C, and (ii) cell 2A (because (i) cells 1B and 1C, and (ii) cell 2A are not served by the same base station).




Figure 2: Soft handover, softer handover, and the active set

7. 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.
8. The user devices, such as mobile phones, cellular data modems in computers and tablets, 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**.
9. Uplink and downlink communications within a cell are usually distinguished because they are transmitted either:
 - 9.1 using different frequencies (i.e. Frequency Division Duplex, FDD) (i.e. one set of frequencies are used in uplink, another in downlink); or
 - 9.2 at a different time (Time Division Duplex, TDD).
10. Communications from different users to a base station, and from different base stations (or, more specifically, different cells), are distinguished using:
 - 10.1 different frequencies (i.e. each user used a different uplink frequency, and different cells used different downlink frequencies), known as Frequency Division Multiple Access (**FDMA**);
 - 10.2 different timeslots, known as Time Division Multiple Access (**TDMA**); and/or
 - 10.3 different codes, known as Code Division Multiple Access (**CDMA**).
11. Real life systems often use combinations of these technologies. At the Priority Date in the UK, 3G networks primarily 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 an **FDD WCDMA** system and generally in this CGK Annex, references are to a FDD system, unless explicitly stated otherwise.
12. Amongst the 3G systems developed is the UMTS system developed by the 3rd Generation Project Partnership (**3GPP**). 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.

13. 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).
14. For example, in Figure 2 above:
 - 14.1 Communications in each of the cells 1A through 3C will use the same pair of frequencies for uplink and downlink communications;
 - 14.2 Each of the cells 1A through 3C will use a different set of scrambling codes; and
 - 14.3 Each mobile phone will use a different scrambling code.
15. It is not necessary to explain in any detail how this is accomplished, however, there are a number of important consequences of this technology:
 - 15.1 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.
 - 15.2 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.
 - 15.3 When operating close to the boundary between two (or more) cells controlled by different base stations, soft handover (or “macrodiversity”) must be used. This is different from previous systems (save for IS-95 CDMA, which was not implemented in the UK), where devices are only connected to a single cell at a time.

- 15.4 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.
16. The remainder of this section contains a description of the aspects of UMTS which are relevant to these proceedings. Many of these aspects of UMTS are also used in HSUPA, with certain modifications, which are explained below.

UMTS System Architecture

17. Figure 3 below shows the overall system architecture of the UMTS system, and its connections to external networks. No relevant modifications were made to the system architecture to implement HSUPA.

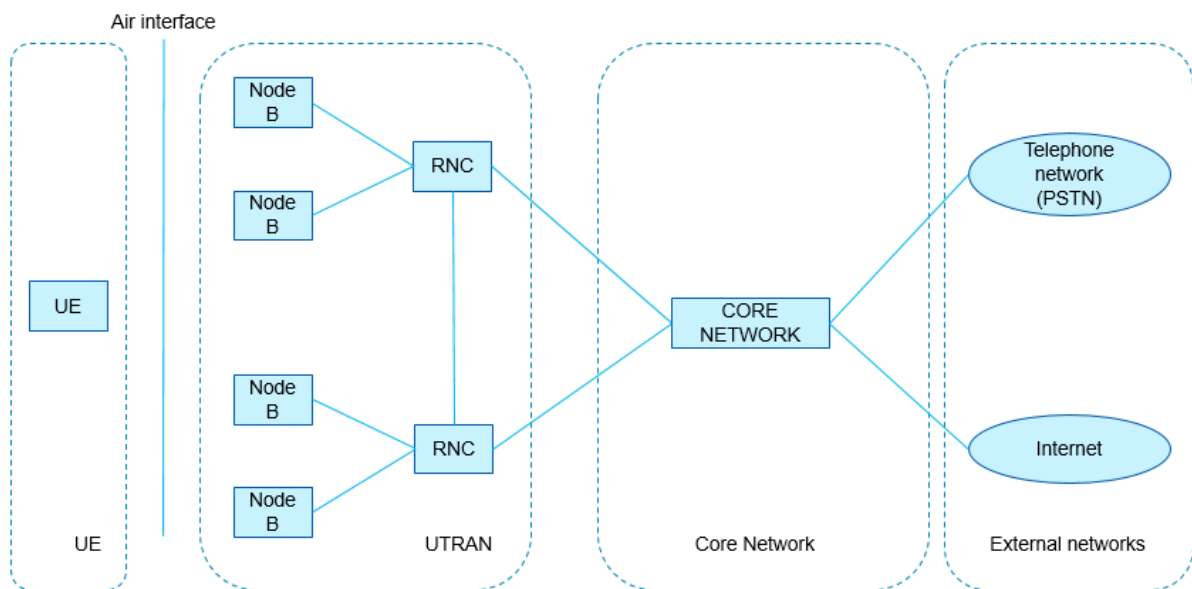


Figure 3: UMTS overall system architecture, functional network elements and logical entities

18. The UMTS system can be thought of as comprising the following functional network elements, as shown in Figure 3 above:

- 18.1 The User Equipment (**UE**), typically a mobile phone;
- 18.2 The UMTS Terrestrial Radio Access Network (**UTRAN**), which handles the radio-related functionality; and

- 18.3 The Core Network, which provides switching and routing of calls and data connections to external networks (shown to the right of the Core Network).
19. 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 the functionality of the Core Network.

Logical entities of the UTRAN

20. As can be seen in Figure 3 above, the UTRAN comprises the following logical entities:
- 20.1 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 4 above; and
- 20.2 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). Each Node B is controlled by a single RNC.
21. As Node Bs are the UMTS-specific version of base stations, the terms “active set”, “soft handover”, and “softer handover” can be applied to Node Bs and cells of those Node Bs in the same way that they were applied to base stations generally.
22. RNCs can serve several different roles simultaneously and are named according to which role is relevant. Thus, if an RNC is described as the Serving RNC (**SRNC**), it means it is the RNC which controls Radio Resource Control (**RRC**) signalling (see further below) for a UE and maintains a connection to the Core Network for that UE. No other roles of the RNC are relevant for the purposes of this trial. RRC signalling from the SRNC is sent through all Node Bs with cells in the active set.
23. 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.

UMTS protocol stack

24. UMTS characterises the functions of the UE, the Node B and the SRNC by specifying the protocols in them.
25. A protocol specifies a set of predefined rules that allow devices to communicate with each other.
26. The relevant UMTS protocols for this report are, referring to the bottom three layers of the OSI model:
 - 26.1 layer 3: Radio Resource Control (**RRC**);
 - 26.2 layer 2:
 - 26.2.1 Radio Link Control (**RLC**); and
 - 26.2.2 Medium Access Control (**MAC**); and
 - 26.3 layer 1: Physical (**PHY**).
27. 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 at paragraph 35 below.
28. 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 – the higher layers are not relevant in this case. Strictly in UMTS there is some physical layer functionality in RNCs and some (very specific) MAC functionality in the Node B, but they are not relevant to this case, so we can proceed on the basis that in the UTRAN in UMTS, 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). Therefore, the Node B and the SRNC combined provide the corresponding functionality (or protocol entities) for all the relevant protocols in the UE.

29. 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 4 below (adapted from Figure 11 of TS 25.301 V6.2.0.) shows some physical layer functionality in the SRNC. This is to deal with the user device being connected to two (or more) cells at once during "macrodiversity" (as mentioned below) and shows the protocols in the UE, the Node B and the SRNC in UMTS. Protocol entities connected with horizontal lines communicate with each other, using lower protocol entities to carry the data.




Figure 4: UMTS data flows through protocols and between logical entities

30. There are a few points to note about this scheme:
- 30.1 Different data will have different originating points and different terminating points. For example:
 - 30.1.1 An email will go from a UE to an endpoint outside the network;
 - 30.1.2 RRC control messages configuring a service will need to go from the SRNC to the UE; and

- 30.1.3 Layer 1 control messages such as power control commands will need to go from the Node B to the UE.
- 30.2 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).
- 31. 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.
- 32. Data which is sent over the UMTS system will either be:
 - 32.1 User data (e.g. a webpage) – also referred to as “user plane data”; or
 - 32.2 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”.

Radio bearers and channels

- 33. Figure 5 below (adapted from Figure 2 of TS 25.301 V6.2.0) shows the protocols, radio bearers and channels relevant to this report. The left hand side of Figure 5 shows flows of control data through the protocol stack. The right hand side shows flows of user data through the protocol stack.

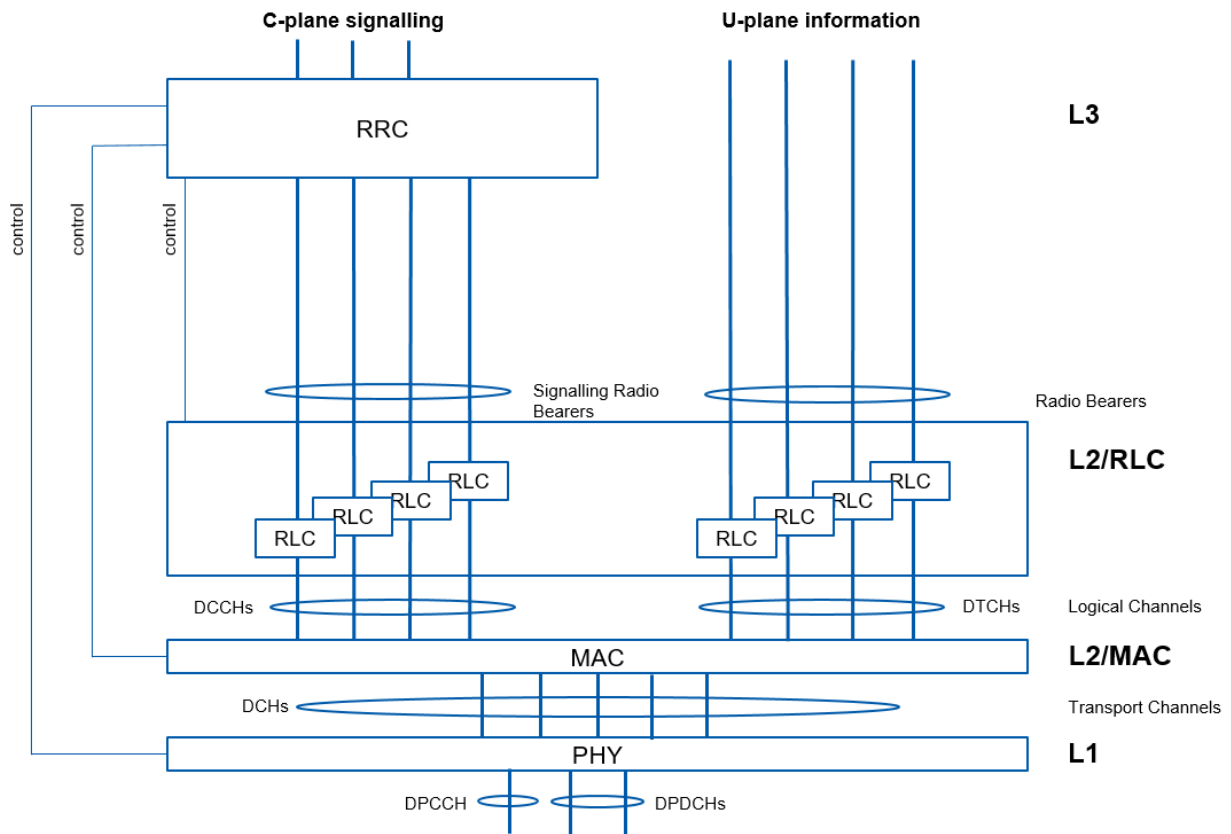


Figure 5: Layers, Protocols, Radio Bearers and Dedicated Channels of UMTS

34. As can be seen in Figure 5 (reading from top to bottom, and left to right):
- 34.1 A “**Signalling Radio Bearer**” is a flow of control data (described as C-plane signalling) between the RRC and RLC protocol entities.
 - 34.2 A “**Radio Bearer**” is a flow of user data (described as U-plane information) into the RLC protocol entity.
 - 34.3 A “**logical channel**” is a flow of data (either user (referred to as “traffic”) or control) between RLC and MAC protocol entities.
 - 34.4 A “**transport channel**” is a flow of data (user, control or both (because of C/T multiplexing)) between MAC and the physical layer protocol entities.
 - 34.5 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).

35. In the UE, for uplink communications, the protocols perform the following functions:

35.1 **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**) provided by lower layers in response to requests from higher layers in the protocol stack. 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 (see further below).

35.2 **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. It stores data received from higher layers in its buffers, maps each flow of data it receives onto its own logical channel (i.e. it takes data from a Radio Bearer, processes it (usually by adding a header) and then puts it on a defined logical channel. If it receives three Bearers, it will output three logical channels).

35.3 **MAC** is responsible for the selection and transmission of data, 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 that deals with **dedicated** transport channels (the **MAC-d**) is relevant. The logical, transport and physical channels defined for UMTS relevant to this case 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 **MAC-d** is responsible for mapping to the **Dedicated** (transport) **CHannel(s)** both the **Dedicated Control CHannel(s)**

(**DCCH**), which are the logical control channel(s), and the **Dedicated Traffic CHannel(s) (DTCH)**, which are the logical user data channel(s). Figure 5 above only shows dedicated logical channels being mapped onto dedicated transport channels, and so the MAC layer shown in it is effectively the MAC-d.

35.4 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. The physical layer multiplexes (**Dedicated**) transport **CHannels (DCH)** which are to be transmitted using the same type of (**Dedicated**) **Physical Data CHannel (DPDCH)** into a **Coded Composite Transport CHannel (CCTrCH)** of a particular type, so that they can be transmitted using the same DPDCH. There can be more than one DPDCH if required due to a sufficiently high data rate, or if there is more than one type of CCTrCH. Control information which originates in the physical layer will be carried on the **Dedicated Physical Control CHannel (DPCCH)**.

36. Set out below are the relevant UMTS channels.

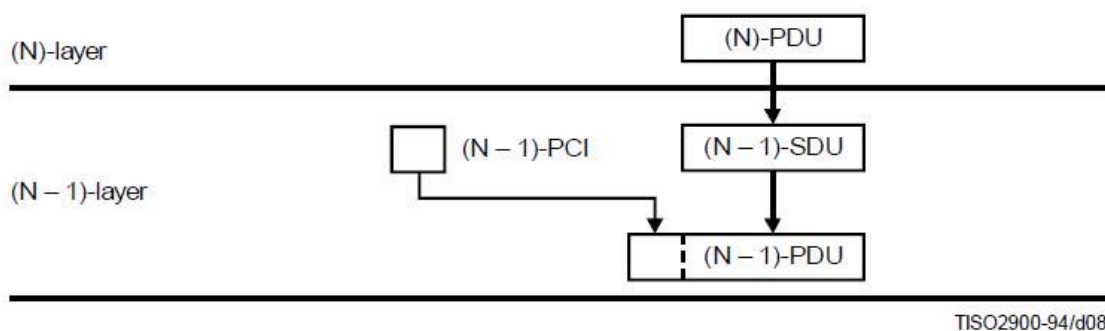
	User Data	Layer 3 Control Data	Physical Layer Control Data
Bearer	Radio Bearer (RB)	Signalling Radio Bearer (SRB)	
Logical Channel	Dedicated Traffic CHannel (DTCH)	Dedicated Control CHannel (DCCH)	
Transport Channel	Dedicated CHannel (DCH)		
Physical Channel	Dedicated Physical Data CHannel (DPDCH)	Dedicated Physical Control CHannel (DPCCH)	

Table 1: Relevant UMTS channels

37. 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 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.

PDUs and SDUs

38. 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”.
39. Figure 6 below (taken from Figure 9 of the **ITU-T X.200** standard (07/94), commonly known as the OSI 7-layer model) 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”.



PCI Protocol-control-information
PDU Protocol-data-unit
SDU Service-data-unit

Figure 6: PDU construction

40. 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 above.

41. Thus:

41.1 RLC receives RLC SDUs and outputs RLC PDUs.

41.2 MAC receives RLC PDUs (as MAC SDUs) and outputs MAC PDUs. In particular, if the MAC PDUs are dedicated MAC PDUs, MAC-d outputs **MAC-d PDUs**.

42. The physical layer receives a MAC PDU as a **Transport Block (TB)**. However, this is merely a difference in terminology, and the Transport Block is in effect the physical layer SDU.

Quality of Service (QoS)

43. 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.

44. 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).

45. There are four QoS traffic classes in UMTS:

45.1 conversational class;

45.2 streaming class;

45.3 interactive class; and

45.4 background class.

46. 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, differences in 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.

47. 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 (such as voice, video and web browsing) for which a QoS traffic class may be used (respectively, conversational class, streaming class and interactive class).
48. Within each class, priority levels may be specified. Using these, the relevant network elements can 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.
49. 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

50. 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 for the SRNC to receive from Node Bs having cells in the active set the UE's transmissions with the desired block error rate (**BLER**). Using the minimum required power means that the transmission from one UE is less likely to block or unnecessarily interfere with transmissions from other UEs using cell(s) in its active set.
51. 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 BLER; the SRNC then sends this SIR target to those Node B(s). The Node B(s) instruct the UE to increase or decrease transmit power to maintain the SIR as close as possible to the target.
52. 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).
53. 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.

54. 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.
55. 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

56. 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.
57. Each application in use will require logical traffic channels (i.e. DTCH) to send the user data, and logical control channels (i.e. DCCH) to manage the channels carrying the user data. So in the example in the preceding paragraph:
- 57.1 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;
- 57.2 RRC would need to:
- 57.2.1 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 in order to send the user data over the air interface; and
- 57.2.2 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).
58. 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-d 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 four-bit 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 demultiplex 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 (TB) by the physical layer.

59. The physical layer can also multiplex transport channels into a single Coded Composite Transport Channel (CCTrCH) for further processing by the physical layer. For example, the three DCH transport channels carrying the user data for the user’s three applications in the example above (i.e. carrying the three DTCH logical traffic channels), would be multiplexed in the physical layer into a single CCTrCH if they are going to be sent over a single dedicated physical channel (or multiple physical channels of the same type, a solution adopted where one physical channel cannot carry the amount of data required). 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.

Sending data

60. Data is transmitted within defined periods known as Transmission Time Intervals (**TTI**). Every TTI (for example, every 10ms) a UE will need to (i) determine which data to send to the Node B and (ii) send it.
61. In UMTS, MAC is responsible for determining which data to send from a UE to a Node B. How it does so is tightly controlled.
62. Each TTI a UE can send a single Transport Format Combination (**TFC**) to the Node B. A TFC is, as the name suggests, a combination of “Transport Formats” that specifies the size of packet from each transport channel and the total size of data sent for each transport channel. When the UE has data to

transmit, it selects a TFC and fills it with data. The skilled person would recognise that the PMS in section 11.4 specifies how this is done in some detail but, since the way it is handled in UMTS is not relevant to this dispute, it is not necessary to set it out.

HSDPA

63. 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. As previously noted, 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. HSDPA was not (and is not) a standalone system: a UE implementing HSDPA is required by the standard to also implement UMTS.

64. By the Priority Date, the standardisation of HSDPA was sufficiently complete that HSDPA was already part of commercially available products.

65. HSDPA introduced various techniques into UMTS, including:

65.1 A shorter (2ms) TTI (compared to the previous shortest TTI of 10ms).

65.2 A more complicated, but more efficient, form of error correction (known as Type II Hybrid Automatic Repeat Request or **HARQ**), being a combination of methods for correcting errors in a transmission.

65.3 Node B-controlled scheduling (i.e. the Node B is 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, HSDPA is more responsive than the legacy channels and thus facilitated the use of shorter TTI and HARQ. In particular, Node B controlled scheduling in HSDPA allows the Node B to decide what resources to allocate to downlink transmissions to a UE in each TTI.

65.4 Higher order modulation.

HSUPA

66. This is a long section and it may assist to set out how it is structured:

66.1 First, how and when HSUPA was introduced into the 3GPP specifications and its status at the Priority Date.

- 66.2 Second a high level overview of how data is sent in HSUPA.
- 66.3 Third, the changes to the UMTS protocol stack and the channel structure which were required to support HSUPA and in particular certain aspects relied upon by Kim.
- 66.4 An overview of the relevant functions of the MAC protocol in HSUPA, this being the part of the HSUPA to which the Patent is addressed.
- 66.5 Fifth, it is necessary to cover in some detail the mechanisms provided in the 3GPP specifications to implement Node B scheduling of uplink data transmission. This is the element of HSUPA in which the problem addressed by the Patent lies and to which the claims of the Patent are addressed.
- 66.6 Sixth, the structure of the MAC-e PDU used in HSUPA which is necessary to understand Kim in context at the Priority Date.

1 The Introduction of HSUPA into the 3GPP Standards

- 67. HSDPA introduced enhanced downlink capabilities, but they were not matched by uplink capability until HSUPA was introduced in Release 6, which was complete by September 2005, with only corrections and clarifications being added after that date. The 3GPP working groups moved on to the development of Release 7 which would include further HSUPA functionality. As indicated by the PMS itself, early Release 7 specifications started to be released shortly before the Priority Date, but Release 7 was in active development at the Priority Date. HSDPA networks were commercially available by the Priority Date but not HSUPA networks (the first commercial HSUPA network was launched in February 2007).
- 68. Like HSDPA, HSUPA was not (and is not) a standalone telecommunications system. A UE implementing HSUPA is required by the standard to also implement both UMTS and HSDPA. Therefore, much of the description of UMTS above is also an accurate description of HSUPA. As was the case in UMTS, the behaviour of the UE in both HSDPA and HSUPA is very tightly specified, while the behaviour of the network is, to the greatest extent possible, left to implementation.
- 69. The aim of HSUPA was to provide enhancements to the uplink similar to those provided to the downlink by HSDPA, in particular in terms of user experience

(throughput and delay) and capacity. It introduced a new “enhanced” uplink transport channel called the **Enhanced (Dedicated) Transport Channel (E-DCH)** - Dedicated because each UE has its own E-DCH data path to the Node B that is continuous and independent from the DCHs and E-DCHs of other UEs. HSUPA also introduced to the uplink techniques similar to those in HSDPA, namely:

- 69.1 the option to choose a shorter 2ms TTI;
- 69.2 fast physical layer HARQ error correction (minor differences not being relevant); and
- 69.3 fast Node B controlled scheduling.

Although unlike HSDPA, HSUPA does not support any higher order modulation schemes.

70. HSUPA also supports the sending of “non-scheduled data”, which is data that is not subject to fast Node B scheduling. The existence of such data is relevant, but the details of its operation are not, so it is not addressed in detail.

2 *High-level overview of HSUPA*

71. The aspect of HSUPA most relevant to this trial is the fast Node B-controlled scheduling. Being physically close to the air interface, the Node B has more instantaneous information about the uplink interference situation than the SNRC and can control uplink data rates in a rapid manner.
72. Put simply, uplink scheduling in HSUPA allows the Node B to control how much data (and power) each UE within its service area is transmitting at any moment in time, both to optimise the user experience of all users and to minimise interference, in turn increasing the overall capacity and throughput of the system. For example, more resources might be allocated to users with a lot of data waiting to be sent, or they might be allocated to users close to the Node B for whom more of the power would be used to send data and less to overcoming the pathloss to the Node B.
73. The scheduling function in HSUPA is split between the Node B (which is responsible for resource allocation) and the UE (which is responsible for making use of the resource allocation to send data).

74. The Node B scheduling is implemented by the UE receiving grant messages from the Node Bs with which it is in soft handover (i.e. where a UE has a connection to multiple cells) and using those to derive a single “Serving Grant” value which will be used to control the amount of scheduled data that can be sent at a time. The Serving Grant is quantified as a relative power level, expressed as a ratio to the power at which the existing UMTS Dedicated Physical Control Channel (DPCCH) is transmitted, but is eventually converted into a number of bits of scheduled data that can be transmitted.
75. There are two types of grant messages used by the Node B to control the maximum amount of transmit power used to send scheduled data. One is an Absolute Grant (which is sent only by the serving E-DCH cell and which is intended to set the Serving Grant to a specific value) and the other is a Relative Grant (which is used to move the Serving Grant to a higher or lower value).
76. The UE provides status information (known as “Scheduling Information”) to the Node B so that the scheduling function receives, among other things, information on the amount of data in the UE’s buffers awaiting transmission and the radio channel conditions or relative power the UE can use for its data transmissions.
77. The 3GPP specifications define rules for when Scheduling Information is to be sent by the UE, for example, the expiration of a timer, which leads to periodic transmission of Scheduling Information. The triggering conditions are discussed in detail below in paragraphs 128 to 135.
78. The Node B scheduler uses the Scheduling Information sent by the UE together with a so-called “happy bit” (which indicates whether the UE is “happy” or “unhappy” with the current data rate) to make decisions on whether to adjust the UE’s Serving Grant.
79. HSUPA can be configured with either a 2ms or 10ms TTI, which means that a single packet of data (called a MAC-e PDU, or, equivalently, an E-DCH transport block) is sent to the physical layer either every 2ms or 10ms.
80. HSUPA uses a “stop-and-wait” HARQ, which involves the transmitter waiting to send a second block of data until it has received an acknowledgement (**ACK**, an indication that a packet has been correctly received) 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, HSUPA calls for multiple HARQ processes. The operation of HARQ is substantially the same for both 2ms and 10ms TTI, but the number of HARQ processes is not. 2ms TTI used 8 HARQ processes, whilst 10ms TTI used 4 HARQ processes. HARQ processes are conventionally numbered from 0 (zero). This allows the HARQ processes to operate in sequence and cyclically (i.e. simply by knowing which TTI is next, the UE can identify the one operational HARQ process for that TTI), such that in the first TTI, HARQ process 0 transmits, then in subsequent TTIs, HARQ processes 1 to 7 transmit whilst HARQ process 0 waits for its ACK or a negative acknowledgment (**NACK**, an indication that an errored packet has been received), then in TTI 9, HARQ process 0 can either transmit the next packet (i.e. the one after the packet sent on HARQ process 7) or retransmit the packet it transmitted last time (in case there has been a NACK (or no ACK)).




Figure 7: Visual representation of HARQ process cycle

81. As part of the initial set-up of HSUPA for a UE, individual HARQ processes can be “restricted” by RRC (meaning that they cannot be used for scheduled data), and as part of the scheduling grant process, individual (or all) HARQ processes can dynamically be activated and deactivated (meaning that they can and cannot be used for scheduled data).

82. HSUPA entails significant changes from UMTS and HSDPA for the sending of data. In each TTI, the UE can send zero or one MAC-e PDU (partly because there is only one uplink transport channel, the new E-DCH). This MAC-e PDU may contain data from multiple 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). If Scheduling Information is to be sent, it is included in the MAC-e PDU. The size of the MAC-e PDU is constrained to match one of 128 pre-defined sizes (each being an E-DCH Transport Format Combination (**E-TFC**)). The E-TFC is the HSUPA equivalent of, but not the same as, the TFC in UMTS.
83. The process of selecting the E-TFC and determining what data is to be included in the MAC-e PDU is somewhat complex, and is described in greater detail later. An overview is set out here.
84. In each TTI, the UE will first update the Serving Grant (according to any Absolute and Relative Grants received), and will then work out what data, if any, it is able to send in that TTI.
85. This involves multiple steps, including:
 - 85.1 Checking whether the current HARQ process is available to send new data (or if it needs to perform a re-transmission of the last E-DCH transport block);
 - 85.2 Checking whether it is required to send Scheduling Information;
 - 85.3 Checking whether scheduled data is allowed to be sent on the current HARQ process (i.e. it is neither deactivated nor “restricted”);
 - 85.4 Determining the maximum size of E-TFC that can be used, given the UE’s available transmit power; and
 - 85.5 Determining the amount of scheduled and non-scheduled data that can be sent (according to the Serving Grant, non-scheduled grants and the maximum supported E-TFC size).
86. The UE will then multiplex together any required Scheduling Information and the maximum possible amount of data, subject to all of these constraints and

the priority ranking of the logical channels, before selecting the smallest E-TFC that will fit the resulting MAC-e PDU.

87. The UE will then transmit the (scheduled and non-scheduled) data (and any Scheduling Information) on one or more of the new E-DCH Dedicated Physical Data Channels (E-DPDCH(s)), and related control information on the new E-DCH Dedicated Physical Control Channels E-DPCCH (see below).

3 *HSUPA protocol stack and relevant channels*

88. HSUPA introduced new channels for sending and controlling the sending of user data on the new Enhanced Dedicated Transport Channel (E-DCH), 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 (in contrast to UMTS, which had no (relevant) MAC functionality in the Node B) and the SRNC to join the new channels back to the pre-existing system.

89. Figure 8 below (adapted from Figure 5.6.10.3-1 of TS 25.301 V7.0.0.) shows the protocols in the UE, the Node B and the SRNC. As with UMTS, protocol entities connected with horizontal lines communicate with each other, using lower protocol entities to carry the data.




Figure 8: HSUPA Protocol Stack in the UE, Node B and SRNC

90. The MAC-d layer protocol entity was retained in HSUPA, and additional entities were introduced:

- 90.1 A MAC-e/es entity in the UE which handles the E-DCH specific functions;

- 90.2 A MAC-e entity in the Node B; and
- 90.3 A MAC-es entity in the SRNC which handles E-DCH specific functionality which was not covered in the MAC-e entity in Node B.
91. As already stated, HSUPA also introduced new uplink channels. The E-DCH carries user plane data and non-physical layer control data. After transport channel processing, the E-DCH maps on to one or more (up to six, to support higher data rates) uplink physical channels, the **E-DCH Dedicated Physical Data Channels (E-DPDCHs)**. The E-DPDCHs require simultaneous transmission of an uplink physical control channel, the **E-DCH Dedicated Physical Control Channel (E-DPCCH)**, to carry physical layer control data to deliver the information the receiver needs in order to know what format of E-DPDCH transmission is being used.
92. Two downlink physical channels were also introduced to facilitate uplink scheduling:
- 92.1 The **E-DCH Absolute Grant Channel (E-AGCH)**, to carry the Absolute Grant scheduling value for the UE; and
- 92.2 The **E-DCH Relative Grant Channel (E-RGCH)**, to carry the Relative Grant step-up/down scheduling commands.
93. Finally, a third downlink physical channel was introduced to carry HARQ acknowledgment information (ACKs and NACKs) to the UE called the **E-DCH HARQ Indicator Channel (E-HICH)**.
94. Table 2 is an updated version of Table 1 above, showing the newly-introduced HSUPA channels (in red), and identifying whether those channels are uplink or downlink channels.

	User Data	Layer 3 Control Data	Physical Layer Control Data			
Bearer	Radio Bearer (RB)	Signalling Radio Bearer (SRB)				
Logical Channel	Dedicated Traffic Channel (DTCH)	Dedicated Control Channel (DCCH)				
	MAC-d flows					
Transport Channel	Enhanced Channel (Uplink)	Dedicated (E-DCH)				
Physical Channel	E-DCH Physical Channel (E-DPDCH) (Uplink)	Dedicated Data	E-DCH Dedicated Physical Control Channel (E- DPCCH) (Uplink)	E-DCH Absolute Grant Channel (E-AGCH) (Downlink)	E-DCH Relative Grant Channel (E-RGCH) (Downlink)	E-DCH HARQ Indicator Channel (E-HICH) (Downlink)

Table 2: Relevant HSUPA channels

4 RRC role in HSUPA

95. RRC configures whether each **MAC-d flow** (a conceptual grouping of logical channels with similar QoS requirements - shown in green in Table 2 above) is either (i) subject to Node B-controlled scheduling or (ii) allocated to a non-scheduled transmission (but not both).
96. Non-scheduled transmissions are conceptually similar to the allocation of a transport format to a logical channel in UMTS – a UE is given the ability to

configure a specific MAC-d flow to have a guaranteed physical layer data rate, effectively disabling Node B scheduling. In contrast, Node B-controlled scheduling allows the UE to use a maximum amount of power to transmit data from any channel configured to utilise Node B-controlled scheduling.

97. As in UMTS, RRC can control the scheduling of uplink data by giving each logical channel a priority between 1 (highest) and 8 (lowest). The UE is required to maximise the transmission of higher priority data, that is to say, data from the highest priority logical channel with data in its buffer.

5 *Relevant MAC functions in HSUPA*

98. The MAC-e/es in the UE is responsible for determining what data, both user plane data and non-physical layer control data, is sent over the new channels each TTI. As explained below (in section 18, beginning on page 104), how the MAC-e/es selects data for transmission is tightly controlled but is based on the allocated resources and the power available for enhanced uplink transmission. The MAC-e/es multiplexes the selected data in order of its logical channel priorities, subject to certain constraints. Data from the highest priority logical channel is concatenated (i.e. inserted one after the other) into a MAC-es PDU and, if allowed, data from lower priority logical channels is concatenated, in priority order, into separate MAC-es PDUs. (Data from a single logical channel might be multiplexed into several MAC-es PDUs, if the MAC-d PDU sizes are different, but this is not relevant to this dispute). One or more MAC-es PDUs is/are multiplexed in accordance with pre-defined limits into a single MAC-e PDU, until the grants or the available power are exhausted. A MAC-e PDU might also include Scheduling Information and padding bits. A MAC-e PDU is carried on the E-DPDCH as a single transport block, the E-TFC.

99. Each E-TFC governs the overall size of data which can be multiplexed into it. The E-TFC used is signalled on the E-DPCCH by means of a seven-bit E-DCH Transport Format Combination Indicator (**E-TFCI**), which indicates the size of the transmitted data by reference to a lookup table of transport block sizes within the **E-TFC Set**. There are four configured transport block tables for the Enhanced Uplink, two for the 10ms TTI, two for the 2ms TTI, in each case optimised for a 336-bit RLC PDU size, or exponentially distributed. They are set out in Annex B to the PMS.

100. The single MAC-e PDU transmitted in a TTI could therefore contain data from a number of logical channels. Some of these logical channels might be grouped together within the MAC-e/es as MAC-d flows.
101. 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 (necessary only 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.

6 *MAC-e/es functions in Node B scheduling*

102. The MAC-e/es:
- 102.1 Uses Scheduling Information to indicate the status (i.e. buffer occupancy and channel conditions) of the UE to the UTRAN, receives scheduling grants from UTRAN, and updates its stored Serving Grant variable to reflect the received scheduling grants;
 - 102.2 Performs E-TFC restriction based on the power available for enhanced uplink transmissions. The UE must have sufficient transmit power to send data as well as having a grant allowing it to transmit that data;
 - 102.3 Generates a MAC-e PDU and selects an appropriate E-TFC to transmit the MAC-e PDU to the Node B in accordance with the outcome of the Serving Grant Update, E-TFC restriction and E-TFC Selection procedures;
 - 102.4 Configures the corresponding control information to be sent on the E-DPCCH;
 - 102.5 Handles any HARQ retransmissions.
103. Aspects of the MAC-e/es functionality are covered in the following sections, finishing with a more detailed discussion of the MAC-e and MAC-es header fields. An understanding of the information carried in these header fields, how they relate to the information sent on the E-DPCCH, and when they are

used, is important for understanding the cited prior art, Kim, and consequently, its relevance to the Patent.

7 First Function: resource allocation and Scheduling Information

104. The circumstances in which a UE can send Scheduling Information to the Node B depend on the resources it has been allocated. It is necessary to understand some terminology relating to handover, the allocation of scheduling grants by the Node Bs in the E-DCH Active Set, and the updating of the Serving Grant by the UE, before turning to the handling and contents of Scheduling Information.

8 Terminology

105. As with UMTS, in HSUPA a UE can be in soft or softer handover with multiple cells at the same time. The set of cells with which a UE is in soft (or softer) handover is known as the **E-DCH active set** (which may be the same as the UMTS active set or a subset of it).

106. Of the Node Bs controlling the cells of the E-DCH active set, one Node B has greater control than the others over the UE. This Node B is known as the **Serving Node B**. The cells of the Serving Node B through which the UE is controlled (and it need not be all cells controlled by the Serving Node B) are known as the **E-DCH Serving Radio Link Set** (or **E-DCH Serving RLS**). The cell of the E-DCH Serving RLS through which the Serving Node B exerts greatest control over the UE is known as the **Serving E-DCH Cell**.

107. Each cell in the E-DCH active set not controlled by the Serving Node B is known as a **Non-serving E-DCH RL**, or a **Non-serving RL**. The Node B(s) which control each Non-serving E-DCH RL have a reduced amount of control over the UE.

108. The E-DCH active set for a UE can be shown using Figure 2, reproduced below as Figure 9. The UE is in soft handover, for example, at instance 3 between cells served by three different Node Bs, and in softer handover at instance 2 between two cells served by the same Node B.




Figure 9: The E-DCH active set

109. Table 3 below shows for the five instances in Figure 9 above the E-DCH active set, the Serving E-DCH Cell (and by implication the Serving Node B), the rest of the E-DCH Serving RLS, and any Non-serving RL(s), and, by column, which cells can transmit which scheduling grants (explained in section 9 below). Note that in instance 5, the Serving E-DCH Cell has changed from 2A to 1B, and Node B1 is the Serving Node B.

	Serving Node B	Serving E-DCH Cell	E-DCH Active Set	Serving E-DCH Cell	E-DCH Serving RLS other than the Serving E-DCH Cell	Non-Serving RL(s)
Scheduling Grant				Absolute Grant / Relative Grant (Up/Down/Hold)	Relative Grant (Up/Down/Hold)	Relative Grant (Down/Hold)
Instance						
1	2	2A	2A	2A	-	-
2	2	2A	2A, 2B	2A	2B	-

3	2	2A	2A, 3C, 1B	2A	-	3C, 1B
4	2	2A	2A, 2B, 3C	2A	2B	3C
5	1	1B	1B, 1C, 2A	1B	1C	2A

Table 3: Functions of cells within the E-DCH active set

9 *Resource allocation*

110. As set out in paragraph 92 above, HSUPA introduced two new downlink physical channels on which grants for the transmission of scheduled data are sent:

110.1 The E-AGCH, on which the Absolute Grant (**AG**) is transmitted; and

110.2 The E-RGCH, on which Relative Grants (**RG**) are transmitted.

111. The Absolute Grant and the Relative Grant are the two ways in which a Node B can modify the resources which a UE will use to transmit scheduled data.

112. The UE maintains a variable, the Serving Grant (**SG**) which is used to store the grant that is derived from the received Absolute Grant and Relative Grant messages. The Serving Grant is defined as a ratio of the E-DPDCH power to the baseline power level of the DPCCH, and represents the maximum power ratio allowed to be used for the transmission of scheduled data on the E-DCH.

10 *The Absolute Grant*

113. The Absolute Grant for a UE can only be transmitted by the Serving Node B through the Serving E-DCH cell (e.g. see Fig 9 above, where in instances 1 to 4, only Node B 2 can transmit an Absolute Grant, and then only through Cell 2A). The Absolute Grant provides the Serving Node B with the greatest amount of control over a UE's Serving Grant, including by specifying the maximum power which can be used for the transmission of scheduled data. The Absolute Grant can vary the Serving Grant the most, although as discussed in paragraph 121 below, the Serving Grant is updated using the Serving Grant Update procedure, and certain values of the RG from a Non-serving RL may override certain values of the Absolute Grant from the Serving E-DCH Cell.

114. The Absolute Grant's functionality varies depending on whether a 10ms or a 2ms TTI is configured, but includes:
- 114.1 For both 2ms and 10ms TTI, the ability to specify the maximum power ratio of the E-DPDCH compared to the DPCCH with reference to a lookup table (Table 16B of TS 25.212 V7.1.0.) for all active HARQ processes, by using the Absolute Grant's "Value" field (which may be a "Zero_Grant" value that specifies that there is no resource available to transmit scheduled data); and
- 114.2 For 2ms TTI only, the ability to deactivate the current, or all, HARQ processes.
115. Scheduled data is not allowed to be sent on an inactive HARQ process (whether deactivated or restricted by RRC). However, as a general rule, both non-scheduled transmissions and Scheduling Information can be sent on deactivated or restricted HARQ processes, so if all HARQ processes are deactivated or restricted, a UE can still send non-scheduled transmissions and Scheduling Information.

11 Relative Grants

116. Relative Grants instruct the UE to increase ("UP"), decrease ("DOWN") or hold ("HOLD") the transmit power previously used to transmit the E-DPDCH in a given HARQ process, with reference to a lookup table set out in the PMS - Table 9.2.5.2.1.1: Scheduling Grant Table (SG-table). It shows the 38 values a Serving Grant can take as the result of the UE implementing a Relative Grant command in the Serving Grant Update procedure, and the index for each such Serving Grant value.
117. Relative Grants can be transmitted by:
- 117.1 The Node B which controls cells in the E-DCH Serving RLS, in which case any value can be transmitted; and
- 117.2 Any Node B(s) which control a cell which is a Non-Serving E-DCH RL, in which case they can only signal "HOLD" or "DOWN".
118. The starting point in interpreting a Relative Grant is to look at the E-DPDCH to DPCCH power ratio used for the E-TFC selected in the previous TTI of the

current HARQ process. In particular, the previous value of the Serving Grant is not used as the starting point.

119. This is then adjusted to account for any non-scheduled data and Scheduling Information that was transmitted. The lookup table is used to convert the power ratio into a baseline index value.
120. A “DOWN” Relative Grant command decreases the baseline index value by one step. An “UP” Relative Grant command increases the baseline index value by either three, two, or one steps.
121. If neither an “UP” nor a “DOWN” Relative Grant command is received, the UE interprets the absence (sometimes referred to as discontinuous transmission or DTX) as a “HOLD”. This is perhaps not the most useful name, as a “HOLD” command does not prevent changes (down or up) to the Serving Grant (for example, because the UE might receive a Relative Grant from another Node B).

12 The Serving Grant Update procedure

122. The skilled person would know that the PMS – in section 11.8.1.3.1 - contains detailed “pseudo-code”, i.e. a step-by-step procedure, which specifies the Serving Grant Update procedure. It provides that, in the event that a UE receives more than one scheduling grant in a given TTI, which may conflict with one another, the new Serving Grant is set to the minimum of the Serving Grant resulting from a “DOWN” Relative Grant (i.e. receipt of multiple “DOWN” Relative Grant in one TTI does not result in multiple steps down) and the resulting Serving Grant from the Serving E-DCH cell’s scheduling grant.

13 Scheduling Information

123. Scheduling Information (sometimes abbreviated in the PMS to **SI**) is part of the uplink control information in HSUPA, the remainder being the 10 bits of control information sent on the E-DPCCH – see section 4.6.11). It provides information to the Node B to allow the Node B to make scheduling decisions for scheduled data transmitted from the UE in the uplink. It is not used to regulate non-scheduled transmissions and contains no information about non-scheduled data.

124. Scheduling Information is specified to include 18 bits of information. Scheduling Information is transmitted in a MAC-e PDU, either alone, or multiplexed with other data, in which case it is included at the end of the MAC-e PDU (but before any padding).
125. Although there is considerable detail about SI in the sections which follow, Lenovo were keen to stress one part of section 9.3.1 in TS25.309 from March 2006 (a document presented to Dr Irvine by Lenovo's solicitors) where it states: *'For the UE to request resources from the Node B(s), Scheduling Requests will be transmitted in the uplink in the form of Scheduling Information and Happy Bit.'* As Lenovo submitted, Mr Townsend agreed that the Skilled Person would be thinking in such terms: i.e. that SI was a means for the UE to request resources.

14 *Contents of Scheduling Information*

126. Scheduling Information comprised four fields, each encoded with a specified number of bits. Those fields are:
- 126.1 **Highest priority Logical channel ID (HLID)**, 4 bits, which identifies the highest priority logical channel with data to transmit, or if more than one, the highest priority logical channel with highest buffer occupancy.
- 126.2 **Total E-DCH Buffer Status (TEBS)**, 5 bits, which indicates a range (Table 9.2.5.3.2-1 in the PMS shows the TEBS values against bytes of data available for transmission) within which is included the amount of data in bytes that is available for transmission (including retransmission in the RLC layer; e.g. if RLC requires an acknowledgement for a packet, has not received one, and is storing that packet, it counts towards TEBS).
- 126.3 **Highest priority Logical channel Buffer Status (HLBS)**, 4 bits, which indicates (as a percentage range (Table 9.2.5.3.2-2 in the PMS) the amount of TEBS which relates to the logical channel identified by HLID.
- 126.4 **UE Power Headroom (UPH)**, 5 bits, which indicates the ratio of the maximum UE transmission power to the DPCH power currently being used. It was commonly known that this could be used to

estimate uplink channel conditions for a UE and its ability to use further uplink resources that might be assigned to it.

127. Figure 10 (taken from Figure 9.2.5.3.2-1 of the PMS) below sets out the format of the Scheduling Information message, containing 18 bits.

UPH (5bits)	TEBS (5bits)	HLBS (4bits)	HLID (4bits)
----------------	-----------------	-----------------	-----------------

Figure 10: Format of the Scheduling Information message

15 *Scheduling Information reporting*

128. Scheduling Information was reported (i.e. transmitted from the UE to the Node B) because either:

128.1 A Scheduling Information report had been triggered. Lenovo characterised this as ‘triggered SI’.

128.2 The E-TFC selected by the UE (in accordance with the process described below at section 18) to be used for sending data had sufficient remaining space to include Scheduling Information (i.e. there were at least 18 bits of space remaining). Lenovo characterised this as ‘non-triggered or opportunistic SI’.

129. Triggers for transmitting Scheduling Information were categorised into two main classes, based on the value of the UE’s Serving Grant.

130. In one of those two classes of triggers (set out in section 11.8.1.6.1 of the PMS), Scheduling Information is triggered by the following events where the Serving Grant is equal to Zero_Grant or all HARQ processes are deactivated (i.e. a UE is not allowed to send scheduled data):

130.1 When TEBS becomes larger than zero (i.e. data arrived in a previously empty buffer).

130.2 When data with higher priority than data already in the buffer arrives (i.e. if a change in HLID needs to be reported (see paragraph 126.1)).

- 130.3 Periodically, if the relevant periodic timer (T_SING, which is described in the PMS in section 11.8.1.6.1 as “*Timer Scheduling Information – “Zero_Grant”*”) is configured, and the timer expires.
131. In the second of those two classes of triggers (set out in section 11.8.1.6.2 of the PMS), Scheduling Information is triggered by the following events where the Serving Grant is not equal to Zero_Grant and at least one HARQ process is activated (i.e. a UE is allowed to send scheduled data):
- 131.1 On change of Serving E-DCH Cell, and the new Serving E-DCH Cell was not part of the previous E-DCH Serving RLS (e.g. in instance 5 in Fig 9).
- 131.2 Periodically, if the relevant periodic timer (“*T_SIG*”, which is described in the PMS in section 11.8.1.6.2 as “*Timer Scheduling Information – different from “Zero_Grant”*”) is configured, and the timer expires.
- 131.3 If the Serving Grant becomes Zero_Grant, or all processes are deactivated (i.e. the UE receives (i) an Absolute Grant = “Zero_Grant” or (ii) Absolute Grant = “INACTIVE”, and 2ms TTI was configured, and either only that HARQ process was previously active or the Absolute Grant Value field was set to “ALL” (ignoring the “secondary E-RNTI”, a concept not relevant for present purposes)), and in either case the UE has data in its buffer.
132. There was also a further trigger (set out in section 11.8.1.6.3 of the PMS) at the Priority Date that was not within either of the preceding two categories. In this trigger, if a MAC-e PDU containing triggered Scheduling Information fails to be delivered to the Serving E-DCH RLS (i.e. either a NACK or no ACK is received from the Serving Node B – which Lenovo characterised as ‘delivery failure SI’), then:
- 132.1 If the Scheduling Information is sent with higher layer data multiplexed (i.e. at least one MAC-es PDU) in the same MAC-e PDU, the sending of a new Scheduling Information shall be triggered.
- 132.2 If the Scheduling Information is sent alone, the transmission of Scheduling Information shall not be triggered.

133. Even if multiple events trigger the transmission of Scheduling Information, by the time it can be transmitted (it may be delayed by a HARQ retransmission), only one Scheduling Information shall be sent in the MAC-e PDU.
134. There is a general prohibition on sending Scheduling Information if TEBS is zero (i.e. the buffer in the UE is empty), even if Scheduling Information is otherwise triggered. The only exception to this prohibition is that Scheduling Information shall be included in a MAC-e PDU if the E-TFC which would otherwise be used has sufficient space to include Scheduling Information (as mentioned above at paragraph 128.2).
135. Furthermore, as previously stated, the UE can send Scheduling Information on HARQ processes which are deactivated or restricted by RRC, i.e. even when scheduled data cannot be sent.

16 *Grants for the transmission of triggered Scheduling Information*

136. If Scheduling Information is triggered, section 11.8.1.6 of the PMS requires that it shall be included in the next available transmission, although no more than one Scheduling Information shall be included in any transmission. A HARQ retransmission would take priority, but subject to that, Scheduling Information is allowed to take place on the HARQ process for the upcoming transmission, even if it is deactivated, or restricted by RRC.
137. Therefore, if Scheduling Information is triggered, section 11.8.1.4 of the PMS requires the UE to assume that there is a non-scheduled grant available for transmission of the Scheduling Information, and that the Scheduling Information has a priority higher than any other logical channel.

17 *Second Function: E-TFC restriction*

138. In each TTI, a UE will only select an E-TFC which it has sufficient power to transmit.
139. In a first step, the UE determines power available for enhanced uplink transmission by estimating the remaining power available on the uplink based on the maximum UE uplink transmission power, after the power required for various other channels has been taken into account. The maximum transmission power is a known parameter at the UE. It is limited either by the physical capabilities of the UE or signalled by the network.

140. In a second step, the UE determines which E-TFCs are supported, and which are blocked, by comparing:

140.1 the power available for uplink transmissions; to

140.2 the power required to transmit each E-TFC taking into account the power offset attribute (which specifies whether a particular MAC-d flow requires more than the power per bit to be used based on the relevant reference E-TFC, and if so, how much more) of the highest priority MAC-d flow to be transmitted in the next available TTI.

141. If the power required is greater than the power available, the E-TFC is in a blocked state, subject to any configured “minimum set”, which specifies E-TFCs that are exempt from such blocking and may be transmitted irrespective of the power available: see section 11.8.1.4 of the PMS.

142. If the power required is not greater than the power available, the E-TFC is in supported state.

143. The process of determining which E-TFCs are in which state is known as **E-TFC restriction**. Only E-TFCs in a supported state are considered for transmission each TTI.

18 Third Function: MAC-e PDU generation and E-TFC Selection

144. Based on the outcome of the E-TFC restriction procedure (taking into account power required for uplink transmissions for UMTS and HSDPA, and the E-DPCCH), the Serving Grant Update procedure, and the Scheduling Information reporting procedure, the UE can determine what data to send in the upcoming TTI.

145. The Prior MAC Specification contains two descriptions of MAC-e PDU generation and E-TFC Selection.

146. The first, is in section 11.8.1.4 titled, E-TFC Selection (the **Normative Description**).

147. The second, in Annex C, is in the form of informative pseudo-code for E-TFC Selection (the **Informative Description**).

- *The Normative Description*

148. By the time the UE comes to generate a MAC-e PDU, it has determined:

- 148.1 Which RLC queue has the highest priority data in its buffer that can be sent in a given TTI.
- 148.2 The number of bits of data in that buffer.
- 148.3 Which E-TFC(s) are supported based on the HARQ profile of the MAC-d flow to which the logical channel serving that RLC queue belongs. For each MAC-d flow, a HARQ profile specifies the maximum number of HARQ transmissions for that MAC-d flow and a power offset attribute.
- 148.4 Which other MAC-d flow(s) can be multiplexed into a MAC-e PDU based on the highest priority MAC-d flows multiplexing list. The RRC can configure for each MAC-d flow a multiplexing list. The multiplexing list identifies for each MAC-d flow, the other MAC-d flows from which data can be multiplexed in the same TTI.
- 148.5 How much data there is in the buffer(s) relating to those allowed MAC-d flow(s).
- 148.6 The Serving Grant as a power ratio based on the Serving Grant Update procedure.
- 148.7 The maximum number of bits of scheduled data to which the Serving Grant corresponds.
- 148.8 Which non-scheduled grant(s) (if any) are applicable in the upcoming TTI.
- 148.9 The maximum amount of data which can be transmitted according to each of those non-scheduled grant(s).
- 148.10 How much data there is in each buffer for each of those non-scheduled grant(s).
- 148.11 Whether Scheduling Information has been triggered, and therefore needs to be transmitted, and is therefore assumed to have a non-scheduled grant available for transmission with higher priority than any logical channel.

As an aside, the UE will also have determined (i) whether a retransmission on this HARQ process is required; and (ii) whether the transmission overlaps with a compressed mode gap, but since neither are relevant (for different reasons) they need not be considered further.

149. The Normative Description lists the requirements which must be met when selecting an E-TFC and the data to be carried in the MAC-e PDU. In particular, the Normative Description states that:

149.1 *“The data allocation shall maximize the transmission of higher priority data”.*

149.2 *“The amount of data from MAC-d flows for which non-scheduled grants were configured shall not exceed the value of the non-scheduled grant”.*

149.3 *“The Scheduling Information is always sent when triggered”.*

149.4 *“Only E-TFCs in supported state shall be considered”.*

149.5 *“The E-TFC resulting in the smallest amount of padding for the selected MAC-es PDUs and corresponding MAC-e/es headers, shall be selected including the case when the Scheduling Information needs to be transmitted”.*

150. The Normative Description also:

150.1 states that if the UE is *“not in a power limited condition”*, there should be a quantization (or adjustment) to the amount of scheduled data which can be multiplexed to match the next smaller supported E-TFC.

150.2 Lists (in the subparagraph of section 11.8.1.4 of the PMS which begins with the words *‘When not in a power limited condition...’*) the factors to be taken into account when converting the Serving Grant from a power ratio to a number of bits. Mr Townsend pointed out that the Normative Description does not specify how this should be done, which is true. Dr Irvine responded by pointing out that in section 11.8.1.4 there is a specific reference to section 5.1.2.5B.2.3 of TS 25.214 Physical Layer Procedures (FDD). He said that if the Skilled Person wanted to know how a Serving Grant value should be converted to a

number of bits, they would turn to this reference in TS 25.214, which details how the UE will be sent reference E-TFCs, in the form of pairings of an E-TFC and a Serving Grant level. The UE then knows the number of bits which can be sent when the Serving Grant matches one of the reference E-TFCs and it calculates the number of bits which can be sent for other Serving Grant levels proportionally.

- *The Informative Description*

151. Although the skilled addressee would know that the Informative Description only “describes one possible implementation” of MAC-e PDU creation and E-TFC Selection, they would also understand that it demonstrated a sensible method of complying with the relevant parts of the standard. Mr Townend was of the view, which Dr Irvine did not dispute and with which I agree, that they would use it to aid their understanding of the Normative Description, especially with respect to how the various steps could be put together in a practical manner.

19 *Fourth function: Configuring control information sent on the E-DPCCH*

152. In order to correctly receive and deconstruct the MAC-e PDU, the Node B will need to receive two sets of control information from the UE:

152.1 The control information on the E-DPCCH; and

152.2 The control information in the MAC-e header (see from paragraph 158 below).

153. The E-DPCCH carries ten bits of information in each TTI, made up of:

153.1 the E-TFCI (E-DCH Transport Format Combination Indicator) (7 bits);

153.2 the RSN (Retransmission Sequence Number) (2 bits); and

153.3 the happy bit (1 bit).

154. The E-TFCI indicates the size of the transmitted data (i.e. together the MAC-es PDUs, the MAC-e header, any Scheduling Information, and any padding) by reference to a lookup table of transport block sizes within the E-TFC Set.

155. The RSN specifies whether the E-TFC contains a retransmission, and if so, whether the first, second or subsequent retransmission. Beyond the second

retransmission, the combination of the RSN and the transmission timing allows the receiver to determine the exact transmission number.

156. The happy bit indicates whether or not the UE is “unhappy”. The UE is “unhappy” if it is transmitting the maximum amount of data currently allowed, could transmit at a higher data rate, and will take more than a configurable period of time to empty its buffer based on the current Serving Grant. It is “happy” in all other circumstances.

20 *Fifth function: HARQ retransmissions*

157. If the previous MAC-e PDU sent on a HARQ process requires retransmission then (subject to the maximum number of HARQ transmissions not being already reached) that MAC-e PDU will be retransmitted using the same E-TFC in the next TTI for which the HARQ process is operational, and so a new MAC-e PDU cannot be transmitted in the TTI.

21 *MAC-e and MAC-es header information*

158. In HSUPA, there is no need for a MAC-d header. The MAC-d merely routes data from RLC to the MAC-e/es in the UE in response to requests from the MAC-e/es to the RLC (and routes packets from the MAC-es to the RLC in the SRNC). This is in contrast to UMTS, where the MAC-d’s functions include C/T multiplexing and the addition of relevant headers (see paragraph 58 above). Therefore, in HSUPA, a MAC-d PDU is the same as: (i) an RLC PDU; (ii) a MAC-d SDU; and (iii) a MAC-es SDU.
159. Each MAC-es PDU carries higher layer data in the form of MAC-d PDUs of the same size and from a single logical channel, and has a MAC-es header. One or more MAC-es PDUs can be multiplexed into a MAC-e PDU, which has a MAC-e header portion per MAC-es PDU. This is shown in Figures 9.1.5.2 and 9.1.5.2a of the PMS, which show a MAC-es PDU and a MAC-e PDU respectively. These are reproduced below as Figure 11 and Figure 12 respectively.

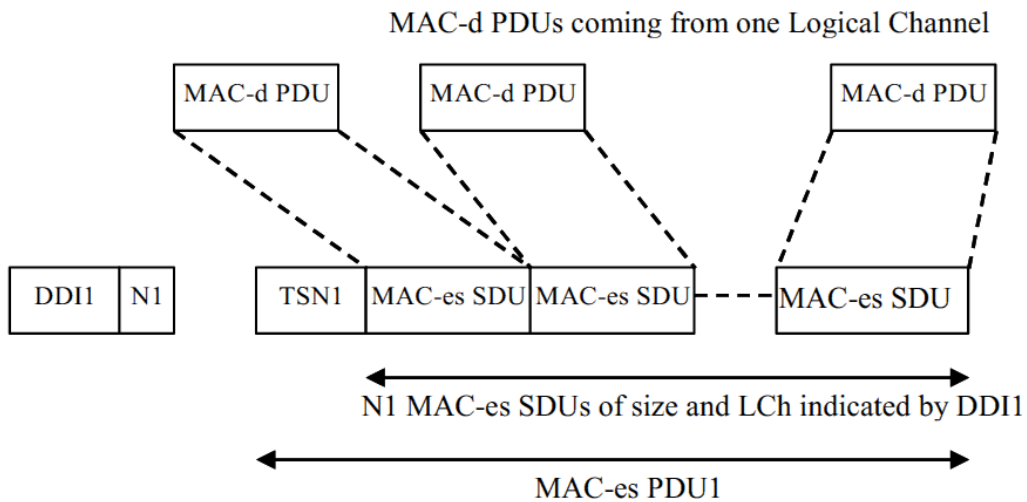


Figure 11: Figure 9.1.5.1 of the PMS (MAC-es PDU)

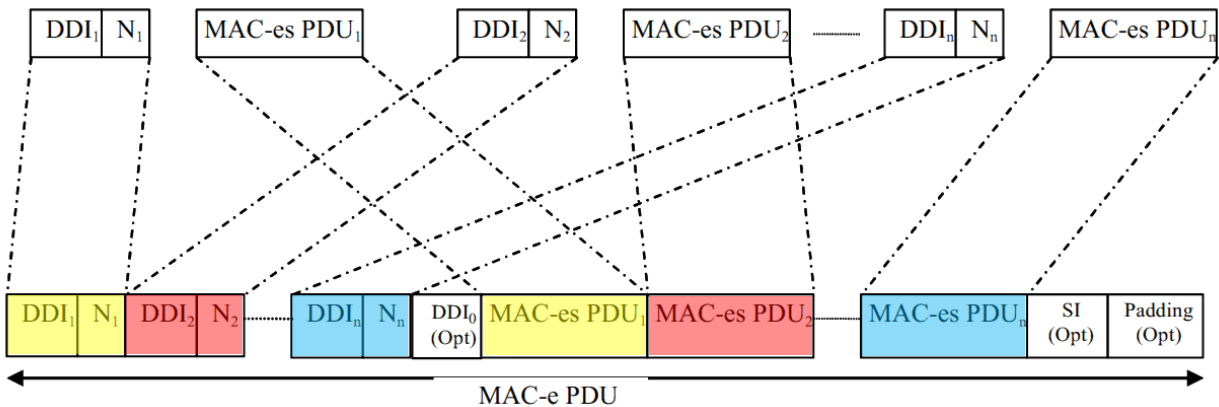


Figure 12: Figure 9.1.5.2a of the PMS (MAC-e PDU)

160. As can be seen from Figure 11, each MAC-es PDU has a single header field, the **TSN**, or Transmission Sequence Number (6 bits) associated with it. The TSN identifies the order of the MAC-es PDU in relation to other MAC-es PDUs containing data from the same logical channel. This allows the MAC-es to reorder MAC-es PDUs so that it can deliver RLC PDUs to the RLC in order.
161. As can be seen in Figure 12 (in which colours have been added for clarity):
- 161.1 for each MAC-es PDU in a MAC-e PDU, there will be two MAC-e header fields associated with it:

- 161.2 a **DDI**, or Data Description Indicator field (6 bits), which specifies the logical channel and MAC-d flow the data in that MAC-es PDU came from, and the RLC PDU size; and
- 161.3 an **N** field, which specifies the number of RLC PDUs in the MAC-es PDU;
- 161.4 if a MAC-e PDU comprises x MAC-es PDU, there will be $18x$ bits of header information (one TSN, one DDI, one N field per MAC-es PDU); and
- 161.5 a MAC-e PDU may include one of three optional elements (indicated as optional by the use of “(Opt)”), namely (from left to right):
 - 161.5.1 DDI_o (also referred to as the “**special DDI**”);
 - 161.5.2 “SI”; and
 - 161.5.3 padding.
- 162. The inclusion of these elements is optional in that they may not be included in each MAC-e PDU. For example, there is no requirement to include Scheduling Information in every MAC-e PDU. It is only included in the circumstances set out in section 15 above. However, it is necessary that the MAC-e PDU sent on the E-DPDCH(s) is the same size as indicated by the E-TFCI on the E-DPCCH. The availability of the DDI_o and padding is therefore necessary to allow transmission and correct reception of a MAC-e PDU, as explained below.
- 163. After including all MAC-es PDUs (and associated headers), how the amount of remaining space in the MAC-e PDU (determined from the E-TFC size) is used is determined by a **three-part rule** which is set out in section 9.2.4.2:
 - 163.1 If it is 24 bits or more, DDI_o (111111) is included, following by Scheduling Information (even if not triggered), and padding (if necessary to fill the PDU);
 - 163.2 If it is between 18 and 24 bits, then Scheduling Information (even if not triggered) and padding are included; and

- 163.3 If it is less than 18 bits, then only padding is included. E-TFC Selection rules ensure that this cannot occur if Scheduling Information has been triggered.
164. These rules taken together with the fixed size header portions (and the content of the header information) allow the Node B to infer where the boundaries between MAC-es PDU are, whether or not Scheduling Information is included, and how many bits of padding are included. In other words, the Node B has sufficient information to split up the MAC-e PDU into its constituent parts.
165. Where Scheduling Information is to be sent alone (i.e. with no MAC-es PDUs), the MAC-e PDU takes a simpler form, as shown in Figure 9.1.5.2b of the PMS and which is reproduced below as Figure 13. In this case, the MAC-e PDU simply comprises the 18 bits of Scheduling Information alone, with no header information or padding.




Figure 13: Figure 9.1.5.2b of the PMS (MAC-e PDU (SI is sent alone))

End of the CGK Annex