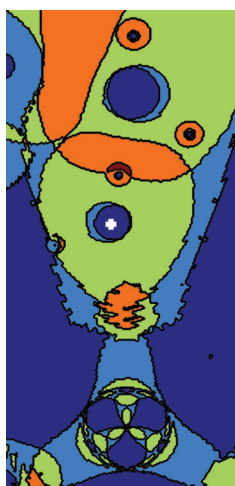# CELL WILTING AND BLOSSOMING FOR ENERGY EFFICIENCY

ALBERTO CONTE AND AFEF FEKI, ALCATEL–LUCENT BELL LABS
LUCA CHIARAVIGLIO, DELIA CIULLO, AND MICHELA MEO, POLITECNICO DI TORINO
MARCO AJMONE MARSAN, POLITECNICO DI TORINO AND INSTITUTE IMDEA NETWORKS

ABSTRACT

In this article, we consider the adoption of sleep modes for the base stations of a cellular access network, focusing on the design of base station sleep and wake-up transients.. We discuss the main issues arising with this approach, and we focus on the design of base station sleep and wake-up transients, also known as cell wilting and blossoming. The performance of the proposed procedures is evaluated in a realistic test scenario, and the results show that sleep and wake-up transients are short, lasting at most 30 seconds.

TRENDS IN LOW-POWER BASE STATIONS

The amount of data exchanged over mobile networks has increased exponentially over the past few years, mainly because of the introduction of a variety of smartphones, and of the iPhone in particular. This trend is not expected to slow down in the near future, as the mobile market is continuously enriched with new services and innovative wireless-enabled devices in the fields of consumer and professional electronics (e.g., photo cameras, e-books), captors and sensors (e.g., fire alarms), and machine-to-machine devices (e.g., cars, parking, vending machines).

A fundamental element of this (r)evolution is the omnipresence of high-capacity connectivity, enabled by the fast development of broadband wireless technologies: today third generation/ Universal Mobile Telecommunications System (3G/UMTS), tomorrow fourth generation (4G), in the form of Long Term Evolution (LTE), as well as the process of densification [1] of wireless networks, with the deployment of short-range low-power base stations (BSs) to create so-called pico- and femtocells. This last evolution is a major trend in the current growth of mobile networks, and is based on a well-known law in wireless systems: the gain resulting from the reuse of the same radio resource in non-contiguous cells drives the increase of the number of BSs deployed within a given area (i.e., their density).

Cell densification is becoming a key technique in wireless networks deployment, as it enables the increase of capacity and energy efficiency in wireless networks [1]. The reason is that in a non-empty space (i.e., in the vast majority of real cases) the power necessary to provide service to a user increases quicker than the distance between the BS and the user's mobile terminal (MT) (mainly due to the attenuation undergone by the transmission). As an example, the total power consumed by four BSs is generally smaller than the consumption of a high-power BS covering an equivalent area, and the capacity is about four times higher. Unfortunately, the first wave of densification is occurring as an additional layer to existing macrocell BSs, aiming at absorbing the peaks of capacity demand, and not as their replacement, an evolution expected to happen in a longer timeframe. Densified wireless networks are thus composed of superposed small (low-power) cells and large (high-power) cells, and are generally referred to as heterogeneous networks.

Currently, heterogeneous networks are natively overdimensioned: when traffic is low, the probability of unused (i.e., unnecessary) BSs increases considerably, particularly when the cell size is small.

NEED FOR SLEEP MODES

Unused BSs are one of the main sources of energy waste in current wireless networks. This is due to the fact that today BSs (regardless of their technology: UMTS, LTE, WiMAX, or WiFi) need to permanently signal their presence and the availability of the cellular service. In addition, the current BSs are also designed to continuously listen to the radio environment in order to detect incoming MTs. Continuous emission and reception are always-running processes, which consume energy even when no MT is using or requesting the BS service (e.g., at night).

With the above mentioned issues in mind, the application of sleep modes to BSs constitutes a promising approach to improve energy efficiency. Sleep modes allow turning off BSs when and where they are not necessary, especially in low-

The authors consider the adoption of sleep modes for the base stations of a cellular access network, focusing on the design of base station sleep and wake-up transients.

load periods, such as nights, weekends, and holidays. Up to 20–40 percent of energy can be saved with this approach [2, 3].

Switching off a portion of the deployed BSs in a given area guarantees energy savings, even when using commonly available chipsets. This technique stands as a complementary approach to the design and development of specific energy-efficient components (e.g., highly efficient power amplifiers). While the latter approach is acceptable for high-end and high-cost macrocell BSs, the possibility to use commonly available electronic components is a fundamental aspect when designing low-cost and low-power BSs. In addition, since the radio emission is interrupted during sleep modes, interference between cells is also reduced, and unnecessary radiation is eliminated. For all these reasons, increasing attention in the research community and telecom industry is paid to such energy-efficient mechanisms.

SLEEP MODE BASICS

The design of a BS sleep mode must satisfy three requirements: minimize the consumed energy, guarantee the availability of wireless access over the service area, and minimize the perceived degradation of the user experience.

These targets drive the definition of the two primary sleep mode mechanisms:
- Sleep mode entrance: specifies when, how, and which BSs enter sleep mode with minimal degradation of the service offered to end users
- Sleep mode exit: specifies when, how, and which BSs wake up to respond to traffic demand increases, with minimal degradation of the service offered to end users

SLEEP MODE ENTRANCE

Sleep mode entrance design must address two aspects: detect when a BS can enter sleep mode, and describe how the transition from active to sleep state is implemented. A simple trigger of sleep mode start can be obtained by monitoring BS activity to detect when it is unused or underutilized. A trivial case is when the cell is empty for a certain period (i.e., no service is active). As we observed, the likelihood of this event increases as the size of the cells diminishes. Further gains in terms of energy saving can be obtained by considering that, in low load conditions, neighboring BSs can take the responsibility of the traffic of a BS that is turning off. Operationally, this technique requires that we:
- Identify the set of neighboring BSs that can absorb the load of the BS to be switched off
- Check that the neighboring BSs can effectively be in charge of the additional load, possibly by applying specific compensation actions, like up-tilting their antennas or slightly increasing their emitted power

Once the sleep mode trigger is confirmed, the transition from active to sleep state can start. This transition must be designed in such a way that any interruption of ongoing calls/sessions is prevented. Indeed, in several systems, such as 3G systems, it is not always possible for the network to order all the MTs to hand over out of a cell that is going to be switched off, because the
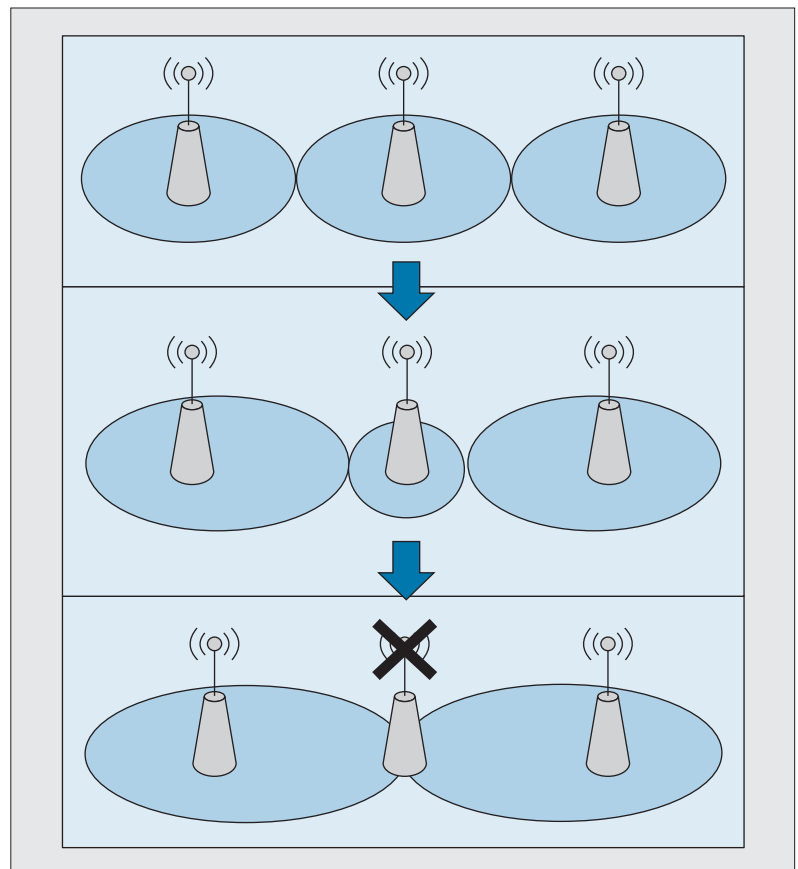

Figure 1. *Example of BS wilting procedure.*

interference among adjacent cells might be too large to allow for handovers; in some cases, an abrupt power-off of a BS can induce call interruptions. Call drops can be avoided by adopting a slow reduction of the BS transmit power. This technique is commonly named BS *wilting* or *progressive switch-off*. Figure 1 illustrates the BS wilting concept: three BSs are initially active, and offer service to the MTs that are camping in their cells. When the central BS starts its switch-off transient, it progressively reduces its transmitted power so that its coverage shrinks while the area of adjacent cells expands.

During this process, MTs served by the central BS can hand over to neighboring BSs. At the end of the transient, the central BS is in sleep mode, and coverage is provided by the neighboring cells.

During BS wilting, if any MT experiences critical degradation without any possibility to migrate to another BS, it can signal to the BS, which in turn can suspend (and even revert) the process. Idle MTs camping under the BS also need to be warned of the ongoing process, and be invited to relocate away. If they fail to relocate, they should alert the BS to stop the wilting process. A cell barred status is usually signaled by the wilting BS, alerting the idle MTs and avoiding new MTs moving in from neighboring cells.

SLEEP MODE EXIT

Sleep mode exit (BS wake-up) should be triggered when the data traffic conditions imply a high risk of overload in neighboring active BSs,

Base station wilting and blossoming transients must be very gradual, to allow mobile terminals to hand over toward the base station that can serve them best after the transient.

or an unacceptable quality of service for end users. An obvious solution is thus to allow the active BSs to alert sleeping neighbors as soon as their load exceeds a given threshold, close to saturation. In this case, the sleeping BSs must receive a wake-up signal on their backhaul interface (i.e., the dedicated link connecting the BS to the core network) and turn on rapidly. In this approach, the challenge resides in choosing a good threshold to wake up the BS, neither too early, to avoid unnecessary interference and energy waste, nor too late, to avoid poor performance. As a general rule of thumb, this kind of parameter is initially evaluated through extensive network simulations, and then adapted to operating conditions at runtime. The transition from sleep to active mode must be implemented so that denial of service and call drops are prevented. This can be achieved through a progressive switch-on process, known as BS blossoming, during which the power of the BS is slowly raised to its nominal target value. Other challenging issues are the criteria and techniques to detect and locate the traffic peaks, to wake up the right BS or set of BSs.

User-Friendly Wilting and Blossoming Transients

Current and future cellular technologies, such as UMTS and LTE, are essentially interference-limited; hence, the quality of the network operations can be jeopardized by abrupt changes in transmitted power levels. For example, an MT that receives a strong signal from a nearby BS typically cannot hear the signal of other BSs, and loses its connection to the network if the nearby BS is switched off too quickly. Conversely, an MT that receives a weak signal from a distant BS loses its connection if a nearby BS is switched on too fast because of the generated strong interference.

This suggests that the BS wilting (a.k.a. wake-up or switch-on) and blossoming (a.k.a. sleep or switch-off) transients must be very gradual, to allow MTs to hand over toward the BS that can serve them best after the transient. However, gradual transients can reduce the effectiveness of a BS sleep mode scheme if they are very long, since the power consumed by the BS during the transients is typically quite close to full operational power.

In this section we consider a small portion of a realistic cellular network, and examine the case of gentle BS wilting and blossoming transients, computing their expected duration with an accurate signal propagation model based on a tool developed by Alcatel Lucent called Wireless System Engineering (WiSE). The tool analyzes propagation in complex environments involving real building configurations (including construction materials) by means of ray tracing techniques [4].

A preliminary analysis of the duration of gentle switch-off transients was originally reported in [5], based on a much less accurate propagation model.

We look at BS X, whose sleep (or wake-up) decision is made at instant $t_0 = 0$, based on some sleep mode start decision.

In the case of a BS sleep transient, we need to derive a *power reduction profile* for X that jointly guarantees three conditions:
• The MTs connected to X can handover to neighboring BSs, and only very few (if any) are dropped
• The handover overload for all BSs is kept small, i.e., the number of simultaneous handovers never exceeds the maximum defined by the operator
• The transient duration is as short as possible, compatibly with the requirements of the network and the users

Similarly, in the case of a BS wake-up transient, we need to derive a *power increase profile* for X that jointly guarantees that the MTs connected to neighboring BSs can handover to X, and only very few (if any) are dropped, as well as the second and third conditions above.

The power reduction/increase profile is defined by the function $P_X(t)$ that specifies the transmitted power of BS X at time t. For the profile we use a step function such that from time t_i to time t_{i+1} the transmission power of X is constant and equal to $P_X(t_i)$, with $i = 0,1, \cdots, M$ and $t_0 < t_1 < \cdots < t_M$. In the sleep transient case, $P_X(t_0) > P_X(t_1) > \cdots > P_X(t_M) = 0$ and the typical normal transmission level of X is $P_X(t_0)$. At time t_M, BS X is switched off. In the switch-on case, $P_X(t) = 0$ for $t < t_0$ and at t_0 the wake-up phase starts, with $P_X(t_0) > 0$, and $P_X(t_0) < P_X(t_1) < \cdots < P_X(t_M)$. At time t_{M-1} the BS X is already operating at its typical transmission power, and at time t_M the transient can be considered completed.

The duration of the wilting and blossoming transient is thus given by t_M in both cases.

In this article, as an example, we examine the case in which the power emitted by the BS being switched off is progressively halved, i.e., $P_X(t_1) = P_X(t_0)/2$, $P_X(t_2) = P_X(t_1)/2$, and so on, up to the minimum transmission power $P_X(t_M-1)$; at the following step the BS is switched off, $P_X(t_M) = 0$. For a BS wake-up transient, we examine the case in which the power emitted by the BS is progressively doubled, i.e., $P_X(t_1) = 2 P_X(t_0)$, $P_X(t_2) = 2 P_X(t_1)$, and so on, up to full transmission power $P_X(t_{M-1})$. Clearly, other power reduction/increase profiles might be decided, according to operator-specific criteria, such as minimizing the transient duration, making constant the expected number of handovers per power level, and so on.

At the beginning of the sleep transient, X is transmitting at its typical value, $P_X(t_0)$. We compute in this situation the number of MTs connected to X that also hear other BSs. Assuming that at time t_0 these MTs start to handover out of X, because X signals that it is about to switch off, we compute the duration of the time interval it takes to complete the handovers. Time t_1 is set equal to this duration. At t_1, X decreases its transmission power to $P_X(t_1) = P_X(t_0)/2$. Given that the power decreased, some MTs that are connected to X start also hearing other BSs and can handover out of X. We again derive how much time it takes for these handovers, and this duration defines the amount $t_2 - t_1$ that X needs to spend at transmission power level $P_X(t1)$. We iterate this process up to the minimum power

level $P_X(t_{M-1})$. The possible remaining MTs that cannot hear other BSs but X at power level $P_X(t_{M-1})$ are dropped when, at t_M, X is completely switched off. The risk of dropping MTs is obviously a negative effect of the sleep transient. The occurrence of dropping, however, can be reduced to an acceptable level by appropriate power reduction profiles with very low values of the transmit power in the last steps of the transient. In addition, when strict dropping avoidance is required, a more conservative approach can be adopted by stopping (and possibly reverting) the wilting process as soon as an MT raises a "losing connectivity" alert. This kind of alert is commonly available on current cellular technology, for example, event 2D in wideband code-division multiple access (WCDMA)/UMTS technology [6] and event A2 in LTE [7].

Instead, at the beginning of the switch-on period, X starts transmitting at power level $P_X(t_0)$. We compute in this situation the number of MTs that receive the signal of X at higher power than for other BSs. At time t_0 these MTs start to hand over toward X, and we compute the duration of the time interval it takes to complete the handovers. Time t_1 is set equal to this duration. At t_1, X doubles its transmission power to $P_X(t_1) = 2$ $P_X(t_0)$. Given that the power is increased, some more MTs start hearing X better than the other BSs and can hand over to it. We again derive how much time it takes for these handovers, and this duration defines the amount of time $t_2 - t_1$ X needs to spend at transmission power level $P_X(t_1)$. We iterate this process up to the maximum power level $P_X(t_{M-1})$; during the time interval from t_{M-1} to t_M, some more MTs hand over toward X, and at time t_M the transient can be considered complete. At each power increase step, the stronger signal generated by X risks causing excessive interference and call dropping, for MTs that are close to X but connected to other BSs. For this reason, the power increase should be small enough to guarantee that MTs connected to other BSs are not disconnected and have the chance to hand over.

A more formal derivation of the average duration of sleep transients is presented in [5].

A Case Study

As a case study, we consider a portion of the central area of the city of Munich, Germany, which corresponds to a square of 800×800 m comprising 1 macrocell, 8 microcells, and 10 femtocells. A map of the considered area, together with a side view, is presented in Fig. 2.

The received power for each location and for each BS is computed adopting a planning tool based on ray tracing with up to four reflections. The granularity of the evaluation is a square of 2 $\times$ 2 m. The evaluation of the received power is made at an elevation of 1 m, the typical height of an MT in the pocket. We assume that the MTs are uniformly distributed in the service area and that their density is ρ. Unless otherwise specified, we set $\rho = 5 \times 10^{-3}$ users/m^2, corresponding to about 3200 users in the considered area (i.e., a value that can correspond to intermediate load).

Let H be the maximum number of handovers

Figure 2. *Case study: Map with cell identifiers (SC = micro cell, F = femto cell). Aerial view (top) and side view (bottom).*

that can occur simultaneously toward a new BS (due to signaling channel limits), and let T_H be the time interval required to complete one handover procedure. Typically, the percentage of handovers in a cell can involve around 30 percent of the active MTs, but the actual number depends on many factors; we therefore let H span over a quite large interval, setting H equal to 2, 10, or 20.

When an MT starts a handover, it needs to identify the new BS. The identification is fast if the carrier frequency of the new BS is the same as that of the old one. A longer time interval is required if the carrier frequency changes, but we do not consider this case. Typical durations for the identification time and handover time are T_I = 0.8 s and T_H = 0.5 s, respectively, as indicated in [8]. Table 1 reports the main parameters for each cell type.

Wilting and Blossoming Transient Durations

First, we consider the scenario reported in Fig. 2 and evaluate the durations of the sleep and wake-up transients for the BSs of microcell SC_1 and femtocell F_2. We then modify the scenario by considering the case in which no macrocell is present.

We consider first SC_1, whose transmission power in normal conditions is $P_{SC1}(t_0) = 1$ W,

and we focus on the sleep transient. The interference map when all the BSs transmit at their maximum power is reported in Fig. 3. BS SC_1's position is indicated with a white marker. The vertical bar in the figure indicates how many BSs an MT can hear in each point.

We assume that the transmission power of SC_1 is halved at each step for a total of $M = 8$ steps, from 1 W to 1/128 W, before going to zero.

Table 2 reports:
• The time t_1 needed to perform, before starting the actual transmit power reduction, all the possible handovers out of X
• The time $t_M - t_1$ necessary to hand over the MTs while the power of SC_1 decreases
• The total transient duration, t_M

The last row of the table indicates the average number of handovers to perform during times t_1, $t_M - t_1$, and t_M. Observe that the procedure allows no user to be dropped (instead of the 72.4 that would have been dropped on average with an uncontrolled and abrupt switch-off of SC_1) at the cost of an additional delay before entering the sleep mode between 10 and 26 s.

The power decrease profile is detailed in Fig. 4. The sleep transient duration is less than 26 s in the worst case ($H = 2$), a negligible amount with respect to the typical periods for which a BS is in sleep mode in low traffic conditions.

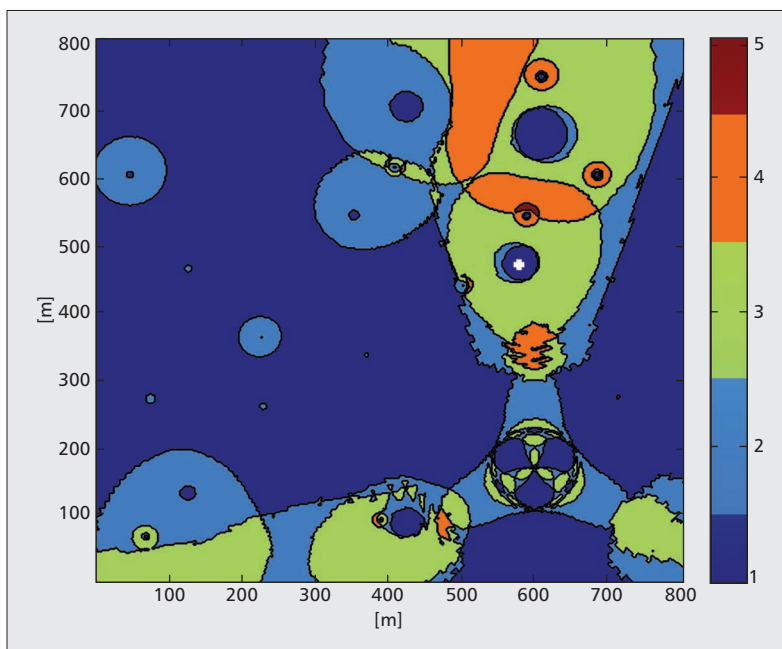| Size	Type	Power [W]	Height [m]	Azimuth	Tilt
Macrocell	Tri-sector	39.81 (sect.)	33	(90,210,330)	–5
Microcell	Isotropic	1	4	0	0
Femtocell	Isotropic	0.02	2	0	0

Table 1. *Cell parameters.*

Figure 3. *BSs interference areas.*

Of course, different assumptions can be made regarding the power reduction profile. As an example, we also computed sleep transient durations for a linear power reduction profile in which at each step the power decreases by 0.1 W; the transient lasts, in this case, between 15 and 29 s.

The computed durations for the cell wilting transients, of at most 30 s, depend on the user density, and decrease if fewer users are camping in the cell area. We can thus confidently assume that the time lapse required for BS switch-off in periods of low traffic is shorter than half a minute.

Consider now the wake-up transient for SC_1, with a power increase profile composed of exponential steps (power is doubled at each step). The durations of the power level steps is shown in Fig. 5. Again, the transient lasts about half a minute, with the assumed parameters and user density.

To study the sleep and wake-up transients for femtocell BSs, we consider the case of F_2. Usually, the number of MTs served by a femtocell can be up to four or five. We observe in Fig. 3 that all MTs connected to F_2 can also hear other BSs. This means that before the switch-off, we can force the MTs connected to F_2 to hand over to another BS. Thus, in this case, we do not need to progressively decrease the transmission power of the femtocell BS, since we can switch off F_2 as soon as the MTs are transferred to other BSs. The average number of users connected to F_2 is $N_h(t_0) = 0.18$. This means that the average time needed to switch off the femtocell BS F_2 is equal to one handover time: $t_1 = t_M = 0.5$ s. Therefore, the switch-off time of a femtocell BS can be assumed to be on the order of a second, or very few seconds at most. Similarly, the wake-up transient does not require a gradual increase of power. When the femtocell BS switches on, some MTs start hearing it and are requested to hand over to it. The duration of this process is again on the order of a few seconds.

No Macro Cell

We now study the same scenario, with only micro- and femtocells, to consider a setting with lower capacity and less uniform coverage. For the sake of brevity, we focus on sleep transients only. Since now no macrocell is present, to take into account the reduced capacity we use a smaller user density, $\rho = 10^{-3}$ users/m^2.

The interference map for this case is reported in Fig. 6. We select again SC_1 as the microcell BS to enter sleep mode.

Table 3 refers to the power reduction profile in which transmit power is halved at each step. The average number of users that are dropped at the end of the transient when the BS is finally switched off is 0.1, which is a negligible value.

Figure 7 reports the power reduction profiles for both the cases in which, at each step, the power is halved or decreased by a constant amount equal to 0.1 W. The transients always last about half a minute in this case as well. Note that with the linear power reduction profile, the last step is much higher than in the exponential case. This implies the risk of disconnecting a non-negligible number of users when the BS power goes to zero.

H	t_1 [s]	$t_M - t_1$ [s]	t_M [s]
2	16	9.3	25.3
10	3.5	7.8	11.3
20	2	7.8	9.8
No. of handovers	63	9.4	72.4

Table 2. *Switch-off transient for SC$_1$: average switch-off times and number of handovers.*

CHALLENGES FOR ADVANCED OPTIMIZATIONS

Cell sleep mode is not a last-minute novelty: rudimental sleep modes have been in use in cellular network management for several years. They are usually based on static configuration of on-off periods, defined through the analysis of traffic history (e.g., night vs. day, or business vs. residential hours), careful network planning, as well as post-treatment of alerts and traces to detect possible problems. Today, research efforts focus on the definition of advanced optimization algorithms to further improve the performance in terms of energy savings and intercell interference reduction. Some recent relevant works in this field consider the switch-off of some BSs [2, 3, 9, 10, 11] or active transmitters [12]. Other approaches, instead, focus on planning with energy and spectral efficiency optimization as the primary design target [13, 14].

One of the major open challenges resides in the possibility to dynamically detect and locate capacity needs, and to turn on/off BSs accordingly in (quasi) real time. This leads to a difficult (and currently unsolved) optimization problem: how to set the best combination of active and sleeping BSs in a given area in order to satisfy the needs of users, while at the same time achieving maximum energy savings.

Another important issue to be solved concerns the difficulty of guaranteeing continuous coverage and the avoidance of coverage holes when switching off a base station. This is indeed very difficult, as radio propagation is strongly affected by the environment and thus poorly predictable. Even the progressive cell wilting process described previously is not error proof. As an example, consider a house where a femtocell BS enters sleep mode after the last MT is switched off. Later, the house owner may wish to switch on his MT in a location where no other BS can be heard. This is clearly a deadlock: the femtocell BS is off and cannot detect the MT, and the MT has no means to signal its problem. To solve this issue, a new approach is emerging [15] that proposes to move from an avoidance problem to a handling problem: coverage holes' existence is accepted, and MTs shall be allowed to transmit a blind wake-up message (a.k.a. reverse-paging message) to ask for any BS to switch on.
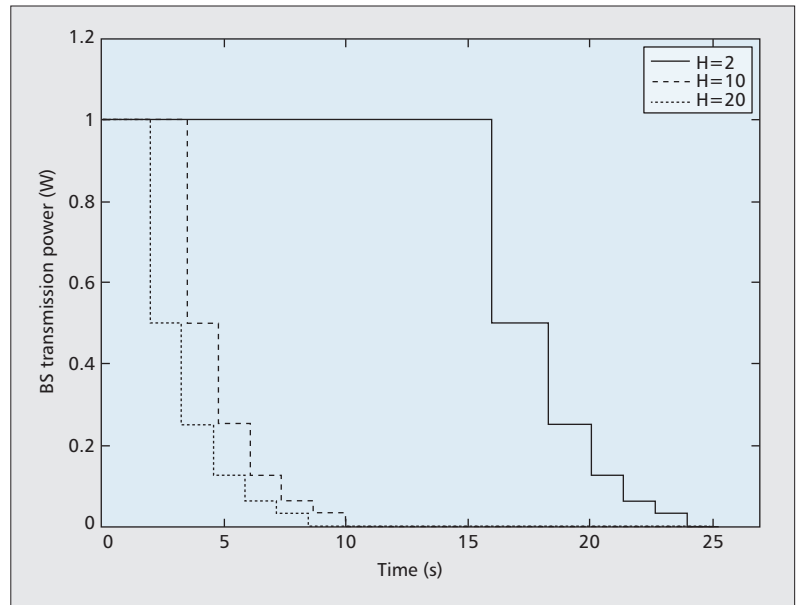
Coverage hole handling (through blind wake-

Figure 4. *Switch-off transient for SC$_1$. Exponential power decrease profile, with H = 2, 10, 20.*
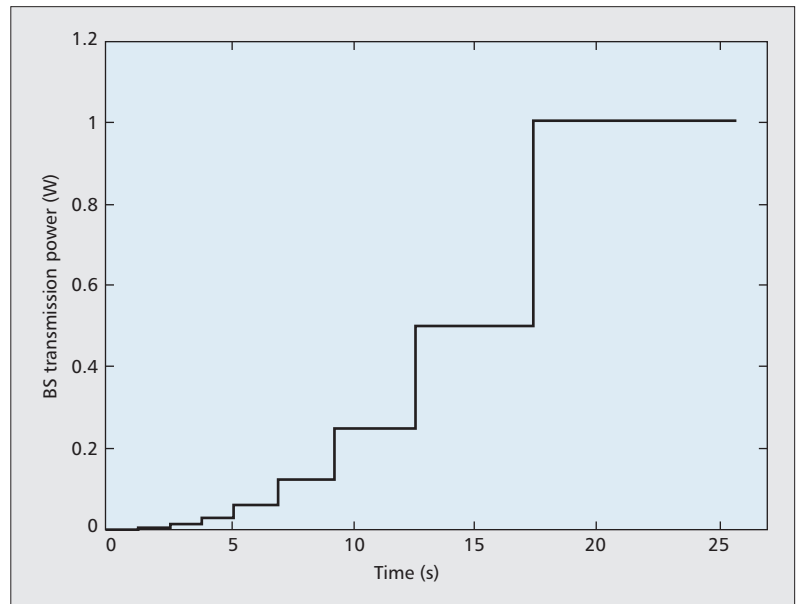

Figure 5. *Switch-on transient for SC$_1$. Exponential power increase profile, with H = 2.*

up) will open the door to a new generation of sleep mode algorithms, applicable also to homogeneous deployments, and capable of achieving very substantial energy savings. Unfortunately, blind wake-up is not available on current cellular systems, and its definition will go through a long standardization cycle before being effectively implemented in real products.

BS sleep modes must also deal with closed access cells, that is, "private" cells that can be used only by a limited and well identified set of users. In this case, the mechanisms previously described must be enhanced with an identification of the MT in order to avoid switching on a private BS to serve an MT that has no right to use it.

H	t_1 [s]	$t_M - t_1$ [s]	t_M [s]
2	16	14.4	30.4
10	10.9	3.5	14.4
20	10.4	2	12.4
No. of handovers	62.33	21.66	83.99

Table 3. *No macro cell: average switch-off times and number of handovers.*
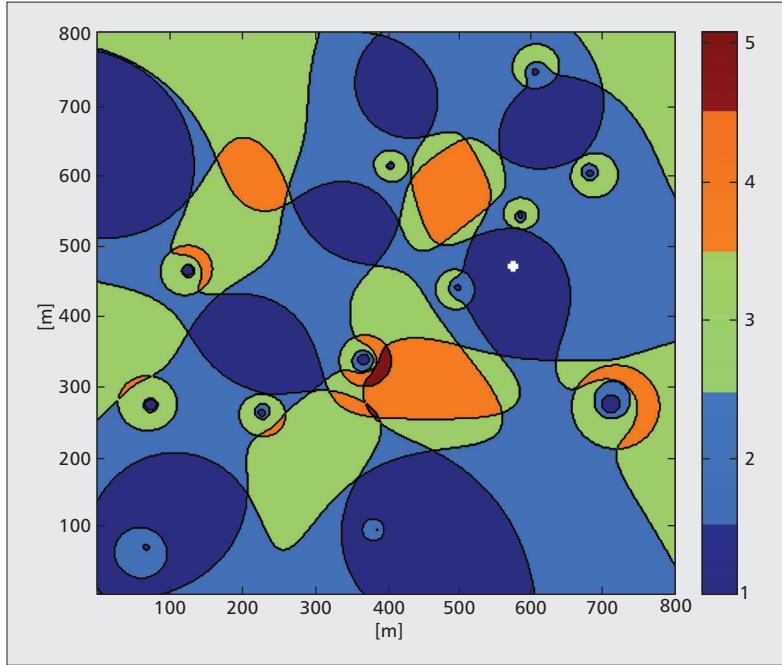

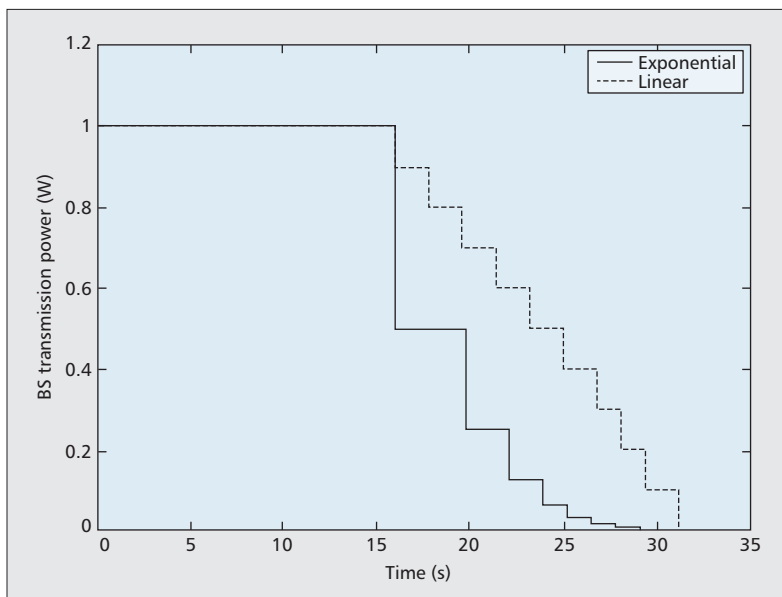**Figure 6.** *No macro cell. BSs interference areas.*

Figure 7. *No macro cell. Switch-off transient for* SC_1. *Exponential and linear power decrease profiles, with* $H = 2$.

CONCLUSIONS

In this article, we have analyzed cell wilting and blossoming techniques that consist in a progressive base stations switch-off and -on, respectively. Specifically, we have considered a realistic cellular network scenario, and we have computed the duration of BS sleep and wake-up transients using an accurate signal propagation model. Our main finding shows that these transients are very short, thus allowing BSs to be switched off and on in short time, with no significant reduction of the energy savings achievable with sleep mode approaches.

This indicates that the use of BS sleep modes can be very effective in reducing energy consumption and might become pervasive in the wireless networks of the future. We can envision that in a not-so-distant future, driving back home from a party, late at night, listening to our favorite web radio or the messages recorded on our answering machine, or email messages being read by a voice synthesizer, over a wireless cellular connection, the BSs in charge of our MT will progressively wake up along our route as we drive by, and go back to sleep after we drive past them. Since traffic is low, late at night, keeping BSs on implies an unreasonable energy waste, and the switch-on and switch-off transients durations are much shorter than the times between cars passing by, so energy savings are not compromised.

REFERENCES

[1] F. Richter and G. Fettweis, "Cellular Mobile Network Densification Utilizing Micro Base Stations," *ICC 2010*, Cape Town, South Africa, May 2010.
[2] M. Ajmone Marsan *et al.*, "Optimal Energy Savings in Cellular Access Networks," *1st Int'l. Wksp. Green Commun.*, Dresden, Germany, June 2009.
[3] M. Ajmone Marsan *et al.*, "Energy-Efficient Management of UMTS Access Networks," *ITC 21 — 21st Int'l. Teletraffic Congress*, Paris, France, Sept. 2009.
[4] S. Fortune *et al.*, "WiSE Design of Indoor Wireless Systems: Practical Computation and Optimization," *Computing in Sci. and Eng.*, vol. 2, no. 1, 1995, pp. 58–68.
[5] M. Ajmone Marsan *et al.*, "Switch-Off Transients in Cellular Access Networks with Sleep Modes," *GreenComm 4*, Kyoto, Japan, June 2011.
[6] TS 25.331, 3GPP; Tech. Spec. Group Radio Access Network; Radio Resource Control (RRC); protocol spec..
[7] TS 36.331, 3GPP; Tech. Spec. Group Radio Access Network; Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC); protocol spec..
[8] H. Holma and A. Toskala, *WCDMA for UMTS*, Wiley, 2004.
[9] E Oh, B. Krishnamachari, "Energy Savings through Dynamic Base Station Switching in Cellular Wireless Access Networks," *IEEE GLOBECOM*, Miami, FL, Dec. 2010.
[10] Z. Niu *et al.*, "Cell Zooming for Cost-Efficient Green Cellular Networks," *IEEE Commun. Mag.*, Nov. 2010, vol. 48, no. 11, pp. 74–79.
[11] K. Dufkova *et al.*, "Energy Savings for Cellular Network with Evaluation of Impact on Data Traffic Performance," *Euro. Wireless 2010*, Lucca, Italy, Apr. 2010.
[12] L. Saker and S. E. Elayoubi, "Sleep Mode Implementation Issues in Green Base Stations," *PIMRC 2010*, Istanbul, Turkey, Sept. 2010.

[13] A. J. Fehske, F. Richter, and G. Fettweis, "Energy Efficiency Improvements Through Micro Sites in Cellular Mobile Radio Networks," *GreenComm '09*, Honolulu, HI, Dec. 2009.

[14] F. Richter *et al.*, "Micro Base Stations in Load Constrained Cellular Mobile Radio Networks," *PIRMC 2010 W-Green Wksp.*, Istanbul, Turkey, Sept. 2010.

[15] I. Haratcherev, C. Balageas, and M. Fiorito, "Low Consumption Home Femto Base Stations," *1st IEEE PIMRC 2009 Wksp. Indoor and Outdoor Femto Cells*, Tokyo, Japan, Sept. 2009.

Biographies

ALBERTO CONTE (alberto.conte@alcatel-lucent.com) is a team manager in the Networks Research Domain at Alcatel-Lucent Bell Labs, Paris, France. After receiving M.S. degrees in telecommunication engineering from the Politecnico di Torino in Italy and Institut Eurécom in Sophia Antipolis, France, he joined the software department department at Bell Labs, France, initially working on embedded systems and VoIP softswitches. He then shifted his research interests to wireless technologies for both home/enterprise (WiFi) and public cellular systems. He and his team pioneered the definition of end-to-end mobile solutions based on WiMAX technology, and he played a key role as Alcatel-Lucent representative at the WiMAX Forum. His current research focuses on 4G networks and their optimizations, including high-density self-x deployments, energy savings solutions, and ubiquitous access enablers to improve end user experience. He is an active lecturer at Ecole Nationale Supérieure des Télécommunications (ENST), INSA-Lyon, and Supelec (France), where he teaches WiMAX technology to M.S. degree students. He is also a member of the Alcatel-Lucent Technical Academy.

AFEF FEKI (afef.feki@alcatel-lucent.com) is a research engineer at Alcatel Lucent Bell Labs, France. She holds an engineering degree in telecommunications as well as an M.Sc. in optics and radio frequency, both from the Polytechnic Institute of Grenoble in France, and she received a Ph.D. in telecommunications from Telecom ParisTech with high honours. Her research interests at Bell Labs include 4G wireless networks, femto and small cells deployment, and green networks, as well as self configuration and management of radio networks. Previously, she worked at Orange Labs R&D France where she was involved in various projects on cognitive radio, dynamic spectrum access issues, and the optimization of radio resource management. She has published over 10 journal and conference papers, and holds several patents. She is co-chair of the International Workshop on Indoor and Outdoor Femto Cells (IOFC), TPC co-chair of WiOpt 2012, and TPC member for other conferences such as WCNC 2012 and VTC 2011.

LUCA CHIARAVIGLIO (chiaraviglio@tlc.polito.it) received his Master's degree in computer engineering and Ph.D. degree in electronics and communications engineering, both from Politecnico di Torino in 2007 and 2011, respectively. Between August 2009 and April 2010 he was a visiting scholar at the WING group of Boston University under the supervision of Prof. Ibrahim Matta. His main research interests are in the field of energy-aware design for backbone and access networks.

DELIA CIULLO (ciullo@tlc.polito.it) received her Master's degree in telecommunications engineering and Ph.D. degree in electronics and communications engineering, both from Politecnico di Torino in 2007 and 2011, respectively. In 2009 she was a visiting student at the CNRG group of MIT under the supervision of Prof. Eytan Modiano. She is currently a postdoctoral researcher at Politecnico di Torino. Her research interests are in the fields of energy-aware networks, scaling properties in wireless networks, and P2P-VOD systems.

MICHELA MEO (meo@tlc.polito.it) received her Laurea degree in electronics engineering in 1993, and her Ph.D. degree in electronic and telecommunications Engineering in 1997, both from the Politecnico di Torino, Italy. Since November 1999 she has been an associate professor at Politecnico di Torino. She has co-authored more than 150 papers and edited six special issues of international journals. Her research interests are in the field of green networking, performance evaluation and modeling, traffic classification and characterization, and P2P.

MARCO AJMONE MARSAN [F] (ajmone@tlc.polito.it) is a full professor in the Electronics Department of Politecnico di Torino and a part-time chief researcher at IMDEA Networks, Madrid, Spain. From September 2002 to March 2009 he was director of the Institute for Electronics, Information and Telecommunications Engineering of the National Research Council. From 2005 to 2009 he was Vice-Rector for Research, Innovation and Technology Transfer at Politecnico di Torino. In 2002 he was awarded a honorary doctoral degree in telecommunications networks from the Budapest University of Technology and Economics. He is the chair of the Italian Group of Telecommunications Professors, and the Italian Delegate to the ICT Committee of the Seventh Framework Programme of the European Union.