

An Overview of Energy-efficient Base Station Management Techniques

Yi Zhang*, Łukasz Budzisz[†], Michela Meo*, Alberto Conte[‡], Ivaylo Haratcherev[‡], George Koutitas[§],
Leandros Tassioulas[§], Marco Ajmone Marsan*, Sofie Lambert[¶]

* Politecnico di Torino, Italy, E-mail:[yi.zhang,michela.meo,ajmone]@polito.it

[†] Technische Universität Berlin, 10587 Berlin, Germany, E-mail: budzisz@tkn.tu-berlin.de

[‡] Alcatel-Lucent Bell Labs, France, E-mail:[alberto.conte,ivaylo.haratcherev]@alcatel-lucent.com

[§] University of Thessaly, 38221 Volos, Greece, E-mail: george.koutitas@gmail.com, leandros@inf.uth.gr

[¶] Ghent University - iMinds, Belgium, E-mail: sofie.lambert@intec.ugent.be

Abstract—Cellular networks have been traditionally dimensioned to fulfill the desired quality of service (QoS) requirements at all times, and consequently their deployment has been planned to meet the expected peak of the user demand. However, with the user demand recently increasing at exponential pace, concerns about the cellular networks energy consumption have been raised. In response, energy-efficient resource management schemes have been proposed, which take into account energy consumption, and control how much of the network infrastructure is actually needed at different times, and how much can be temporarily powered off to cut energy consumption. Since most of the energy consumed in cellular networks is used by base stations (BSs), algorithms for managing BSs seem to be the most urgent development to achieve energy-efficient operation. This paper provides a quick overview of the BS management techniques that were recently proposed for cellular networks. In addition, an outlook on real implementation aspects, including current commercial products, and trends in the development of energy-efficient hardware is also given.

Index Terms—Energy-efficient networks, cellular networks, base stations, femtocells

I. INTRODUCTION

Energy efficiency is gaining increasing attention in various areas of research, both academic and industrial. This attention has also targeted communication networks, which have become a significant contributor to the total electricity consumption in the world, with their share increasing from 1.3% in 2007 to 1.8% in 2012 [1], mostly because of the rapid growth of user traffic demand.

Considering cellular access networks, we can observe that they have traditionally been designed targeting performance maximization at full load. Indeed, the full load working condition is the most critical and challenging, in terms of efficient use of the available resources, and, consequently, in terms of capital expenditures (CAPEX). However, most of the time, networks work at low to medium load. This is especially true for wireless access networks, in which the user aggregation level is very limited: cells cover limited areas where users tend to exhibit similar behaviors, so that the load profile exhibits large variations between peak and off-peak values, with long periods of low load. As an example, Fig. 1 shows traffic measured in one of the operating Italian mobile networks; solid lines refer to a cell in a business area; dashed

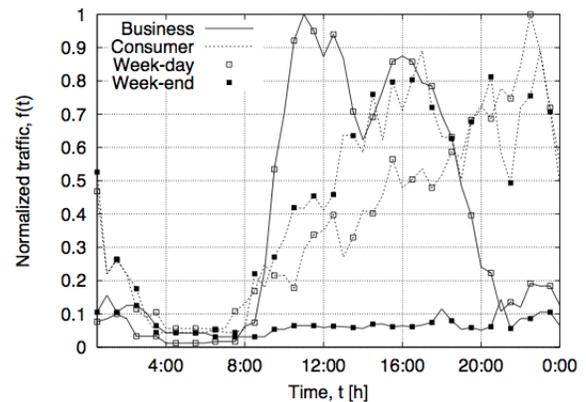


Fig. 1. Daily traffic profiles for a cell in a business area and a cell in a consumer area, week-day and week-end profiles measured in a network in operation.

lines refer to a cell in a consumer area; the empty markers identify the profile of a week-day; the solid markers refer to a week-end day. Traffic values are obtained by averaging and normalizing the measurements (at 15-minute intervals) collected during a week. From the figure we see that traffic fluctuates significantly during a day, and that periods of low activity are long. This would not be an issue if the power consumption of BSs were *load proportional*, i.e., if a direct proportionality existed between carried traffic and consumed power. Unfortunately, this is not true: the power versus load profiles of base stations, and of the entire network, exhibit very limited *load proportionality*. The values of the power consumption at zero or low load are large fractions, typically about 60-80%, of the consumption at full load. The limited load proportionality of cellular access equipment can be a serious obstacle to achieve the objective of energy efficiency.

The energy consumption of cellular networks can be reduced by making the consumed power more load proportional. The research efforts in this field have taken two main directions. On the one hand, manufacturers are focusing on designing devices that consume less power, and whose consumption is more load proportional; on the other hand, new network architectures are being proposed to make the entire network consume less, and in a more load proportional

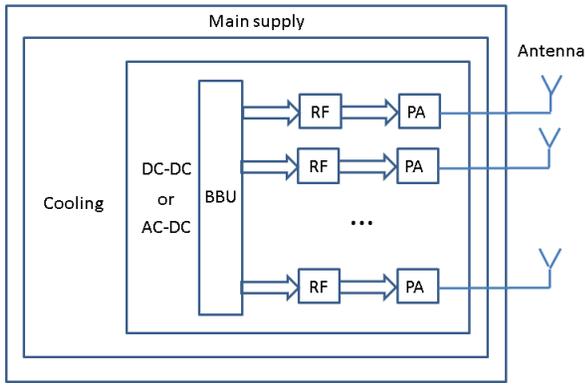


Fig. 2. Components of a BS.

fashion. Due to the fact that base stations (BSs) are the main energy consumers in cellular access networks, this paper overviews the issue of BS management to achieve energy efficiency (load proportionality) in cellular access networks.

The paper is further organized as follows: Section II investigates BS power consumption in cellular networks using LTE as an example; Section III advocates the introduction of sleep modes to save energy in cellular networks; Section IV surveys the existing BS management algorithms; Section V summarizes the current state of implementation of BS management schemes; and finally, Section VI concludes the paper.

II. ENERGY CONSUMPTION OF BASE STATIONS

In cellular wireless access networks, end users connect to BSs through wireless channels, and each BS, in turn, is connected to some other network element through either a wired or a wireless point-to-point link, which is part of the back-haul network. Several generations of cellular wireless access technologies exist, the most recent being LTE, also commonly called 4G. LTE is an energy-efficient technology: assuming as a relevant power consumption metric for cellular wireless access networks the power necessary to serve each user [W/user], in [2] it is claimed that, in urban areas, with a typical user density of 300 users/km², LTE requires 18 W/user, compared to 27 W/user of WiMAX, and 68 W/user of HSPA.

In an urban area, an LTE macro BS can cover an area of about 0.22 km² with a range of about 500 m. In suburban/rural environments, the covered area can grow to 2.6 km² with the same transmission power. Considering that the power consumption of the macro BS is around 1 kW, the power consumption per unit area is approximately 4500 W/km² for the urban area, and 400 W/km² for the suburban/rural area, respectively [2].

As shown in Fig. 2, a BS is comprised of a baseband unit (BBU) and one or more transceivers, each one of which contains a radio frequency part (RF), a power amplifier (PA), and an antenna, connected through a feeder. In addition, a BS, which normally uses an input voltage of 48 V, also contains an AC-DC converter, a DC-DC converter, and a cooling system (sometimes just a fan) [3].

In traditional deployments, the BS equipment is located far from the antenna, so that long feeder cables are necessary,

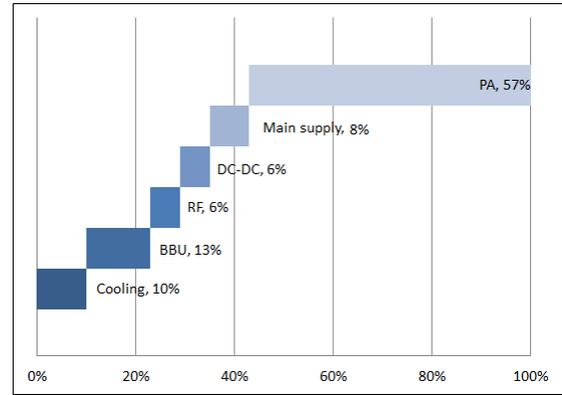


Fig. 3. Percentage power consumptions of various components of a BS.

and power loss occurs. In the case of LTE BSs, quite often the RF and PA components can be located close to the antenna, so as to eliminate the feeder cable losses. This layout is called Remote Radio Unit (RRU) or Remote Radio Head (RRH). The RF and PA components can even be integrated into the antenna. Besides allowing fewer losses, an additional advantage of the RRU layout is that in some cases cooling becomes unnecessary. In an LTE BS, the most power hungry element is the PA. LTE adopts an OFDM modulation scheme with high Peak to Average Power Ratio (PAPR), forcing the amplifier to operate in a linear region between 6 and 12 dB lower than the saturation point. This reduces the Adjacent Channel Interference (ACI), but increases power consumption. The chart in Fig. 3 shows the percentage contribution to power consumption of the various components of an LTE macro BS [3].

The peak power consumption of an LTE macro BS is about 1350 W, in the case of 3 sectors, 2 antennas/PA per sector, one carrier with 10 MHz bandwidth, and 2 x 2 MIMO. With a RRU configuration, the peak value decreases to about 800 W, thanks to a reduction of the energy consumed by the PA and by cooling. The actual instantaneous power consumption of a BS depends on the PA load, which in turn depends on the traffic carried by the BS. The relation between emitted power and the load ρ can be expressed by a linear relation:

$$P(\rho) = a + b\rho \quad 0 \leq \rho \leq 1, \quad (1)$$

where a is the consumption when the BS is active, but carries no traffic, and b is the load proportionality parameter. In an LTE macro BS, the parameters a and b depend on the number of antennas according to the following expression:

$$P(\rho) = N_{TX} (P_0 + \Delta_p P_{max} \rho) \quad 0 \leq \rho \leq 1, \quad (2)$$

where N_{TX} is the number of antennas, P_{max} is the maximum power out of the PA, P_0 is the power consumed when the RF output is null, Δ_p is the slope of the emission-dependent consumption. By properly optimizing the BS, the top devices can achieve a load proportionality of about 40%, meaning that at zero load the consumption is about 60% of that at full load, and, between zero and full load, power grows linearly with the load.

TABLE I
PARAMETERS OF THE CONSUMPTION MODEL FOR LTE BSs

Macro BS type	N_{TX}	P_{max} [W]	P_0 [W]	Δ_p
No RRU	6	20	130	4.7
With RRU	6	20	84	2.8

Table I reports the typical values of the parameters for an LTE macro BS with and without RRU.

III. SLEEP MODES OF BASE STATIONS

Usually, most of the components and subsystems of a BS admit various sleep modes that differ for the degree of functionality that is inhibited when the sleep mode is entered (except for the case that parts of TX/RX chain cannot be switched off separately). Each sleep mode has a specific energy consumption value that depends on how deep the sleep mode is, where deeper sleep modes are associated with higher degrees of reduction of the BS functionality, as well as to lower consumption. Each sleep mode requires a given wake-up time, defined as the time needed to restore the normal full functionality of a component or subsystem, starting from a given sleep mode. The deeper the sleep mode is, the longer the wake-up time is. Typical values range from tens of seconds to a couple of minutes for small cells, and up to 10-15 minutes for macro BSs.

Wake-up times have a strong impact on the feasibility of BS management algorithms. One of the major requirements of the operators for the adoption of sleep modes is the fact that they should not affect the quality of service provided to users, and that the network is guaranteed to operate smoothly and continuously. Thus, when using sleep modes, at the occurrence of unexpected traffic increases or bursts, i.e., when capacity is needed, sleeping BSs should be quickly powered on and return to the fully operational mode.

The definition of sleep modes, that include energy consumption values and wake-up times, is necessary to properly design the *BS management algorithm*. The BS management algorithm is the core mechanism to exploit sleep modes and consists of a strategy with which BSs in a given area enter or exit sleep modes; in other terms, the algorithm defines, at any instant of time, the network *configuration* that consists of the state of each BS.

While the most recent BSs have already been designed for frequently entering sleep modes, still most of the BSs that are deployed today were designed foreseeing only occasional switch-on and switch-off, so that the use of sleep modes for these devices might be limited. In particular, frequent changes in powering state might threaten the robustness of some components of the BS and finally lead to higher failure rates of the device, or parts of it. Depending on different kinds of BS management algorithm, the BS might need to enter a number of sleep modes that varies between one per day (with switches between two possible network configurations) to frequencies of the order of more than one per hour, when decisions about network configurations are taken after dynamic traffic estimates or measurements. A quite interesting issue is, thus, the assessment of the robustness of a BS to frequent

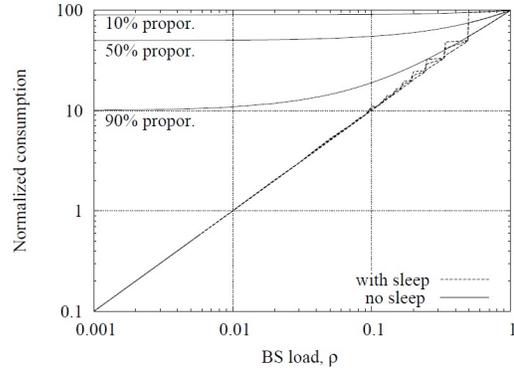


Fig. 4. Effect of the use of sleep modes, for different degrees of load proportionality

powering state changes, where with robustness we mean the lifetime of the equipment, the number and type of maintenance and repair events. For an evaluation of the feasibility of a BS management algorithm, these values should be related to the normal lifetime of a technology, that is the typical time between the technology deployment and the next generation of devices to be designed, produced and deployed in operative networks.

The practical implementation of a BS management algorithm requires that the BS can be remotely controlled, so that a decider engine can, in any instant, command the proper BS powering mode. This, in turn, implies some additional control processes of the BS hardware, some software or hardware interfaces to manage the BS powering state changes. Similarly, other parameters of the BS might need to be remotely controlled. For example, the RF transmission power or the antenna tilting, might need to be adjusted depending on the configuration of the network, i.e., depending on the powering state of the BS and of the adjacent BSs.

In order to quantify the relevance of sleep modes for power saving, consider a very simple and idealized scenario, where all cells in a given area are identical, and traffic is the same in all cells. Assume also that, when a given fraction of BSs is put to sleep, the other BSs can provide coverage. When the aggregate traffic of k cells is so low that one cell only can carry all of it, $k-1$ BSs can be put to sleep and one only can remain active. Then, given that the power versus load model is given by Eq. (1), and assuming that the consumption in sleep mode is given by P_{sleep} , the average energy consumption of the k BSs is:

$$P_S(\rho) = \frac{a + kb\rho}{k} + \frac{(k-1)P_{sleep}}{k} \quad \rho < \frac{1}{k} \quad (3)$$

Indeed, if the fraction of load in a cell is less than $\frac{1}{k}$, then just one BS out of k is capable of carrying the traffic, and the other $k-1$ can be put to sleep. Sleeping BSs consume a power P_{sleep} , while the only BS that remains powered on consumes the energy resulting from the traffic of k BSs. Assuming that we can always keep on the smallest number of BSs, Eq. (3) holds as far as $\frac{1}{k+1} < \rho < \frac{1}{k}$: when the load goes below $\frac{1}{k+1}$, a number k of BSs can be switched off out of $k+1$.

Figure 4 shows the behavior of the average consumption, $P_S(\rho)$, from Eq. (3), versus the load per cell, ρ . Different degrees of load proportionality are considered, where with load proportionality we intend the fraction of consumption that is proportional to load, that is $b/(a+b)$ according to the notation in Eq. (1). We consider 3 values: i) 10% of load proportionality, as is the case of most of the BSs that are deployed today, ii) 50% of load proportionality, as in the top devices that are under development, and iii) 90% of load proportionality, that is not realistic with today technologies and knowledge, but we consider it as a reference extreme case. Sleep modes have beneficial effects even when the devices are highly performing in terms of load proportionality.

IV. EXISTING BS MANAGEMENT SCHEMES

This section provides a short survey of BS management schemes which can provide significant benefits in terms of energy savings. The main goal is to keep QoS at acceptable levels and provide a self-balanced and self-organized operation of the critical nodes, while saving energy and providing load-proportional power consumption of the network. The control scheme applied by the proposed solutions depends on the architecture of the network and the different application scenarios. To this end, two different case studies have been identified. The first refers to a flat network architecture, where a single technology and a single type of BSs is assumed. In the second case, we explore a multi-tier architecture, where multiple access technologies or types of BSs can cooperate and provide, e.g., traffic offloading opportunities.

A. Flat Network

Flat network architectures are cellular network architectures consisting of one type of cells, e.g., either only macrocells or only microcells, or a combination of the two. In this case we assume a single technology that is under the administrative domain of one mobile operator. This means that the electricity costs related to the operation of the BSs are assigned to the operating expenditures (OPEX) of the operator. We distinguish two different types of sub-architectures, i.e., the *non-overlapping* architecture and the *overlapping* architecture.

Non-overlapping architectures assume that the coverage areas of the BSs do not overlap. This is mainly the case in microcell- or macrocell-only deployments. In that case BS switching on/off schemes can be achieved by X/Y switch-off patterns. Parameter X denotes the number of BSs that remain powered on, while Y is the number of BSs that can be switched off (set in sleep mode) [4], [5]. Since there is no overlapping between BSs, the BSs that are on should increase their cell range to provide additional coverage. The following two options exist. In the first case the BSs surrounding the candidate to be switched off slightly increase their cell range so as to fill the coverage hole created by the powered-off BS. In the second case, the peripheral BSs around the BS that remains on are switched off, while the central BS needs to largely increase its range, so as to fill the coverage hole resulting from the switch-off of the surrounding stations [6].

The decision on which approach to follow is a function of the architecture of the network and the possible cell range. To be more precise, the second approach, where the central BS covers the hole created by the switch-off of the surrounding BSs, requires an antenna tilt together with the increase of the RF output power. In general, when $X > Y$, a slight increase of the RF output power of the BSs that remain on is enough to provide coverage. In the other case (when $X \leq Y$), an increase of the RF output power of the BS and a reduction of the antenna tilt is required to provide coverage. In the *non-overlapping* architecture, the most commonly used control scheme that decides on the state of operation of the BS (on/off) or the antenna tilt, can be centralized or distributed. In the centralized approach, a central controller which has under its administrative domain all BSs, sends the required commands to the BSs, according to traffic/load criteria in each cell of the network. In the distributed case, each X/Y group of BSs in the network decides to change their state of operation. The grouping of BSs in the X/Y scheme is made by the central controller of the network, or by the mobile network operator [6], [7]. For real application scenarios, a centralized control scheme is usually preferred, so as to eliminate probabilities of coverage holes in the network. Of course, this increases the complexity of the algorithm, but guarantees reliable QoS.

Overlapping network architectures assume that the BS coverage areas may overlap. This is a typical case of macro-micro arrangements over non-uniform traffic areas, where micro stations are used to provide the required capacity under the coverage umbrella of macro stations. In such topology, two types of BSs are distinguished: the *critical* stations, usually the macrocell, which cannot be put into sleep mode, due to coverage issues, and the *flexible* stations, which can be set in sleep mode. The flexible stations are usually microcells that are deployed under the macrocell coverage. Since there is an overlap between critical and flexible stations, there is no need for an increase of the cell range of the BSs that remain on. This makes the procedure of energy management much easier and more reliable for the QoS of the network, compared to flat architectures. The reason is that coverage holes will never appear in the network. The control scheme that can be implemented in this case can be centralized, distributed or pseudo-distributed. Due to the negligible probabilities of creating a coverage hole, centralized control is avoided, to keep complexity at low levels. In the distributed case, each flexible BS decides on the state of its operation (on/off), independently of the conditions of the neighboring cells. In the pseudo-distributed case, flexible stations are assigned to critical stations, usually under cell overlap criteria [8]. Self-organization can also be considered, as presented in [9]. The critical stations, considering traffic load criteria within their cells, but also according to traffic conditions in the flexible cells, decide on the state of operation of the flexible stations.

B. Multi-Tier Network or Offloading

Multi-Tier Network architectures exist when the cellular macro-micro network co-exists and can cooperate with other

technologies, e.g., small cell networks [10]. A typical example is provided by the femtocell or WiFi layer. In both cases, the user can connect or roam between the multi-tier networks under user association rules. The main objective is to provide an online user association algorithm (or offloading solution), which will satisfy objectives that are related to energy efficiency and load control of the network, which falls under the administrative domain of the mobile operator. This network is the macro-micro architecture which is responsible for the very high OPEX of the operator. Based on current femtocell and WiFi technologies, the power consumption of small cell networks is almost flat, and independent of the traffic served at the node. For this reason, such an offloading scheme is fair in terms of electricity cost 'migration' and the only issue to be addressed is the QoS at both sides. The major benefit for the mobile operator is longer switch-off time periods of the BSs, but also the support of switch-off deactivation even in the case of peak traffic.

For the case of femtocell layer cooperation, the main constraint is that femtocells should be in open access, and thus be able to absorb traffic from the cellular network [11]. Open access can be managed by the mobile operator (since femtocells are connected to his network) or it can be left on the decision of the femtocell owner. The second issue that must be addressed is the user-to-femtocell association rule. There are two strategies to follow. The *no priority* user association rule treats BSs and femtocells as equal priority stations, and users are assigned according to best server criteria. The second case, *femto priority*, assumes that the user is always connected to a femtocell in case there is adequate coverage, independent of the channel condition from the macro-micro layer. With this approach, more users are connected to femtocells, and thus greater savings are expected. The most practical control and coordination scheme is the pseudo-distributed case, where critical stations should be able to control the flexible stations, but also the status of the access of the femtocells that will be assigned under its administrative domain [8]. In this case, the mobile operator is relieved of managing a large number of femtocells. The controller of the critical stations decides when to initiate open access at femtocell layer, according to external conditions that can be related to RES capacity, traffic loads, electricity prices, etc.

For the WiFi network the situation is different. Mobile operators usually do not have the administrative rights to the WiFi APs. The cooperation between the macro-micro cellular network and WiFi APs is a user association rule in multi-Radio Access Technology (multi-RAT). To support such connectivity, mobile devices that support both WiFi and cellular technologies can automatically, or under mobile operator command flow, migrate data in the preferred network. Furthermore, with the development of multipath-TCP techniques [12], it is feasible to transmit data in both networks, equally or unequally weighted along the two paths (WiFi and cellular). Similar to the femtocell case, offloading is achieved at the administrative domain of the mobile operator, with a negligible increase of electricity cost for WiFi owners. Regarding the control

scheme, if the mobile operator would like to manage the route of data, then broadcasted signals to mobile terminals that will set the preferred radio access technology should be transmitted. Following a similar procedure as with the femtocell layer cooperation, *other RAT priority* or *no priority* should be implemented, so as to decide the user-to-technology association.

V. IMPLEMENTATION OF MANAGEMENT SCHEMES

Equipments and networks require the ability to exploit existing enablers and to deal with considerable constraints and limitations, which risk reducing the expected energy savings. In this section, the limitations and constraints of the integration of the green approaches with existing systems are addressed.

A. Limitations and Constraints for Integration With Existing BS System

BS management solutions are impacted by real network architectures and deployments. The most important limitation concerns algorithms aiming at turning off base stations in single-layer deployments. In fact, in this scenario, it is extremely difficult to guarantee that coverage holes do not occur somewhere in the area around a BS entering sleep mode; the problem is particularly critical for indoor or deep indoor environments. The same consideration applies also in the case of home femtocells and access points, which may be responsible for providing coverage (and not only capacity) in some specific places, e.g., basements, urban canyons, etc. Such risk is unacceptable by mobile operators, and this drastically reduces the chances to see such algorithms deployed in commercial networks. Cell-compensation mechanisms can be a solution only in case of very high guarantees against coverage failures. On the contrary, algorithms exploiting the co-existence of several layers are expected to be the more and more utilized in future networks, such as the HetNets (Heterogeneous Networks, with large umbrella cells and several small cells under their coverage); similarly, multi-RAT approaches gain momentum, as a valid way to increase the mobile network capacity.

Besides, BS management algorithms need to take into account real equipments hardware and software limitations. The most important of them concern the possibility of a base station to actually enter into some energy-saving mode, the amount of saving achievable by these modes, and the transient time for moving from these modes to a fully operational state (active state). Currently deployed base stations (of whatever technology) do not implement such low-power-states, and can usually work into two states only: fully active or complete switch-off (apart for the backhaul network interface, which remains active to receive switch-on commands). Algorithms based on long time-scales (i.e., tolerating off to on transients in the order 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 stations, using the aforementioned OA&M interface. Algorithms requiring shorter wake-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 fully active mode depends on the number (and type) of components that are switched off (see, e.g., [13] for an analysis on 3G femtocell, and [14] for WLAN APs). A cold-restart requires several minutes to complete, with the delay increasing with the size of the base station. Modern base station power profiles (e.g., [15]) usually include such limitations, and should be used when validating the algorithms performance. When considering faster algorithms, the new generation of power amplifiers implementing dynamic gating and drain voltage control enable symbol-time-scale algorithms (1 ms and below) to work correctly and efficiently.

B. Issues to be Addressed for Future Equipments

With respect to the predictions of the evolution of technology, equipment, and components, some risks of limiting the effectiveness of energy-efficient BS management solutions can also be identified. For example, the tendency to integrate components onto a single chip, such as in the case of Radio-on-Chip (RoC), Digital-on-Chip (DoC) and System-on-Chip (SoC), presents the risk of limiting the flexibility of turning off different parts of such components. In fact, while allowing more efficient circuits, integration of several components implies the inability to turn off some parts of them without turning off the others. This may result in components that can be no longer turned off for energy savings. Similar problems come from the centralization of the baseband processing, like in the BBU pooling concept. In such cases the granularity of a BBU component can be the one of several cells, making it more difficult to turn off some of the cells in a fine-tuned way. The impact of longer term evolutions, like the advent of Switched Mode Power Amplifiers (SMPA), requires further investigation.

Additional attention must be brought to the internal hardware architecture of a BS. For example, modern multi-carrier and multi-technology BSs (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.

VI. SUMMARY AND CONCLUSIONS

Cellular networks, originally dimensioned to meet the peak of the user demand, are now facing the forecasted traffic explosion, and striving for the deployment of more infrastructure. Therefore, energy consumption starts to play an increasingly important role in the overall operation expenditures, and mobile network operators are keen on finding viable solutions to cut their energy bill. It has been shown in this paper that developing more energy-efficient hardware is only partially solving this problem, and thus there is a need to look for better alternatives to more effectively cut the overall energy consumption. As most of the energy consumed in cellular

networks is attributed to the BS operation, the development of energy-efficient BS management schemes seems to be of utmost importance. To this end, this paper has identified the most prominent examples of BS management algorithms, which can be applied in different network architectures that are currently present in the cellular network landscape. Energy-saving approaches applied in flat, multi-Tier networks and offloading technology have been discussed. To complement this picture, the aspects relevant to the implementation of BS management schemes in current and future infrastructures have been identified, to indicate the key drivers that could facilitate the adoption of the proposed energy saving schemes.

ACKNOWLEDGMENT

This work has been supported by the TREND project (Towards Real Energy-efficient Network Design, grant agreement n. 257740), a Network of Excellence funded by the European Community's Seventh Framework Programme.

REFERENCES

- [1] S. Lambert, W. V. Heddeghem, W. Vereecken, B. Lannoo, D. Colle, and M. Pickavet, "Worldwide electricity consumption of communication networks," *Optics Express*, vol. 20, no. 26, pp. B513–B524, Dec. 2012.
- [2] W. Vereecken, W. Van Heddeghem, M. Deruyck, B. Puype, B. Lannoo, W. Joseph, D. Colle, L. Martens, and P. Demeester, "Power consumption in telecommunication networks: overview and reduction strategies," *IEEE Communications Magazine*, vol. 49, no. 6, pp. 62–69, 2011.
- [3] EARTH (Energy Aware Radio and network technologies) project, "Energy efficiency analysis of the reference systems, areas of improvements and target breakdown," 2012.
- [4] M. A. Marsan, L. Chiaraviglio, D. Ciullo, and M. Meo, "Optimal energy savings in cellular access networks," in *IEEE International Conference on Communications (ICC '09) Workshops*, Jun. 2009, pp. 1–5.
- [5] —, "Multiple daily base station switch-offs in cellular networks," in *4th International Conference on Communications and Electronics (ICCE '12)*, Aug. 2012, pp. 245–250.
- [6] Z. Niu, Y. Wu, J. Gong, and Z. Yang, "Cell zooming for cost-efficient green cellular networks," *IEEE Communications Magazine*, vol. 48, no. 11, pp. 74–79, 2010.
- [7] L. Chiaraviglio, D. Ciullo, G. Koutitas, M. Meo, and L. Tassiulas, "Energy efficient planning and management of cellular networks," in *IEEE Wireless on-demand Network Systems and Services (WONS '12)*, Jan. 2012, pp. 159–166.
- [8] S. Kokkinogonis and G. Koutitas, "Dynamic and static base station management schemes for cellular networks," in *IEEE Global Communications Conference (GlobeComm '12)*, Dec. 2012.
- [9] K. Sambanis, D. Kutscher, and M. Bruneer, "Self organized energy efficient networks," in *IEEE Personal Indoor and Mobile Radio Communication (PIMRC '10)*, Sep. 2010, pp. 1665–1670.
- [10] S. Bhaumik, G. Narlikar, S. Chattopadhyay, and S. Kanugovi, "Breathe to stay cool: adjusting cell sizes to reduce energy consumption," in *The 1st ACM SIGCOMM workshop on Green Networking*, 2010, pp. 41–46.
- [11] A. De Domenico, R. Gupta, and E. Calvanese Strinati, "Dynamic traffic management for green open access femtocell networks," in *IEEE 75th Vehicular Technology Conference (VTC '12-Spring)*, May 2012, pp. 1–6.
- [12] C. Pluntke, L. Eggert, and N. Kiukkonen, "Saving mobile device energy with multipath tcp," in *Proceedings of the sixth international workshop on MobiArch*, 2011, pp. 1–6.
- [13] I. Haratcherev and A. Conte, "Practical energy-saving in 3g femtocells," in *IEEE Green Broadband Access (GBA) workshop, in conjunction with ICC 2013*, Jun. 2013.
- [14] I. Haratcherev, M. Fiorito, and C. Balageas, "Low-power sleep mode and out-of-band wake-up for indoor access points," in *GLOBECOM Workshops, 2009 IEEE*, 2009, pp. 1–6.
- [15] NTT DOCOMO, Alcatel-Lucent, Alcatel-Lucent Shanghai Bell, Ericsson, Telecom Italia, "[r1-113495], base station power model." [Online]. Available: http://www.3gpp.org/ftp/tsg_ran/wg1_r11/TSGR1_66b/Docs/R1-113495.zip