

Energy Efficient Management of two Cellular Access Networks

Marco Ajmone Marsan

1) Electronics Department
Politecnico di Torino, Italy

2) IMDEA Networks, Madrid, Spain
marco.ajmone@polito.it

Michela Meo

Electronics Department
Politecnico di Torino, Italy

michela.meo@polito.it

ABSTRACT

In this paper we evaluate the energy saving that can be achieved with the energy-aware cooperative management of the cellular access networks of two operators offering service over the same area. We evaluate the amount of energy that can be saved by using both networks in high traffic conditions, but switching off one of the two during the periods when traffic is so low that the desired quality of service can be obtained with just one network. When one of the two networks is off, its customers are allowed to roam over the one that is on. Several alternatives are studied, as regards the switch-off pattern: the one that balances the switch-off frequencies, the one that balances roaming costs, the one that balances energy savings, and the one that maximizes the amount of saved energy. Our results indicate that a huge amount of energy can be saved, and suggest that, to reduce energy consumption, new cooperative attitudes of the operators should be encouraged with appropriate incentives, or even enforced by regulation authorities.

1. INTRODUCTION

The number of cellular network accesses worldwide has recently surpassed the staggering number of 4 billions, with an average of more than 6 subscriptions every 10 people, and with a maximum in Italy, where, on average, each person (whether newborn or retired) subscribes over one and a half wireless network contracts [1, 2]. This huge number of wireless terminals, together with the access network equipment necessary to serve them, consume an enormous amount of energy. This is why network operators are paying huge energy bills. For example, in Italy, the incumbent telecommunication network operator electricity bill amounts to over 2 TWh, second only to that of the national railway system [3]. The impact of such energy cost on the end user telephone bill is non-marginal. Quite similar considerations apply to other nations, for example Japan [4]. The reduction of those energy bills can provide a significant reduction of OPEX (operational expenditure), which can offer relief for companies in times of recession like those we are living, but can also be, at least partially, transferred to end users to reduce the unit cost of network access. This situation has caught the interest of networking researchers, and has stimulated the birth of an innovative research line, often called "Green Networking".

Energy efficiency has traditionally been an issue tackled by hardware designers and equipment manufacturers. The first area of networking that paid attention to energy consumption was represented by sensor networks, where the peculiarity of the network nodes made energy quite a significant element of the network design space. The first work introducing the energy issue in the mainstream networking area appeared at SIGCOMM 2003, and refers to the greening of the Internet [5]. The attention to energy issues in networking has been drastically rising in the last two years, when some specific meetings were organized to discuss the problem. As an example of recent works on green networking, see [6, 7, 8].

In our previous works [9, 10] we tackled the issue of the energy-aware management of individual cellular access networks, estimating the amount of energy that can be saved by an operator that tries to optimize its energy consumption by reducing the number of active cells in its own access network during the periods when they are not necessary, because traffic is low. The periodic reduction of the traffic in some portions of a cellular access network is due to both the typical day-night behavior of users, and the daily swarming of users carrying their mobile terminals from residential areas to office districts and back, resulting in the need for large capacity in both areas at peak usage times, but in reduced requirements during the periods in which the area is lightly populated (day for residential areas and night for office districts). The assumption that some cells in the access network can be switched off when traffic is low implies that radio coverage and service provisioning can be taken care of by the cells that remain active, which requires a, possibly small, increase in the emitted power, and some adjustment in other network parameters, such as antenna tilting; moreover, some switch-off patterns result unfeasible due to specific site positioning that require some cells to be always on to provide full coverage.

In this paper we explore a different, possibly more viable, approach, which considers the fact that metropolitan areas are normally served by several competing operators, all providing full coverage of the metropolitan area, and dimension networks according to their number of subscribers. When traffic is high, the resources of each operator are exploited at capacity, resulting in the quality of service (QoS) used as a target design objective. When traffic is low, the resources of each operator become redundant, and at some point just one of the existing networks would be able to carry all the traffic in the area. Thus, if the operators cooperate, they can switch off their network in turn, and save energy.

In this paper we assume that just two operators, say A and B, are offering service in a metropolitan area (but the extension to the case of more than two operators is trivial), and that they are will-

ing to cooperate in order to save energy, by accepting the competitor's subscribers as roaming customers while their home network is switched off. In other words, operator A periodically switches off its access network, and when the network of operator A is off, operator B accepts the roaming traffic of operator A, and vice-versa. Of course, by doing so, part of the energy required to power one of the two access networks is saved, at the price of transferring some customers from their home network to another operator. This implies some increased technical complexity, as well as some subtle costs, deriving for example by the information gathered by one operator about the behavior of the competitor's customers; but, as our results will show, the saved energy can be huge. We claim that the approach we propose can be implemented with minimal impact on existing networks, since it boils down to one operator becoming for some time periods a virtual operator, which offers service through the other operator's infrastructure.

In this paper we study this scenario with simple analytical models, and we quantify the benefits that can be achieved with several different switch-off patterns, corresponding to the cases when: i) the switch-off frequencies are balanced, so that the two operators alternate in switching off their network, ii) roaming costs are balanced, so that no roaming cost must be evaluated, iii) the energy savings of the two operators are equal, iv) the amount of saved energy is maximized.

2. THE MODEL

We consider an area served by two operators, A and B, whose access networks (termed *network A* and *network B*, respectively) are dimensioned according to the traffic demands of their customers, but fully cover the area. We denote by N_A and N_B the numbers of customers of the two operators, and by $f_A(t)$, $f_B(t)$, with $t \in \{0, T\}$ spanning over 24 hours, the daily traffic profiles of the two networks. We assume that the daily traffic profile repeats periodically (i.e., we neglect the weekend effect on traffic, which is a conservative assumption as regards energy savings), and that the average per-user traffic in the two access networks is the same, so that the overall traffic in the networks is proportional to the respective number of users:

$$\frac{f_A(t)}{N_A} = \frac{f_B(t)}{N_B} \quad (1)$$

Thus, letting $\alpha = N_B/N_A$, we get $f_B(t) = \alpha f_A(t)$.

In the next section we will introduce specific traffic profiles for the derivation of numerical results, but it can be noted that quite often real traffic patterns exhibit a periodical structure, with highs around mid-morning and mid-afternoon, and lows at night. In Figure 1 we report a typical shape of traffic profiles, setting at $t = 0$ the peak traffic value. For the sake of simplicity, we mostly consider traffic profiles that are symmetrical with respect to $T/2$, but the extension to the asymmetrical case only requires some heavier notation. In the results section we will also discuss one asymmetrical case.

With no loss of generality, we assume $N_A > N_B$, so that also $f_A(t) > f_B(t)$, and $\alpha < 1$. Let f_{max} identify the maximum traffic level reached in either of the two access networks, which, for the previous assumption is network A.

Since each access network is dimensioned so as to carry its peak traffic with the desired level of QoS (which we assume to be equal for the two networks), when the traffic reduces to a level such that just one network is able to carry the traffic of both without violating

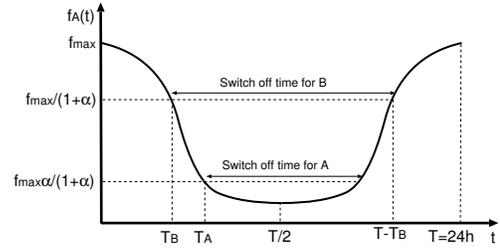


Figure 1: Typical daily traffic profile for network A, $f_A(t)$, and possible switch-off periods for networks A and B (only one at a time can be switched off).

QoS constraints, the other network can be switched off.

In particular, network B can be switched off from time T_B until time $T - T_B$ (because of the symmetry assumption), if:

$$f_A(T_B) + f_B(T_B) = f_{max} \quad (2)$$

But, since $f_B(T_B) = \alpha f_A(T_B)$, this means:

$$f_A(T_B) = \frac{f_{max}}{1 + \alpha} \quad (3)$$

so that T_B can be obtained as:

$$T_B = f_A^{-1} \left(\frac{f_{max}}{1 + \alpha} \right) \quad (4)$$

Similarly, network A can be switched off from time T_A to time $T - T_A$, when:

$$f_A(T_A) + f_B(T_A) = \alpha f_{max} \quad (5)$$

since αf_{max} is the maximum traffic in network B, for which such network was dimensioned. T_A can be obtained as:

$$T_A = f_A^{-1} \left(\frac{\alpha f_{max}}{1 + \alpha} \right) \quad (6)$$

From our assumption that $N_A > N_B$, it follows that $T_B < T_A$, so that network B, which carries less traffic and requires less capacity, can be switched off for longer periods of time, as expected.

In what follows, we consider different switch-off patterns.

2.1 Balanced switch-off frequencies

Assuming that during each 24-hour period just one of the two networks is turned off, we can define the switch-off frequencies P_A and P_B as the fraction of days in which networks A and B are switched off. If every day one of the two networks is switched off,

$$P_A + P_B = 1 \quad (7)$$

Balancing the switch-off frequencies implies $P_A = P_B = 0.5$.

2.2 Balanced roaming costs

When one of the two access networks (say network B) is switched off, its subscribers are accepted as roaming customers by operator A, whose access network remains on. This implies some roaming costs, corresponding to the traffic generated by customers of operator B that is carried by network A. We can quantify the amount of

roaming traffic carried by network A when network B is switched off as:

$$\int_{T_B}^{T-T_B} f_B(t)dt = \int_{T_B}^{T-T_B} \alpha f_A(t)dt \quad (8)$$

On the contrary, the amount of roaming traffic carried by network B when network A is switched off is:

$$\int_{T_A}^{T-T_A} f_A(t)dt \quad (9)$$

Assuming that the unit roaming costs for the two operators are equal, in order to balance costs, so that it is not necessary to transfer money from one operator to the other to pay for the traffic carried when the home network is switched off, we must set:

$$P_A \int_{T_A}^{T-T_A} f_A(t)dt = P_B \int_{T_B}^{T-T_B} f_B(t)dt \quad (10)$$

which leads to

$$\frac{P_A}{P_B} = \alpha \cdot \frac{\int_{T_B}^{T-T_B} f_A(t)dt}{\int_{T_A}^{T-T_A} f_A(t)dt} \quad (11)$$

2.3 Balanced energy savings

The total cost of operating each of the two networks can be roughly estimated as a cost per user C_U (which represents the total network CAPEX + OPEX – capital and operational expenditures – divided by the number of subscribers; we assume C_U to be the same for the two operators) multiplied by the number of customers, minus the energy saved by switching off the network, which is computed as an energy cost per unit time, C_e , multiplied by the switch-off period duration, and by the switch-off frequency. The cost for networks A and B, are, respectively:

$$C_A = C_U N_A - C_e(A)(T - 2T_A)P_A \quad (12)$$

$$C_B = C_U N_B - C_e(B)(T - 2T_B)P_B \quad (13)$$

The energy cost for each network is a function of many parameters. We consider as two extreme cases the situation where the energy cost does not depend on the number of network subscribers, which implies $C_e(A) = C_e(B) = C_e$, and the situation where a direct proportionality exists between the number of subscribers and the energy cost, so that $C_e(A) = C_e N_A$ and $C_e(B) = C_e N_B$, that is $C_e(B) = \alpha C_e(A)$. Real energy costs are probably neither independent on the number of customers, nor proportional, but somewhere in between. In the next section we will consider both cases as possible bounds to the real situation. To balance the energy savings of the two network operators, it is necessary to set:

$$C_e(A)(T - 2T_A)P_A = C_e(B)(T - 2T_B)P_B \quad (14)$$

which, under the assumption (7) that every day one of the two networks is switched off, leads to:

$$P_A = \begin{cases} \frac{T-2T_B}{2(T-T_B-T_A)} & \text{if } C_e(B) = C_e(A) \\ \frac{\alpha(T-2T_B)}{T-2T_A+\alpha(T-2T_B)} & \text{if } C_e(B) = \alpha C_e(A) \end{cases} \quad (15)$$

and P_B is obtained from $P_B = 1 - P_A$.

2.4 Maximum saved energy

The total energy saved can be expressed as:

$$C_e(A)(T - 2T_A)P_A + C_e(B)(T - 2T_B)P_B \quad (16)$$

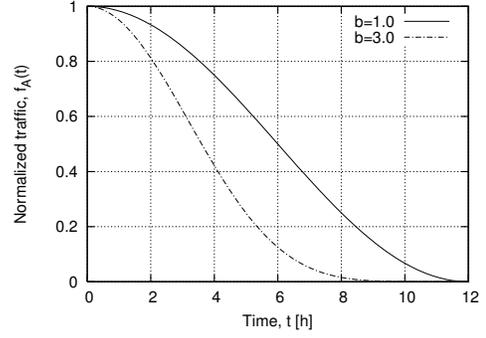


Figure 2: Sinusoidal-like traffic profile.

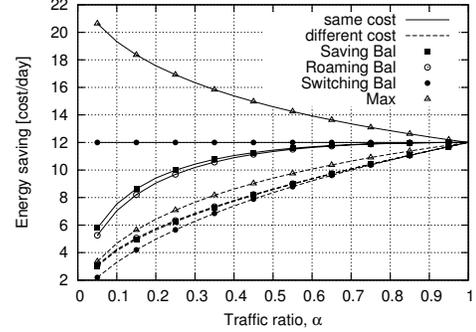


Figure 3: Energy savings under variable cost models and switch-off pattern; traffic shape parameter $b = 1$.

If $C_e(A) = C_e(B) = C_e$, since the assumptions we introduced are such that $T_B < T_A$, in order to maximize the amount of saved energy we must set $P_B = 1$ (which implies $P_A = 0$). In other words, the maximization of the saved energy in this case requires that network B (the smaller network) is switched off every day. This is quite intuitive, since the unit energy costs are equal for the two networks, but the possible switch-off duration is longer for the smaller network.

On the contrary, if $C_e(A) = C_e N_A$ and $C_e(B) = C_e N_B$, we must maximize

$$(T - 2T_A)P_A + \alpha(T - 2T_B)(1 - P_A) \quad (17)$$

that leads to $P_A = 1$, if $(T - 2T_A) > \alpha(T - 2T_B)$, and $P_A = 0$ otherwise.

3. NUMERICAL RESULTS

In the derivation of numerical results we consider two different scenarios, characterized by their daily traffic profile.

In the first scenario, the daily traffic profiles follow a sinusoidal-like periodic behavior expressed by:

$$f_A(t) = \frac{1}{2^b} [1 + \sin(\pi t/12 + \pi/2)]^b \quad (18)$$

with $b = 1, 3$. The traffic profiles in the 12-hour period following the peak are shown in Fig. 2, assuming $f_{max} = 1$. It must be noted that with $b = 3$ the curve has steeper slope, and the average traffic is lower.

The energy savings that can be achieved with the proposed switch-off patterns are reported in Figs. 3 and 4, for $b = 1$ and $b = 3$, respectively. Results are reported as a function of α , the traffic unbalance between the two networks. Solid lines refer to the case

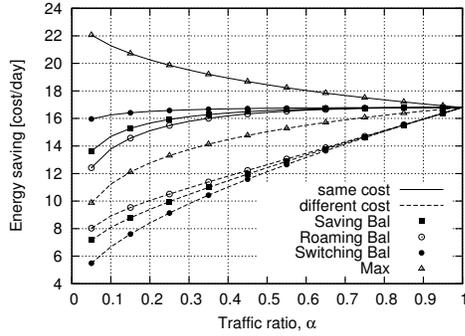


Figure 4: Energy savings under variable cost models and switch-off pattern; traffic shape parameter $b = 3$.

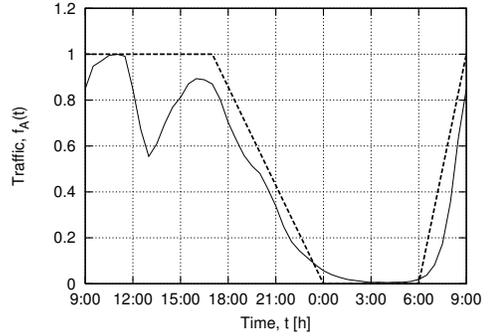


Figure 6: The real normalized traffic profile, and its approximation.

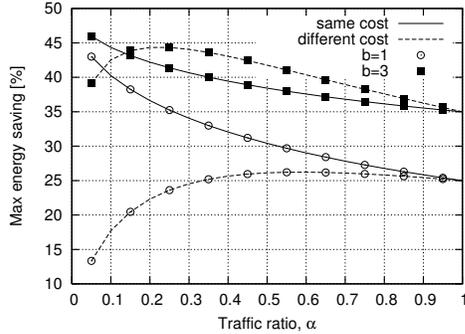


Figure 5: Maximum energy savings as percentage of the total energy cost; variable cost models, and $b = 1, 3$.

$C_e(A) = C_e(B)$, labeled as "same cost" in the figure; dashed lines refer to the case $C_e(B) = \alpha C_e(A)$, labeled as "different cost". Savings are expressed as the average amount of energy saved per day, measured in units corresponding to the energy cost of 1 hour for operator A, i.e., $C_e(A)$. For example, with *same cost* for the two operators, a saving equal to $12C_e(A)$ means that 12 hours of energy per day can be saved out of the total daily cost $2 \cdot 24C_e(A) = 48C_e(A)$; this corresponds to a 25% saving in energy. In the case of *different cost*, $C_e(B) = \alpha C_e(A)$, one hour of energy saved switching off operator B's network corresponds to a saving equal to $\alpha C_e(A)$; the total daily cost is, in this case, $24(1 + \alpha)C_e(A)$. Thus, while the total cost varies with α in the case of *different cost*, it is constant for *equal cost*.

Consider first the *same cost* case with $b = 1$ (solid lines in Fig. 3). When traffic is balanced ($\alpha = 1$), obviously, the same saving is achieved, independently of which network is switched off; because of the shape of the sinusoidal traffic profile with $b = 1$, one of the two networks can be switched off for half of the time, resulting in 12 hours (25%) saving. From the curves we see that more than 25% savings can be achieved for all values of α , in the case in which the amount of saved energy is maximized. In this case, network B is always switched off, since costs are equal, but the period in which network B can be turned off is longer. Exactly 25% savings is achieved for all values of α , in the case in which networks are switched off in alternating days, again because of the shape of the traffic profile. Progressively lower savings for decreasing values of α are observed in the cases when roaming costs and energy savings are balanced. Differences with respect to the maximum become substantial for low values of α . The roaming balance pattern is slightly less effective than saving balance.

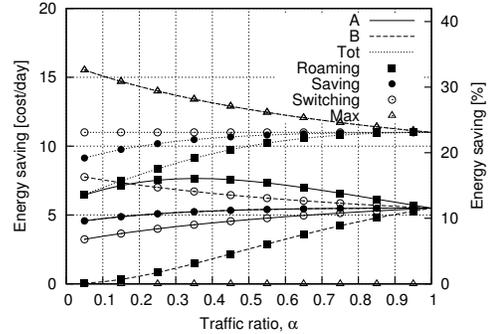


Figure 7: Realistic operator traffic. Same cost model. Total daily savings, and average daily saving for each operator under different switch-off patterns.

Differences among switch-off patterns are much smaller in the case of *different cost*, and savings become smaller for decreasing values of α , but here we must consider that also the total energy cost is decreasing for decreasing α . The maximum energy saving in this case is obtained by always switching off network A, due to its higher energy cost.

Similar observations can be drawn from Fig. 4 that corresponds to the case $b = 3$. In this case, larger savings are possible due to the steeper slope of the traffic pattern, and the lighter overall traffic. The maximum achievable savings as a percentage of the total energy cost are shown in Fig. 5 for both $b = 1$ and $b = 3$. Savings in the case $b = 3$ are larger due to the reasons mentioned before (steeper decrease and lower average traffic), and become quite substantial, being larger than 35% for all values of α .

As a second scenario, we now consider a more realistic traffic profile, derived from true traffic data, collected in the network of an Italian operator. The solid line in Fig. 6 reports the measured daily traffic profile of the operator, normalized so that $f_{max} = 1$. The dashed line shows the simplified model we use instead of $f_A(t)$ to generate numerical results.

In Fig. 7 we focus on two operators with *same cost*, and we consider the usual switch-off patterns. The figure reports the total daily energy savings (dotted lines), as well as the savings that are actually achieved by the two operators by switching off network A (solid lines), and network B (dashed lines), respectively. Savings are again reported taking as unit the hourly costs of energy, $C_e(A)$ or $C_e(B)$. The scale on the right reports percentages of the total energy cost. The maximum total energy saving can be achieved by always switching off network B, that can be switched off for

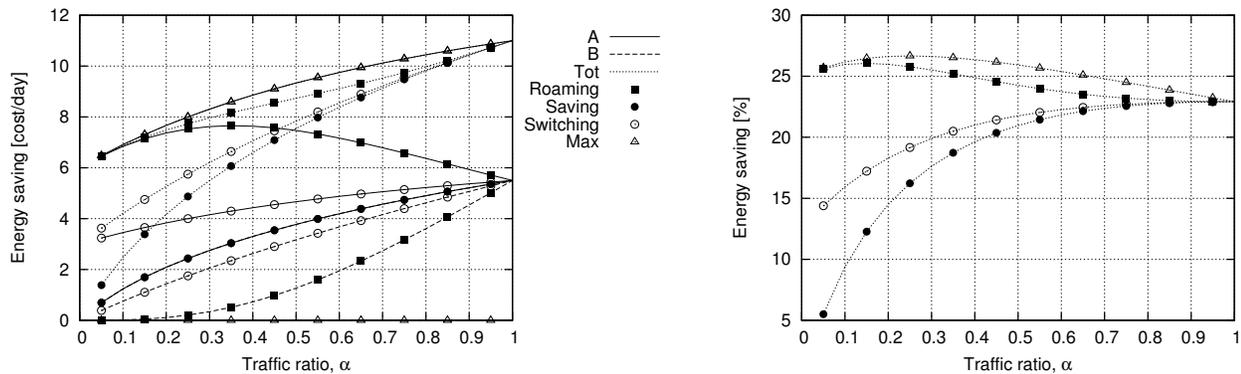


Figure 8: Realistic operator traffic. Different cost models. In the left plot, total daily savings, and average daily saving for each operator under different switch-off patterns. In the right plot, percentage savings with respect to the total energy cost of running the two networks.

longer periods of time than network A, because of its lower traffic load; the largest the traffic unbalance is (low values of α), the more energy is saved. The switch-off pattern that produces the smallest total energy savings is roaming balance; however, notice that a saving of 12% can be achieved in the worst case (very low values of α). Again, when there is no traffic unbalance, all switch-off patterns yield equal results; in this case the total energy saving reaches $11C_e(A)$, which corresponds to almost 25% of the total energy consumed by the two networks.

Fig. 8 shows results for the case in which the energy costs are proportional to the number of subscribers, i.e., $C_e(B) = \alpha C_e(A)$; in this case, the total daily cost of running the two networks is equal to $24(1 + \alpha)C_e(A)$, i.e., it increases with α . The left plot reports energy saving in cost unit equal to $C_e(A)$, the right plot the percentage saving computed over the total daily cost. Contrary to the previous case (*same cost* for A and B), the roaming balance switch-off pattern is now closest to the maximum achievable saving. Interestingly, the maximum is achieved by always switching off network A (remember that it was the opposite under *same cost*). Indeed, even if network A can be switched off for shorter periods than network B, having A off translates into larger energy savings. These results suggest that the actual cost model, as well as the traffic unbalance degree α , can lead to different conclusions about the most convenient policy to adopt. However, it is worth noticing that the actual savings that can be achieved are always remarkable (typical values are above 15-20%, and can grow well above).

Probably, the most interesting observation that can be drawn from these results is that the coexistence of two operators in a competitive market is not energy efficient. A new, green approach to networking requires a new cooperative attitude between operators.

4. CONCLUSIONS

Our study started with the observation that metropolitan areas are normally served by a few competing cellular network operators, which provide 24/7 full coverage, each dimensioning its network according to peak traffic, but providing redundant resources when traffic is low. It becomes thus possible to save energy by switching off some networks when their resources are not necessary, provided that operators are willing to cooperate, by accepting the competitor's subscribers as roaming customers while their home network is switched off, i.e., by temporarily becoming virtual operators.

With very simple mathematics, in the case of just two operators,

we quantified the energy savings with several different switch-off patterns, showing that the amount of energy that is possible to save with this approach can be substantial (of the order of 20%, and above), thus also proving that the coexistence of two operators in a competitive market is not energy efficient, unless a cooperative approach is used. Conversely, our results indicate that an energy-efficient approach to cellular networking is coherent with the present trend, which consists in several virtual operators sharing one physical infrastructure.

These preliminary results suggest that many parameters influence the actual amount of energy that can be saved, leading to different conclusions about the most convenient policy to adopt. The next step in our research will therefore consist in the development of more detailed models to better quantify the effectiveness of cooperative attitudes of cellular network operators for a green approach to networking.

5. REFERENCES

- [1] ITU, "Measuring the Information Society - The ICT Development Index", 2009.
- [2] EC, "Progress Report on the Single European Electronic Communications Market 2008 (14th Report)", Mar. 24, 2009.
- [3] S. Pileri, "Energy and communication: engine of the human progress", Intelc Opening Keynote Speech, 2007, http://www.ega.it/intelec2007/img/Stefano_Pileri.pdf.
- [4] M. Etoh et al., "Energy Consumption Issues on Mobile Network Systems", IEEE SAINT 2008, Turku, Finland, July 2008.
- [5] M. Gupta, S. Singh, "Greening of the Internet", ACM SIGCOMM, Karlsruhe, Germany, August 2003.
- [6] M. Pickavet et al., "Worldwide Energy Needs for ICT: the Rise of Power-Aware Networking", IEEE ANTS Conference, Bombay, India, December 2008.
- [7] M. Liebsch, X. Perez-Costa, "Utilization of the IEEE 802.11 Power Save Mode with IP Paging", ICC 2005.
- [8] A. Jardosh, K. Papagiannaki, E. Belding, K. Almeroth, G. Iannaccone, B. Vinnakota, "Green WLANs: On-Demand WLAN Infrastructure", UCSB Rep. ID: 2008-03, March 2008.
- [9] L. Chiaraviglio, D. Ciullo, M. Meo, M. Ajmone Marsan, "Energy-Aware UMTS Access Networks", W-GREEN 2008, Lapland, September 2008.
- [10] M. Ajmone Marsan, L. Chiaraviglio, D. Ciullo, M. Meo, *Optimal Energy Savings in Cellular Access Networks*, GreenComm 2009, Dresden, Germany, June 2009.