5

Energy-Efficient Base Stations

Alberto Conte

Alcatel-Lucent Bell Labs, Centre de Villarceaux, Nozay, France

5.1 Introduction

With the explosion of mobile Internet applications and the subsequent exponential increase of wireless data traffic, the energy consumption of cellular networks has rapidly caught the attention of the entire telecommunication community: industrials, operators, academics and government institutions. One of the first actions taken has been to monitor and understand where and by which cellular equipments the energy is consumed. Several studies have been conducted in parallel (e.g. [1–3]), and while the figures may slightly differs, all come to the same conclusion: whatever the technology is used (UMTS, HSPA and LTE), the major part of the energy ($\sim$50–60%) of a mobile network is consumed by the radio access network (RAN), and in particular by the set of base stations, followed by the core network ($\sim$30%), and data centres ($\sim$10%). The impact of the base stations comes from the combination of the power consumption of the equipment itself (up to 1500 W for a nowadays macro base station) multiplied by the number of deployed sites in a commercial network (e.g. more than 12000 in the United Kingdom for a single operator [3]). In order to effectively improve the energy efficiency of the future mobile networks, it is thus important to focus the attention on the base station. This chapter aims a providing a survey on the base stations functions and architectures, their energy consumption at component level, their possible improvements and the major problems that must be faced in order to make such improvement effective. When possible, examples from the 3GPP LTE technology are provided, in order to help the reader to map generic concepts to real-life networks. However, it is not the intention of this chapter to present in detail the LTE technology. The reader interested in an extensive presentation of this technology can refer to Ref. [4].

5.2 BS Architecture

5.2.1 *Generic Cellular Network Architecture*

The mobile networks are generically decomposed into core network (CN) and RAN. As concrete example, Figure 5.1 shows the functional architecture of LTE networks as specified by 3GPP; here, RAN is called evolved universal terrestrial radio access network (eUTRAN, [5]) and the core network is called Evolved Packet Core (EPC, [6, 7]).

The core network interfaces the mobile network with the "external" world, allowing the mobile terminals to access the public Internet, specific applications (e.g. VoIP), and private networks (e.g. corporate networks). At transport layer, the CN is in charge of providing the IP anchoring for the user equipment (and thus its IP address) such that the UE peers (e.g. a video server) are not impacted by the terminal mobility; all the data received from or destined to the UE goes through the anchor point (PDN Gateway in 3GPP LTE), whatever the position and speed of the UE. It is the job of the core network to track the movements of the UE and to forward its traffic to right position, in general using some type of IP tunnelling (e.g. 3GPP uses GTP [8]). Thanks to the anchoring function, the core network is also well suited to perform traffic control and shaping, allowing or denying the access to specific services in accordance with the subscriber profile. At control layer (MME, HSS and PCRF functions in 3GPP LTE), the core network in charge of managing the subscriber profiles and to perform authentication, authorization and accounting (AAA) of each user.

The core network does not handle any RF aspects specifics to the radio channel, its action stops at IP level (layer 3). It is the role of the RAN to take in charge all the radio specificities and to hide them from the core network. It handles the signalling and data exchanges with the mobile terminal (or User Equipment, UE) over the radio channel on one side, and interfaces with the core network, through the backhaul, on the other side. In modern architectures (e.g. LTE), the RAN is composed essentially by the base stations, as single equipment; other architectures, based on multiple elements, are also possible and have been used in the previous generation of cellular networks, for example in 3G/UMTS networks (with NodeB and RNC)
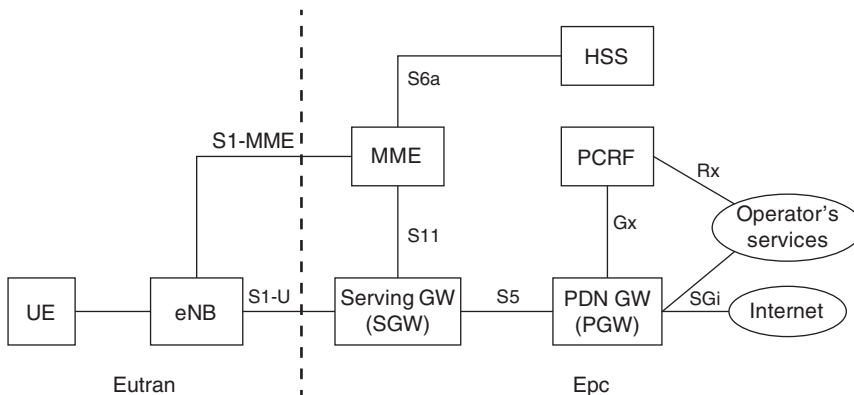

Figure 5.1 3GPP LTE architecture

or in 2G/GSM networks (with BTS and BSC). In the following only the single-element case is considered, the extension to the multi-elements case being straightforward.

5.2.2   Base Station Functions

The identification of the functions implemented by the base station is important in order to understand its hardware and software architecture, introduced in the next sections.

The role of the base station is primarily to take in charge the connection with the user equipments through the last wireless hop, providing all the necessary functions which include the data delivery, the mobility management, the overall radio resource management and the assistance to the UEs for a wide range of operations, such as network discovery, attachment and paging.

In the user plane, the role of the BS is to forward to the right destination (UE) the data received from the core network in the downlink direction (DL: BS → UE) and to perform the reverse operation on the uplink (UL: UE → BS). More specifically, the base station user plane comprises of the following:

- L1 or PHY (physical) layer, including both digital and analogue signal processing, such as symbol and carrier modulation and demodulation, signal amplification and transmission and reception through the antenna system – in 3GPP LTE, this corresponds to implementing the OFDMA (downlink) and SC-FDMA (uplink) PHY layers, whereas in 3G/UMTS this corresponds to implementing the WCDMA PHY layer.
- The medium access control (MAC) layer, in charge of regulating (scheduling) the access to the radio resources (time and frequency spectrum) among the multiple users, QoS levels and flows, also including signalling and control messages – in 3GPP LTE, this layer also includes the hybrid-ARQ (HARQ) functionality and is described in Ref. [9].
- The (radio) link control, including segmentation and reassembly, to adapt the upper-layer packet size to the PHY layer, and retransmissions for error-free transmissions – in 3GPP LTE, this is called the radio link control (RLC) [10].
- Additional PHY agnostic functions, such as security (integrity and ciphering) and header compression – in 3GPP LTE, this is performed by the Packet Data Convergence Protocol (PDCP) [11].
- In single-element RANs, there is a need for packet buffering, forwarding and reordering between peer BSs during mobility – in 3GPP LTE, this is also performed by the PDCP [11].
- On the network side, the backhaul interface (e.g. Ethernet or microwave, depending on the deployment scenario) and the IP (tunnelling) stack (GTP in 3GPP LTE).

In the control plane, the base station relies on the user plane stack (L1, MAC, RLC and PDCP in 3GPP LTE) for the transport of signalling and control messages over the air interface. The BS control plane is in charge of a wide range of tasks, among which broadcast of system information, assistance to network detection and selection (e.g. synchronization), connection control (including security key generation), channel quality measurements configuration and reporting (of serving and neighbour cells), radio admission control, mobility control (of both idle and active UEs), paging, radio link failure detection and recovery, and so on. In addition, the base station is also in charge of transparently transporting the control plane exchanges
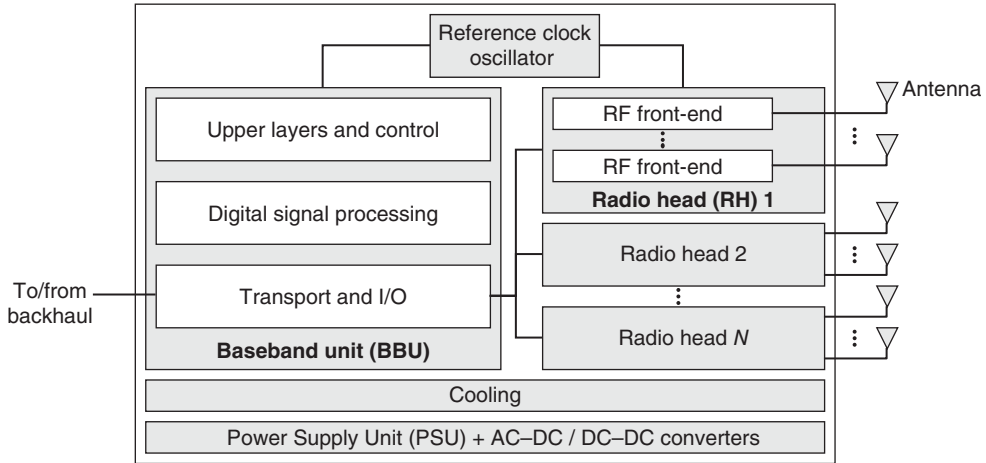
Figure 5.2 Schematic BS architecture

that take place between the UEs and the core network, such as the UE authentication and authorization, or the IP address assignment.

5.2.3 Generic BS Internal Architecture

A very large set of base stations exist in the market, differing by the size, form and internal architecture. This is due by the (co)existence of several cellular technologies, form factors (e.g. macros, micros and femtos; see Section 5.2.4) and telecom equipment makers, each one using its own design. It is thus impossible to provide a unique yet detailed description of a cellular base station; the technical choices to build them are too varied. In order to have a tractable analysis, this section proposes a simplified BS internal architecture, generic enough to be applicable (i.e. adaptable) to describe all kind of base station (Figure 5.2). It will be used in the following of this chapter to analyze the possible choices at component level, the obtained power consumption, and to study the possible evolutions towards more energy-efficient equipments.

The proposed architecture describes the base station in terms of the following five main building blocks:

- The baseband unit (BBU), in charge of all the digital operations
- One or more radio heads, in charge of the analogue operations and working on one or more antennas (e.g. for multi-antenna schemes, such as MIMO)
- A controlled oscillator locked on a GPS source or equivalent in charge of providing a very precise reference clock to the full system, BBU and RHs. For indoor small-size BS deployments (like femtos), where GPS signal is not available and backhaul connectivity not reliable enough to transport a stable clock reference, a stabilized oscillator (e.g. a temperature compensated crystal oscillator, TCXO) can be used
- The power supply unit, also including the different AC/DC and DC/DC converters (base stations normally use an input voltage of 48 V)

- The cooling system, which may vary from air conditioning to fan and even passive cooling in case of very low power base stations (such as femtocells).

The *Baseband Unit* (BBU) is the main block of the digital part of the base station. Referring to the previous section, the BBU is in charge of all the functions from L1/PHY (excluding the analogue part) up to the MAC and network layer, in both user and control plane. It is also in charge of the overall base station control, including O&M. The BBU elementary building blocks are given as follows:

- Transport and I/O layer supporting the different transport protocol stacks (such as IP and IP-tunnelling, GTP in LTE) and including:
    - The backhaul/backbone interface, which can vary from a simple Ethernet (e.g. for femtocells) to optical or microwave link. It is in charge of receiving/sending the IP packets (data and control) from/to the core network and also between the base stations in case of peering (such as in LTE, through the X2 interface).
    - An internal switch in charge of all the I/O that take place during the data/signal processing. Modern base stations abandon the legacy bus architectures to move towards more performing and flexible switches.
    - The in-phase and quadrature (I/Q) interface, used to exchange the digital samples of the baseband signal between the BBU and the radio head(s). This interface can be fully internal (e.g. directly exploiting the I/O switch) or external, e.g. in case of remote radio Head architectures (see Section 5.2.4). Open Base Station Architecture Initiative (OBSAI, [12]) and Common Public Radio Interface (CPRI, [13]) over an optical link are widely used standards for this interface.
- The digital signal processing (or baseband processing), performing the L1/PHY (digital) operations. In the downlink direction, it modulates the MAC-scheduled IP packets into the final baseband signal, described as I/Q digital samples. The digital samples of the baseband signal are then passed to the RH in order to create the final analogue signal that will be transmitted over the air. In the uplink direction, it demodulates the digital samples of the received baseband signal (passed by the RH) in order to extract all the transported information (data and control). More specifically, the most common operations of this block are (non-exhaustive list): FFT/iFFT, channel coding/decoding (turbo/Viterbi), channel estimation, synchronization, MIMO pre-coding, signal detection, digital filtering, up/down sampling and equalization. These operations are very demanding in terms of processing power, and specific hardware accelerators are commonly used for the most critical of them, such as turbo coding/decoding and FTT/iFFT.
- The upper layers and control, including:
    - Shared Memory (and controllers, e.g. DDR 2/3), used by the other blocks to temporary store the data during the different BBU operations. This also includes the buffering of the received data (IP packets) waiting to be sent to the UEs.
    - The "MAC and above" processing (upper layers), including link control (RLC in LTE), security and header compression (PDCP in LTE). In particular, the MAC is in charge of scheduling, that is, selecting which of the packets waiting in the internal buffer(s) will be transmitted in the next transmission slot (scheduling). Packet handling, scheduling and encryption are the most processing greedy operations of this block, with encryption often requiring a specific hardware accelerator.

○ The control and O&M part is in charge of the overall platform control, as well as the support to the remote base station management systems, allowing monitoring and configuring the equipment from the operators O&M platform.

The *Radio Head* (RH) is the analogue part of the base station. Each BS can have one or more radio heads, depending on the number of supported sectors (typically one radio head per sector), carriers and technologies. Each radio head in turn can have one or more antennas, depending on the supported MIMO modes. For example in LTE, MIMO 2×2, requiring two antennas per RH, is currently the default configuration, but higher antenna diversity schemes are supported. Figure 5.3 depicts a generic internal architecture of a radio head, generically composed of the following:

- I/Q interface, used to exchange I/Q digital samples with the BBU
- D↔A Converters, also including radio-head-related digital processing, such as the digital pre-distortion (DPD)
- RF transceiver (RF up/down conversion)
- RF amplifiers (power amplifier, PA, in downlink and low-noise amplifier, LNA, in uplink)
- Antenna interface, including duplexer and high-frequency filters (to split and isolate transmitted and received signals over the same antenna and RF cable), and the RF cable connection to the antenna (feeder or jumper)
- Antenna(s)

In the downlink direction, the digital-to-analogue converter (DAC) is in charge of constructing the analogue baseband signal starting from the digital I/Q samples received from the BBU. The baseband analogue signal is then passed to the transceiver, in charge of translating (up-converting) it into the passband signal, centred around the targeted high-frequency carrier (e.g. 2.6 GHz). The high-frequency passband signal is then amplified by the power amplifier and transmitted to the antenna via the specific interface. From energy consumption point of view, the power amplifier is the most consuming element of a radio head.

The analogue DL chain is not exempt from introducing signal distortions. In particular, high-power amplifiers can have non-linear in–out transfer. For this reason, a loopback (dashed
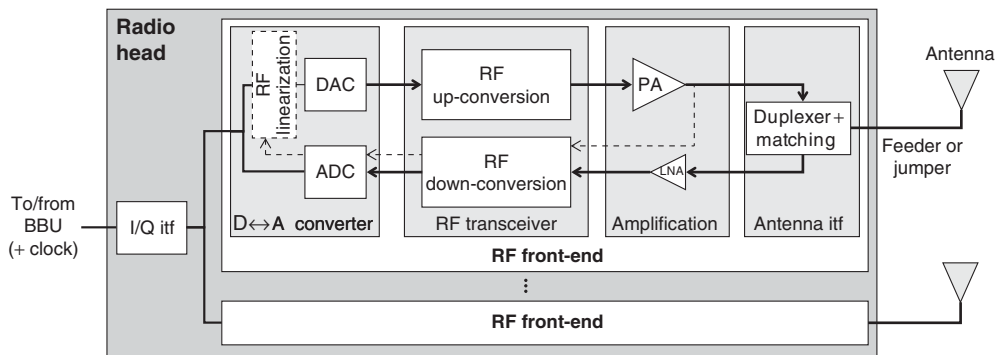

Figure 5.3 Radio head schematic architecture

line in Figure 5.3) is sometime introduced (especially in high-end / high-power macro-base stations); the PA output signal is sent back (through the uplink chain) in order to detect and correct unpredictable distortions. This is done by a (optional) RF-specific digital processing (RF linearization), functionally just before the DAC conversion, and usually including crest factor reduction for decreasing the input signal peak-to-average power ratio (PAPR) and DPD for increasing the PA linearity.

In the uplink direction, the received high-frequency passband signal is filtered and amplified by the low-noise amplifier, before being passed to the transceiver for down-conversion to baseband signal. Several down-conversion techniques can be used, the most common being the superheterodyne and the direct-conversion.

In superheterodyne architectures, the high-frequency signal is converted to a lower intermediate frequency (several successive IF stages are also possible) before obtaining the final baseband signal. Superheterodyne is the preferred choice for high/mid-power base stations (macro and micro BSs). Direct-conversion [14] (or zero-IF) is a simpler technique where the baseband signal is immediately obtained with a single step. Direct-conversion is commonly used in low-power base stations (pico and femto BSs) [15].

In the last stage of UL analogue processing, the baseband signal is passed to the analogue-to-digital converter (ADC), where it is sampled and quantized in order to obtain I/Q samples describing it in the digital domain. The I/Q samples are then sent to the BBU for information extraction (demodulation).

5.2.4   Types of Base Station

As said, there exist a wide set of base stations on the market, differing by the size (i.e. emitted power), spatial-reuse strategy (sectorization), supported carriers and technologies and internal architecture.

Regarding the emitted RF power, base stations are often classified into (see Table 5.1) *large cells* (macro and micro base stations), usually deployed at high heights, such as on top of building roofs or telecom towers, and *small cells* (pico and femto base stations), deployed at ground level (e.g. at home or on top of street lamps).

Macro base stations are deployed to provide coverage over wide areas, several hundreds of metres to few kilometres, depending on the deployment environment (e.g. urban vs. rural), whereas micros are typically used in urban environments to fill smaller spots (few hundreds of meters) or to provide additional capacity in crowded environments (e.g. stadium and large shopping malls). Pico base stations can serve a very limited amount of users (10–20

Table 5.1   Macro/micro/pico/femto.

BS type	N_{sec}	P_{max} (W)	P_{max} (dBm)
Macro	3	40.0	46
Micro	1	5.0	37
Pico	1	1.0	30
	1	0.25	24
Femto	1	0.1	20

at maximum) over a short range area (few tens of meters, maximum ∼150 m), whereas femtocells are usually deployed in indoor/in-house environments, serving less than 10 users.

The architecture of large cells is driven by the technical performances (e.g. emitted power, number of parallel active users, multi-antenna schemes), resulting in relatively big and expensive equipments. Due to the high emitted RF power, the requirements on power amplifiers are very challenging to meet; high power over a large linearity region (especially in the case of orthogonal frequency-division multiplexing (OFDM) systems, presenting a high crest factor or PAPR) requires top-class components assisted by DPD and crest factor reduction processing. Even if some alternatives exist, (multi-stage) superheterodyne architectures are the consensual choice for the transceiver. In current large cells, the digital part (BBU) is usually build around several components or even cards, tailored to the specific tasks, such as hardware accelerators and DSPs for digital processing and ciphering, and more generic processors and cards for control and upper layers. However, new architectures tend to move towards system-on-chip solutions, integrating all the BBU functions into a single component. High emitted power translates also into high consumed power and thus high temperatures to dissipate, requiring expensive (and energy consuming) cooling systems, such as air conditioning or at least fan-based free-air circulation (in micro cells). Large cells can support complex multi-antenna diversity schemes, such as MIMO 2×2, 4×4 and in future even more. This implies 2, 4 or more antennas per radio head. Since the antennas of large cells are commonly mounted at high heights, a feeder cable is usually used to connect them to the BS site (e.g. in the basement of the building).

The architecture of small cells is primarily driven by size and cost; since the working conditions are easier (limited RF power ranges), the power amplifiers are often simpler, cheaper and less efficient than in the large cells case and usually do not require any linearization pre-processing (and thus, no loopback in the RH). Transceivers are based on simple direct-conversion solutions, such as zero-IF. System-on-chip solutions are commonly used for the digital part, integrating all the components of the BBU. Cooling is performed by small free-air fans, or by simple passive dissipation in the smallest equipments (femtocells). Small cells have a very limited number of antennas: one or maximum two for MIMO 2×2 schemes. Since the base station site and the antenna are co-located, the antenna connection is usually a loss-free jumper, or even a direct connector.

Base stations also differ by the number of supported sectors per site (Figure 5.4). Macro cells are typically multi-sectors, mostly 3-sectors with roughly 120° angular width. This implies three radio heads for each base station, each one serving a single sector with directional antennas. Micro, pico and femto base stations are typically single sector with omnidirectional antennas.
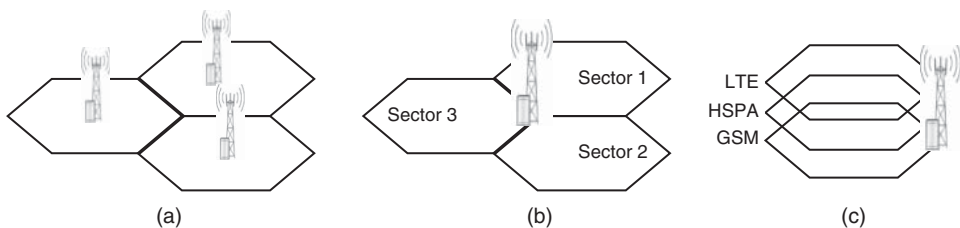

| (a) | (b) | (c) |

Figure 5.4   Types of base stations (BS) (sectors and carriers). Single-sector (a) versus tri-sector (b) versus multi-carrier or multi-technology (c) BS
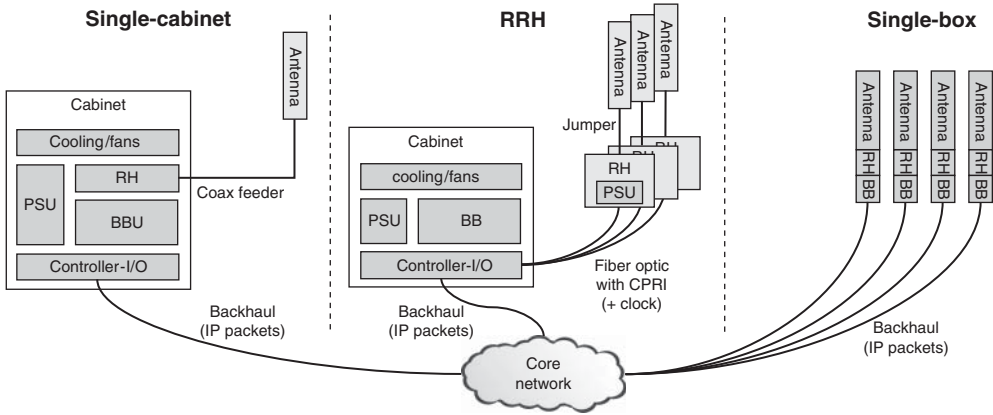
Figure 5.5 Types of BS (single cabinet versus RRH versus in-a-box)

Base stations (and in particular macros and micros) can also support multiple carriers, as well as multiple standards and technologies, such as GSM + HSPA + LTE (Figure 5.4(c)). This allows BS site reuse, reducing the need of acquiring new sites when deploying a new technology (e.g. LTE). Modern base stations are able to support multiple carriers and technologies on a single radio head (per sector).

Finally, three main base station architectures can be identified nowadays (Figure 5.5): the classic single-cabinet architecture, the remote radio head (RRH) and the single-box (or all-in-one) small architecture.

The *single-cabinet* architecture is the classic large base station architecture: all the elements of the BS (except the antennas) are packed into a single cabinet, on one or several blades, at the BS site (e.g. the basement of building). The antennas in turn are mounted at high height (e.g. rooftop) and connected to the base station site through a feeder cable. The drawback is that long feeder cables introduce important signal attenuation, commonly approximated at −3 dB, meaning that half of the signal power is lost in the cable.

Remote Radio Head (RRH) architectures have been introduced to mitigate the problem. In RRH, the radio head is separated from the BBU and mounted close to the antenna. This drastically reduces the distance between RRH and antenna and thus the signal attenuation. The connection between BBU and RRH is based on loss-free optical fibres and CPRI standard [13]. Remote radio head is a first example of "green" base station architecture, since its primary goal is to reduce energy wasting. Its draw-back is a diminished accessibility for maintenance and higher wind load for the antenna tower or mast.

Single-box (or all-in-one) architecture is the natural choice for small base stations. In this architecture, all the BS components, including the antennas are collocated in the same limited-size box (in a way very similar to WiFi access points).

5.3 Base Station Energy Consumption

Several studies have been recently carried on in order to better understand the power consumption of base stations at component level and the impact of traffic load variations on it [14–16].

This section summarizes the main results of these studies and proposes an analytical power model, useful when studying and simulating energy-related aspects of cellular BS.

5.3.1   Analysis of Energy Consumption at Component Level

All the detailed analyses (e.g. [14, 15]) of the power consumed by a base station at component level show similar results: the most consuming blocks (see Section 5.2.3) are the power amplifier, the signal processing (digital and analogue), the power supply unit (especially the AD/DC and DC/DC converters) and the cooling system. However, there exist huge energy consumption differences across the different components. Moreover, such differences vary depending on the type (size) of the base station (macro, micro and femto/pico, as shown in Figure 5.6.

   In macro base station, the power amplifier is by far the most consuming component. It participates for approximately 60% to the total BS energy consumption. The remaining 40% are shared between the signal processing (with the digital part consuming two times more than the analogue one), the cooling system (requiring air conditioning) and the power supply (including AC/DC and DC/DC conversion). It must be noted that in some cases, the air conditioning can account for up to 18% of the total BS power consumption. The reason why the power amplifier is the most consuming component in macro BS comes from the need of high RF output power combined with the intrinsic low efficiency of power conversion. Moreover, since macro base
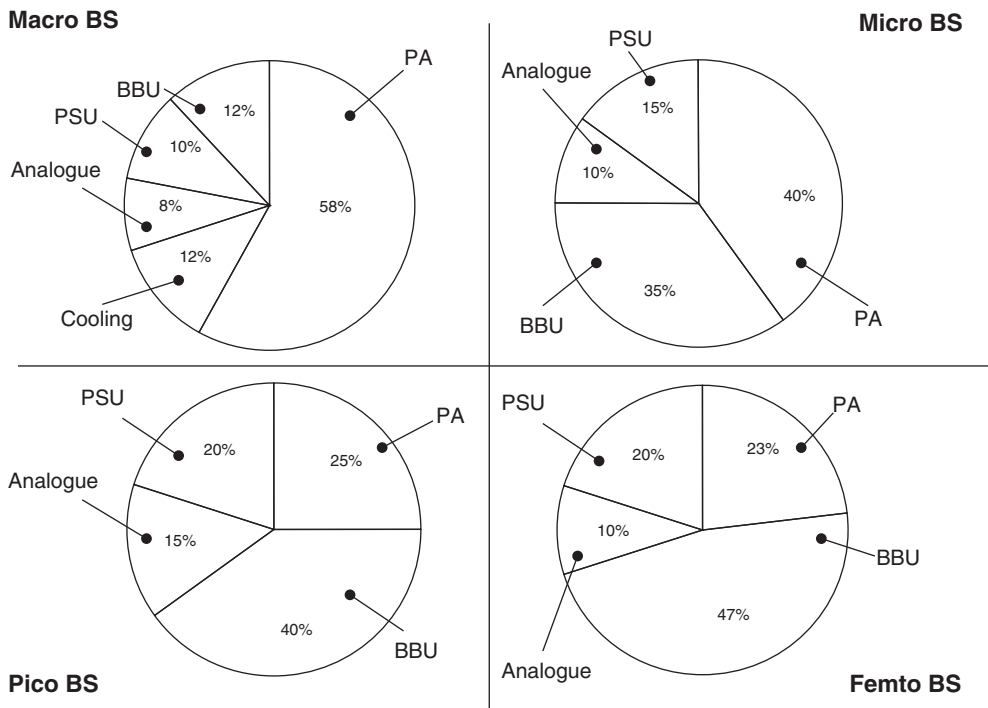

Figure 5.6   Per component energy consumption breakdown for different types of BSs

stations commonly have several sectors (and antennas per sector), several power amplifiers are necessary, one per RF chain (RF front end).

In micro BSs, the required output power and the number of power amplifiers is lower than in the macro case. As a consequence, the power amplifier still accounts for the largest share of energy consumption (~40%), but the digital processing (BBU) is now at an equivalent level. Since a simple fan is usually sufficient, the cooling impact is largely reduced. The remaining share (~25%) is shared between the power supply unit (~15%) and the analogue part (~10%).

It appears evident that to reduce the energy consumption of large cells (macro and micro) the improvement efforts must be focused on the power amplifier and digital signal processing. For macro base stations efforts in improving heating dissipation and cooling techniques are also very important, with the ambitious objective of replacing highly-consuming air conditioning with much greener fan-based systems.

In small base stations (pico and femto), the power consumption sharing is radically different. The digital processing (BBU) becomes the most consuming block (40–50%), followed by the power amplifier (20–25%). The PSU (power supply unit) accounts for around 20%, with the remaining 10–15% consumed in the analogue part. Cooling (passive) does not impact the power consumption repartition. To reduce the energy consumption of small cell, the efforts must be focused on reducing the impact of BBU and, on a minor part, of PA. Power supply also requires an attentive engineering. The major challenge however comes from the fact that small cells are subject to strong price pressure, making impossible to use high-end high-efficient components.

Finally, it must be noted that in LTE base stations (whatever their size, large and small) the operating conditions of the power amplifier are worsened by the OFDM modulation scheme which has a high PAPR. This forces the amplifier to operate in a linear region between 6 and 12 dB lower than the saturation point (which is in general the optimal efficiency point of a PA) [5]. This reduces the adjacent channel interference (ACI), but increases the PA inefficiency, and thus its power consumption.

5.3.2   *Impact of Load Variations*

It has been shown that the data traffic load in a cellular network varies during the day, depending on the time [17–20]. For example, in a business (office) area, the data traffic is very high during the day, but very low during the evening and night period. On the contrary, mobile traffic in residential areas tends to be concentrated in the evening and early night hours. However, whatever the deployment area (residential, business, urban, rural, etc.), it is possible to approximate the daily load variation with simple yet realistic models, summarized in Figure 5.7 (100% = average load).

It is thus important to analyze the power consumption of a base station depending on its non-constant load. The best behaviour (from energy consumption point of view) would be a perfect linear one, with zero energy consumption at 0% load (i.e. when the base station is neither sending nor receiving any data from users) and maximum at 100% load. However, the reality is far from this. As an example, Figure 5.8 shows the power consumption of a real LTE macro base station (without any energy-efficiency feature), as function of its load, defined as the number of used resource elements (RE) with respect to the maximum REs available (600) per TTI (here: 600 REs per 1 ms). Since the load can vary per sector, the figure shows the energy consumption for one single sector. In order to keep the result generic, the power
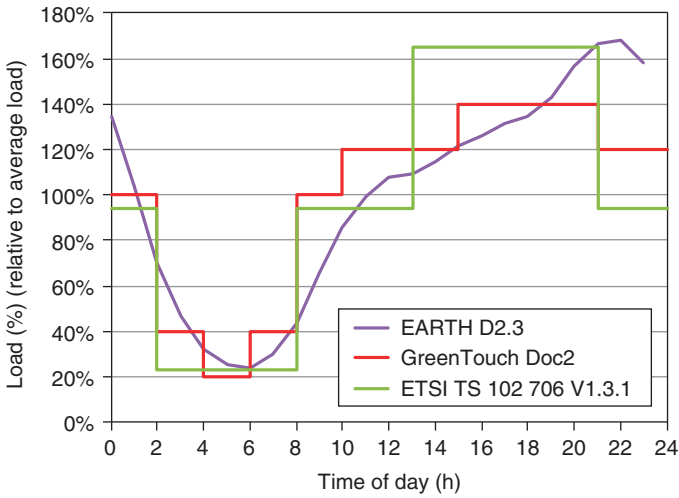
Figure 5.7   Mobile traffic variation profiles from [17] (label EARTH D2.3), [18, 20] (label GreenTouch Doc2), and [19] (label ETSI TS 102 706 V1.3.1)
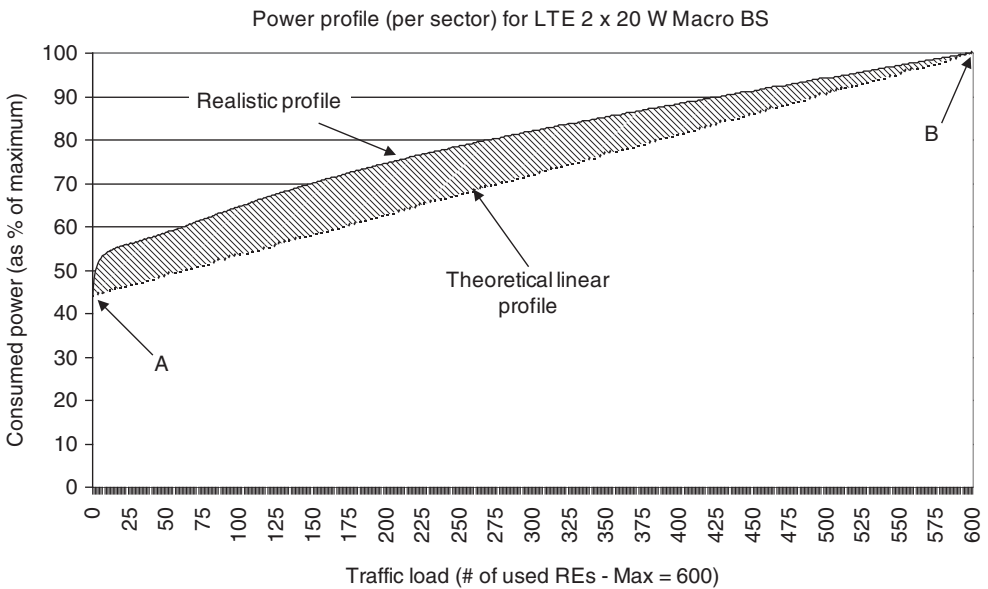

Figure 5.8   Example of BS power consumption as function of traffic load

consumption is normalized with respect to the maximum consumption (per sector). As an example, the consumption of a tri-sector, 2×20 W (MIMO) LTE macro BS is approximately 1500 W, with consumption close to 500 W per sector.

By analyzing Figure 5.8, several important observations can be done on the power consumption of a macro cell (note that the following does not apply in the same way to micro, pico and femto cells):

- The energy consumption of the macro BS effectively grows with the traffic load.
- The growth is not linear; there is an initial step when moving out from 0%. In addition, the slope is steeper at low load than at high loads.
- There is a huge offset at zero load; the energy consumed by a macro BS not serving any user is approximately half of the energy consumed at full load. This translates into very low-energy efficiency at low loads.

The non-linearity of the of the slope implies that it is better to either transmit at full load/power (point A on Figure 5.8) or to not transmit any data (point B). In fact, any other operating point (i.e. the real curve) is always higher (more consuming) than the theoretical linear curve (dashed line). This observation is (one of) the reason at the base of energy-aware scheduling techniques presented in Section 5.4

Even more important, from an energy consumption point of view, is the offset at 0% load (point A). There are several causes to this offset. First, as any real electronic equipment, the base station consumes energy just for the reason of being switched-on: circuits must be power supplied and thus energy is consumed. It is impossible to have ideal equipments that consume absolutely no power when turned-on, even if they *do nothing*. In addition, the BS is always active, even a 0% load. In fact, cellular technologies require that the BSs always broadcast signalling and control channels to assist mobile terminals in important operations, such as detection of service availability, acquisition of system characteristics (e.g. used MIMO scheme), mobility and paging. The mandatory emission of signalling (even when no users are present) is often called *coverage tax*, as it can be seen as the cost of constantly advertising the service availability. It is estimated that approximately 10% of the total RF power is constantly spent in signalling. Figure 5.9 shows the problem on an LTE frame structure; the resource element (squares) occupied by signalling and control channels (even when no user is present) is greyed. It can be observed that an LTE base station is obliged to transmit signalling (e.g. pilots) four times per millisecond (i.e. every 250 us) at least.

The power consumption profile as function of the traffic load drastically varies depending on the type of base station [17, 21]. This is summarized in Figure 5.10, where it can be observed that the dependency from the traffic load is high on large cells but almost non-existent in small BSs (the femtocell power consumption varies from 11 to 9 W approximately). This can be explained by the lower impact of power amplifier on the total consumption of smaller cells. In addition, as the power amplifiers on small cells must be cheap, their performances are often poor, with a flat response in term of power consumption. On the other hand, and whatever the type of BS, it can be noted that the energy consumption of all the other components poorly scale with load. Only the BBU is slightly impacted by the load, but this impact is today very limited. This provides an additional investigation direction for improving the energy efficiency of future base stations. Finally, a (small) non-linearity of the power profile is visible in macro- (as expected) and micro base station. In small cell the profile is linear (and almost flat).
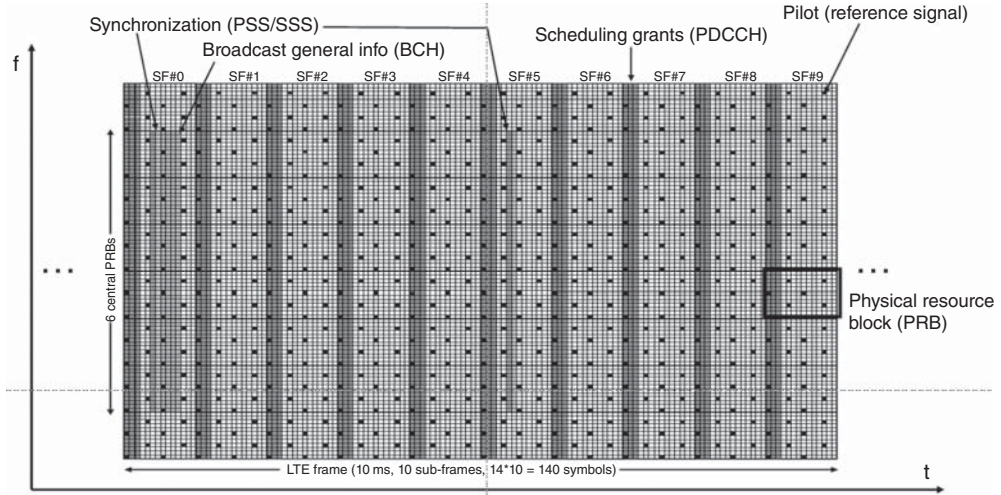
Figure 5.9   Coverage tax: impact of signalling and control channels on LTE frame

5.3.3   Power Models

The goal of the BS power model is to provide a simple and easy-to-use mathematical model of the power consumption of a BS, as function of the traffic load. This is often useful when analyzing and comparing different green wireless solutions. In this chapter, the power model used is a simplified version of the one described in [22], which is an evolution of the model proposed in Ref. [17, 21]. This model allows estimating the base station consumption (P_{in}) as function of the output RF power (P_{out}), the traffic load of the sector (λ) and the type of BS (Table 5.2): number of sectors (N_{sec}), maximum RF power (P_{max}), consumed power at zero load (P_0) and slope (Δ_p). It also models a sleep mode (see Section 5.4), with a specific power consumption (P_{sleep}), that can be activated only at zero load. The model does not distinguish the number of antennas per sector; a 2×20 W macro BS corresponds to a 40 W macro.

$$P_{in} = \begin{cases} N_{sec} \cdot P_{out}, & 0 < \lambda \leq 1 \\ N_{sec} \cdot P_{sleep}, & \lambda = 0 \end{cases}$$

$$P_{out} = P_0 + \Delta_p \cdot P_{max} \cdot \lambda$$

Figure 5.11 present the traffic profiles issued by the proposed model in three cases: macro (3 sectors, 2×20 W per sector), micro (5 W) and femto (100 mW).

To conclude this section, it is worth to be noted that other power models have been proposed in the literature in the recent years, as for example [23].

5.4   Evolutions Towards Green Base Stations

The recent attraction drawn by the telecom research and industrial world on green aspects has allowed identifying several promising techniques to reduce the energy consumption of
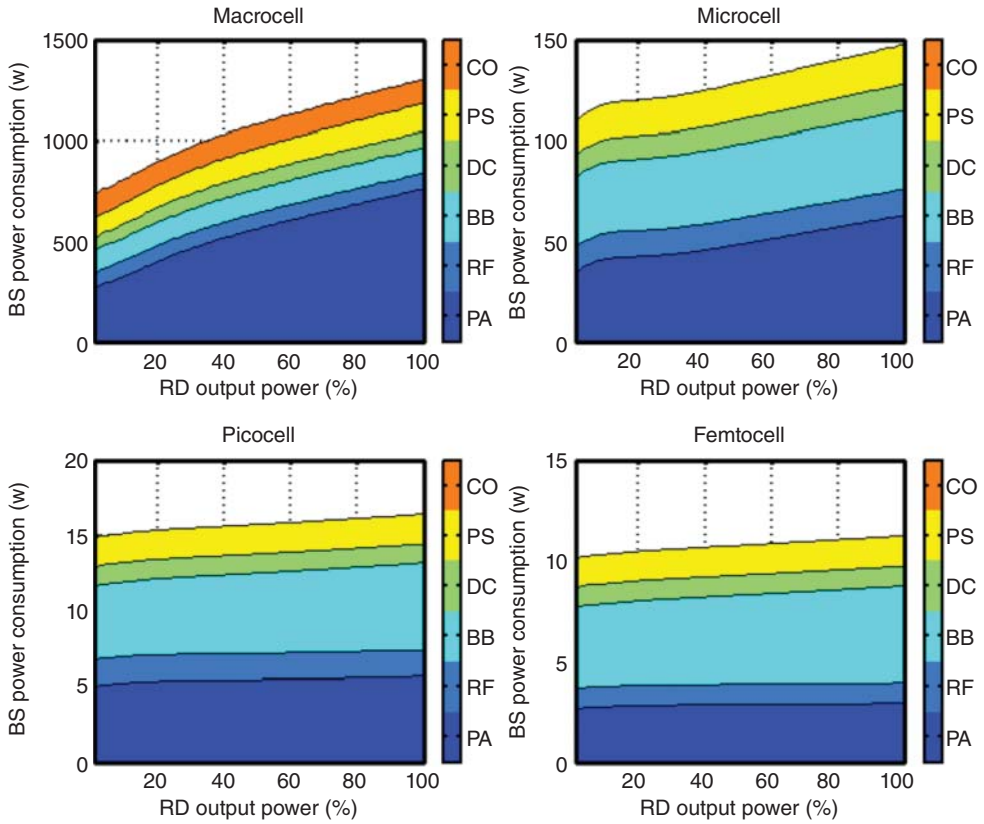
Figure 5.10 Power consumption profiles for different types of base station (year 2010 values) – Courtesy of European Community's 7th Framework Program FP7 project EARTH – Legend: CO = Cooling, PS = power supply, DC = DC–DC converters, BB = baseband processing, RF = RF transceiver (w/o PA), PA = power amplifier

Table 5.2 BS power profile parameters.

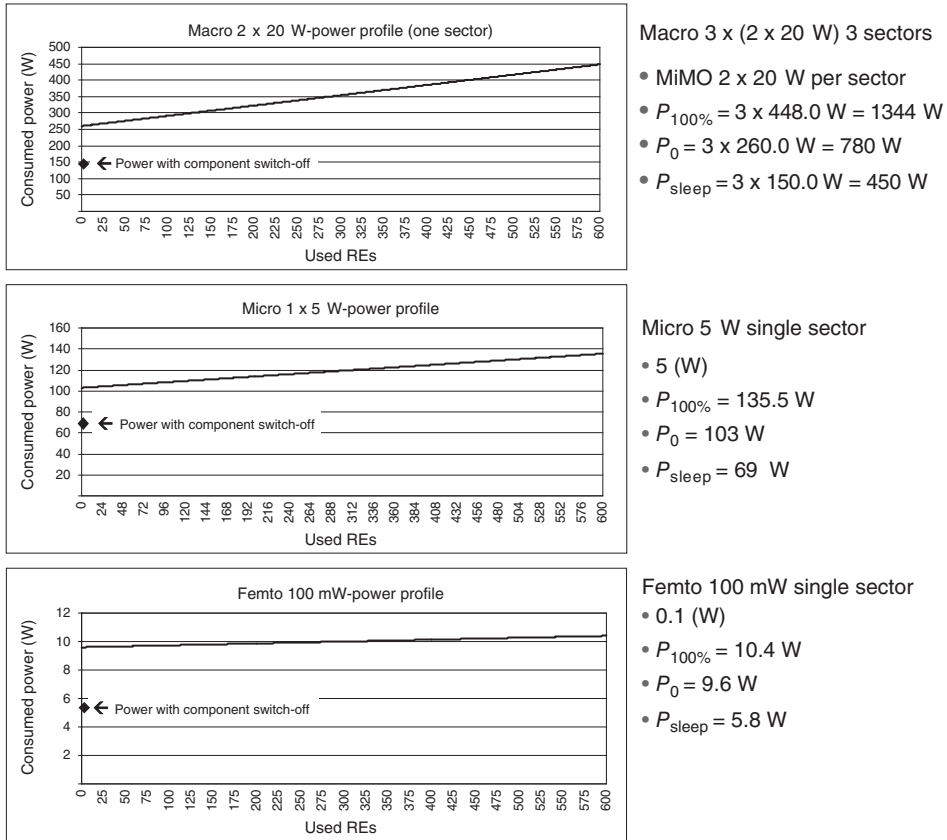| BS type	N_{sec}	P_{max} (W)	P_{max} (dBm)	P_0 (W)	Δ_p	P_{sleep} (W)
Macro single cabinet	3	40.0	46	260.0	4.7	150.0
Macro RRH	3	40.0	46	168.0	2.8	112.0
Micro	1	5.0	37	103.0	6.5	69.0
Pico	1	1.0	30	96.2	1.5	62.0
	1	0.25	24	13.6	4.0	8.6
Femto	1	0.1	20	9.6	8.0	5.8

Figure 5.11   Power model profiles for different BS types (macro, micro and femto)

cellular base stations. This section focuses on the evolutions at base station level (component, architecture and operation) from which the higher energy savings are expected, namely, improvements at components level (both hardware and operating-point), implementation of BS stand-by modes, traffic-load adaptation and evolution of BBU architecture, including the cloud approach. When possible, this section also points out the possible limitation and blocking points of the described evolutions.

5.4.1   Component Level Evolutions

New materials and new components architectures constitute the most basic level how wireless BS can effectively benefit of the improvements developed for the worldwide electronic industry. They have an impact on all the BS components, and particularly on the most consuming ones, such as power amplifier, and baseband processing. Indirectly, cooling can also benefit of such improvements; the improved efficiency of the other components usually implies a reduction of the heat to be dissipated.
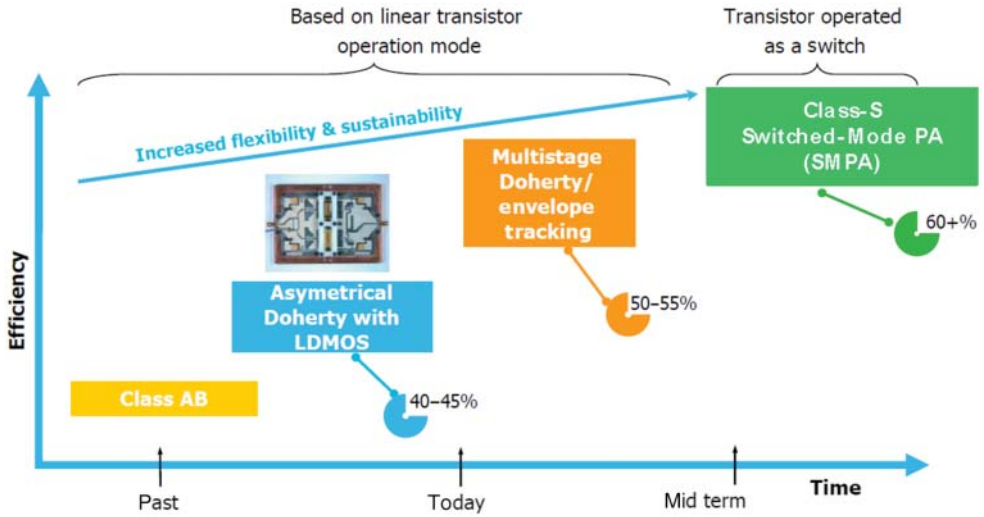
Figure 5.12 Principle of envelope tracking power amplifier (ETPA)

5.4.1.1 New Power Amplifiers architectures

As shown in Section 5.3.1, the PA is responsible for a very significant part of the BS power consumptions, especially in the case of large macro cells. Therefore, high-efficiency PAs have become the subject of intensive research resulting in continuous improvement of performances ([24, 25]).

Several solutions focused on improving the power amplifier efficiency under high PAPR (and backed off) conditions, for example in LTE. Signal conditioning algorithms, such as crest factor reduction for decreasing the PAPR and digital pre-distortion for increasing the PA linearity, enable the PA operation closer to saturation, corresponding to the maximum efficiency working point. These techniques usually come as complement to more radical PA architecture evolutions (Figure 5.12).

The current state-of-the-art is represented by Doherty power amplifiers (DPA), which contain one main (carrier) amplifier always active and an auxiliary one (peaking) active only when signal peaks occurs ([26, 27]). This aspect makes DPA particularly adapted to signals with high PAPR. DPAs show an efficiency of around 40–45%.

Efficiency can be further improved by two new architectures currently under intensive study: envelope tracking PAs (ETPA, [28]) and switch mode power amplifier (SMPA, [29]). In ETPA, the DC supply voltage is constantly adjusted so that to track the variations of the transmitted input. Since the difference between the DC supply voltage and the transmitted envelope voltage is dissipated as heat, envelope tracking PAs allow to drastically reduce this dissipation, and thus to increase the overall efficiency to around 50–55%. EPTA is better view for small cells systems where the instantaneous required pass band is lower.

Class-S SMPA architecture promises very high theoretical efficiency. This comes from the fact that the active component (transistor) is operated in an ON/OFF way: when ON, the transistor acts as a very low resistance (ideally short circuit, such as a closed switch), when OFF

it acts as an open circuit. Since the transistor shows a zero ON resistance (thus, no voltage) and an infinite OFF resistance (thus no current), a 100% efficiency is theoretically achieved. In reality, efficiencies of around 60% are expected.

5.4.1.2 Signal-Aware Power Amplifiers

Another family of PA evolutions deals with the variations of traffic (and thus requested output power) over the day, and even in real time. Section 5.3.2 showed that such variations can be very high, and in very low-load conditions the PA works in an operating region far from maximum efficiency (saturation). It is thus interesting to defined mechanisms to adapt the PA working point to the actual traffic load. This is usually obtained by modifying the gate voltage while keeping constant the drain voltage of the transistor. However, this technique requires updating the digital-pre-distortion so that the operating point modification is not interpreted as a distortion. This implies that such modification cannot be performed too frequently. Another approach consists in using only two operating points: the nominal one and one with gate voltage to zero (PA deactivation), used only when there is no signal to transmit (i.e. neither data nor signalling). In this way, the digital-pre-distortion is not impacted (when the gate voltage is at zero there is no signal to pre-distort), allowing very rapid activation/deactivation, to track to real-time (less than 1 ms) traffic variations. During the deactivation period (gain voltage to zero), the power consumption of the PA can be reduced by a factor approaching 10 times with respect to the power consumed at maximum power.

5.4.1.3 Improvements of BBU

The Baseband Unit is currently undergoing important evolutions, brought in particular by important improvements from components and architecture. From an energy point of view, the use of low-consumption DSP and ASICs architectures must be preferred to more flexible but also more energy greedy FPGA architectures. This is today in line with the mainstream trend of moving away from multiple-cards architectures and to adopt integrated solutions, often based on system-on-chip (SoC) architectures. For large base stations, multiple-SoC architectures are hitting the market, where several generic cores are integrated with memory facilities and specific hardware accelerators for computationally expensive operations (such as turbo coding/decoding and ciphering). This integration brings a natural reduction of power consumption of the BBU, and the opportunity for BBU processing scaling, where the internal cores are activated/deactivated according the BS load.

BBU processing scaling requires specific design in terms of both hardware and software. It becomes extremely important to design the BBU parallel processes in a way that processing flexibility can be achieved. As an example, dedicating each processing unit to different parallel operations is not the good approach, because it makes impossible to turn off some processing unit at low load. A good design is in turn to concentrate always-on processes on some processing units and to execute all the load-depending operations together on a pool of processors; at low load, a single processing unit is needed whilst new units can be activated when the load increases.

5.4.2   BS Operation Improvements

The improvements of the base station operation target to improve the proportionality of the power consumed with the traffic load, in particular by exploiting the periods of time when the base station is poorly load (or even not serving any user). Two main techniques have been identified: real-time adaptation to traffic variation and BS stand-by modes.

5.4.2.1   Smart Load Adaptation to Traffic Load Variations

Smart load adaptation techniques apply to active base stations and aim at adapting the BS operating mode to the variation of traffic load. Such solutions combine component-level improvements of the power amplifiers (i.e. signal-aware PAs presented in Section 5.4.1.2 with smart scheduling strategies exploiting such improvements. Two main approaches can be identified [30]: scheduling policies adapting the actual used bandwidth, combined with power amplifiers with adaptive operating point, and scheduling policies creating micro-sleeps periods, using power amplifiers with deactivation mode [31, 32]. The second technique is also known as micro- or cell-discontinuous transmission (DTX) because during the micro-sleep periods the BS suspends its RF transmission. It must be noted that these solutions apply only to OFDM-based networks, such as LTE and LTE-A, and cannot be used in, for example, WCDMA base stations.

Smart load adaptation techniques gains have been estimated [33] to approximately 30% reduction of BS energy consumption.

5.4.2.2   Activation/Deactivation of RF Resources

In this case, the adaptation to the traffic load conditions is obtained by simply activating/deactivating the BS radio resources.

It is possible, for example, to adapt of the number of transmit antennas depending on the traffic load. For example, modern base stations (BSs) use up to eight antennas. When the traffic demand is low, it is possible to switch off some of the RF front ends and use a less spectral efficient transmit mode.

Another approach is to reduce the transmit band. For example, in LTE by using narrower carriers at low loads (e.g. fall-back mode from 20 to 5 MHz). When several carriers are active in the same sector (e.g. a 3-carrier per sector WCDMA Base Station), it is possible to implement a fallback mode which uses only one carrier. A similar approach is possible with multi-technology BS (multi-RAT, e.g. GSM + WCDMA + LTE), by switching off some of the technologies.

In all these cases, a particular attention must be brought to the internal hardware BS architecture. For example, modern multi-carrier and multi-technology base stations (recently hitting the market) tend to use the same RF head for several carriers or several technologies (e.g. one RF head for one sector, offering two 3G carriers and one LTE carrier). In that case, algorithms based on carrier or RAT switch off do not bring significant savings, as the power amplifier remains active for serving the remaining carriers/technologies.

5.4.2.3   Base Station Sleep Modes

During low-load periods, it can happen that a base station becomes unnecessary, either because no users are using it or because the remaining users can be absorbed by neighbour base stations. In these cases, in order to avoid energy wasting, it is interesting to move the base station into some low-consumption mode (sleep or stand-by). Unfortunately currently, deployed base stations (whatever technology) do not implement such sleep modes and can usually work into only two states: full active or complete switch off. Algorithms based on long timescales (i.e. tolerating off–on transients of the order of tens of minutes and above) are not affected by this problem; they can be implemented by turning completely off the base station. This can be achieved even on already deployed base station, using the OAM interface. Algorithms requiring shorter wakes-up delays cannot be based on a complete switch off of the base station. In fact, the delay necessary to move from a power-saving mode to the full active mode depends on the number (and type) of components that are switched off (see, e.g. [34] for an analysis on 3G femtocell). A cold-restart requires several minutes to complete, with the delay increasing with the size of the base station.

New sleep modes must thus be defined. These modes are characterized by the BS wake-up delay (i.e. the time required by the BS to move from the stand-by to the full active mode): the longer is the acceptable wake-up delay, the higher is the number of internal components that can be turned off, and thus the higher are the energy savings. Two main cases can be identified (note that the ultra-fast real-time on/off case corresponds to the cell-DTX presented in Section 5.4.2.1: fast wake-up (few seconds) and longer wake-up time (several tens of seconds to few minutes).

For fast wake-up, very few components can be turned off: the radio head (some care must be brought to digital-pre-distortion stabilization delay), and possible some cooling (in case of single-cabinet, otherwise the RH switch off does not lead to temperature reduction in the BBU cabinet). The other components must remain on, especially the BBU that requires a considerable amount of time to reactivate, incompatible with fast wake-up. It can be concluded that theoretically, a well-designed fast-wake-up sleep mode can target to save the energy consumed by the RH (e.g. ~60% for a macro cell, see Figure 5.6).

If longer wake-up time can be accepted, additional components can be turned off, including the oscillator, the digital signal processing and the upper layers. The backhaul interface and some control process must be kept in order to be able to receive the wake-up request from the neighbour cells or the management system. Cooling can be turned off, while power supply shall be designed to be partially turned off, otherwise it must remain on to supply the backhaul interface and the remaining control processes. A rough estimation of potential energy gains is around 80–90%.

5.4.3   BS Architecture Evolutions

With respect to energy efficiency, two major architecture evolutions are currently being analyzed and defined by the telecom industry (Figure 5.13): massive-MIMO (i.e. the massive increase of the number of active antennas on the radio side) and cloud-RAN (i.e. execution of baseband processes in remotely located farms of generic servers). The two evolutions are not antagonistic and will probably combine together: the centralized high-processing power
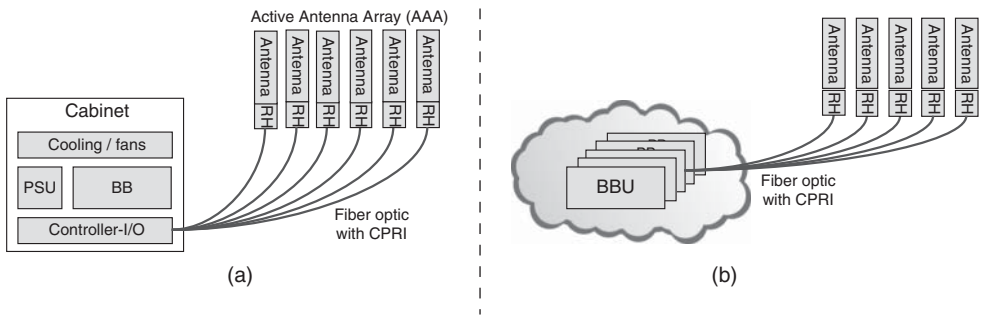
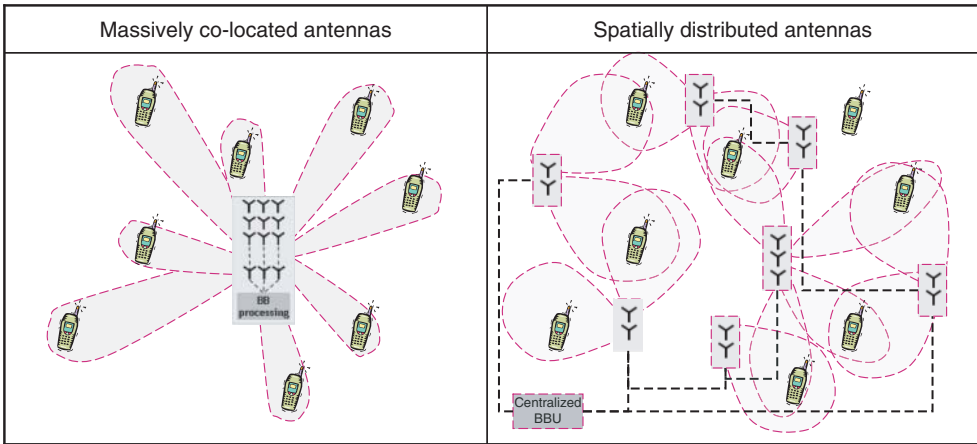**Figure 5.13** BS architecture evolutions: massive-MIMO and Cloud-RAN

Figure 5.14 Distributed versus centralized massive-MIMO

of cloud-RAN is in fact well adapted to the huge amount of signal processing required to synchronize the emission of massive-MIMO systems.

5.4.3.1 Massive-MIMO Architecture

Massive-MIMO aims at increasing the spectral efficiency through spatial multiplexing ([35–37]). The impact at BS level comes from the multiplication of radio heads and RF front end (hundreds of antennas), centralized on a single site or distributed over a specific region (Figure 5.14). Energy gains come from the fact that the RF power transmitted per antenna is much less than in traditional architecture because the radiated energy can be focused on the receiver. However, such gains risk to be diminished when the high-spectral efficiency of massive-MIMO is not fully exploited, i.e. at low traffic loads. The engineering of such systems is thus challenging ([38]); the used front end must be low power yet efficient enough

in order to effectively reduce the total consumption. Dynamic activation/deactivation of the used front ends can also be considered, during low-load periods.

5.4.3.2 Cloud-RAN Architecture

In the cloud-RAN approach ([39]), the BBU of several base stations are centrally located in a remote location, executed on server farms. The centralized BBU are connected to the distributed radios via a high-capacity backhaul networks, usually based on fibre. The idea is that such a distributed network could be less costly to deploy for mobile operators because it could run on smaller cell sites. In addition, centralized processing offers high computational capacity and ease multi base stations cooperation, for example, distributed MIMO (e.g. CoMP, coordinated multi-point transmission/reception defined by 3GPP) or massive MIMO. From the energy efficiency point of view, the centralization of BBUs (also known as BBU-pooling) promises the possibility to scale the processing resource with the access network load. This is similar to the BBU processing scaling presented in Section 5.4.1.3, but extended to the multi-BS case, with a coarser granularity.

References

[1] A. Fehske; G. Fettweis; J. Malmodin; G. Biczok; "The global footprint of mobile communications: the ecological and economic perspective," IEEE Commun. Mag., vol. 49, no. 8, pp. 55–62, 2011.

[2] S. Vadgama; M. Hunukumbure; "Trends in green wireless access networks," 2011 IEEE International Conference on Communications Workshops (ICC), pp. 1–5, 2011.

[3] P. J. Chung; "Green Radio—the case for more efficient cellular base stations," (slides) UK-Taiwan ICT Workshop: Smart & Green Communications, University of Taiwan, 2010).

[4] S. Sesia; I. Toufik; M. Baker; LTE – The UMTS Long Term Evolution: from Theory to Practice, New York: John Wiley & Sons, 2009.

[5] 3GPP Technical Specification 36.300, "Evolved universal terrestrial radio access (E-UTRA) and evolved universal terrestrial radio access network (E-UTRAN); overall description; stage 2," www.3gpp.org/ftp/Specs/html-info/36300.htm.

[6] 3GPP Technical Specification 23.401; "General packet radio service (GPRS) enhancements for evolved universal terrestrial radio access network (E-UTRAN) access," www.3gpp.org/ftp/Specs/html-info/23401.htm.

[7] 3GPP Technical Specification 23.402, "Architecture enhancements for non-3GPP accesses," www.3gpp.org/ftp/Specs/html-info/23402.htm.

[8] 3GPP Technical Specification 29.060, "General packet radio service (GPRS); GPRS tunnelling protocol (GTP) across the Gn and Gp interface," www.3gpp.org/ftp/Specs/html-info/29060.htm.

[9] 3GPP Technical Specification 36.321, "Evolved universal terrestrial radio access (E-UTRA); medium access control (MAC) protocol specification," www.3gpp.org/ftp/Specs/html-info/36321.htm.

[10] 3GPP Technical Specification 36.322, "Evolved universal terrestrial radio access (E-UTRA); radio link control (RLC) protocol specification," www.3gpp.org/ftp/Specs/html-info/36322.htm.

[11] 3GPP Technical Specification 36.323, "Evolved universal terrestrial radio access (E-UTRA); packet data convergence protocol (PDCP) specification," www.3gpp.org/ftp/Specs/html-info/36323.htm.

[12] OBSAI, Open Base Station Architecture Initiative, www.obsai.com.

[13] CIPRI, Common Public Radio Interface, www.cpri.info.

[14] Abidi, A.A.; "Direct-conversion radio transceivers for digital communications," IEEE J. Solid-State Circuits, vol. 30, no. 12, pp. 1399–1410, 1995.

[15] B. Debaillie; A. Giry; M. J. Gonzalez; L. Dussopt; M. Li; D. Ferling; V. Giannini; "Opportunities for energy savings in pico/femto-cell base-stations," Future Netw. Mobile Summit (FutureNetw), 15–17, pp. 1–8, 2011.

[16] Desset, C.; Debaillie, B.; Giannini, V.; Fehske, A.; Auer, G.; Holtkamp, H.; Wajda, W.; Sabella, D.; Richter, F.; Gonzalez, M.J.; Gódor, I.; Olsson, M.; Imran, M.; Ambrosy, A.; Blume, O.; "Flexible power modeling of LTE base stations" IEEE WCNC 2012, Paris, 2012.

[17] EARTH Project, "Deliverable D2.3 - energy efficiency analysis of the reference systems, areas of improvements and target breakdown," available at: www.ict-earth.eu/publications/deliverables/deliverables.html, 2012.

[18] GreenTouch, www.greentouch.org.

[19] ETSI TS 102 706 V1.3.1 (2013-07): "Environmental Engineering (EE) Measurement Method for Energy Efficiency of Wireless Access Network Equipment"

[20] GreenTouch, "GreenTouch Green Meter Research Study: Reducing the Net Energy Consumption in Communications Networks by up to 90% by 2020," A GreenTouch White Paper, Version 1.0, June 26, 2013

[21] G. Auer; V. Giannini; C. Desset; I. Godor; P. Skillermark; M. Olsson; M. A. Imran; D. Sabella; M. J. Gonzalez; O. Blume; A. Fehske; "How much energy is needed to run a wireless network?," IEEE Wireless Commun., vol. 18, no. 5, pp. 40–49, 2011.

[22] NTT DOCOMO, Alcatel-Lucent, Alcatel-Lucent Shanghai Bell, Ericsson, Telecom Italia, "[r1-113495], base station power model." [Online]. Available: http://www.3gpp.org/ftp/tsgran/wg1rl1/TSGR166b/Docs/R1-113495.zip.

[23] Debaillie B.; Desset, C.; Louagie F.; "A flexible and future-proof power model for cellular base stations," in Vehicular Technology Conference (VTC), Scotland, May 2015

[24] S. C. Cripps; RF Power Amplifiers For Wireless Communications, 2nd edn, Artech House Microwave Library, 2006.

[25] J. Joung; C. K. Ho; S. Sun; "Green wireless communications: a power amplifier perspective," 2012 Asia-Pacific Signal & Information Processing Association Annual Summit and Conference (APSIPA ASC), pp. 1–8, 2012.

[26] W. H. Doherty; "A new high efficiency power amplifier for modulated waves," Proc. Inst. Radio Eng., vol. 24, no. 9, pp. 1163–1182, 1936.

[27] B. Kim; J. Kim; I. Kim; J. Cha; "The Doherty power amplifier," IEEE Microw. Mag., vol. 7, no. 5, pp. 42–50, 2006.

[28] J. Kim; J. Son; S. Jee; S. Kim; B. Kim; "Optimization of envelope tracking power amplifier for base-station applications," IEEE Trans. Microwave Theory Tech., vol. 61, no. 4, pp. 1620–1627, 2013.

[29] A. Wentzel; C. Meliani; W. Heinrich; "RF class-S power amplifiers: state-of-the-art results and potential," 2010 IEEE MTT-S International Microwave Symposium Digest (MTT), pp. 812–815, 2010.

[30] A. Ambrosy; O. Blume; H. Klessig; and W. Wajda; "Energy saving potential of integrated hardware and resource management solutions for wireless base stations," in 2011 IEEE 22nd International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC), pp. 2418–2423, 2011.

[31] P. Frenger; P. Moberg; J. Malmodin; Y. Jading; and I. Godor; "Reducing energy consumption in LTE with cell DTX," in 2011 IEEE 73rd Vehicular Technology Conference (VTC Spring), pp. 1–5, 2011.

[32] R. Gupta; E. Calvanese Strinati; and D. Ktenas; "Energy efficient joint DTX and MIMO in cloud radio access networks," in 2012 IEEE 1st International Conference on Cloud Networking (CLOUDNET), pp. 191–196, 2012.

[33] Ambrosy, A.; Blume, O.; Ferling, D.; Jueschke, P.; Wilhelm, M.; Xin Yu; "Energy savings in LTE macro base stations," 7th IFIP Wireless and Mobile Networking Conference (WMNC), 2014, pp. 1–8, 2014.

[34] I. Haratcherev and A. Conte; "Practical energy-saving in 3G femtocells," in IEEE Green Broadband Access (GBA) workshop, in conjunction with ICC 2013, 2013.

[35] T. L. Marzetta; "Noncooperative Cellular Wireless with Unlimited Numbers of Base Station Antennas," IEEE Trans. Wireless Commun., vol. 9, no. 11, pp. 3590–3600, 2010.

[36] H. Q. Ngo; E. G. Larsson; T. L. Marzetta; "Energy and spectral efficiency of very large multiuser MIMO systems," IEEE Trans. Commun., vol. 61, no. 4, pp. 1436–1449, 2013.

[37] J. Hoydis; S. ten Brink; M. Debbah; "Massive MIMO in the UL/DL of cellular networks: how many antennas do we need?," IEEE J. Sel. Areas Commun., vol. 31, no. 2, pp. 160–171, 2013.

[38] E. G. Larsson; F. Tufvesson; O. Edfors; T. L. Marzetta; "Massive MIMO for next generation wireless systems," eprint arXiv:1304.6690, 2013.

[39] Lin, Y.; Shao, L.; Zhu, Z.; Wang, Q.; Sabhikhi, R. K.; "Wireless network cloud: architecture and system requirements," IBM J. Res. Dev., vol. 54, no. 1, pp. 4:1–4:12, 2010.