CHAPTER 7

Wideband CDMA Air Interface: Protocol Stack

7.1 General Points

The unifying principle in the UTRAN development work has been to keep the mobility management (MM) and connection management (CM) layers independent of the air interface radio technology. This idea has been realized as the access stratum (AS) and nonaccess stratum (NAS) concepts (Figure 7.1). The AS is a functional entity that includes radio access protocols[1] between the UE and the UTRAN. These protocols terminate in the UTRAN. The NAS includes core network (CN) protocols between the UE and the CN itself. These protocols are not terminated in the UTRAN, but in the CN; the UTRAN is transparent to the NAS. The MM and CM protocols are GSM CN protocols; GPRS Mobility Management (GMM) and Session Management (SM) are GPRS CN protocols. Just as the NAS tries to be independent of the underlying radio techniques, so also have the MM, CM, GMM, and SM protocols tried to remain independent of their underlying radio technologies. The apparent dependence of these higher-layer protocols on the radio access protocols will be clarified later in this chapter.

The NAS protocols can be kept the same, at least in theory, regardless of the radio access specification that carries them. Thus, it should be possible to connect any 3G radio access network (RAN) to any 3G CN. This is a nice principle and a worthy goal, but in practice, its implementation is not simple. True independence among the layers in a protocol stack is difficult and expensive to implement. Time and budget constraints usually conspire to allow short cuts to appear in signaling implementations, which causes some interdependence to work itself into the details. The idea of separate access strata is, nevertheless, helpful in understanding the mechanisms and reduces development and testing costs.

One practical result of this concept is that GSM's MM and CM resources are used almost unchanged in 3G NAS. More precisely, the NAS

1 Protocol: "System of rules governing formal occasions" (Oxford English Dictionary). This is in fact a rather good description of a communications protocol. The UE and the network entities must have strict rules in their communication so that both entities know exactly what should be done in each occasion.

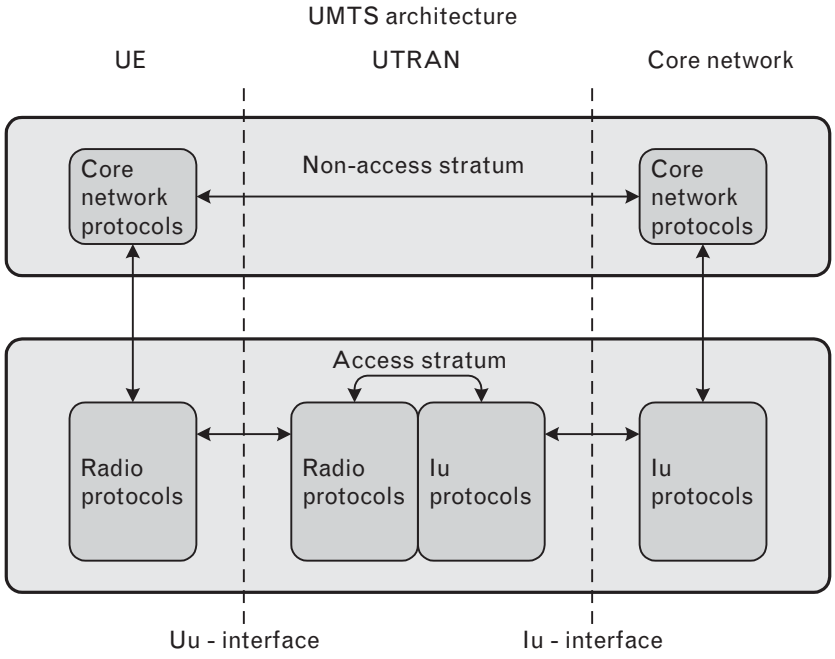

FIGURE 7.1
Stratum model.

layers will be similar to the future GSM MM and CM layers. The reader should understand that some changes have to be made to the legacy GSM CN protocols to meet the future GSM requirements. The upgrades to the current GSM CN will allow support for both the GSM and the UMTS RANs; GSM must be transformed into one of the UMTS radio modes. Present-day GSM operators would not accept any other solution; they want access to 3G. Therefore, it is interesting to notice that future GSM enhancements are being specified in GSM/EDGE Radio Access Network (GERAN) working groups that are part of the 3GPP organization.

Because the CN protocols already exist in GSM and are hardly new developments for 3G as such, they are not discussed thoroughly here. If necessary, they can be easily studied from numerous GSM and GPRS books (e.g., [1–5]). A short overview of each of the tasks is, however, offered here for continuity.

The lower layers (from the AS) are, as the reader can imagine, quite different from GSM. The radio access technology (RAT) used in the UTRAN is CDMA, but in the GSM it is TDMA. From this difference it follows that the protocols used are also very different. The packet-based GPRS protocol stack is also quite different from the UTRAN because of the difference in the RAT, even though they both are packet-based techniques. The widely stated claim about GPRS being a steppingstone to 3G is actually true, but only from the network and marketing points of view. GPRS actually provides a halfway step to a UMTS solution for the network infrastructure because the GPRS CN components can, in many cases, be

reused in a 3G network. On the mobile station side, however, the truth is quite different, as a GSM/GPRS mobile and a WCDMA mobile don't have much in common in their protocol stacks, at least not in the AS parts of them. GPRS is also an important marketing test, allowing the wireless industry to see how subscribers accept new nonvoice content and applications. Think of UMTS as a GPRS network with an advanced and highly adaptive radio interface.

7.2 Control Plane

Radio interface protocols can be divided into two categories: horizontal layers and vertical planes (Figure 7.2).

There are three protocol layers in the AS: physical layer (L1), data-link layer (L2), and network layer (L3). The data-link layer can be further divided into several sublayers: medium access control (MAC), radio link control (RLC), broadcast/multicast control (BMC), and packet data convergence protocol (PDCP). The network layer also includes several sublayers, but among these only the radio resource control (RRC) belongs to the AS. The other sublayers within the network layer are part of the NAS (CN) protocols; these appear above the dotted line in Figure 7.2.

There are also two vertical planes; the control (C) and user (U) planes. The MAC and RLC layers exist in both the C and U-planes. The RRC is

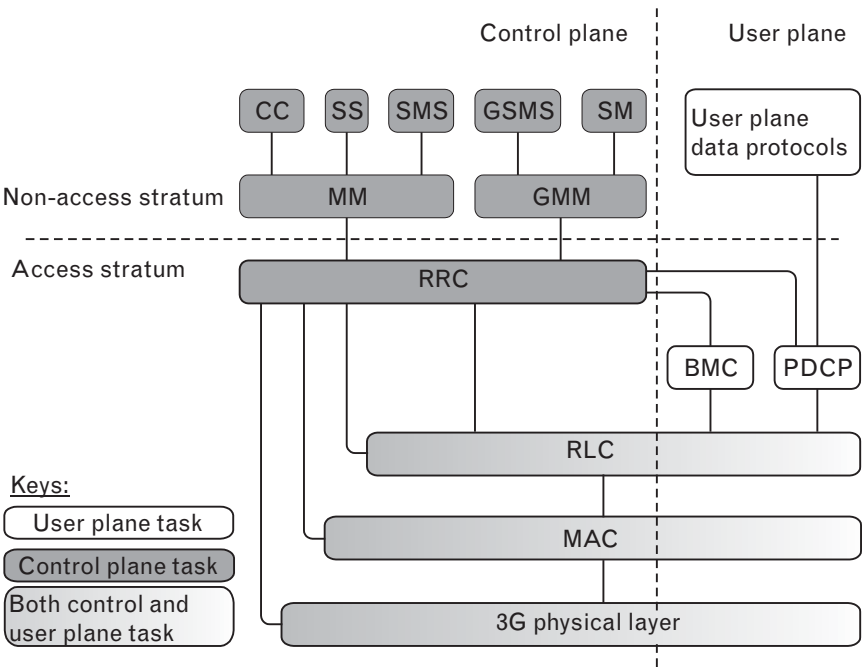

FIGURE 7.2
Protocol tasks in the UTRAN air interface.

found only in the C-plane (i.e., RRC = Radio Resource Control), and the BMC and PDCP are found only in the U-plane. The C-plane carries control data, information that is needed by the protocol tasks to run the system. The U-plane, on the other hand, carries data that is generated by the user, or by a user application. The U-plane data is typically digitally coded voice, but increasingly also other forms of data.

All of these aspects are explained in subsequent chapters: first the C-plane (Figure 7.3) and then the U-plane protocols. The RRC gets the most thorough treatment in this book because it also manages the other protocol layers in the AS. Understanding the RRC is an essential prerequisite to understanding the air interface's inner workings. Each protocol layer performs strictly defined functions, possibly exchanging information with other layers via protocol interfaces.

In the following sections, try to notice the differences between the concepts of function and service. A function is something a protocol does for itself. This may require communication with its peer task and exploitation of the services provided to it by the layers below it. A service is something that is provided to higher protocol layers as a result of the functions of the protocol task itself. It is thus quite possible that the same process can be classified as both a function and a service.

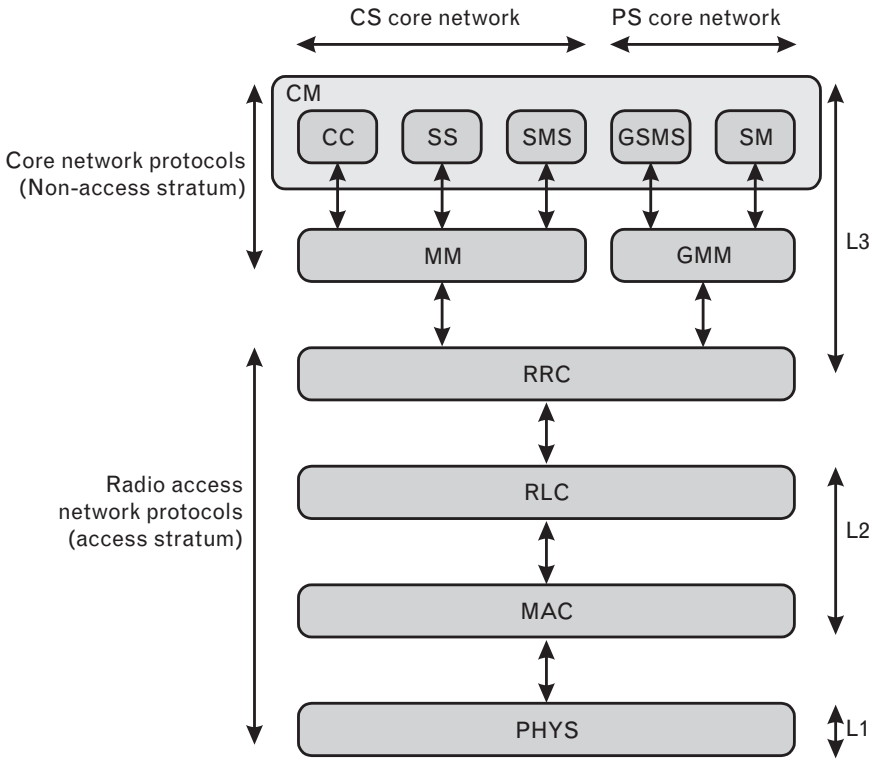

FIGURE 7.3
WCDMA C-plane protocol stack.

7.3 MAC

The UTRAN MAC is not the same protocol as the GPRS MAC, even though they both have similar names and handle similar tasks in similar ways. The UTRAN MAC can even contain different functionalities depending on whether it supports FDD, TDD, or both modes.

The MAC is not a symmetric protocol; the entities in the UE and in the UTRAN are different. A MAC task contains several different functional entities that are depicted in Figure 7.4. Note that this figure depicts the UE MAC. The UTRAN MAC is slightly different from the UE MAC and more complex.

- MAC-b handles the broadcast channel (BCH). The UTRAN has one MAC-b for each cell; the UE may have one or multiple MAC-b's, depending on the implementation. Several MAC-b's may be used for receiving neighbor cell BCHs. This entity is active in the downlink direction only. Note that in the UE this entity will be very simple.
- MAC-c/sh deals with common and shared channels except the HS-DSCH. It handles the paging channel (PCH), the forward access channel (FACH), the random access channel (RACH), and the downlink shared channels (DSCH). The uplink common packet channel (CPCH) in the FDD mode and the uplink shared channel (USCH) in the TDD mode are also handled by this entity. One MAC-c/sh exists in each UE and one exists in the UTRAN for each cell.
- MAC-d handles dedicated logical channels and the dedicated transport channels. The UE has one MAC-d, and the UTRAN has one MAC-d for each UE with assigned DCHs.
- MAC-hs handles the HSDPA functionality. The HS-DSCH is a high-speed downlink shared channel. The UE has one MAC-hs if it is HSDPA-capable; the UTRAN has one MAC-hs for each cell that supports HS-DSCH. Note that a UE does not have to support HS-DSCH and DSCH reception simultaneously. MAC-hs is a bit of a special case among other functional entities because it works with 2-ms subframes, whereas the other entities use 10-ms frames. This tight timing constraint also means that especially the HARQ function control cannot be handled via higher-layer protocols as usual, but must be handled directly from layer 1. In Figure 7.4 this is depicted as the associated downlink signaling. This data flow comes from an HS-SCCH physical channel. Correspondingly, the associated uplink signaling is mapped into an HS-DPCCH physical channel. From a

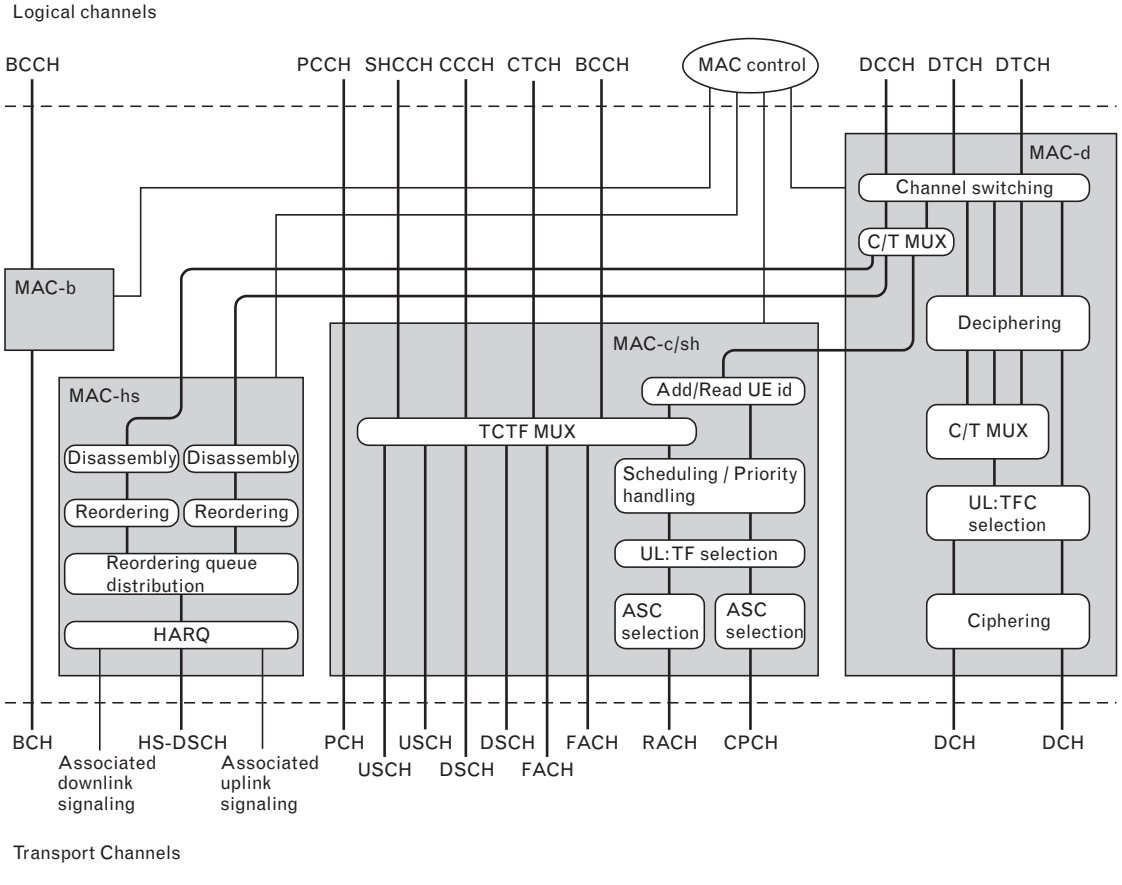

FIGURE 7.4 *MAC protocol layer functional entities.*

protocol stack architecture perspective, it would have been clearer to name these data flows as new transport channels.

The MAC operates on transport channels (see Section 3.2.2) between the MAC and layer 1. The logical channels are described between the MAC and RLC in Section 3.2.1. The internal configuration of the MAC is controlled by the RRC. The MAC is the lowest sublayer in layer 2; it has a thorough understanding of how to manipulate the physical layer on behalf of the layers above it.

7.3.1 MAC Services

The services MAC provides to the upper layers include the following:

- Data transfer;
- Reallocation of radio resources and MAC parameters;
- Reporting of measurements to RRC.

7.3.2 MAC Functions

MAC functions include the following:

- Mapping between logical channels and transport channels;
- Selection of the appropriate transport format for each transport channel depending on the instantaneous source rate;
- Priority handling between data flows of one UE;
- Priority handling between UEs by means of dynamic scheduling;
- Identification of UEs on common transport channels;
- Multiplexing/demultiplexing of higher-layer PDUs into/from transport blocks delivered to/from the physical layer on common transport channels;
- Multiplexing/demultiplexing of higher-layer PDUs into/from transport block sets delivered to/from the physical layer on dedicated transport channels;
- Traffic-volume monitoring;
- Transport-channel type switching;
- Ciphering for transparent RLC;
- Access service class selection for RACH and CPCH transmission;

- HARQ functionality for HS-DSCH transmission;
- In-sequence delivery and assembly/disassemby of higher layer PDUs on HS-DSCH.

Some of these functions are further explained in the following sections.

7.3.2.1 Priority Handling Between Data Flows of One UE

The priority of a data flow is used when the MAC layer chooses suitable transport format combinations (TFCs) for uplink data flows. Higher-priority data flows can be given higher bit rate combinations, and low-priority flows may have to use low bit rate combinations. A low bit rate can also mean a zero bit rate.

Note that there is not a single priority parameter attached to a data flow, but MAC has to derive it from at least two sources: the buffer occupancy parameter received from the RLC and the MAC logical channel priority received from the RRC.

At radio bearer setup/reconfiguration time, each logical channel involved is assigned a MAC logical channel priority (MLP) in the range 1, …, 8 by the RRC. The details of the TFC selection algorithm are not defined in the MAC specification. Rather, the specification gives a list of constraints the algorithm implementation has to fulfill.

7.3.2.2 Identification of UEs on Common Transport Channels

If a UE is addressed on a common downlink channel or it uses the RACH, the UE is identified by the MAC layer. There is a UE identification field in the MAC PDU header for this purpose. If the message was addressed to this UE, it is routed further to the RLC, and from there, either to the RRC, the BMC, or the PDCP. Other messages are trashed.

7.3.2.3 Traffic-Volume Monitoring

The UTRAN-RRC layer performs dynamic radio access bearer (RAB) control. Think of the RRC as a kind of mediator between the network and the radio interface. The MAC layer is obliged to eventually react in some appropriate way to the RRC's needs. Based on the required traffic volume, the RRC can decrease or increase the allocated capacity. The task of monitoring the traffic volume is allotted to the MAC. The UE-MAC layer monitors the uplink transmit buffer, and the UTRAN-MAC layer does the same for the downlink buffer. If the queue in either entity goes out of range, the corresponding RRC is notified. The UE-RRC must further notify the UTRAN-RRC. It is the UTRAN-RRC that has to make decisions about

radio resource allocations because only the UTRAN-RRC knows the total load situation of the whole system.

The monitoring of the traffic volume is controlled by the RRC. It may command the MAC to perform either periodic or event-triggered monitoring. In the case of periodic monitoring, the MAC sends a new report periodically after a timer has expired. In the case of event-triggered monitoring, the RRC gives a range of allowed buffer values, and once the transmission queue goes out of range, an alarm indication is sent back to the RRC (see Figure 7.5).

The purpose of the traffic-volume monitoring procedure is to allow for efficient radio resource usage. If the allocated resources are not sufficient for the generated traffic, the UTRAN may reconfigure itself and add resources. This may mean allocating a DCH instead of a shared channel or simply reducing the SF on a particular channel. Similarly, if the traffic-volume monitoring shows that the allocated resources are underutilized,

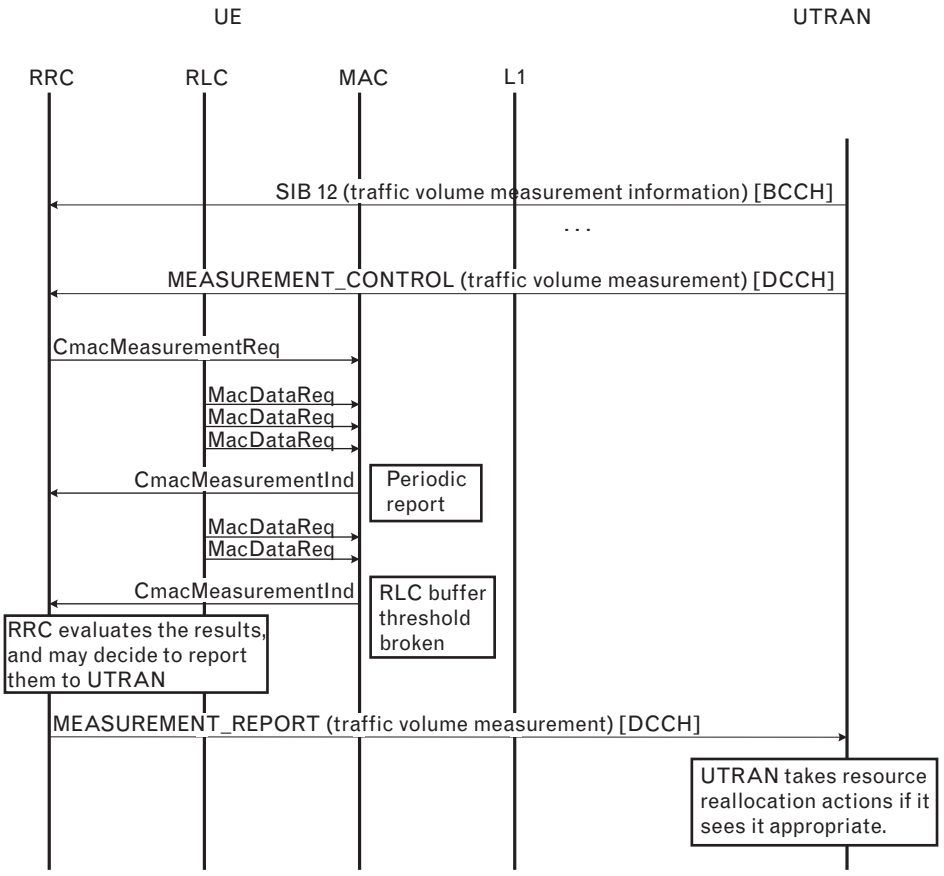

FIGURE 7.5 *Traffic-volume monitoring.*

the UTRAN may reconfigure the connection from a dedicated resource to a shared resource, or increase the current spreading factor.

Note that the transmission buffer to be monitored is actually in the RLC layer, but the buffer occupancy information is relayed down to the MAC layer with each MacDataReq signal. See Section 7 of [6], for a description of the traffic-volume-monitoring interlayer procedure.

7.3.2.4 Transport-Channel Type Switching

The MAC executes the switching between common and dedicated transport channels based on a switching decision made by the RRC. In the UE the dynamic-transport-channel-type switching function maps and multiplexes the DCHs (DCCH and DTCH) into logical channels. Note that in 3GPP jargon, the function *mapping between logical channels and transport channels* refers to a functionality in the MAC-c/sh, which has a more static nature, and *transport-channel type switching* refers to the more dynamic functionality in MAC-d.

7.3.2.5 Ciphering for Transparent RLC

If the RLC layer is in transparent mode (i.e., it is just a "pipe" between the PDCP and the MAC), then ciphering must be done in the MAC layer. Otherwise, it will be performed in the RLC layer. Ciphering prevents the unauthorized interception of data. The ciphering algorithm to be used is the same in the MAC and in the RLC; that is, the UE does not have to use more than one ciphering algorithm at a time. However, the ciphering algorithm may be changed according to commands from the UTRAN. See Section 7.5.2.16 for a further description of ciphering.

7.3.2.6 Access Service Class Selection for RACH Transmission

The MAC gets a set of access service classes (ASCs) from the RRC, and it chooses one of them to be used for the RACH transmission. These classes define the parameters used in a RACH procedure, including access slots and preamble signatures. The algorithm itself uses two variables: MAC logical channel priorities (MLP) and the maximum number of ASCs (NumASC). The MinMLP parameter is set as the highest logical channel priority assigned to the logical channel in question (note that a smaller MLP number means a higher priority; the range is from 1 to 8). The ASC number is obtained as follows:

If all the transport blocks in a transport block set have the same MLP, then select: ASC = min(NumASC, MLP).

If the transport blocks in a transport block set have different MLPs, then select: ASC = min(NumASC, MinMLP).

The ASC enumeration is such that it corresponds to the order of priority (ASC 0 = the highest priority; ASC 7 = the lowest priority). ASC 0 would only be used for very important reasons, such as emergency calls.

7.3.2.7 Hybrid ARQ Functionality for HS-DSCH Transmission

Hybrid ARQ (HARQ) is an acknowledged retransmission scheme that is employed on the HS-DSCH channel. In the UE this functionality is relatively simple as the HARQ entity only has to check the correctness of the received packet and send either a positive or a negative acknowledgment back to the peer HARQ entity in UTRAN. However, the UTRAN HARQ has more complex duties. It has to take care of at least the following tasks:

- *Scheduling of data.* There are probably several active HSDPA UEs in the cell, and their data transmission processes may have different priorities. The scheduler has to select which data is sent first. Note that retransmitted data is probably of higher priority than data that is transmitted for the first time.

- *Buffering of data.* Because HS-DSCH employs acknowledged data transmission, the transmitting entity cannot discard the data as soon as it has been transmitted, but it must be buffered until a positive acknowledgment has been received

- *Retransmission functionality*. If a negative (or no) acknowledgment is received for a packet, the packet must be retransmitted. Because HARQ employs link adaptation, the retransmission may use a different modulation scheme, and a different redundancy version. This is to increase the probability of a successful transmission. Note that the UTRAN cannot select these quantities at will, but the result must comply with the allowed transport format combinations. In the case of HS-DSCH, the transport format selection is different from other channels because here the dynamic part of a transport format includes also the modulation scheme and the redundancy version.

Note that a failed packet will not be retransmitted forever in MAC-HARQ. If the data is important enough, it will also be protected by a higher layer (RLC) retransmission protocol, which takes care of the problem if it does not receive a positive acknowledgment in time.

7.3.2.8 In-Sequence Delivery and Assembly/Disassembly of Higher-Layer PDUs on HS-DSCH

Because of the HARQ retransmission protocol, it is possible that the UE receives the data packets via HS-DSCH in an order other than that in which they were originally transmitted. Thus, there has to be a reordering buffer in the UE. Assembly and disassembly functions are needed because the data packets in the HS-DSCH are probably a different size than in the RLC layer buffers. HS-DSCH is optimized for very high-speed data transfer; thus, the packets on this channel are typically very large.

7.3.3 TFC Selection

TFC selection is in fact a process that combines several functions from the earlier list. First some definitions:

- *Transport format* (TF) defines what kind of data and how much is sent on each transport channel in each transport time interval (TTI). TTI length is equal to the duration of a radio frame or a multiple of it.
- *Transport format combination* (TFC) is a set of TFs that are sent simultaneously (within the same TTI) on different active transport channels to or from the same UE. Indirectly, TFC gives the data rate used.
- *Transport format combination set* (TFCS) is a set of TFCs. The UE has to select one TFC from a set of allowed TFCs for data transmission in each TTI. The TFCS to be used is signaled to the UE via RRC signaling, but this set can be limited later by several different network procedures. As a result, only some TFCs from the original set are allowed TFCs at a given time.

The MAC layer has to choose a set of TFs, so that given the current channel conditions, the maximum amount of highest-priority data could be transmitted over the air interface. This is not a simple task. The MAC layer itself knows from the configuration data which transport formats and which combinations of transport formats are valid. However, not all such combinations are usable all the time. The current channel conditions could impose limitations on what TFCs can be used. Those combinations that could carry the highest amount of data also need the highest transmit power in the physical layer. This could be more than the maximum allowed transmit (TX) power the transmitter can use. In a CDMA system, more data basically means more power. And the more noise there is in the radio interface, the higher the transmitting power must be. Thus, it is quite conceivable that especially in a noisy environment, only some of the TFCs can be used. The UTRAN can signal a temporary TFC limitation to a UE via RRC layer

signaling. But still, this does not remove the need to monitor the required TX power and limit the TFCs further if necessary.

On the other hand, the data to be transmitted is in the data buffers in the RLC layer. MAC has to get the buffer occupancy information from RLC, as well as the priority of the data in those buffers. It has to send as much data as possible, and at as high a priority as possible. And obviously, the MAC layer cannot send more data than there is in RLC buffers. Note that the data in RLC buffers is not a stream of bits, but a group of PDUs, and those cannot be divided or combined at will. Moreover, the MAC layer is not allowed to choose TFCs that require the RLC layer to add padding bits to its PDUs to make them match with the chosen TFC (i.e., to choose too-large TFCs). Further complexity is caused by the compressed mode as this will cause the transmitter to either send less data or use more power, in which case there is again the danger of exceeding the maximum allowed TX power. Furthermore, many applications employ variable rate adaptive codecs, and the MAC layer has to cooperate with them so that the produced bit rate exactly matches with some allowed TFC. And to make all this a bit more challenging, the reader must remember that the TFC selection must be made once every 10 ms, that is, the length of the radio frame. In fact, the selection frequency is equal to the length of the shortest configured TTI duration, so it could be 10, 20, 40, or 80 ms. But still, the algorithm must be based on the worst case, that is, 10 ms.

The TFC selection algorithm is not, and will not be, specified by the 3GPP. Only the TFC selection criteria are given in the MAC specification [7], and it can be implemented in a more or less efficient way. TF selection must be done on all DCHs, and also on RACH and CPCH channels.

The MAC protocol is specified in [7]. TFs were discussed in this book in Section 3.5.

7.4 RLC

One RLC task contains several different functional entities. For bearers using the transparent mode service or the unacknowledged mode (UM) service, there is one transmitting and one receiving entity for each bearer. For bearers using the acknowledged mode (AM) service, there is only one combined transmitting and receiving entity for each bearer. Different modes are used for different types of data. If the data is of an important nature, it needs lots of protection and AM service. On the other hand, some data is not suitable for AM service. For example, it is no good to use AM for voice. The AM retransmission protocol could guarantee that a voice packet does get through eventually, but a retransmitted voice packet cannot be used anymore because of the additional delay. A voice packet must be received in time without delays or it is worthless.

In general, the RLC layer is in charge of the actual data packet (containing either control or user data) transmission over the air interface. It makes sure that the data to be sent over the radio interface is packed into suitably sized packets. The RLC task maintains a retransmission buffer, performs ciphering, and routes the incoming data packets to the right destination task (RRC, BMC, PDCP, or voice codec).

The transparent mode is used for the BCCH, PCCH, SHCCH, DCCH, DTCH, and CCCH channels. For the CCCH and SHCCH, the transparent mode is used only in the uplink direction. Transparent mode means that very little processing is done to the data in the RLC. It contains transmission and receiver buffers and also, in some cases, segmentation and reassembly functions. Note that no RLC header is added to data units in the transparent mode (see Figure 7.6).

Despite this rather limited functionality, one instance of an RLC transparent entity is needed per direction and per bearer—one for the uplink and one for the downlink.

UM is used for the DCCH, DTCH, CTCH, and the downlink SHCCH and CCCH channels. The RLC adds a header to the PDU and ciphers/deciphers it. As in transparent mode, one instance is needed per direction and per bearer (see Figure 7.7).

AM can be used for the DCCH and DTCH channels. The SDUs are segmented or concatenated onto the PDUs of fixed length.

The multiplexer (MUX) chooses the PDUs and decides when they are delivered to the MAC. The MUX may, for example, send RLC control

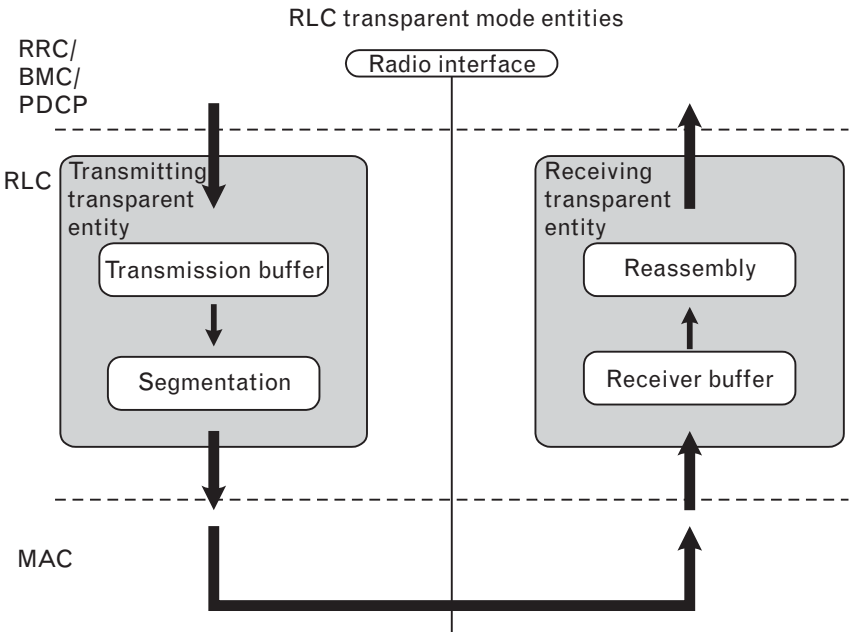

FIGURE 7.6
Transparent entities in RLC.

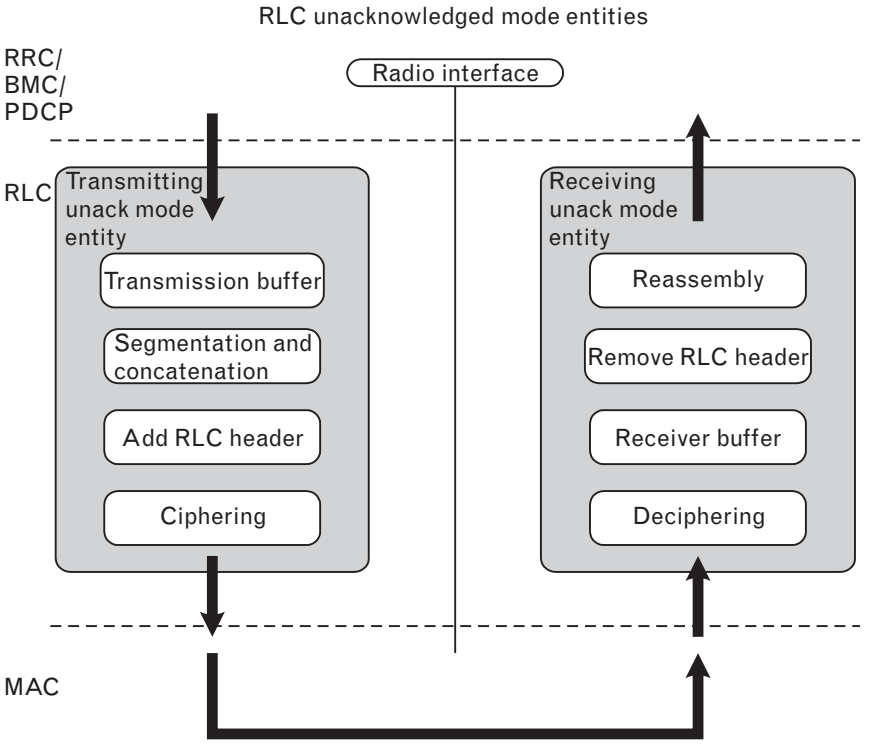

FIGURE 7.7
UM entities in RLC.

PDUs on one logical channel and data PDUs on another logical channel, or it may send everything via one logical channel.

If the data in AM mode does not fill the whole PDU, then padding is used to fill the rest of the PDU. This padding can be replaced with piggy-backed control information in order to increase the transmission efficiency.

There is only one AM entity per bearer in the UE that is common to both the uplink and the downlink (see Figure 7.8).

7.4.1 RLC Services

The following are services provided to upper layers:

Transparent Data Transfer Service

- Segmentation and reassembly;
- Transfer of user data;
- SDU discard.

Unacknowledged Data Transfer Service

- Segmentation and reassembly;

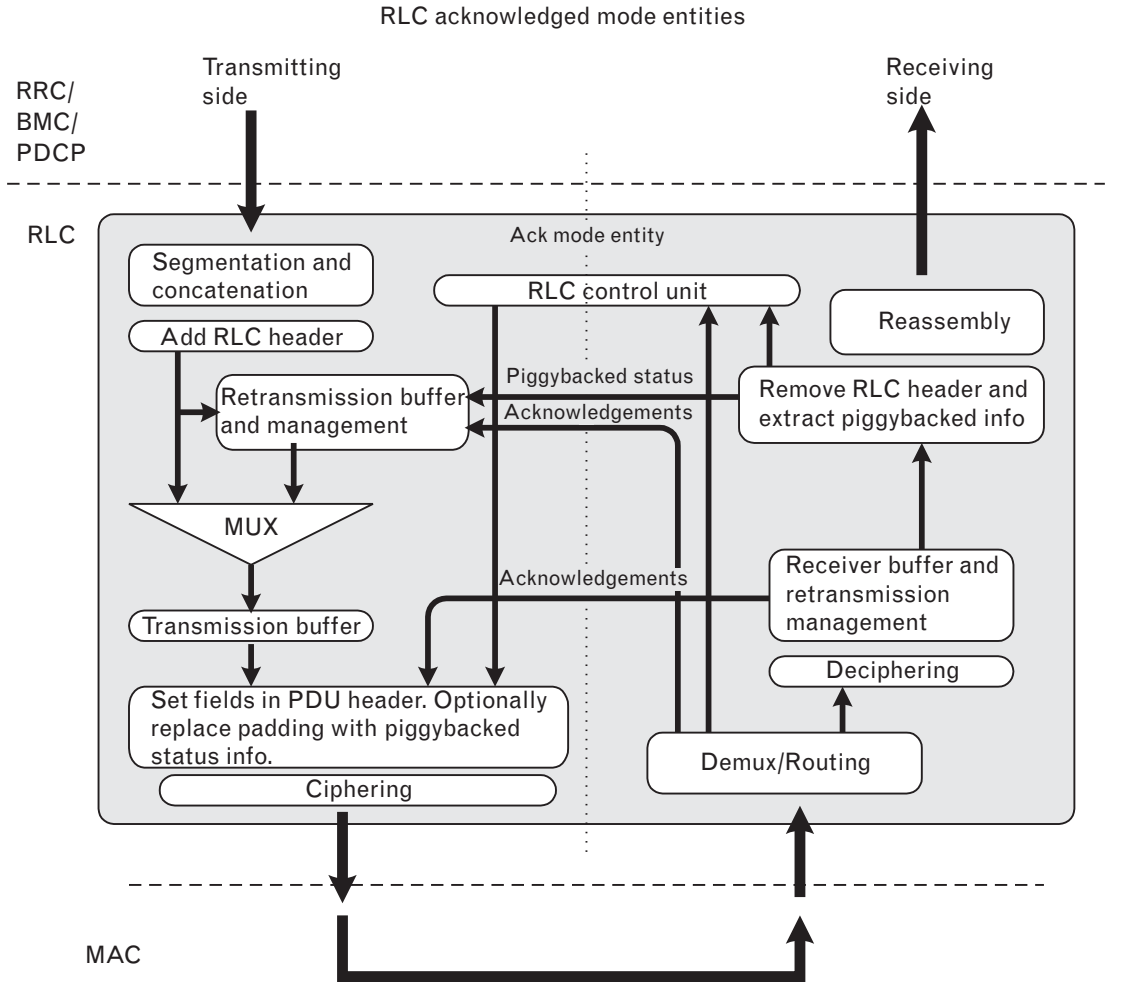

FIGURE 7.8 *AM entity in RLC.*

- Concatenation;
- Padding;
- Transfer of user data;
- Ciphering;
- Sequence number check;
- SDU discard.

Acknowledged Data Transfer Service

- Segmentation and reassembly;

- Concatenation;
- Padding;
- Transfer of user data;
- Error correction;
- In-sequence delivery of higher-layer PDUs;
- Duplicate detection;
- Flow control;
- Protocol error detection and recovery;
- Ciphering;
- SDU discard.

Maintenance of Quality of Service (QoS) as Defined by Upper Layers

Notification of Unrecoverable Errors

7.4.2 RLC Functions

The following functions are supported by the RLC:

- Segmentation and reassembly of higher-layer PDUs into/from smaller RLC payload units;
- Concatenation (RLC SDUs may be concatenated so that they will fill the RLC PUs);
- Padding;
- Transfer of user data;
- Error correction;
- In-sequence delivery of higher-layer PDUs;
- Duplicate detection;
- Flow control;
- Sequence number check (in unacknowledged data transfer mode);
- Protocol error-detection and recovery;
- Ciphering (in UM and AM modes);
- Suspend/resume function.

The RLC protocol is defined in [8].

7.5 RRC

We turn now to the most important subject of this chapter: the RRC. The RRC controls the configuration of the lower layers in the protocol stack, and it has control interfaces to each of the lower layers (PDCP, BMC, RLC, MAC, and layer 1). It is the conductor of the protocol stack orchestra.

7.5.1 RRC Services

The RRC provides the following services to the upper layers:

- *General control*. This is an information broadcast service. The information transferred is unacknowledged, and it is broadcast to all mobiles within a certain area.
- *Notification*. This includes paging and notification broadcast services. The paging service broadcasts paging information in a certain geographical area, but it is addressed to a specific UE or UEs. The notification broadcast service is defined to provide information broadcast to all UEs in a cell or cells. Note that the notification broadcast service seems to be quite similar to the general control service.
- *Dedicated control*. This service includes the establishment and release of a connection and the transfer of messages using this connection. These connections can be both point-to-point and group connections. Message transfers are acknowledged.

7.5.2 RRC Functions

The RRC functions include the following:

- Initial cell selection and cell reselection (includes preparatory measurements);
- Broadcast of information (system information blocks [SIBs]);
- Reception of paging messages;
- Establishment, maintenance, and release of RRC connection;
- Establishment, reconfiguration, and release of radio bearers;
- Assignment, reconfiguration, and release of radio resources for the RRC connection, which includes such things as the assignment of codes and CPCH channels;
- Handovers (HOs), which include the preparation and execution of HOs and intersystem HOs;

- Measurement control;
- Outer-loop power control;
- Security mode control (ciphering control, integrity protection, counter check);
- Routing of higher-layer PDUs (direct transfer);
- Control of requested QoS;
- Support for DRAC (fast allocation of radio resources on the uplink DCH);
- Contention resolution in the TDD mode;
- Timing advance in the TDD mode;
- Management of the CBS service (the service itself is implemented in BMC);

These functions are explained in the following sections. Some are quite similar to the corresponding GSM procedures, but there are also some new functions. When we compare the RRC to the GSM-RR functions, the most important changes include SHOs, intersystem HOs, and more flexible channel configuration management, which yields a more efficient usage of the available resources.

Some of these functions are depicted in the signaling flow diagrams in Chapter 11. The RRC protocol is specified in [9].

7.5.2.1 Initial Cell Selection

The initial cell selection, as well as other cell-evaluation procedures, is quite different from the GSM cell-selection procedures. Cell-selection procedures are discussed in [10], which offers a rather cryptic presentation. The purpose of the initial cell-selection procedure is to find a cell, not necessarily the best cell, but a usable cell, for the UE to camp on after power-on.

In the UTRAN, the number of carrier frequencies is quite small. One operator typically operates only on two or three frequency carriers. In the first phase of UMTS in Europe, the frequency allocation for UMTS-FDD is 2×60 MHz (uplink/downlink), which means that there can be, at most, only 12 carrier frequencies of 5-MHz bandwidth each. These carriers are then divided between up to six operators. Each carrier will only support one operator. This obviously forces the operators to coordinate their network-planning activities near national borders because the same frequency can be used by different operators in adjacent countries.

The specifications do not accurately dictate how the initial cell-selection procedure should be implemented; it is left for the UE

manufacturers to decide. Most of the functionality, however, has to be in the physical layer, and the RRC layer has only a management role. The initial cell-selection procedure is performed on one carrier frequency at a time until a suitable cell is found. In principle the process includes the following:

1. Search for primary synchronization channels (P-SCHs);
2. Once such a channel is found, acquire time-slot synchronization from it;
3. Acquire frame synchronization from the corresponding S-SCH;
4. Acquire the primary scrambling code from the corresponding CPICH;
5. Decode system information from the cell to check whether it is a suitable cell for camping (i.e., it contains the right PLMN code and access to it is allowed).

The synchronization process in the physical layer is explained in Section 3.1.9. Here the issue is considered from the whole UE point of view.

All P-SCHs have the same fixed primary synchronization code. The search procedure should yield a set of P-SCHs in the area. Because the P-SCH is only transmitted during the first 256 chips of each time slot, the beginning of its transmission also indicates the start of a time slot in the corresponding cell.

In the second phase of the process, the received signal is correlated with all possible secondary synchronization code (S-SCH) words on the S-SCH. There are 16 different SSCs, and these can be combined into 64 different code words, each with a length of 15 SSCs. Once the right code word is found, this gives the UE the frame synchronization and the code group identity, which indicates eight possible primary scrambling codes for the control channels.

The third phase of the procedure consists of finding the right primary scrambling code for this cell. Each candidate cell's primary scrambling code (there are eight of them as shown in the second phase) is applied, in turn, to the common pilot channel (CPICH) of that cell. Because the CPICH carries a predefined bit/symbol sequence, the UE knows when it has found the correct primary scrambling code. The resolved primary scrambling code can then be used to detect the CCPCH, which carries the BCH, which contains the system information the UE is seeking. There are various ways to optimize this procedure to make it quicker.

Note that phase five actually contains another major procedure, PLMN (i.e., the operator) selection. PLMN is identified by a PLMN code, a number that is transmitted on the BCCH channel of that network. A UE tries to find its home PLMN, the operator it has a contract with. In principle, a UE should first scan through all UTRAN frequencies until a good

PLMN is found, and then start an initial cell-selection process on that frequency. Note that one frequency can only be used by one operator (except in areas near country borders). However, while looking for the right PLMN code, the UE has already obtained all the necessary information for camping on a suitable cell, and no new scanning procedure is necessary once the correct PLMN is found. The situation is different if the UE is roaming abroad, and the home PLMN is not found. In that case RRC has to report all available PLMNs to NAS and wait for its selection decision, which can be either automatic or manual (user selection). This is time consuming, and many readers may have noticed this phenomenon when arriving at an airport in a new country and switching their GSM phones on. It may take a very long time before the phone registers to a network, especially if the phone is a multimode model with several frequency bands to scan. PLMN selection process is probably best described in [11].

The initial cell-selection process is repeated as many times as necessary until the first suitable cell is found for camping. Once the UE has managed to camp on a cell, it decodes the system information from it, including the neighbor cell list. This information can be used to help the UE find the best cell to camp onto. Note that the initial cell-selection procedure only found a cell to camp on (the first possible cell). It is possible that this cell will not be the best possible cell. For example, there could have been other frequencies including better cells for this particular UE that had not yet been scanned.

The neighbor cell list immediately tells the UE which frequencies and neighbor cells should be checked while the best possible cell is being searched for. The list includes additional information that can be used to optimize the cell-synchronization procedure, information such as the primary scrambling codes and timing information (optional, relative to the serving cell). With this information it should be possible to quickly descramble the CPICH from a neighbor cell.

From the CPICH it is possible to calculate the received chip energy-to-noise ratio (Rx Ec/No) for this cell. This measurement is acquired for each neighbor cell in the list. Based on this information, the UE can determine whether there are better cells available. From a possible candidate cell, the UE must decode the system information to check that it is not barred for access.

If the neighbor cell list contains cells from another RAT—for example, GSM cells—and the serving cell quality level is worse than the Ssearch parameter, then the GSM cells must be taken into consideration in the cell-reselection procedure.

The initial cell-selection procedure described here is to be used in case there is no information on the current environment stored in the UE. However, normally the UE starts the cell selection with a stored information cell-selection procedure. The UE may have stored the necessary

information of the cell it was previously camped on, such as frequency and scrambling code. The UE may first try to synchronize into that cell, and if it fails, it may trigger the initial cell selection.

7.5.2.2 Cell Reselection

The cell-reselection procedure, or as the 3GPP calls it, the cell reselection evaluation process, is performed in idle mode to keep the UE camped on a best cell. If the UE moves or the network conditions change, it may be necessary for the UE to change the cell it is camped on. This procedure checks that the UE is still camped on the best cell, or at least on a cell that is good enough for the UE's needs.

In normal idle mode, the UE has to monitor paging information and system information and perform cell measurements. The cell-reselection procedure will be triggered if the measurements indicate that a better cell has been found, or if the system information of the current cell indicates that new cell access restrictions are applied to the cell in question, such as cell barred. [2]

System information block 3 (SIB3) is an important message here because it tells the UE the quality parameter to measure, and also all the parameters for the cell-reselection evaluation algorithm.

The neighbor cells to be measured are given in the neighbor cell list (SIB11). The results of these measurements are evaluated periodically.

System information (SIB3) may also contain various optional threshold parameters that define when to perform various measurements (S_x is the measured quality parameter of the serving cell):

- If $S_{\text{intrasearch}}$ is given and $S_x \leq S_{\text{intrasearch}}$, then the UE must perform intrafrequency measurements.
- If $S_{\text{intersearch}}$ is given and $S_x \leq S_{\text{intersearch}}$, then the UE must perform interfrequency measurements.
- If $S_{\text{searchRAT}\ n}$ is given and $S_x \leq S_{\text{searchRAT}\ n}$, then the UE must perform inter-RAT measurements

If a threshold parameter is not given, then the UE must always perform the corresponding measurements.

Based on these measurements the UE periodically evaluates the best-cell status. If it seems that there is a better cell available, it will trigger a cell-reselection procedure.

2 An operator can bar a cell, for example, during maintenance or testing. Operators may have special test mobiles that can ignore the cell bar-flag and these can be used to test the base station functionality while access is refused for other mobiles.

7.5.2.3 Cell Reselection and Random Access

There is a potential problem with cell-reselection and network-access attempts. In the idle mode, the UE cannot monitor its environment continuously. That would quickly wear the battery down. Instead, the monitoring process is periodic, and the periodicity is set by the network. The longer the period, the bigger is the danger that the UE may not be camped on the best possible cell. Normally, this would not be a problem as the timers used here are pretty short and the next measurement process will fix the problem by triggering a reselection.

However, problems may be caused if the UE decides to launch an access attempt while it is not attached to the best cell. An access attempt includes the transmission of RACH bursts to a serving Node B. A WCDMA system is sensitive to interference, and sending RACH bursts to the "wrong" Node B could introduce severe and unnecessary interference to any Node B close to the UE. The UE ramps up the RACH transmission power until it receives a response from the base station, and if this RACH burst is addressed to a Node B far away from the UE, the received signal power in the nearby Node Bs will be unacceptably high. Notice that the RACH signal is interference for all other Node Bs except for the one to which it is addressed. A possible scenario could involve a user who walks around a corner and immediately starts a call. The mobile may still be camped on the old cell, although there might be a much better one available on the new street. If the RACHing process is done on the old cell, the new cell could be easily blocked because of it.

The very first 3GPP specifications tried to solve this problem with a procedure called the *immediate cell evaluation*. It was a procedure that was to be performed just prior to a random-access procedure. The purpose for this procedure was to make sure that the UE is camped on the best possible cell before it starts to send RACH bursts to the network. The problem with the immediate cell evaluation procedure was the delay it caused. The sending of RACH bursts should be started quickly, especially if the call setup procedure was triggered by a paging message. Immediate cell evaluation took some time, especially if as a result of the evaluation a cell reselection was required. The UE would then have to decode at least part of the SIBs in the new cell before it could start to send RACH bursts there because it needs to know the random-access parameters to be applied in the new cell.

Because of these problems, immediate cell evaluation was removed from the specifications. Now, if the random-access procedure fails, the UE is required to trigger a cell-reselection procedure immediately. However, this solution does not remove the problem described earlier. We can only hope that the presented scenario is rare enough.

7.5.2.4 Broadcast of System Information

Broadcast information (or system information, as it is also known) is information about the system and the serving cell that is sent by the network in a point-to-multipoint manner; the information is broadcast to all UEs. It is typically information that is common to all mobiles in a cell; thus, it can be sent using a broadcast service.

Broadcast information consists of messages called system information blocks (SIBs). A SIB contains system information elements of the same nature. There are 18 different blocks, named SIBs 1 through 18. In addition to the SIBs, there is a master information block (MIB) and up to two scheduling blocks (SBs).

Broadcast information is sent via the BCCH logical channel, which is mapped into the BCH in the idle mode and in the CELL_FACH, CELL_PCH, and URA_PCH in connected mode substates. However, in the CELL_FACH substate, the UE can also receive SIB 10 via the FACH.

SIBs are sent according to a certain schedule. The blocks that are more important than others are sent more often, and the blocks of lesser importance are sent less often. The schedule is not fixed, but it can be adjusted by the UTRAN according to the current loading situation. This provides a great deal of flexibility for air interface management.

A mobile station must find out the schedule of various SIBs so that it can wake up and receive only those blocks it needs and skip reception of the others. This is possible because the blocks are arranged as a tree. This tree always starts from a MIB, which must be received and decoded first. The location of a master block is easy to determine because in the FDD mode, it has a predefined repetition rate (8), and a position (0) within the repetition cycle. This means that once a mobile knows the current frame number (it is sent in every block), it can compute the cell system frame number (SFN) mod 8, and find out the position of this block within the 8-block rotation. In the TDD mode the MIB repetition cycle can be 8, 16, or 32 frames. The value that the UTRAN is using is not signaled; the UE must determine it by trial and error.

The MIB indicates the identity and the schedule of a number of other SIBs. These SIBs contain the actual system information the UE is looking for. The MIB may optionally also contain reference and scheduling information for one or two SBs, which give references and scheduling information for additional SIBs. Scheduling information for a SIB may be included only in the MIB or one of the SBs. This is depicted in Figure 7.9. The mobile must maintain this tree in its memory, so that it can decode only those blocks it needs and skip the rest. Note that this arrangement saves power and also provides the UTRAN the possibility to add new types of SIBs to the protocol if such are needed later. Based on the experiences with GSM, new system information messages are not just possible, but likely.

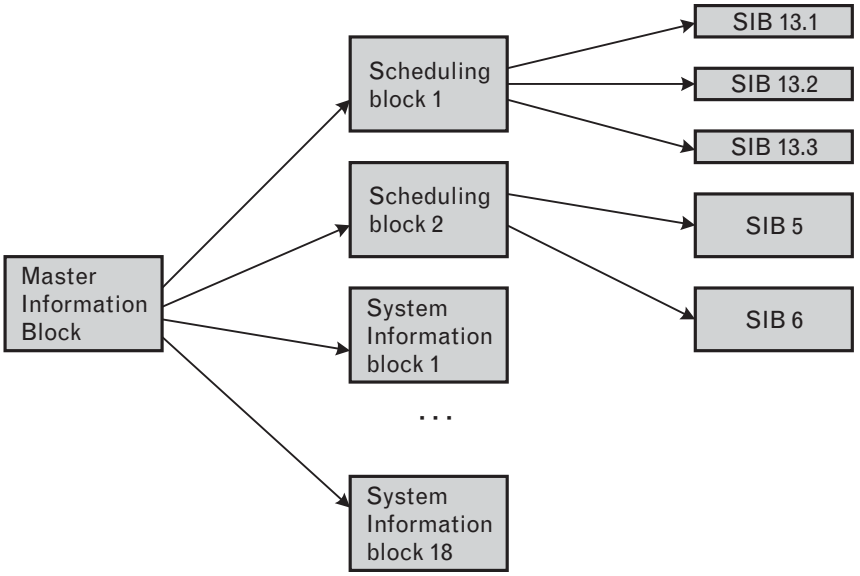

FIGURE 7.9
SIB scheduling tree.

They can be added in later phases of the system as new services and functions are needed. If a mobile finds schedules of blocks it does not recognize, it simply ignores them. Other mobiles with updated protocol software can, however, use these. If a mobile notices that the schedule in its memory does not match the schedule used by the UTRAN, it must delete the stored schedule and start building the scheduling tree again beginning from the MIB.

The network may indicate that some information in a SIB has changed by setting the update flag (value tag) in a higher block; that is, in the same block that contains the schedule for this block. Once this tag changes, the mobile knows that it should recover the corresponding system information again.

Because of the tree structure of the scheduling information, the update flag scheme is always reflected to the value tag of the MIB; that is, if any SIB changes, then the MIB also changes. To keep the mobile from decoding the MIB continuously just to find out whether any value tag has changed, the value tag of the master block itself is sent on the paging channel. This information is sent in the BCCH modification information element within the paging type 1 message. The mobile has to monitor the paging channel in any case, or it couldn't receive incoming calls. If any SIB with a value tag is changed, this will be noticed from the changed value of the MIB value tag sent via the paging channel. The listening (camped) mobiles must then receive the MIB itself and examine which value tags have been changed, and further decode those blocks.

The values of some information blocks change too frequently for this value tag scheme to be practical. For these blocks without a value tag, a

timer is used instead. Every time this timer expires, the corresponding SIB is decoded and the timer is started again.

Some modifications to SIBs are of greater importance than others. For example, if channels are being reconfigured, and mobile stations don't know about this at once, several malfunctions may take place. Therefore, a scheme has been devised that makes it possible for the mobiles to decode the new information immediately after the modification. In such a case, the BCCH modification information in the paging channel contains both the time of the change and the new value of the master block value tag after the change has occurred. The receiving mobile must start a timer using the given value as a time-out value, and once the timer expires, it has to decode the MIB, as well as all the changed SIBs.

Because the paging channel is only monitored when a mobile is in its idle CELL_PCH and URA_PCH states, it is also necessary to transmit the MIB value tag on the FACH, so that all mobiles in the CELL_FACH state can receive this information. This information is added to a system information change indication message on the FACH.

There are altogether 18 different SIBs plus the MIB, and they are briefly explained in the following paragraphs. For a more thorough description, consult the RRC specification [9]. Most probably, new SIBs will be added with the new specification releases.

MIB

- Contains PLMN identity;
- Includes references to other SIBs.

SB

- Contains scheduling information of number of SIBs.

SIB 1

- Contains NAS system information;
- Includes UE timers and counters to be used in idle mode and in connected mode.

SIB 2

- Contains a list of URA identities.

SIB 3

- Contains parameters for cell selection and reselection.

SIB 4

- Contains parameters for cell selection and reselection;
- To be used in connected mode only.

SIB 5

- Contains parameters for the configuration of the common physical channels (PhyCHs) in the cell.

SIB 6

- Contains parameters for the configuration of the common and shared PhyCHs in the cell;
- To be used in connected mode only.

SIB 7

- Contains the fast-changing parameters UL interference and dynamic persistence level;
- Changes so often, its decoding is controlled by a timer.

SIB 8

- Contains static CPCH information to be used in the cell;
- Used in FDD mode only;
- To be used in connected mode only.

SIB 9

- Contains CPCH information to be used in the cell;
- Used in FDD mode only;
- To be used in connected mode only;
- Changes so often, its decoding is controlled by a timer.

SIB 10

- Contains information to be used by UEs having their DCH controlled by a DRAC procedure;
- Used in FDD mode only;
- To be used in CELL_DCH state only;
- Changes so often, its decoding is controlled by a timer.

SIB 11

- Contains measurement control information to be used in the cell.

SIB 12

- Contains measurement control information to be used in the cell;
- To be used in connected mode only.

SIB 13

- Contains ANSI-41 system information;
- Includes four associated SIBs 13.1–13.4;
- Contains references (schedules) of the subblocks;
- To be used only when the CN of the system is ANSI-41.

SIB 14

- Contains parameters for common and dedicated physical channel (DPCH) uplink outer-loop power control information;
- Used in TDD mode only.
- Changes so often, its decoding is controlled by a timer.

SIB 15

- Contains assistance information for UE positioning methods;
- Allows the UE-based positioning methods to perform positioning without dedicated signaling;
- Allows the UE-assisted positioning methods to use reduced signaling;
- Includes five associated SIBs 15.1–15.5.

SIB 16

- Contains predefined channel configurations to be used during handover to UTRAN;
- Includes radio bearer, transport channel, and physical channel parameters to be stored by UE in idle and connected mode;
- There may be several different occurences of SIB 16 in each cell, but the UE is not required to read all of them before initiating RRC connection establishment.

SIB 17

- Contains fast-changing parameters for the configuration of the shared physical channels;
- Information becomes invalid after time specified by the repetition period (SIB REP) for this SIB;
- To be used in connected mode only;
- Used in TDD mode only.

SIB 18

- Contains the PLMN identities of the neighboring cells;
- To be used in shared access networks to help with the cell reselection process (see Chapter 12).

7.5.2.5 Paging

Paging is a procedure that is used by the UTRAN to tell a mobile that there is an incoming call waiting. Establishing the radio connection in the UTRAN is always initiated by the UE; thus, this procedure is needed to inform the UE that an establishment should, in fact, be attempted.

The paging information in the idle mode is carried by paging type 1 messages. One message may contain several paging records, each containing a paging request for a different mobile. It is also possible that a paging message does not contain any paging records at all, but only a BCCH modification information element, which contains the value tag for the MIB. Once a mobile receives a modified value tag, it knows that the MIB must be read from the BCCH.

In principle, the mobile must monitor the PCH continuously to make sure that it does not miss any paging messages and therefore lose incoming calls. However, a continuous reception scheme would soon wear down

the battery; thus, a mechanism called discontinuous reception (DRX) is employed.

The DRX scheme is based on the fact that each UE (actually each USIM card) has a unique IMSI, and from that IMSI and the IE "CN domain specific cycle length coefficient" (received in SIB 1), it is possible to compute the paging occasions for this UE. These are frame numbers, and the network makes sure it will deliver paging messages to certain UEs only during the said frame numbers; the mobile knows when it is safe to fall asleep, confident that the network will hold to its paging schedule. Further enhancement (and substantial power savings) is achieved by introducing a paging indication channel (PICH). The UE actually listens only for this PICH periodically, and when a positive indication appears, then the UE knows to listen for the actual PCH as it may only now contain a message addressed to this mobile. There may also be several PCH/PICH pairs in one cell. This is indicated in SIB 5. The UE selects the one to be used based on its IMSI.

Several mobiles may listen for the same paging occasion. Thus, the UE (i.e., the UE's RRC layer) must check whether any of the paging identities of the received paging records matches its own identity. If a match occurs, the paging indication is forwarded to the MM function, which triggers a call-establishment procedure.

Note that there is a trade-off between mobile standby times and call setup times. If the discontinuous reception procedure uses a long DRX cycle (i.e., the UTRAN can send paging messages relatively seldom), then UEs do not have to listen for the PICH so often, which saves power. This results in longer UE standby times. However, the drawback is longer call setup time with mobile-terminated calls. These parameters can be set by the operator, and they can also be changed dynamically because they are continuously sent via the BCCH.

The paging description discussed so far has concentrated on idle-mode paging. It is also possible that the UTRAN pages the UE in the connected mode. In the CELL_PCH and URA_PCH states, the paging request triggers a UE state change. The UTRAN uses this procedure when it has some additional downlink data to be sent to the UE. The DRX cycle length to be used may be different from the idle mode in the CELL_PCH and URA_PCH states. The UE must use the shortest cycle length of any CN domain it is connected to, or the UTRAN DRX cycle length, whichever is shorter.

The actions the UE takes once a paging message is received in the RRC depend on the RRC state. In idle mode the UE shall react thus:

- If the IE "paging originator" is the CN, compare the included identities of type "CN UE identity" with all of its allocated CN UE identities.

- For each match, forward the identity and paging cause to the upper-layer entity indicated by the IE "CN domain identity."

- If the IE "paging originator" is the UTRAN, ignore that paging record.

In connected mode the UE shall behave thus:

- If the IE "paging originator" is the UTRAN, compare the included identities of type "UTRAN originator" with its allocated U RNTI.
  - For each match, the UE will enter CELL FACH state and perform a cell-update procedure with cause "paging response."
- If the IE "paging originator" is the CN, ignore that paging record.

So, if the paging originator is the CN, a possible paging response is initiated by the MM (RRC connection request, establishment cause = paging response). If the paging originator is the UTRAN, then a possible response is initiated by the RRC (cell update, cause = paging response).

Dedicated Paging
Dedicated paging is a paging message (paging type 2) that is sent in connected-mode states CELL DCH and CELL FACH. It is used, for example, to establish a signaling connection. Note that there is already an existing signaling connection in these states, but the dedicated paging procedure can be used in cases in which a CN other than the current serving CN wants to originate a dialogue with a UE. The RRC receives and decodes the message and forwards the paging identity and the cause to the NAS entity indicated by the CN domain identity. Note that for the access stratum, this is just another signaling message on a dedicated connection.

7.5.2.6 RRC Connection Establishment

The UMTS separates the concepts of a radio connection from a radio bearer (RB). The UE requests an RRC connection, but the network requests an RB setup. Experienced readers may recall that the bearer capability was attached to the radio channel in GSM. Once a resource was allocated it was not possible to change the bearer capabilities. The bearer capability can be changed dynamically during the radio connection in the UMTS provided that the network has the necessary resources. This is a very important enhancement.

An analogy to the relationship between a radio connection and an RB might be a train system, where an RRC connection is the track and a bearer connection is the railway carriages on the track. The track makes it possible

to transfer goods, but the carriages define the kinds of goods delivered and in what quantity they can be delivered.

The RRC connection establishment procedure is quite simple in the RRC level (Figure 7.10), but note that a rather complex RACHing procedure (see Section 11.5) must take place in the lower layers before the RRC connection request message is received in a Node B.

Note that the various forms of this procedure are similar regardless of who initiated the connection: mobile or network. A network-originated connection is triggered by a paging procedure. Only the establishment cause field in the RRC connection request message indicates the reason for the connection establishment.

The RRC must store the value of "initial UE identity" encapsulated in the RRC connection request message, and check that it receives the same value back in the RRC connection setup message. It is possible that several UEs have sent RACH messages at the same time; thus, the response message may not be addressed to this particular UE. This procedure is also known as contention resolution, and it is required because the channels used for RRC connection request and setup messages are sent on common channels. If the received initial UE identity is the same as the sent value, then the UE can continue with the connection setup. Otherwise, the connection establishment must be started again.

The RRC connection setup message is a large one containing significant information to be used in the L1/L2 configuration. This includes transport format sets and transport channel information. The final message (RRC connection setup complete) is sent via the new DCH.

The network can also reject the connection setup and respond with an RRC connection reject instead. Note that an RRC connection is not the same thing as a dedicated connection. An RRC connection may also be mapped onto a common or shared channel.

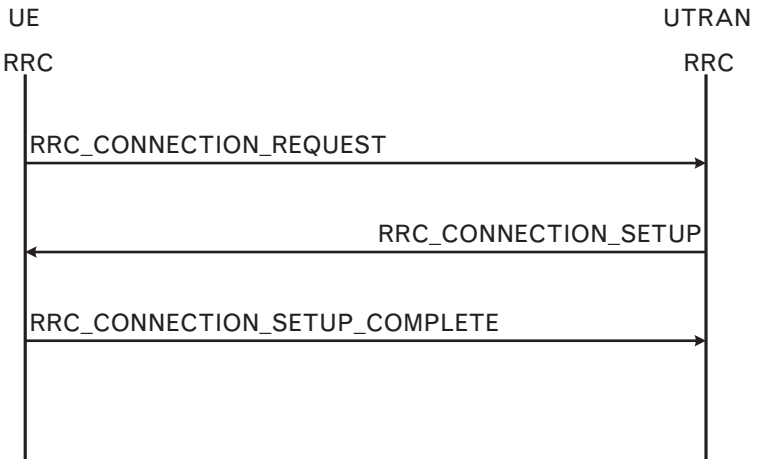

FIGURE 7.10 *RRC connection establishment.*

7.5.2.7 RRC Connection Release

Whereas an RRC connection is always initiated by the UE, its release is always initiated from the network side. There are two versions of this procedure depending on whether a DPCH exists or not. The RRC layer signaling is similar in both cases; the network sends an RRC connection release message, and the UE responds with an RRC connection release complete (Figure 7.11).

If a dedicated physical connection exists, then it is used for the transport of these messages. Both messages are sent using unacknowledged mode in the lower layers. The UE may send its response message several times (the number is defined by variable V308), and after the last transmission, it releases all its radio connections. Similarly, the UTRAN releases all UE-dedicated resources after receiving the complete message. In fact, this is done once a corresponding timer expires, even if it doesn't receive the connection release complete message.

If there is not a dedicated connection, then the downlink RRC connection release message must be sent via the FACH, and the uplink RRC connection release complete message via the RACH. Only one occurrence of both messages is sent, and the acknowledged mode is used. This means that a data acknowledgment message is sent as an acknowledgment in the RLC protocol level to the RRC connection release complete message via the downlink FACH.

7.5.2.8 RRC Connection Reestablishment

If a UE loses the radio connection suddenly, it may attempt a connection reestablishment. This requires a quick cell reselection and sending a request for reestablishment message to the UTRAN.

Once the UE detects that it has lost the connection, it starts timers T301 and T314/T315. The reestablishment attempt must succeed while these timers are still running. On their expiration a possible ongoing reestablishment is abandoned and the UE returns to its idle mode.

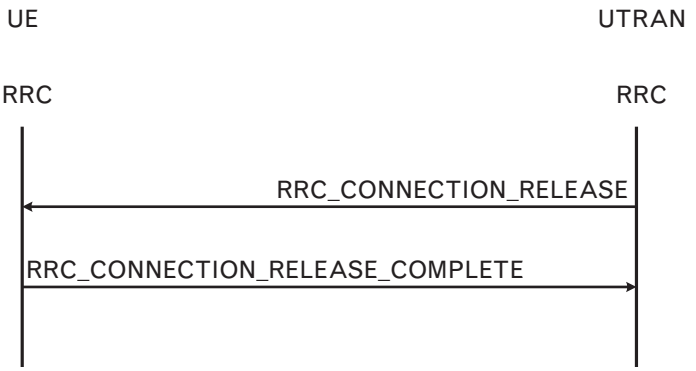

FIGURE 7.11
RRC connection release.

The reestablishment is requested by the NAS (MM) in the UE. The UE must find a suitable cell for reestablishment and do so quickly. The old serving cell will not be accepted, as it is still regarded as unusable.

In the new chosen candidate cell, an RRC connection reestablishment request is sent on the uplink RACH. If the UTRAN can reconnect the old connection, it responds with an RRC connection reestablishment message on the FACH. The UE configures its L1 according to the information obtained from this message, gains synchronization, and then responds with an RRC connection reestablishment complete message on the new DCCH (see Figure 7.12).

Note that a connection reestablishment may require a considerable number of tasks to be carried out by the network; the call may have to be rerouted via new switches and base stations. The probability of these being successful diminishes quickly with time, so a reestablishment procedure must be performed as quickly as possible.

7.5.2.9 Radio Bearer Establishment

As previously explained in the discussion of RRC connection establishment, radio connections and RBs are two separate concepts in UMTS.

A radio connection is a static concept. It is established once, and it exists until it is released. There is only one radio connection per (typical) terminal. On the other hand, the RB defines what kind of properties this radio connection has. There may be several RBs on one radio connection, each having different capabilities for data transfer. RBs are also dynamic; they can be reconfigured as necessary. This is not to say that radio connections cannot be configured. They will be reconfigured in many ways. The most common radio-connection-reconfiguration procedure is probably the HO procedure.

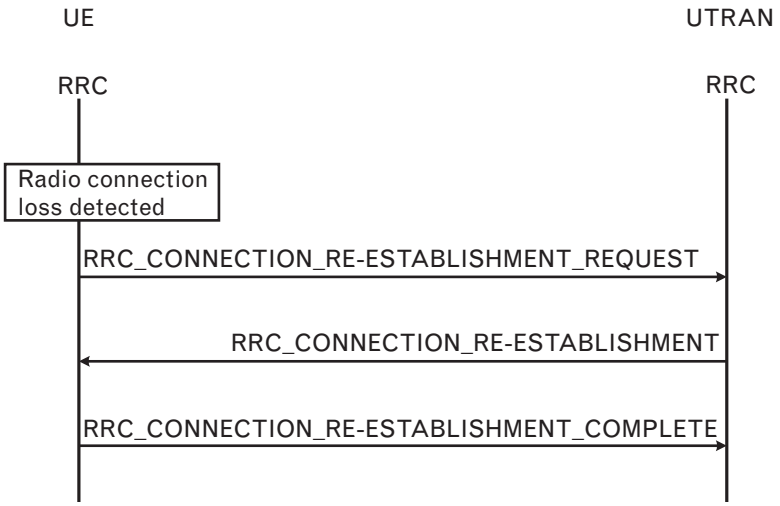

FIGURE 7.12
RRC connection reestablishment.

Indeed, it is also possible to have an RB without a dedicated radio connection. Circuit-switched bearers or bearers using real-time services typically need dedicated radio channels to meet their delay requirements. Packet-switched bearers or bearers using non-real-time services, however, often do not need a permanent association to a dedicated radio resource; they can use shared or common channels.

RB establishment is always initiated by the UTRAN. This is because an RB uses a certain amount of radio resources from the network and these resources are quite limited. Only the network knows what kind of resources it can grant to a UE.

The UTRAN RRC gets a request for a new bearer from the higher layers in the NAS. Down at the RRC level, the signaling is simple: the UTRAN sends a radio bearer setup message, and the UE responds with a radio bearer setup complete (Figure 7.13). However, the interlayer signaling to lower layers can be quite different depending on the requested QoS parameters and whether there already exists a suitable physical channel. These variations are discussed in Section 11.2.

7.5.2.10 Radio Bearer Reconfiguration

This procedure is used to reconfigure an RB if its QoS parameters have been changed or traffic-volume measurements indicate that more or fewer resources are required. Also, this procedure can either be synchronized or unsynchronized, depending on whether the old and new RB setups are compatible. The synchronized bearer-reconfiguration procedure includes an activation time parameter, which is used to ensure that the change takes place at the same time in both the UE and the UTRAN. At the end of the procedure, the old configuration, if such exists, must be released from Node B(s) (see Figure 7.14).

Note that the coexistence of the old and the new configuration in unsynchronized change does not mean that a UE has to maintain two

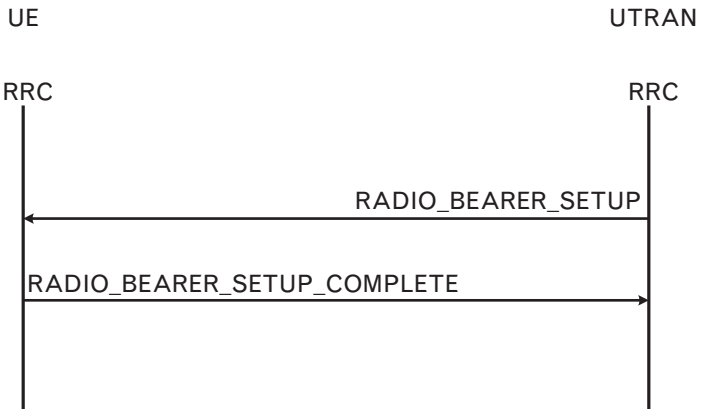

FIGURE 7.13
Radio bearer establishment.

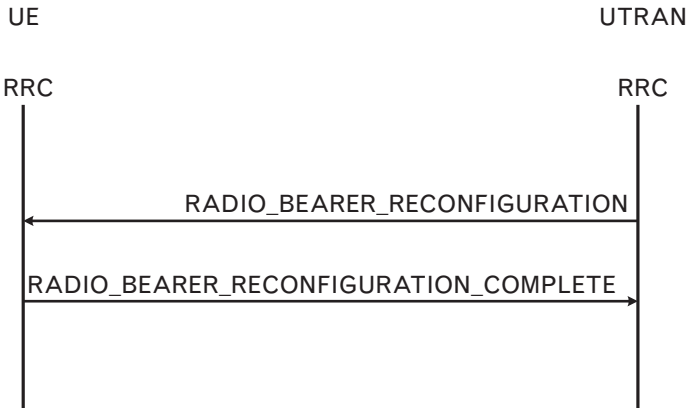

FIGURE 7.14
Radio bearer reconfiguration.

configurations at the same time. It means that a UE can use one configuration (the new one) and the network another (the old one) temporarily during the reconfiguration procedure, and they can still continue to communicate. A synchronized change is used when the old and the new configurations are incompatible; thus, the changes have to take place exactly at the same time.

7.5.2.11 Radio Bearer Release

This procedure releases one or more RB(s) (Figure 7.15). Again this procedure has two variations, synchronized and unsynchronized. The unsynchronized procedure is simpler because the old and new configurations can coexist and therefore the change does not have to take place simultaneously in the UE and in the UTRAN. The synchronized procedure, on the other hand, must use a time parameter to ensure synchronization.

The RB release procedure may include a physical channel modification or deactivation, depending on the requirements of the new situation (i.e., what kind of bearers still exist after this release).

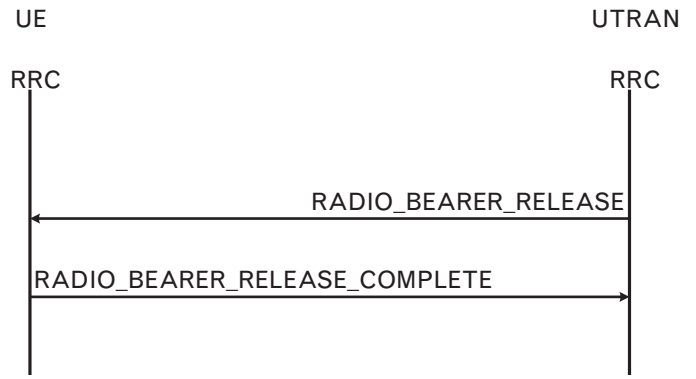

FIGURE 7.15
Radio bearer release.

7.5.2.12 Management of the RRC Connection

This service includes assignment, reconfiguration, and release of radio resources for the RRC connection, for example, assignment of codes and CPCH channels.

The RRC layer handles the assignment of the radio resources (codes and CPCH channels) needed for the RRC connection, including the needs of both the control and the user planes. The RRC layer may reconfigure radio resources during an established RRC connection. This function includes coordination of the radio resource allocation between multiple RBs related to the same RRC connection. The RRC controls the radio resources in the uplink and downlink, such that the UE and the UTRAN can communicate using unbalanced radio resources (asymmetric uplink and downlink). The RRC signals the UE to indicate resource allocations for the purposes of an HO to GSM or other radio systems.

7.5.2.13 HOs

Being a mobile device, a UE may move during a connection. It is the responsibility of the RRC layer to maintain the connection if a UE moves from one cell area to another. In the case of a circuit-switched connection, this procedure is called an HO. There is no radio connection in packet-switched data transfer; thus, there is no need to maintain it. In this case a UE makes a normal cell reselection, and the data packets are then rerouted via the new cell.

An HO decision is done in the UTRAN RRC, and it is based on, among other things, the measurements done by the UE.

An SHO is a procedure in which a UE maintains its connection with the UTRAN via two or more Node Bs simultaneously. An SHO is a CDMA-specific procedure, made possible by the fact that all base stations use the same frequency; thus, it is relatively easy for the UE to receive several of them at the same time. The SHO procedure is a common condition in the life of a UE. A UE can be in an SHO state for most of the time it is in a call. A user cannot notice any kind of change in the voice quality while the UE enters or leaves the SHO state, except that the voice quality may actually be better in an SHO state than in a normal state just before entering the SHO state.

The base stations to which a UE is connected are said to belong to the UE's active set. SHOs are managed with active set update messages sent by the network (Figure 7.16). Note that the UE should not update its active set by itself, but only according to these messages. A UE in an SHO always consumes more network resources than a UE with a normal single connection to the network. Therefore, it must be the network side that decides which UEs need the additional gain from an SHO, and which UEs can do without.

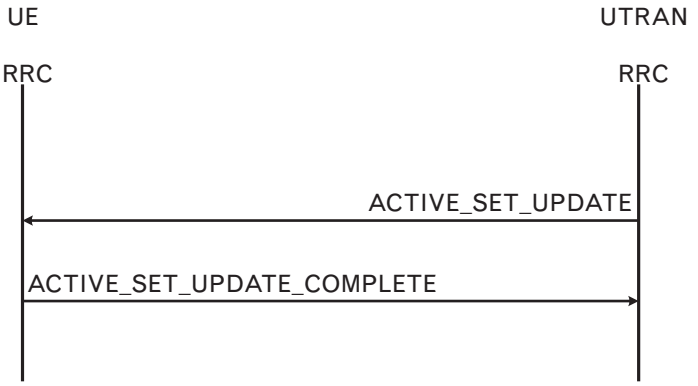

FIGURE 7.16
SHO management.

A hard handover (HHO) corresponds to the normal GSM HO procedure, which is always hard. The term *hard handover* indicates that the old connection is first released before the new one is set up. This may result in an audible break in a speech connection. Within UTRAN an HHO will be performed when the radio frequency channel changes. The UMTS also contains dedicated procedures for intersystem HOs, which are a form of HHO. In practice, these HOs will occur only between UMTS and GSM/GPRS networks at the first phase of UMTS.

Note that there is no special HHO procedure in the UTRAN air interface. An HHO can result from other procedures that reconfigure the air interface. If such a procedure changes the radio frequency of the connection, then an HHO takes place. Procedures that may trigger an HHO include a physical channel reconfiguration, a radio bearer establishment, a radio bearer reconfiguration, a radio bearer release, and a transport channel reconfiguration.

The relocation procedure is a UTRAN internal HO procedure. It reroutes the connection in the UTRAN, but does not affect the radio connection in the air interface, even as there are some implications for higher layers in the UE (e.g., the PDCP may have to support a lossless relocation).

HO processes are further discussed in Section 2.5. Many HO procedures are depicted in Chapter 11.

7.5.2.14 Measurement Control

The measurements performed by the UE are controlled by the RRC layer, which decides what to measure, when to measure, and how to report the results, including both the UMTS air interface and other radio systems. The RRC layer also constructs reports of the measurement results from the UE to the network.

In the idle mode, the measurements are usually made to support cell-reselection procedures, which means they are internal to the UE. It is also

possible that some of the measurement results must be reported to the network when the UE is in the connected mode.

The list of possible measurements the UTRAN may ask the UE to perform is a long one. The UE measurements can be grouped into seven categories or types:

1. Intrafrequency measurements;
2. Interfrequency measurements;
3. Intersystem (or Inter-RAT) measurements;
4. UE positioning measurements;
5. Traffic-volume measurements;
6. Quality measurements;
7. UE internal measurements.

This same grouping is also used in the measurement control message that is sent to the UE to ask it to perform measurements. Each control message can set up, modify, or release a measurement procedure. There may be several parallel measurement procedures running at the same time, so the control message also contains an identity number, which indicates the measurement process this message applies to. Each control message may contain control information for only one measurement type at a time.

The exact rules for how to process various measurement requests and how to perform the actual measurements in each state are given in Section 8.4 of [9].

The results of the measurements can be sent back to UTRAN in measurement report messages. This can happen either periodically, after some triggering event, or perhaps using a combination of these: the first report is sent after the triggering event, and the following messages are sent periodically, while the same triggering condition still applies.

Typically, the actual radio measurements will be performed in layer 1 according to instructions from the RRC layer. However, traffic volume measurements will be performed in the MAC layer. They are based on the RLC buffer occupancy information, which is regularly reported to the MAC layer.

The measurement model for air interface measurements is given in Figure 7.17.

A. *Measurements performed in L1.* These are filtered in L1 so that not all measurement samples will be sent to the RRC. The specification does not give exact filtering rules, but merely the performance objectives and the reporting rate at point B. L1 typically has to collect

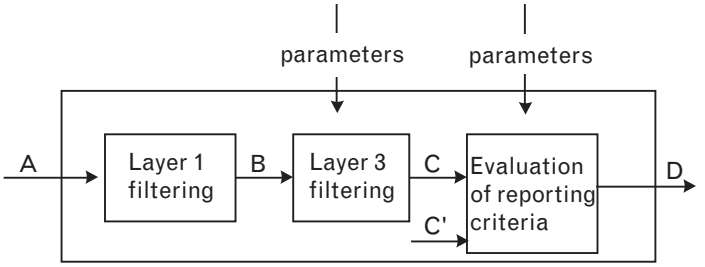

Figure 7.17
Measurement model. (Source: [12]).

several samples from the same measurement process, average them, and send the results to the RRC.

B. *Measurement reports sent to the RRC.* L3 must further filter these reports based on the rules received from the UTRAN in a measurement control message.

C. *A measurement after processing in the RRC filter.* The specification states that the rate of reporting in C is the same as in B. These reports are then evaluated to find out whether they have to be sent further.

C'. *Reporting thresholds.*

D. *Measurement report sent over the air or the Iub interface.*

There is a long list of events that may trigger a measurement report. Normally, a UE does not have to report all these events. They form a pool of events from which the UTRAN can choose the events to be reported. In some cases the event triggers the first measurement report, which is then followed by periodic reporting. The list of possible events is given next.

Intrafrequency reporting events for the FDD mode:

1A. A primary CPICH enters the reporting range.

1B. A primary CPICH leaves the reporting range.

1C. A nonactive primary CPICH becomes better than an active primary CPICH.

1D. Change of best cell.

1E. A primary CPICH becomes better than an absolute threshold.

1F. A primary CPICH becomes worse than an absolute threshold.

Additionally for the TDD mode:

1G. Change of best cell.

1H. Time slot interference on signal code power (ISCP) below a certain threshold.

1I. Time slot ISCP above a certain threshold.

Interfrequency reporting events (for both FDD and TDD modes):

2A. Change of best frequency.

2B. The estimated quality of the currently used frequency is below a certain threshold, and the estimated quality of a nonused frequency is above a certain threshold.

2C. The estimated quality of a nonused frequency is above a certain threshold.

2D. The estimated quality of the currently used frequency is below a certain threshold.

2E. The estimated quality of a nonused frequency is below a certain threshold.

2F. The estimated quality of the currently used frequency is above a certain threshold.

Intersystem reporting events (for both the FDD and TDD modes):

3A. The estimated quality of the currently used UTRAN frequency is below a certain threshold and the estimated quality of the other system is above a certain threshold.

3B. The estimated quality of the other system is below a certain threshold.

3C. The estimated quality of the other system is above a certain threshold.

3D. Change of best cell in other system.

Traffic-volume reporting events:

4A. Transport channel traffic volume exceeds an absolute threshold.

4B. Transport channel traffic volume becomes smaller than an absolute threshold.

Quality reporting events:

5A. A predefined number of bad CRCs is exceeded.

UE internal measurement reporting events:

6A. The UE Tx power becomes larger than an absolute threshold.

6B. The UE Tx power becomes less than an absolute threshold.

6C. The UE Tx power reaches its minimum value.

6D. The UE Tx power reaches its maximum value.

6E. The UE RSSI reaches the UE's dynamic receiver range.

6F. The UE Rx-Tx time difference for an RL included in the active set becomes larger than an absolute threshold (for 1.28 Mcps: The time difference indicated by T_{ADV} becomes larger than an absolute threshold).

6G. The UE Rx-Tx time difference for an RL included in the active set becomes less than an absolute threshold.

UE positioning reporting events:

7A. The UE position changes more than an absolute threshold.

7B. SFN-SFN measurement changes more than an absolute threshold.

7C. GPS time and SFN time have drifted apart more than an absolute threshold.

7.5.2.15 Outer-Loop Power Control

The RRC layer controls the setting of the target of the closed-loop power control mechanism.

The closed loop is the power control method that is used during the UTRAN connection after the initial setup phase includes two subprocesses: the inner and the outer closed-loop power control. The inner-loop power control is an L1 internal procedure, and the outer-loop control also includes the RRC layer.

The inner-loop control uses the SIR_{target} value to adjust the transmission power levels in the air interface. L1 measures the received SIR and compares it to SIR_{target}. If the value is worse, a power increase request is sent to the base station; otherwise, a lower transmission power is requested. CDMA systems are always looking for ways to lower transmitter power. The problem here is that the SIR_{target} value cannot be kept constant. Different connections will have varying QoS targets, mobile terminals will have different speeds, and SHOs will also have an effect on the QoS (i.e., even if the set SIR_{target} is fulfilled, the result may still have too many errors for an application that requires high QoS). Therefore, an outer-loop power control is needed in the UE. This monitors the quality of the received signal and adjusts the SIR_{target} value accordingly; that is, if the received quality is too low, SIR_{target} is increased, and vice versa. A similar outer-loop power control scheme is used both in the UE and in the RNC. In the UE it adjusts the downlink quality, while in the RNC it adjusts the uplink quality.

Once an RB has been established, the UTRAN sends a radio bearer setup message with IE "added or reconfigured DL TrCH information." This IE contains a block error rate (BLER) value for each coded composite

transport channel (CCTrCH) used. This is the quality target the UE attempts to maintain. If the received BLER is lower than the target BLER, then SIR_{target} is increased. If the quality exceeds the minimum required, SIR_{target} is reduced. The quality target control loop is run so that the quality requirement is met for each transport channel. The target BLER value can be dynamically changed by the UTRAN via a radio bearer reconfiguration message.

For the CPCH the quality target is set as the BER of the DL DPCCH as signaled by the UTRAN. This value is signaled to UEs in SIB 8 broadcasts. Similarly, the UE runs a quality target control loop such that the quality requirement is met for each CPCH transport channel that has been assigned a DL DPCCH BER target.

The UE sets the SIR_{target} when the physical channel has been set up or reconfigured. It will not increase the SIR_{target} value before the power control has converged on the current value. The UE may estimate whether the power control has converged on the current value by comparing the averaged measured SIR to the SIR_{target} value.

If the UE has received a DL outer-loop control message from the UTRAN indicating that the SIR_{target} value should not be increased above the current value, it should record the current value as the maximum allowed value for the power control function. Once the RRC receives a new DL outer-loop control message from UTRAN indicating that the restriction is removed, it is then free to continue using the standard power control algorithm.

7.5.2.16 Security Mode Control

The security mode control procedure is used by the UTRAN to control either the ciphering or the integrity protection processes. In the case of ciphering control, this procedure can either start the ciphering, or change the ciphering key. This section also discusses the counter check procedure, as it is closely related to the UE security (see Figure 7.18).

Ciphering Control

This is what is known in GSM as *ciphering key control*. It can trigger the start of ciphering or command the change of the cipher key.

This procedure is always initiated by the UTRAN. It sends a security mode command message to the UE containing the new ciphering key and the activation time. The activation time will be given as an RLC send sequence number (for AM and UM bearers) or as a CFN (for TM bearers). Once the activation time elapses, the UE starts to use the new configuration and sends a security mode complete message back to the UTRAN. Note that this message is still sent using the old ciphering configuration.

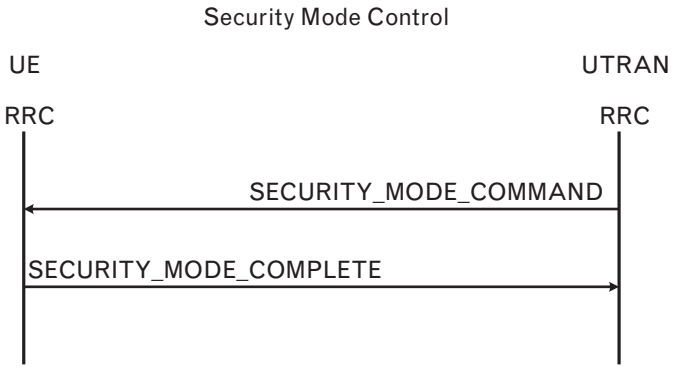

FIGURE 7.18
Security mode control.

Several different ciphering algorithms will be defined by the 3GPP, but the UE only has to support one algorithm at a time (i.e., the same algorithm will be used both in the MAC and in RLC). However, the UTRAN may change the active algorithm at any time during the call. Although the algorithm is the same in the MAC and in the RLC, the used ciphering key sequence numbers will differ. There are three versions of the ciphering key sequence number depicted in Figure 7.19. In 3GPP jargon, the ciphering key sequence number is called *COUNT-C*.

The RLC layer uses the RLC sequence number as a part of the COUNT-C, but in transparent mode this cannot be used because in this case there are no sequence numbers in the RLC, so an eight-bit ciphering frame number from the MAC is used instead.

The actual ciphering process is depicted in Figure 7.20. Note that the deciphering process is similar, except that the input stream is ciphertext, and the output is plaintext.

See [13] for further information on ciphering.

Integrity Protection

Integrity protection is a scheme that guards the signaling traffic in the air interface against unauthorized attacks. An intruder could try to modify the message sequences (e.g., by means of a man-in-the-middle attack). Integrity

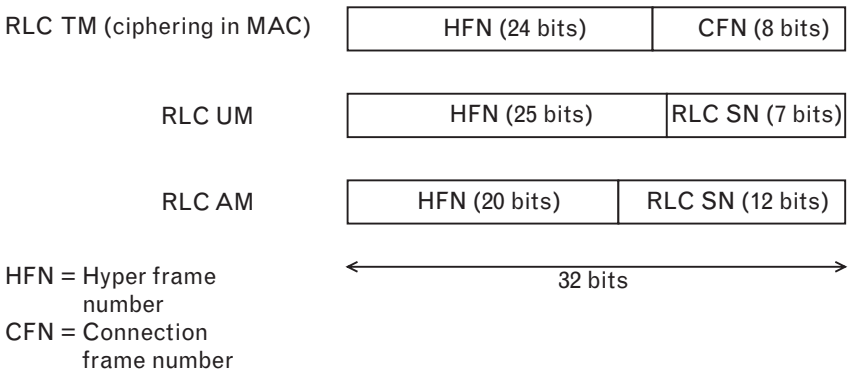

FIGURE 7.19
Ciphering key sequence numbers (COUNT-C).

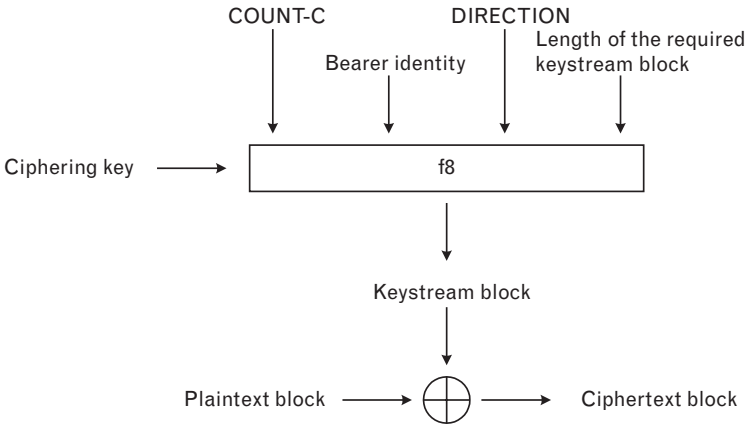

FIGURE 7.20
Ciphering process.

protection ensures that the signaling procedures are not tampered with (or at least makes it very difficult to break the security).

The integrity-protection process is started (and restarted) by the security mode procedure. The same procedure is also used for ciphering control.

To start or reconfigure the integrity protection, the UTRAN sends a security mode command message on the downlink DCCH in AM RLC using the present integrity-protection configuration.

Integrity protection is performed on all RRC messages except:

- HANDOVER TO UTRAN COMPLETE;
- PAGING TYPE 1;
- PUSCH CAPACITY REQUEST;
- PHYSICAL SHARED CHANNEL ALLOCATION;
- RRC CONNECTION REQUEST;
- RRC CONNECTION SETUP;
- RRC CONNECTION SETUP COMPLETE;
- RRC CONNECTION REJECT;
- RRC CONNECTION RELEASE (on CCCH only);
- SYSTEM INFORMATION;
- SYSTEM INFORMATION CHANGE INDICATION;
- TRANSPORT FORMAT COMBINATION CONTROL.

For the CCCH and for each signaling RB, two integrity-protection hyperframe numbers are used (both 28 bits):

1. Uplink HFN;

2. Downlink HFN.

And two message sequence numbers are used (both 4 bits):

1. Uplink RRC message sequence number;
2. Downlink RRC message sequence number.

By combining these numbers, we get two 32-bit integrity sequence numbers, COUNT-I, one for uplink and one for downlink, for each signaling radio bearer (RB 0–4). Once a UE receives a downlink signaling message, it calculates a message authentication code (MAC) based on the stored COUNT-I information and the received message. The algorithm for MAC calculation is given in Section 8.5.10.3 of [9] and in [13]. The calculated MAC must match with the received MAC, otherwise the message has been tampered with and must be discarded.

In the uplink direction, the UE calculates a MAC value and attaches it to the uplink signaling message. The UTRAN has to perform the integrity check in the same way as the UE.

Integrity protection is described in Section 8.5.10 of [9].

Counter Check

Counter check is yet another security procedure in the air interface. It is used to check that the amount of data sent in the air interface is similar in both the UE and in the UTRAN.

The UE must maintain a ciphering sequence number (COUNT-C) for each radio bearer. The UTRAN can query this number from time to time to ensure that there are no intruders taking part in the communication. It sends a counter check message to the UE and includes COUNT-C values for each RB. The UE must compare the received COUNT-C values with the actual COUNT-C values used, and include the number of mismatches in the response message (counter check response) (see Figure 7.21).

Once the UTRAN receives this message, it checks the number of mismatches to find out whether there is anything suspicious going on in the air interface. Note that counter check is used in the user plane, and integrity protection in the control plane. Thus, both procedures are needed, and they complement each other. Indeed, notice that counter check signal exchange is protected by the integrity protection.

The counter check procedure is defined in Section 8.1.15 of [9].

7.5.2.17 Direct Transfer

NAS messages (higher-layer messages) are relayed over the air interface within direct transfer messages. There are separate messages for the uplink and downlink directions, the uplink direct transfer and the downlink direct

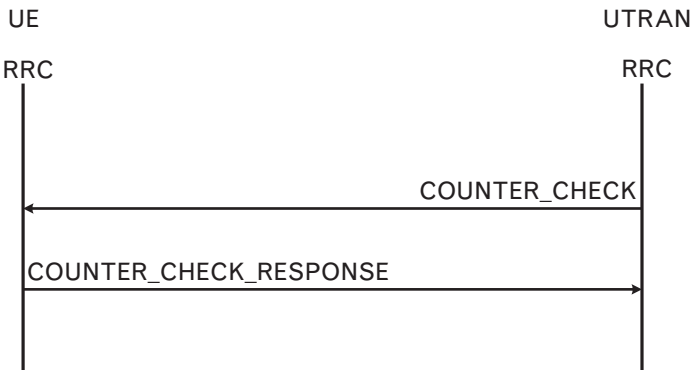

FIGURE 7.21
Counter check.

transfer, respectively. However, if there is not yet an existing signaling connection, then the initial direct transfer message must be used.

The initial direct transfer procedure is typically used when the upper layers request an initialization of a new session. This request also includes the initial NAS message. If there is no RRC connection, the RRC must establish it using the normal RRC connection-establishment procedure. After that the direct transfer message is transferred using AM. A successful signaling-connection establishment is also confirmed to UE-NAS.

In the RLC layer the direct transfer message should be transmitted using either radio bearer 2 (RB2) or RB3. Low-priority messages, such as GSM-SAPI3, should use RB3, and high-priority messages, such as GSM-SAPI0, should use RB2. Note that this priority is indicated by the NAS to the RRC.

Downlink direct transfer messages are routed to higher layers based on the value of the CN domain identity element in the message (CS domain/PS domain). This routing decision is done in the RRC.

Direct transfer is discussed in [9] in Sections 8.1.8, 8.1.9, and 8.1.10.

7.5.2.18 Control of Requested QoS

This function ensures that the QoS requested for the RBs can be met, which includes the allocation of a sufficient number of radio resources for an application. Different applications have different demands for the data transmission services they require. The QoS parameters may include the required bandwidth, the allowed transmission delay, and the allowed data error tolerance. QoS classes in UTRAN are discussed in Section 13.6.5.

The UTRAN air interface is a very flexible one, allowing for the dynamic allocation of system resources. The UTRAN allocates a minimum number of resources for each UE so that the negotiated QoS can be maintained. This chapter has in several earlier sections explained some of the individual procedures related to the QoS and the control of the system's resources, but the overall picture was not explained; it was not shown how

the individual procedures are related to each other. This is the purpose of this section.

The individual procedures presented here include physical channel reconfiguration, transport channel reconfiguration, radio bearer reconfiguration, and traffic-volume measurements.

Increased Data

We begin by considering what happens when the amount of data transmitted increases. The UTRAN controls this procedure. The initial resource allocation (radio bearer setup) was discussed earlier. Once the data transmission is under way, the UE's MAC layer monitors the amount of unsent data in the RLC transmission buffer. If the buffer fills over a predefined threshold, then a measurement report is sent to the UTRAN. This message acts as a request message for additional resources.

The UTRAN may allocate additional resources to the UE depending on the availability of the resources and the negotiated QoS of the UE. There are several ways it can do this.

Small amounts of data are typically, but not necessarily, transferred over common transport channels. Even small amounts of data could be sent via dedicated channels if the data has strict delay requirements, for one cannot easily guarantee the delays in common channels. If the RLC's buffer fills up while the common channels are used, then the UTRAN may assign a dedicated channel to the UE (i.e., the case of RACH/FACH to DCH/DCH). This will be done using the PhyCH reconfiguration procedure; see Section 11.2.5.

If the UE already has a DCH assigned to its application, then the increase in data transmit capacity must be handled by reconfiguring the existing channels. For this purpose there are three possible procedures that could be employed. First among these is the physical channel reconfiguration procedure, which is called on when there are already suitable TFs and TFCs defined for the UE, and these can be used in the new configuration. Second, if the TFs and TFCs must be reconfigured, then the transport channel reconfiguration procedure (Section 11.2.4) must be used. Third, and the most powerful procedure of these is the radio bearer reconfiguration procedure, which can change QoS parameters, change the multiplexing of logical channels in the MAC layer, and reconfigure the transport channels and the physical channels. Note that the UTRAN decides which of the three procedures to use. The UE simply has to follow the orders given to it via the RRC signaling link.

Here we list the procedures in the order of the amount of change they make:

1. Radio bearer reconfiguration;
2. Transport channel reconfiguration;

3. Physical channel reconfiguration.

The radio bearer reconfiguration procedure can be used to change more parameters than the physical channel reconfiguration. A higher-order procedure may also include all the configuration parameters of a lower-order procedure.

It is also possible that new RBs can be allocated or the old RBs can be released, but these actions are not due to the traffic measurements in the RLC, but because a new application/service needs one. Then, of course, a release would mean that the UE no longer needs an RB.

The issues explained above also apply to downlink data, except that the transmission buffer to be monitored is now in the UTRAN.

Decreased Data

If the amount of transmitted data decreases, then the underutilized resources should be deallocated to make them available for other users. The triggering event is, as was the case with increasing data, the transmission buffer in the RLC. If the buffer contents remain less than a lower threshold for a certain time, then the traffic-volume-monitoring process sends a measurement report message to the UTRAN.

It is up to UTRAN to decide what kind of actions (if any) will be taken. As in the increased data case, there are three procedures to choose from: physical channel reconfiguration, transport channel reconfiguration, and radio bearer reconfiguration. If the amount of transmitted data is low enough, the UTRAN may command the UE to release its DCHs and use common channels instead (i.e., the case from the DCH/DCH combination to the RACH/FACH combination). This change can be included in all three reconfiguration procedures mentioned above.

Other Observations

Note that the UTRAN must consider the RLC buffer's content in both the UE and the UTRAN when it decides which procedure to use. For example, even if there is very little uplink traffic (i.e., the buffer in the UE-RLC is empty), it cannot move the UE from a DCH/DCH to a RACH/FACH if there is plenty of downlink traffic. The UE cannot have a DCH in only one direction (i.e., combinations of RACH/DCH and DCH/FACH are not allowed). DCHs must use fast power control, and it is impossible to implement an efficient power control loop using common channel resources in the other direction.

The UTRAN can also "fine-tune" the UE's data transfer resources with the transport format combination control procedure, which merely modifies the allowed uplink transport format combinations within the transport format combination set. Also note that even large changes in the amount of transmitted data do not necessarily trigger any of the

reconfiguration procedures described earlier. The reader must remember that the transmitting entity has a set of TFCs it can choose the suitable one for the current situation. If there is lots of data in RLC buffers, a high-capacity TFC could be used, for example. On the other hand, if common channels are used in the physical layer, the configured TFCS most probably cannot include high-capacity TFCs, as those would be difficult to send over the air interface.

7.5.2.19 Support for DRAC

Many studies have shown that within a typical WCDMA cell, the capacity is uplink limited; that is, if the traffic is increased in a cell, it is the uplink that gets overloaded first. However, this conclusion may not be true in a cell with multimedia users, which will generate much more downlink than uplink traffic; a cell with significant asymmetric traffic.

In the uplink, the spreading codes are not orthogonal, but pseudorandom; thus, the user signals appear as interference to each other. To ease the situation in the uplink, the 3GPP has defined a scheme called dynamic resource allocation control (DRAC). It is a very fast method to spread the load in the uplink DCH while avoiding peaks in the interference level. The UTRAN may assign uplink DCHs with DRAC information elements indicating that the UE must use DRAC in that uplink DCH.

The UTRAN transmits the DRAC parameters regularly via the SIB 10 broadcast message. As this message must be received by the UE in the connected mode, it is mapped into a FACH (and further to an S-CCPCH) channel. These parameters indicate the allowed subset of TFCS according to the given maximum bit rate:

$$\sum_{DCHi\_Controlled\_by\_DRAC} TBSsize_i / TTI_i < MaximumBitRate \tag{7.1}$$

After the first SIB 10 has been received, the UE starts the following process:

1. At the start of the next TTI, the UE will randomly select p, $0 \leq p \leq 1$.
2. If $p <$ Transmission_Probability parameter, then the UE will transmit on the DCH controlled by DRAC during T_{validity} frames using the last stored allowed subset of TFCS, and then returns to step 1. Otherwise the UE will stop transmission on the DCH during T_{retry} frames and then return to step 1.

Transmission time validity (T_{validity}) and time duration before retry (T_{retry}) are indicated to the UE at the establishment of a DCH controlled by this procedure, which may be changed through radio bearer or transport

channel reconfiguration. The UE will always use the latest received DRAC static parameters.

Most probably this scheme is used with bearers that do not have very strict real-time requirements. It is an ideal method to be used with relatively high data rate services with lax delay requirements. The UTRAN can use DRAC bearers to "fill" the free capacity in the uplink. It can measure the uplink data rates and interference, and once there seems to be unused capacity, it can immediately fill that using DRAC managed bearers. The DRAC scheme can modify the bearer data rate even in every frame if necessary. The allowed values for the T_{validity} parameter are 1 to 256 frames.

Note that DRAC support requires simultaneous reception of the SCCPCH and the DPCH channels. DRAC is only applicable when the UE is in the CELL_DCH state. DRAC is also only applicable for FDD systems.

See also Section 14.8 in [9], Chapter 8 in [14], and Section 6.2.5 in [6]. However, I do not believe that DRAC will be the first feature to be implemented in 3GPP networks because most probably the uplink will not be the bottleneck that needs enhancing.

7.5.2.20 Contention Resolution

The RRC must store the value of initial UE identity encapsulated in the RRC connection request message and check that it receives the same value back in the RRC connection setup. It is possible that several UEs have sent on the RACH at the same time; thus, the response message might not be addressed to this UE. This procedure is known as *contention resolution*. If the initial UE identity is the same as the sent value, the UE can continue with the connection setup. Otherwise, the connection establishment must be started again.

7.5.2.21 Timing Advance

This functionality is only supported in the TDD mode. It is used to avoid large variations in signal arrival times at the Node B. Each radio frame in the TDD mode is divided into 15 time slots, which are further allocated either to the uplink or to the downlink. The transmissions from UEs in the cell must arrive at the Node B during their respective time slots. This means that UEs further away from the Node B have to transmit earlier than the nearby UEs so that both transmissions arrive at the Node B at the expected times.

The UTRAN may adjust the UE's transmission timing with timing advance. The initial value for timing advance will be acquired by the UTRAN from the timing of the PRACH transmission. The required timing advance will be given as a 6-bit number (the max value is 63) being the multiple of 4 chips, which is nearest to the required timing advance.

The RRC controls the timing-advance mechanisms. The UTRAN can send the IE UL timing advance with various configuration messages or within the uplink physical channel control message. The UTRAN will continuously measure the timing of transmissions from the UE and send the necessary timing-advance value as an adjustment. With the receipt of this IE, the UE adjusts the timing of its transmissions accordingly in steps of ±4 chips.

The timing-advance mechanism is also needed in the TDD mode's HO procedure. If a TDD-to-TDD HO takes place, the UE transmision in the new cell is adjusted by the relative timing difference, Δt, between the new and the old cell:

$$TA_{new} = TA_{old} + 2\Delta t \tag{7.2}$$

If UL synchronization is used (its support is optional for the UE), the timing advance is subchip granular and with high accuracy in order to enable synchronous CDMA in the UL. The functionality is otherwise similar to the nonsynchronized case, except that the units in the adjustment command represent multiples of ¼ chips.

7.5.2.22 Support for Cell Broadcast Service

Cell broadcast messages are text messages that are broadcast to everybody in a cell. These messages can be received by all mobiles capable of receiving cell broadcast service (CBS) and that are either in idle, CELL_PCH, or URA_PCH states. The user can choose which types of messages will be displayed and which will be discarded based on the message class type.

CBS is described in Section 7.11. Most of the CBS functionality will be put into the BMC task, but the RRC task must handle some configuration and allocation functions for the CBS. This service is a broadcast service; that is, only the downlink direction is used for message delivery. Thus, the supporting functionality in the RRC is different in the UE and the UTRAN.

Initial Configuration for CBS

This function performs the initial configuration of the BMC sublayer. The configuration is delivered to the UE via the broadcast system information; that is, it is received by the UE's RRC. This information is then used to configure the BMC and the L1 so that they can handle received CBS messages.

Allocation of Radio Resources for CBS

This functionality belongs to the UTRAN-RRC only. It allocates radio resources for CBS based on traffic volume requirements indicated by the BMC. The more queued CBS messages in the BMC buffers waiting to be

sent, the more resources should be allocated for the CBS. The radio resource allocation set by the RRC (i.e., the schedule for mapping of the CTCH onto the FACH/S-CCPCH), is indicated to the BMC to enable the generation of schedule messages. The resource allocation for CBS is broadcast as system information.

Configuration for CBS Discontinuous Reception
CBS messages can only be received while there is no active communication between the UE and the UTRAN; that is, in the idle, CELL_PCH, and URA_PCH states. In these states it is important that the UE saves as much power as possible to increase its standby time.

Therefore, the UTRAN only sends CBS messages during predefined times. The UE knows this schedule, and it only has to be prepared to receive CBS messages during those times. This function configures the lower layers of the UE when it will listen to the resources allocated for CBS based on scheduling information received from the BMC.

7.5.2.23 Capability Information

There will likely and hopefully be a wide variety of different types of UEs; multimedia applications imply a wide variety of terminals and appliances. They will have different capabilities, and the network must know the capabilities of a given UE before it can decide what kind of services and resources it can offer to any particular UE.

This information is typically sent to the network if the capabilities of a UE change. This may occur, for example, when a handheld UE is connected to a car kit; thus, its power class changes. The network may also require the UE to send along its capability information during the RRC connection setup procedure. The request can be sent in the RRC connection setup message and the response (the capability information) is included in the RRC connection setup complete message.

A typical message flow in this procedure consists of a UE capability information sent by the UE and a UE capability information confirm sent back to the UE from the network. The procedure can also be initiated by the network by sending a UE capability enquiry. This message triggers the UE to send its UE capability information (see Figure 7.22).

See also Sections 8.1.6 and 8.1.7 in [9] and Section 6.7.1 in [6].

7.6 RRC Protocol States

In GSM as in many other 2G systems, the radio resource protocol states were generally divided into two groups: the idle and the connected states. In

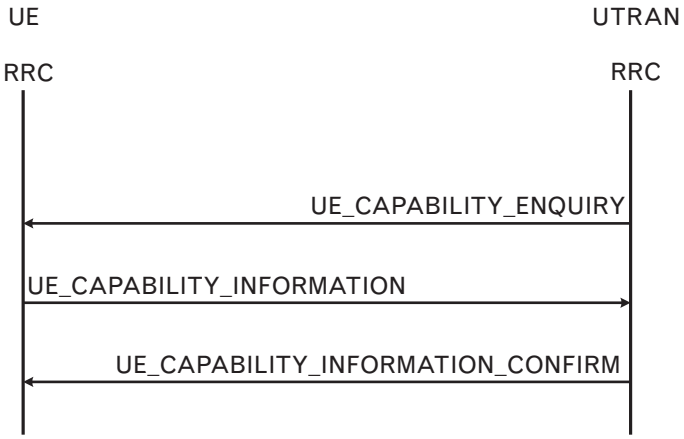

FIGURE 7.22
UE capability information.

the idle state no dedicated radio resources existed between the UE and the base station. We should observe, however, that the idle state is a rather poor name, as the mobile station is far from being "idle." There are several idle-mode tasks it must handle, tasks such as neighbor cell monitoring, cell reselection, paging channel reception, and broadcast data reception. In the connected state, however, a duplex radio connection is in place. The boundary between the idle and the connected mode is pretty clear; it is the existence of a dedicated radio resource. But in the new UTRAN system, this division is blurred.

The idle state in UMTS is similar to GSM, as well as to those we find in other 2G systems: There is no uplink connection whatsoever. The UE has to monitor its radio environment regularly and, when necessary, perform a cell-reselection task. The reception of the broadcast system information and paging messages belong to the UE's idle-mode tasks.

The connected state is different from the corresponding state in circuit-switched 2G systems, but it has similarities with the packet-switched GPRS system. The connected mode is divided into four states (see Figure 7.23). In the connected state there exists a logical RRC level connection between the UE and the UTRAN, but not necessarily a dedicated physical connection.

1. CELL_DCH is a state in which a dedicated connection exists in both directions. This state is entered while an RRC connection is established with dedicated channels, and it is abandoned when the connection is released. This state is comparable to dedicated mode in the 2G circuit-switched networks.

2. CELL_FACH is a state in which there are no dedicated connections, but data can still be transferred. This data transmission is done via common channels. This feature is very useful if the amount of data transferred is small or it is bursty. The use of a common channel

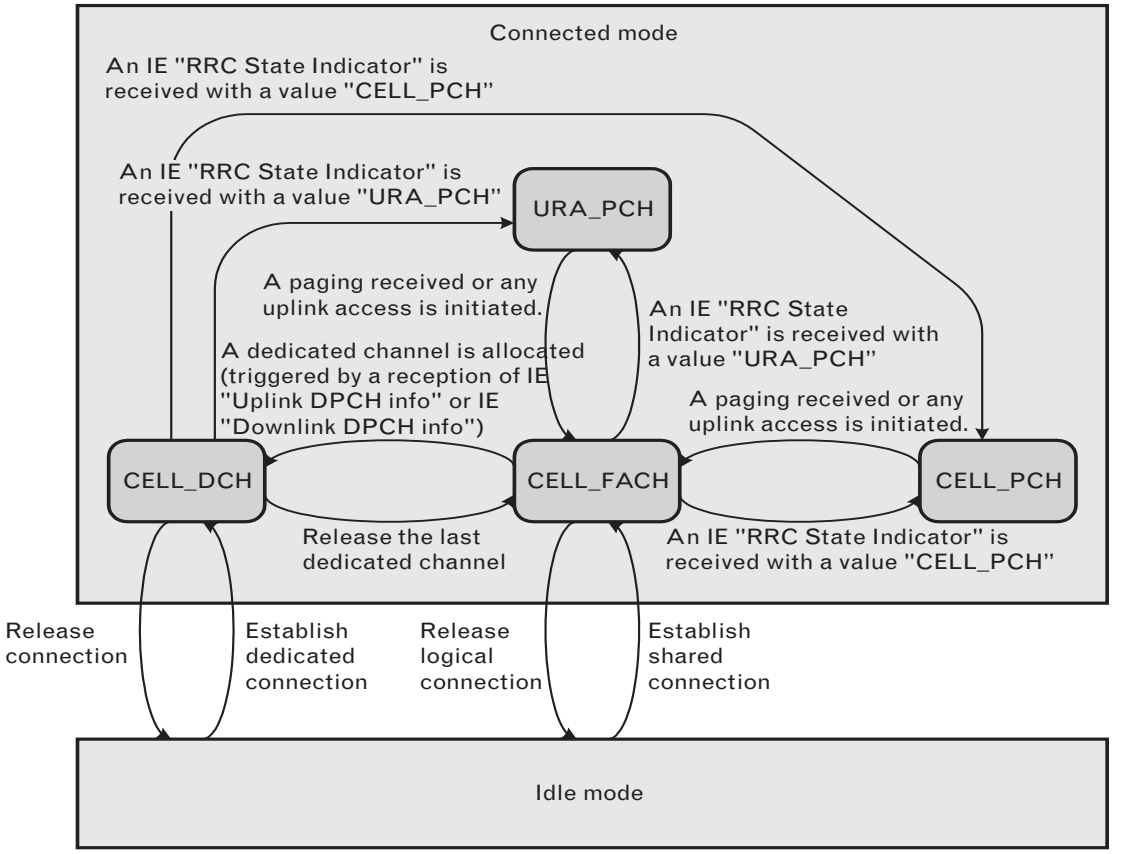

FIGURE 7.23 *RRC protocol states.*

preserves the radio resources in the cell. In the uplink direction, small data packets and control signals can be sent on a RACH or on a CPCH. In the downlink direction, the FACH can be used. In the TDD mode, the USCH and the DSCH can be used. The amount of data transmitted is monitored, and if necessary, dedicated resources can be allocated followed by a state change to the CELL_DCH state.

Because the CELL_FACH state requires the mobile to monitor the FACH channel, it consumes power, which is a scarce resource in handheld equipment. Therefore, if there is no data-transmission activity for a certain time, the RRC moves from the CELL_FACH state to the CELL_PCH state.

3. CELL_PCH state is much like the idle mode because only the PICH is monitored regularly. The broadcast data (i.e., the system

information and cell broadcast messages) are also received. The difference is that the RRC connection still exists logically in the CELL_ PCH state. The RRC moves back to the CELL_FACH state if any uplink access is initiated, or if a paging message is received.

Note that in order for the RRC to move from the CELL_PCH to the idle mode, it must first go to the CELL_FACH state so that connection release messages can be exchanged. If the UE makes a cell reselection while in the CELL_PCH state, it must inform the UTRAN about this. This also requires a temporary cell change to the CELL_FACH state. No uplink activity is possible in the CELL_PCH state itself.

4. URA_PCH is quite similar to the CELL_PCH state, except that every cell change does not trigger a cell-update procedure. In this state an update procedure is only initiated if a UTRAN registration area changes, which is not done with every cell reselection. A state change to this state is requested by the UTRAN if it sees that the activity level of the UE is very low. The purpose of this state is to reduce the signaling activity because of cell updates. The drawback of this arrangement is that if the UTRAN wants to initiate data transmission while the RRC is in this state, it has to expand the paging area from one cell to several cells, possibly to the whole registration area because the location of the UE is not known with great accuracy.

Note that the UTRAN registration area (URA) is a different concept from that of the CN GPRS routing area. The various location concepts in 3G (both in the core network and in the UTRAN) are further discussed in Section 7.7.

We should notice that the CELL_PCH state is actually a subset of the URA_PCH state. As discussed in Section 7.7, it is possible to define overlapping URAs to be used in the URA_PCH state. Thus, the UTRAN operator could define that each cell is a separate URA in addition to other larger URAs. Then the operator could assign small one-cell URAs for slow-moving mobiles, and larger URAs for mobiles with greater mobility. The small URAs could nicely perform the task of the CELL_PCH state. However, it has been decided to keep these states separate.

Generally, the state changes between these states are controlled by the UTRAN, but not always. For example, if a UE that is in the CELL_PCH or URA_PCH state wants to initiate a mobile-originated call, it moves to the CELL_FACH state before initiating the RACH procedure.

The higher layers in NAS, or applications, do not need to know about these states, as they are internal to RRC. The NAS only needs to known if

the AS is in connected or idle mode. If the AS is connected, the NAS can always send data through it. If the internal RRC state happens to be, for instance, CELL_PCH, then RRC will move to CELL_FACH state and send the data via common channels or set up a dedicated connection.

7.7 Location Management in UTRAN

The cell-reselection procedure may also trigger the cell-update procedure. This procedure is used to update the UTRAN registers about the location of the UE. Note that in 3G, the mobility concept is handled separately in both the UTRAN and the CN. This is because the UTRAN and the CN are two separate logical entities, and at least in principle, the CN does not know what kind of access network is connected to it and vice versa. They cannot make use of the location concepts of each other because they do not know about them.

The CN level of mobility is handled with MM tasks. There are two location concepts at the CN level. Location areas are used by the circuit-switched network and routing areas by the packet-switched network. The location area of a UE is stored in the MSC/VLR, where it is used to route the paging messages to the right area. The routing area information of a UE is stored in the SGSN, where it is used for packet-switched paging. If a UE crosses a location area/routing area border, it initiates a location area/routing area update toward the CN. Optionally, the network may also demand that these registrations be done periodically. A successful registration will be acknowledged by the network, which may at the same time issue a new temporary mobile subscriber identity (TMSI) to the UE. If the CN has to page the UE, it will usually do so by referring to its recently assigned alias: the TMSI.

At the UTRAN level, the location concepts are the registration area and the cell area, which are independent of the CN location concepts. The UTRAN location concepts are only valid and used when the UE is in the RRC connected mode. They are also only visible within the UTRAN (see Figure 7.24).

As stated, the UTRAN-level location concepts are only maintained and used while the UE is in the RRC connected mode. For those only accustomed to circuit-switched networks, this may sound a bit strange. Why do we need location management if there is already a connection? In the WCDMA packet-based radio interface, the concept of "connected" is a bit different from that of the circuit-switched world. Here the UE may be in a connected state, even though it does not have a dedicated channel. Common transport channels may be used for data transfer if the data is sporadic or low in volume. As there is no dedicated channel, the network does not

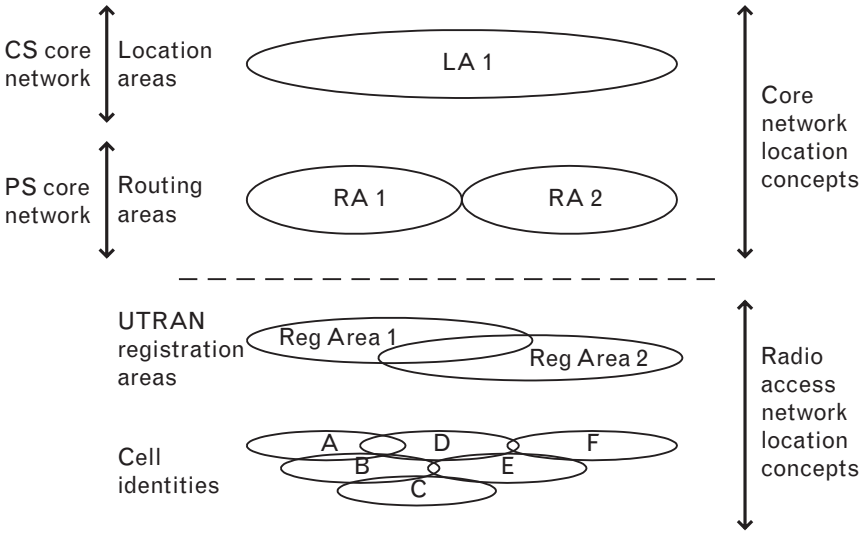

FIGURE 7.24
Location concepts.

directly know about the UE's movements. Therefore, the UE must inform the UTRAN if its location changes while it is in the RRC connected mode without a DCH. If the UE is in the CELL_FACH or CELL_PCH substates of the connected mode, it must inform the UTRAN of every cell change. This procedure is called the cell update (see Figure 7.25). If the UE has higher mobility, then the UTRAN may order the UE to the URA_PCH substate, in which the UE initiates a location registration only when it moves to a cell that belongs to another URA. This reduces the signaling overhead caused by the cell-area updates. The drawback is of course that a paging message may have to be sent to the whole URA, that is, to several cells. Both the cell-area and URA-update procedures can also be done periodically if the UTRAN so orders. Periodic updates are good in that they will remove the "ghost" users from the register of active users. Normally, when a UE is being switched-off, it will inform the UTRAN about this so that UTRAN can remove the UE from the register of active users, and, for

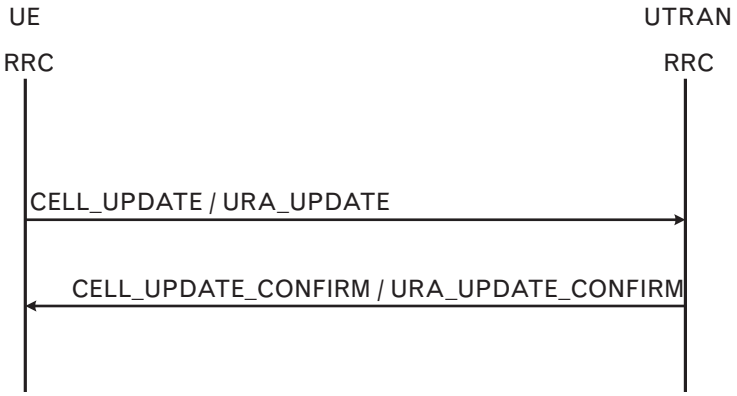

FIGURE 7.25
Cell/URA update.

example, in case of an incoming call, no paging is attempted. However, if the UE is outside the network coverage (for example, inside a thick-walled buiding) when it is switched off, then it cannot send the detach indication, and the UTRAN still assumes it to be active. Periodic updates remove this problem: If a UE does not send its scheduled location update message, it can be assumed to be switched-off.

The URAs can be overlapping or even hierarchical. The same cell may belong to several different URAs, and the UEs in that cell may have been registered to different URAs. SIB 2 contains a list of URA identities indicating which URAs this cell belongs to. This arrangement is done to further reduce the amount of location update signaling because now the UEs moving back and forth in the boundary area of two URAs do not have to update their URA location information if the boundary cells do belong to both URAs.

For example, in Figure 7.26 (left) users A, B, and C constantly cross the URA boundary, triggering URA update procedures. However, if URA areas are overlapping, such as in Figure 7.26 (right), then no update procedures are needed, if user A is assigned URA ID = 44 (or 45), user B is given URA ID = 45, and user C is given URA ID = 46.

As mentioned earlier, in addition to overlapping URA areas, there can also be hierarchical URA structures. One cell can have up to eight different URA identities.

The UTRAN location-registration procedures may also include the reallocation of the temporary identity of the UE. This identity can be included in the cell/URA update confirm message, and it is called the radio

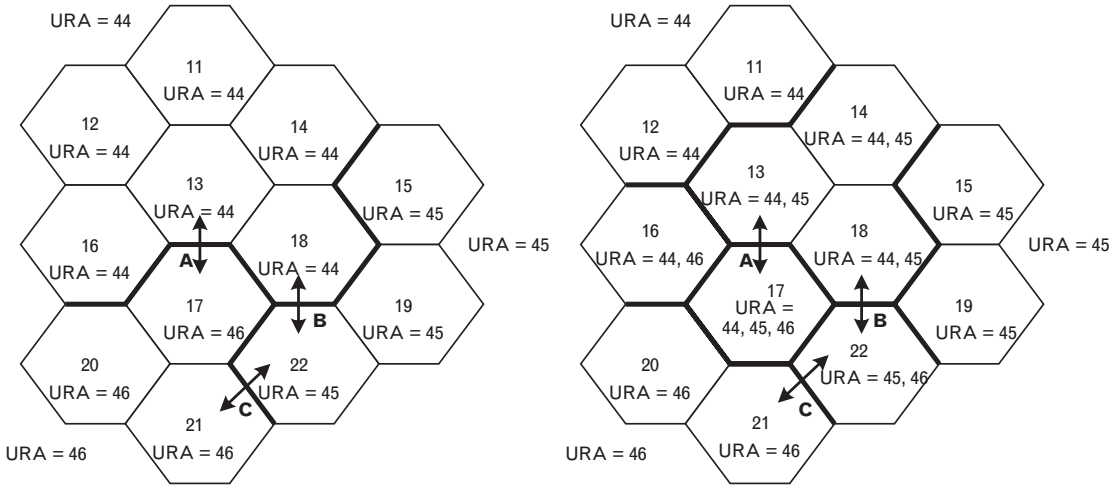

FIGURE 7.26 *Overlapping URA areas.*

network temporary identifier (RNTI). The RNTI is used to address the UE on common transport channels. If dedicated channels are used, then there is no need for an RNTI. The paging messages received in the RRC connected mode normally refer to the UE by using its RNTI. There are two variations of the RNTI: (1) cell RNTI (C-RNTI) and (2) UTRAN RNTI (U-RNTI).

The C-RNTI identifies a UE within a cell, so it can only be used in paging messages when the UE's location is known (i.e., it must be in the CELL_PCH state). This also implies that a new C-RNTI will be allocated to a UE every time it moves to a new cell and conducts a cell-update procedure. The U-RNTI is a UTRAN-wide identity that is used by the UTRAN for paging if it knows that the UE is in the URA-PCH substate (i.e., its location is only known at the URA level). Note that U-RNTI is not a URA-specific identity, but it identifies a UE within the whole UTRAN. Therefore, it can be optionally used for CN-originated paging [15].

The location concepts in the CN and in the UTRAN are not connected in any way. The operator can freely define them independently. Of course, this means that a routing area and a registration area could be defined to be the same, but this is not the intention of the specifications.

7.8 Core Network Protocols in the Air Interface

Only a short overview will be given of the CN protocols in the air interface. This is because these protocols already exist in the GSM/ GPRS system and they can be studied in other sources. A classic GSM reference book is [2]. Other useful GSM publications include [3–5].

7.8.1 Circuit-Switched Core Network

7.8.1.1 Mobility Management

As the name of this task states, one of the main functions of the MM task is location management. But the MM task has also been assigned network-registration and security functions. Typically, the MM procedures can be divided into three groups:

1. MM common procedures;
2. MM specific procedures;
3. MM connection-management procedures.

The MM common procedures can always be initiated while an RRC connection exists. The procedures belonging to this type fall into two classes determined by what entity initiates them:

Initiated by the Network:

- TMSI reallocation procedure;
- Authentication procedure;
- Identification procedure;
- MM information procedure;
- Abort procedure.

Initiated by the Mobile Station:

- IMSI detach procedure.

An MM-specific procedure can only be initiated if no other MM-specific procedure is running or no MM connection exists. The procedures belonging to this type include:

- Normal location updating procedure;
- Periodic updating procedure;
- IMSI attach procedure.

The MM connection-management procedures are used to establish, maintain, and release an MM connection between the mobile station and the network over which an entity of the upper CM layer can exchange information with its peer. An MM connection establishment can only be performed if no MM-specific procedure is running. More than one MM connection may be active at the same time.

In the following a short description of each procedure from the previous list is given.

The TMSI reallocation procedure provides identity confidentiality, that is, protects a user against being identified and located by an intruder. TMSI stands for temporary mobile subscriber identity. It can be used instead of globally unique IMSI to identify a user to the network. Usually the TMSI reallocation is performed at least at each change of a location area.

The authentication procedure permits the network to check whether the identity provided by the mobile station is acceptable or not. The network sends the UE two parameters, RAND and AUTN, and from these and its own secret parameters, it calculates a response parameter RES. The details of this procedure are explained in [13]. Authentication procedure

also provides parameters enabling the mobile station to calculate new UMTS ciphering and integrity keys. The UMTS authentication procedure is always initiated and controlled by the network.

The identification procedure is used by the network to request a mobile station to provide specific identification parameters to the network. This parameter can be IMSI, IMEI, IMEISV, or TMSI.

MM information procedure can be used to provide the mobile station with subscriber specific information. In practice this information includes various time-zone and clock information.

The abort procedure may be invoked by the network to abort any ongoing MM connection establishment or already established MM connection.

The IMSI detach procedure is performed by the UE if it is being deactivated or if the SIM is detached from the UE. The detach procedure is optional, and a flag (ATT) broadcast in SIB 1 message on the BCCH is used by the network to indicate whether the detach procedure is required. The procedure causes the mobile station to be indicated as inactive in the network.

The normal location updating procedure is used to update the location area (LA) information in the network registers. This procedure is triggered when in the idle mode the UE selects a cell that belongs to a different LA than the previous cell. Periodic updating procedure is similar to the normal location updating procedure, except that it will be triggered periodically by a timer. It is used to notify the network about the availability of the UE.

The IMSI attach procedure is used to indicate the UE as active in the network. Typically, this happens when the UE is switched-on. IMSI attach is an optional procedure used whenever IMSI detach is used. This is indicated by a SIB 1 message. IMSI attach employs the normal location update procedure signaling; only the location updating type field in the message indicates that this is an IMSI attach.

The MM protocol is defined in [16].

7.8.1.2 GPRS Mobility Management

The GPRS mobility management (GMM) sublayer provides services to the session management (SM) entity and to the SMS (SMS) support entity for message transfer.

Depending on how they can be initiated, two types of GMM procedures can be distinguished:

1. GMM common procedures:

2. GMM-specific procedures.

GMM common procedures are initiated by the network when a GMM context has been established:

- P-TMSI (re)allocation;
- GPRS authentication and ciphering;
- GPRS identification;
- GPRS information.

GMM-specific procedures can be initiated either by the network:

- GPRS detach.

Or they can be initiated by the UE:

- GPRS attach and combined GPRS attach;
- GPRS detach and combined GPRS detach;
- Normal routing-area updating and combined routing-area updating;
- Periodic routing-area updating;
- Service request.

These procedures are very similar to the corresponding MM procedures described earlier. The main difference is that whereas MM procedures are used between a UE and a circuit-switched core network, GMM procedures are used between a UE and a packet-switched core network.

The GMM protocol is defined in [16]. In practice this protocol is considered to be an extension of the MM protocol and can be implemented within the same protocol entity.

7.8.1.3 Call Control

The call control (CC) protocol is one of several protocols in the connection management (CM) sublayer. This protocol includes the control functions for the call establishment and release.

A CC entity must support the following elementary procedures:

- Call-establishment procedures;
- Call-clearing procedures;
- Call-information-phase procedures;
- Miscellaneous procedures.

A call can be either a mobile-originated call (MOC) or a mobile-terminated call (MTC); that is, it can be initiated by either the mobile or by the network. Optionally the UE can also support a network-initiated MOC. This functionality can be used with the completion of calls to busy subscriber (CCBS) supplementary service.

The call-clearing procedure can be initiated either by the UE or by the network. Note, however, that this means the logical CC-level connection clearing. The actual radio connection (RRC level) is always released by the UE. A radio connection and a CC connection are separate concepts. One can use the radio connection for many other things besides the circuit-switched call, such as SMS and for packet-data applications. Therefore, releasing a call connection does not necessarily mean that the radio connection should also be released. There may be other applications that still need the radio connection.

While the call is active, the CC can perform various procedures. The user-notification procedure informs the user about call-related events, such as user suspension or resume. Support of multimedia calls will be an important procedure especially in UMTS. The dual-tone multifrequency (DTMF) control procedure enables the user to send DTMF tones toward the network. Key presses in the UE containing digit values (0–9, A, B, C, D, *, #) are signaled over the air interface to the MSC, which converts them into DTMF tones and sends them onward to the remote user. Typical applications of the DTMF include various automated information services (e.g., telephone banking: "Press 1 if you want to hear your bank account balance; Press 2 if you want to settle your bills; Press 3 if you want to talk to the operator," and so forth). The support of DTMF is described in [17]. The support for the in-call modification procedure is optional for the UE. This procedure means that the same connection can be used for different kinds of information transfer during the same call, but not at the same time. In practice, this procedure is used for alternating the call between speech and fax services or between speech and data.

Miscellaneous CC procedures include in-band tones and announcements, status inquiry, and call reestablishment. The in-band tones and announcements procedure is used when the network wants to make the mobile station attach the user connection (e.g., in order to provide in-band tones/announcement) before the UE has reached the "active" state of a call. In this case, the network may include a progress indicator (IE) indicating user attachment in a suitable CC message. The status-inquiry procedure can be used to inquire about the status of the peer entity CC. This is a useful procedure in error handling. The call-reestablishment procedure is mostly an RRC-layer matter, as it involves setting up a new radio connection in place of the lost one. Within the CC level, however, this procedure includes provisions for the UE to make a decision as to whether a reestablishment should be attempted. The network-side CC must also identify and

resolve any call states or an auxiliary state mismatch between the network and the UE.

The CC protocol is defined in [16].

7.8.1.4 Supplementary Services

Supplementary services (SS) are value-added services that may or may not be provided by the network operator. The list of various GSM supplementary services is long and ever increasing. These include, for example, the advice of charge (AoC), call forwarding (CF), and call waiting (CW) supplementary services. Because these services belong to the NAS, they are applicable to both GSM and the UMTS. It is likely that later on there will also be UMTS-only supplementary services.

One generic protocol is defined for the control of SS at the radio interface. It is based on the use of the facility information element or the facility message. The exact functionality triggered by this information element or message depends on the information it contains.

SSs are discussed further in Section 13.4. The SS protocol is defined in [18] and in 3G TS 24.08x and the TS 24.09x series of specifications.

7.8.1.5 Short Message Service

The purpose of the SMS is to provide a means to transfer short text messages between a UE and a short message service center (SMSC). These messages are sent using the control signaling resources, and their maximum length can be only 160 characters.[3] SMS is a non-real-time service; a store-and-forward service in which messages can be stored on the SMSC and delivered when the destination UE is available.

The term *SMS-MO* refers to a mobile-originated SMS message; *SMS-MT* refers to a mobile-terminated SMS message.

Note that a UTRAN 3G network will also include an enhanced version of the SMS called the multimedia messaging service (MMS); see Section 12.4.

The SMS protocol is defined in [19].

7.8.2 Packet-Switched Core Network

7.8.2.1 Session Management

The main function of the SM protocol is to support packet data protocol (PDP) context handling of the user terminal. Note that there is no "connection" concept in a (IP) packet-switched system as we know it in a circuit-

3 Nowadays users can send longer than 160-character text messages with their mobiles, but technically those are divided into several SMS messages, each with a maximum of 160 characters and sent separately over the air interface.

switched system. However, the communicating entities do need to know about the characteristics of the data to be transferred. This task is performed by the PDP context-activation procedure. Other functions this task must perform include PDP deactivation and PDP modification.

The SM procedures for identified access can only be performed if a GMM context has already been established between the UE and the network. If no GMM context has been established, the MM sublayer must initiate the establishment of a GMM context by use of the GMM procedures. After GMM context establishment, the SM uses services offered by GMM. Ongoing SM procedures are suspended during GMM procedure execution.

The SM protocol is defined in [16].

7.8.2.2 GPRS Short Message Service Support

The GPRS Short Message Service (GSMS) protocol task handles the SMS service while the UE is attached to the PS CN; that is, to the GPRS system. In practice this protocol is an extension of the circuit-switched SMS protocol, and both will typically be implemented within one protocol task entity. See [19] for further information.

7.9 User Plane

The lower layers of the U-plane are exactly the same as those of the C-plane (MAC and RLC). PDCP and BMC, however, exist only in the U-plane. The control of all the AS U-plane tasks is handled by the RRC. The U-plane is responsible for the transfer of user data, such as voice or application data, whereas the C-plane handles the control signaling and the overall resource management (see Figure 7.27).

7.10 Packet Data Convergence Protocol

As the name implies, the PDCP task is a convergence layer between the actual data protocol in the NAS and the radio access protocols in layer 2 (Figure 7.28). The PDCP itself is an AS protocol. This protocol entity is only used in the U-plane. The required control signaling for the PDCP is handled by the RRC. The PDCP handles the same functionality in the UTRAN as the SNDCP task does in the GPRS system.

The network layer in the NAS can accommodate several different data protocols. These current (and future) protocols must be transferred transparently over the UTRAN. This is the task of the PDCP, which must

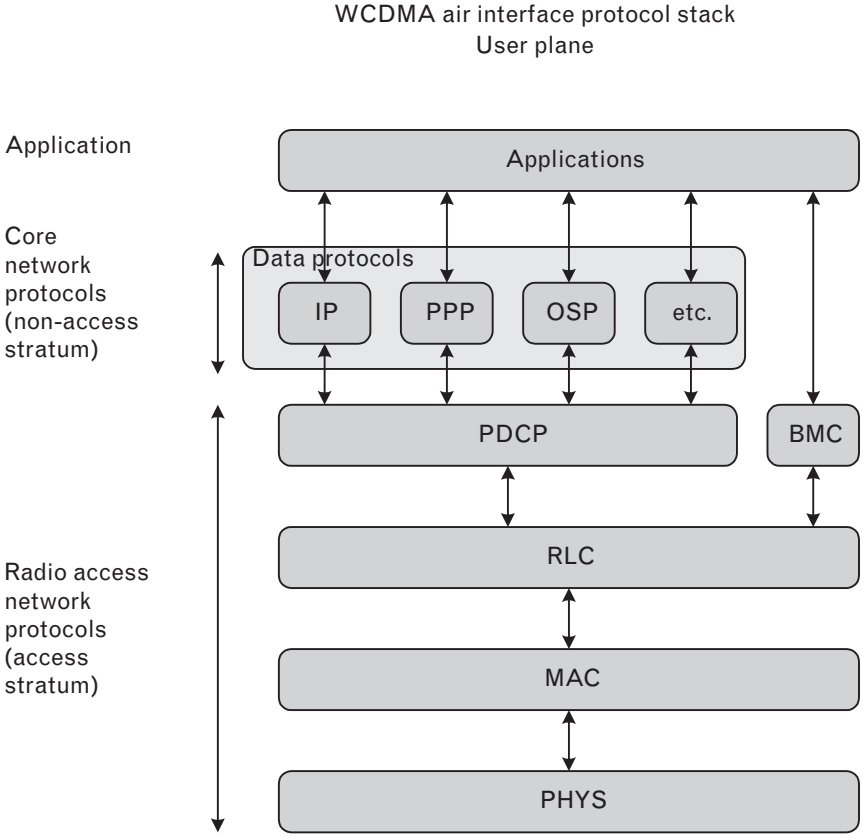

Figure 7.27
WCDMA U-plane protocol stack.

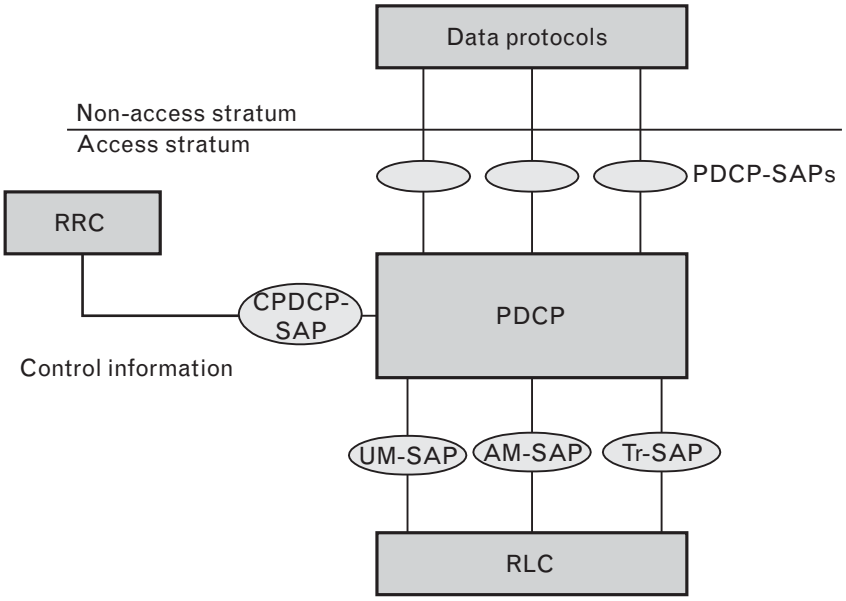

Figure 7.28
PDCP model.

hide the particularities of each protocol from the UTRAN. The packets from all of these protocols will be conveyed over the UTRAN without any changes to the UTRAN protocols.

Therefore, the functions the PDCP shall perform include the following:

- Header compression and decompression of IP data streams;
- Transfer of user data;
- Maintenance of PDCP sequence numbering;

Header compression and decompression are performed by the PDCP to optimize the channel efficiency in the radio interface. The network data protocols are not especially designed for wireless environments; thus, they may have unnecessarily large header fields in their data packets. It is the task of the PDCP to compress these headers to more compact representations. Each data protocol has its own header format, so the PDCP must accommodate different compression algorithms. The particular algorithms and parameters are negotiated by the RRC protocol task, which indicates the result to the PDCP.

The transfer of user data includes forwarding the NAS data to the RLC layer and vice versa. Note that if acknowledged transfer mode is used in the RLC, then buffering of N-PDUs received from NAS is needed. They must be stored until the peer entity RLC acknowledges that they have been successfully sent.

The maintenance of PDCP sequence numbering is used in the SRNS relocation procedure only. If the SRNS changes in the UTRAN and AM is used in RLC, then an orderly continuation of data transfer requires that the PDCP sequence numbering information must be are exchanged between the UE and the UTRAN. This is called a "lossless SRNS relocation." The RRC may also command that a lossless SRNS relocation will not be supported.

The PDCP protocol is specified in [20].

7.11 Broadcast/Multicast Control

Broadcast/multicast control is a layer 2 sublayer that exists only in the U-plane. The necessary control information is received from the RRC, just as in the PDCP sublayer. This layer handles only downlink broadcast/multicast transmission (see Figure 7.29).

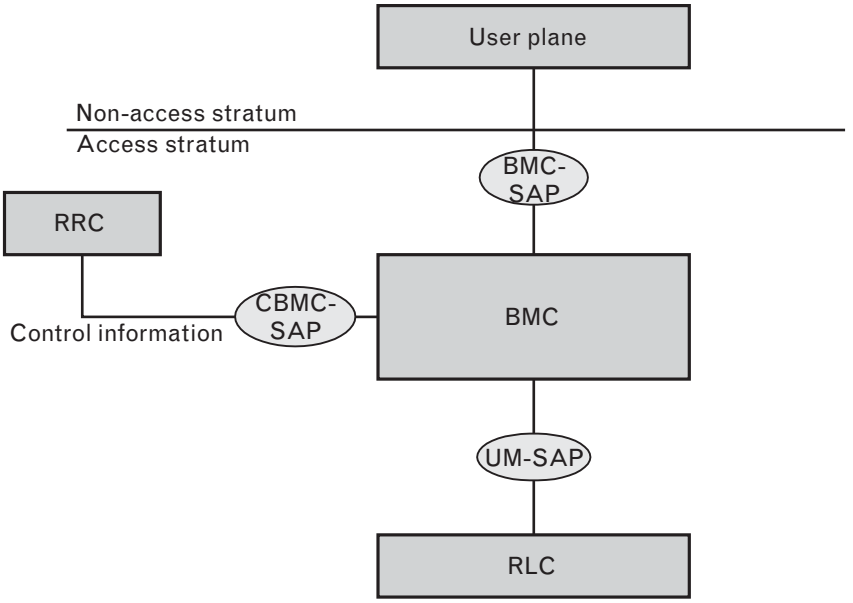

FIGURE 7.29
BMC protocol model.

The BMC task implements the transfer of cell broadcast messages. Cell broadcast service (CBS) is not a UMTS-only service because it is also used in GSM, though not yet very widely. This service is specified for UMTS in [21]. Cell broadcast messages are SMS text messages (although they can be much longer than the normal SMS messages) that are broadcast to everybody in a cell (or in a set of cells). A CBS message consists of CBS pages. One page contains 82 octets, and if 7-bit characters are used, it is possible to send 93 characters in a page. A CBS message may contain up to 15 pages, which gives a maximum size of 1,395 characters for CBS messages. These messages can be received by all mobiles capable of receiving CBS. But the mobiles have to be in either the idle, the CELL_PCH, or the URA_PCH states. Cell broadcast messages are assigned a message class type, which can be used by the UE to filter and receive only those messages that are of interest to it. The categories to be subscribed to could include contents, such as news, traffic information, and weather forecasts. The user can, of course, reject all cell broadcast messages.

The functions of BMC are specified in [22]. These include the following:

- Storage of cell broadcast messages;
- Traffic-volume-monitoring and radio resource requests for CBS;
- Scheduling of BMC messages;
- Transmission of BMC messages to UEs;
- Delivery of cell broadcast messages to the upper layer (NAS).

The BMC entity in the RNC is responsible for storing the cell broadcast messages to be sent. They cannot usually be sent further right after they have been received from the core network. Their transmission to the UE must be scheduled, and also typically be repeated several times. Thus, the UTRAN-BMC needs some storage space.

The BMC in the RNC must also estimate the expected amount of traffic volume that is required for transmission of queued CB messages. This is indicated to the RRC so that it can allocate the necessary radio resources.

The CB messages are scheduled to enhance the performance of the UEs receiving them. A UE can listen for dedicated CB scheduling messages, and from those, extract the scheduling of the actual information bearing CB messages type by type. Thus, a UE does not have to receive all CB messages but only those that it knows belong to categories it has subscribed to. Outside the scheduled message sending times, the UE can enter DRX mode to save power. The BMC entity in the RNC is responsible for building the schedule and sending this information in schedule messages. The BMC in the UE must receive these messages and then inform the RRC so that it knows when to listen for the actual CBS messages. The configuration of layer 1 is done by the RRC, not by the BMC.

The actual transmission of the CB messages is done according to the defined schedule, and the BMC in the UE should forward to upper layers only those messages belonging to subscribed groups. The BMC also has to compare the message IDs and serial numbers of the received messages to the IDs and the numbers already received and stored. If they are identical, then the received message can be discarded.

We can see from this description that the BMC task in the UE is rather simple, but more involved in the RNC, which has many more functions to handle.

The broadcast/multicast protocol specification is in [22]. Broadcast services are further discussed in [21, 23].

7.12 Data Protocols

The PDCP layer connects to the standard data protocols in the NAS. Because these protocols are not specific to the 3G system, they are not discussed further here. The PDCP layer handles the header compression for these protocols because in some cases the size of the header would consume too large a share of the available transmission bandwidth.

The point-to-point protocol (PPP) is defined in [24, 25]. Both IPv4 and IPv6 will be supported by the PDCP. IPv4 is specified in RFC 791 and IPv6 in RFC 2460.

7.13 Dual-System Protocol Stack in UE

Figure 7.30 describes one possible implementation of the dual-system (GSM-3G) protocol stack in the UE.

As can be seen, the 3G UTRAN protocol stack is quite separate from the GSM-GPRS protocol stack in the AS level. Code reuse in the radio access protocols is not possible, but on the other hand, this kind of separation makes the implementation of new protocol tasks easier; the difficult dual-system issues do not have to be addressed in every line of new code. Once again, notice that the RLC/MAC in 3G and the RLC/MAC in GPRS are not the same protocols.

The NAS [i.e., core network protocols (MM and CM)] are, however, similar in both systems and they can be reused. Both the MM and CM layers from the GSM system will require some small modifications to accommodate the 3G radio access protocols below them. The GSM core network protocols will be upgraded to also support the UMTS features, at the same

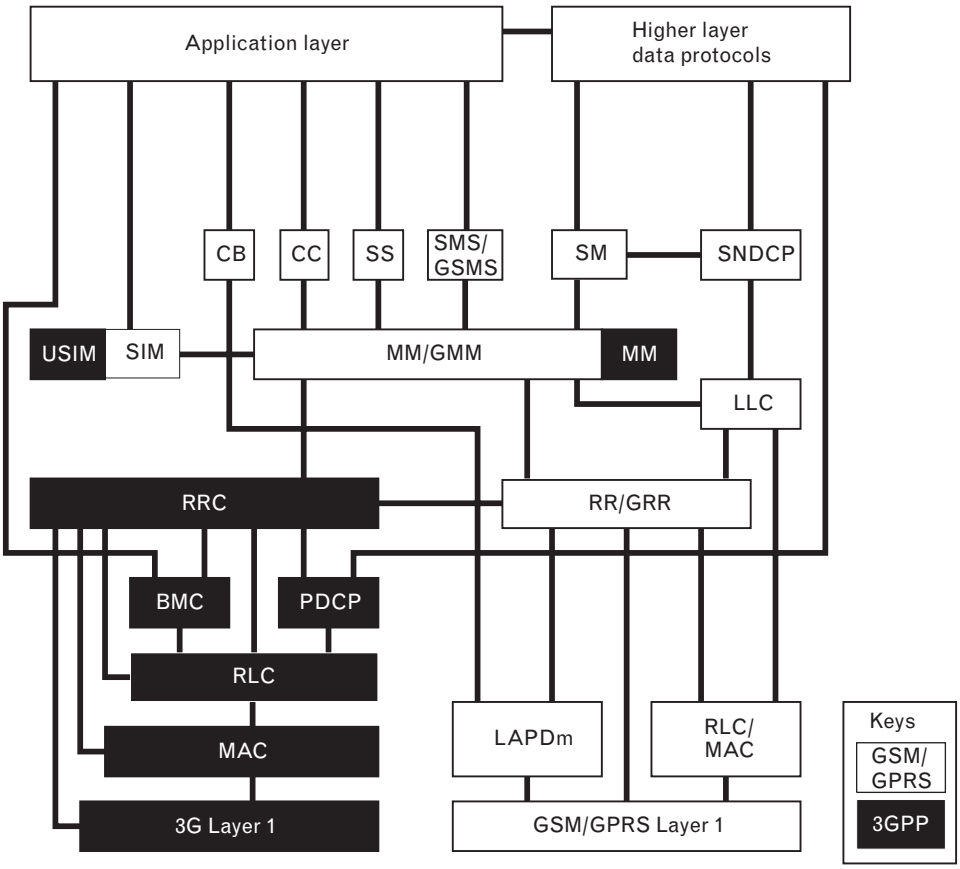

FIGURE 7.30 *Dual-system GSM/GPRS-3G (UTRAN) protocol stack.*

time retaining their GSM compatibility. The same core network can then support both the GSM and UMTS RANs. This means that these core network protocols must also be similar in the mobile station. The 3GPP will also handle the future specification work for the GSM core network protocols, so this makes the goal of common NAS protocols for both GSM and 3G an easier task.

References

[1] Heine, G., *GPRS from A–Z*, Norwood, MA: Artech House, 2000.

[2] Mouly, M., and M.-B. Pautet, *The GSM System for Mobile Communications*, published by the authors, 1992.

[3] Walke, B., *Mobile Radio Networks*, New York: Wiley, 1999.

[4] Mehrotra, A., *GSM System Engineering*, Norwood, MA: Artech House, 1997.

[5] Redl, S., M. Weber, and M. Oliphant, *An Introduction to GSM*, Norwood, MA: Artech House, 1995.

[6] 3GPP TS 25.303, v 5.0.0, Interlayer Procedures in Connected Mode, 2002.

[7] 3GPP TS 25.321, v 5.0.0, MAC Protocol Specification, 2002.

[8] 3GPP TS 25.322, v 5.0.0, RLC Protocol Specification, 2002.

[9] 3GPP TS 25.331, v 5.0.0, RRC Protocol Specification, 2002.

[10] 3GPP TS 25.304, v 5.0.0, UE Procedures in Idle Mode and Procedures for Cell Reselection in Connected Mode, 2002.

[11] 3GPP TS 23.122, v 4.1.0, NAS Functions Related to Mobile Station (MS) In Idle Mode, 2001.

[12] 3GPP TS 25.302, v 5.0.0, Services Provided by the Physical Layer, 2002.

[13] 3GPP TS 33.102, v 4.3.0, 3G Security; Security Architecture, 2001.

[14] 3GPP TR 25.922, v 5.0.0, Radio Resource Management Strategies, 2002.

[15] 3GPP TS 25.301, v 5.0.0, Radio Interface Protocol Architecture, 2002.

[16] 3GPP TS 24.008, v 5.3.0, Mobile Radio Interface Layer 3 Specification; Core Network Protocols-Stage 3, 2002.

[17] 3GPP TS 23.014, v 4.0.0, Support of Dual Tone Multi-Frequency (DTMF) Signaling, 2001.

[18] 3GPP TS 24.010, v 4.2.0, Mobile Radio Interface Layer 3; Supplementary Services Specification; General Aspects, 2001.

[19] 3GPP TS 24.011, v 4.1.0, Point-to-Point (PP) Short Message Service (SMS) Support on Mobile Radio Interface, 2002.

[20] 3GPP TS 25.323, v 5.0.0, Packet Data Convergence Protocol (PDCP) Specification, 2002.

[21] 3GPP TS 23.041, v 4.2.0, Technical Realization of Cell Broadcast Service (CBS), 2001.

[22] 3GPP TS 25.324, v 5.0.0, Broadcast/Multicast Control BMC, 2002.

[23] 3GPP TR 25.925, v 3.4.0, Radio Interface for Broadcast/Multicast Services, 2001.

[24] IETF RFC 1661 "The Point-to-Point Protocol (PPP)," W. Simpson (ed.), July 1994.

[25] IETF RFC 1662 "PPP in HDLC-Like Framing," W. Simpson (ed.), July 1994.