

TREND: Toward Real Energy-efficient Network Design

Marco Ajmone Marsan
Politecnico di Torino / Institute IMDEA Networks
ajmone@polito.it

Stefano Buzzi
CNIT / University of Cassino
buzzi@unicas.it

Luca Chiaraviglio, Michela Meo
Politecnico di Torino
chiaraviglio@tlc.polito.it, michela.meo@polito.it

Carmen Guerrero
Universidad Carlos III de Madrid
guerrero@it.uc3m.es

Filip Idzikowski
Technische Universität Berlin
filip.idzikowski@tu-berlin.de

Yabin Ye, Jorge López Vizcaíno
Huawei Technologies Düsseldorf GmbH
{yabin.ye, jorge.vizcaino}@huawei.com

Abstract—This paper briefly describes the objectives of the TREND (Toward Real Energy-efficient Network Design) Network of Excellence of the European Commission 7th Framework Programme, and outlines some of the main results obtained so far within the project, looking at wireless access networks, core networks, and content distribution issues.

Keywords – energy-efficient networking; green networking; FP7

I. INTRODUCTION

What are the means to best reduce the energy consumption of today's networks without compromising requirements in network and service performance? What are the best suited engineering criteria and principles to actively support energy efficiency along the sequence of network design, planning, and operation? What changes in the design of network equipment are necessary in the short and long term in order to obtain the largest possible energy saving? Which communication and management paradigms and protocols will be able to mediate and ensure the most effective distributed energy control? What are the most promising and sustainable long-term approaches to energy efficient networking, assuming that a clean-slate network design is possible, and what are potential migration strategies to achieve this? What kind of mutually beneficial incentives can be proposed to network operators, service providers, and users, in order to maximize energy efficiency?

The TREND Network of Excellence is generating the knowledge that will help answering these questions, by integrating the efforts of partners with complementary expertise into a virtual centre of excellence on energy-efficient networking. The TREND achievements include: i)

collecting data to assess the power consumption of terminals, devices and infrastructures; ii) identifying energy-friendly devices, technologies, protocols and architectures, and investigating how they can be introduced into operational networks; iii) defining new energy-aware network design criteria; iv) experimentally proving the effectiveness of the proposed approaches. A holistic approach is taken, considering all network segments, from user terminals, to access, to backbone, to data centers.

TREND is funded by the European Commission 7th Framework Programme as the only Network of Excellence in the field of energy-efficient (or green) networking. The acronym TREND stands for "Toward Real Energy-efficient Network Design". The TREND partners are two leading equipment manufacturers, three important European operators, five universities, and two research centers of academic origin. Details can be found on the TREND web site: <http://www.fp7-trend.eu/>. TREND started at the beginning of September 2010, and funding is expected to last three years. Since the program start, several research groups have decided to cooperate with TREND, and have been affiliated to the project as Collaborating Institutions. The current list of TREND CIs can be found at <http://www.fp7-trend.eu/content/collaborating-institutions>. TREND has also established links with the GreenTouch Consortium (<http://www.greentouch.org/>), and is considered a Cooperative GreenTouch Project.

The first year of the TREND NoE has been struck by the tragic loss of its coordinator, Fabio Neri, who tragically and unexpectedly passed away on April 16th, 2011. In spite of this tragic event, the TREND activities have progressed as planned, with several significant achievements, some of which are outlined in the next sections.

The TREND research activities are organized in six work-packages (WPs). The technical WPs address the following

topics: Assessment of power consumption in ICT; Energy efficiency in access and home networks; Energy efficiency in core networks; Instrumenting the network for power saving. The first three WPs structure research activities and coordinate integration efforts in different aspects of energy efficiency in networking. Each WP organizes the research by identifying domains and specifying tasks. Integration, exploitation of complementary expertise, and rationalization of research efforts, take place also through Integrated Research Actions (IRAs), which involve several partners on carefully selected and well-defined technical topics. WP4 coordinates experimental activities, aiming at both providing an effective framework for integration among partners, and at technically proving the effectiveness of the proposed approaches.

The synergistic work and dissemination activities of the involved Institutions is raising the consciousness on the environmental impact of network infrastructures, and provides guidelines for an energy-aware evolution of the Internet, thereby reinforcing the leading European position in designing a sustainable future.

The rest of this paper summarizes some of the main achievements of TREND in this initial research period. Sections 2 and 3 report work on energy efficiency in the wireless access network, looking at transmission aspects, and at Base Station sleep modes, respectively. Sections 4 and 5 consider the core network, looking at routing and protection. Finally, Section 6 addresses the issues related to energy efficiency in content distribution.

More details and results can be found in the TREND publications, which are listed in the TREND web site at <http://www.fp7-trend.eu/content/publications>

II. ENERGY EFFICIENCY AT THE PHYSICAL LAYER

Although the majority of the TREND activities are focused on the upper layers of data networks, some research also addresses the problem of energy efficiency at the physical layer. Here, energy efficiency may be defined as the amount of information bits that are successfully (i.e., with no errors) delivered at the receiver, for each Joule of energy used for the transmitting device operation. A reasonable expression for such a metric is thus the following:

$$EE = R \frac{f(\gamma)}{p_t + p_c}, \quad [\text{bit/J}] \quad (1)$$

where R is the transmission rate, γ is the Signal-to-Interference plus Noise-Ratio (SINR) at the receiver, $f(\gamma)$ (the efficiency function) approximates the probability of correct data reception, while p_t and p_c are the transmitted power, and the power consumed in the electronic circuitry of the transmitting device. A possible expression for the efficiency function may be $f(\gamma) = (1 - \exp(-\gamma))$, which is a good approximation for the probability of correct data reception for several digital modulation schemes in the AWGN

(Additive White Gaussian Noise) channel, but any increasing and S-shaped function ranging between 0 and 1 may be good as well. Energy efficiency in (1) is measured in bit/J, and, for data transmission, it is of interest to tune the transmit power such that EE achieves its maximum value, so that, for instance, a mobile device is capable of transferring the largest possible amount of data given the energy stored in its battery.

Since the received SINR γ is usually proportional to the transmit power, it is readily seen that EE in (1) approaches zero both for the cases in which $p_t \rightarrow 0$ (corresponding to a long duration of the battery but close-to-zero throughput due to decoding errors in the receiver) and $p_t \rightarrow \infty$ (corresponding to the case of error-free reception but close-to-zero battery duration). Accordingly, it descends that an optimal value for the transmit power does exist somewhere in the range $[0, \infty]$, and it can be found through a maximization of EE. While things are smooth in an interference-free scenario, in a cellular system, the problem is actually more complex, due to the fact that mobile users mutually interfere with each other, thus implying that a change of the transmit power of a certain user has an effect on the received SINR (hence, on the energy efficiency) of all the neighbors using the same frequency band. Otherwise stated, the transmission parameters of each user have an impact on the energy efficiency achieved by the other users in the network. Two possible approaches are thus possible: (a) the transmit powers of all users are decided by the network, i.e. there is a centralized control; or (b) each user may independently decide its transmission power, i.e. a decentralized or non-cooperative strategy is pursued. While approach (a) is the usual one, currently implemented in cellular data networks, TREND researchers have investigated approach (b), using game-theoretic tools.

As an example, in [1,2] a multi-cell uplink orthogonal frequency division multiple access (OFDMA) system is considered, and the problem of non-cooperative subcarrier allocation and transmit power control is approached for energy-efficiency maximization. Resource allocation procedures based on the implementation of non-cooperative games are proposed and analyzed, showing that there are instances where the non-centralized approach, with each user selfishly tuning its transmit power so as to maximize its own energy efficiency and disregarding what happens to other users, reaches a steady-state, a situation known as Nash equilibrium, where no user is interested in changing its strategy (for instance, transmit power), provided that the other users do not change theirs. In particular, our study here first focuses on the problem of power control only (assuming that subcarrier allocation has already taken place), and, then, analyzes the problem of joint power control and subcarrier choice, where each user has to independently choose a given number, N say, of the L available subcarriers. Interestingly, it was found that, even though no constraint is put on the subcarrier choice, frequency reuse is automatically implemented, in the sense

that each user, to maximize its own EE value, autonomously avoids choosing subcarriers already employed by nearby users. For a toy system with 4 access points, $N=10$ subcarriers and $L=3$ carriers to be selected by each user, results in [1] show that, for the case in which a total of 10 users are active (so that each carrier is used on average by 3 users), the joint power control and subcarrier allocation algorithm at equilibrium achieves an average EE that is about 60% of the EE achieved in an ideal interference-free environment, where no interference has been considered.

III. ENERGY-AWARE CELLULAR NETWORKS

Energy consumption in wireless cellular access networks has emerged as a major problem in recent years, due to the constant increase in the number of subscribers and to the explosion in the amount of data exchanged by smartphone users. Base Stations (BSs) are the main energy consumers of mobile operators (they account for about 80-90% of the energy bill), and the cost associated with their overall energy consumption is comparable to the Operational Expenses (OPEX) of the operator.

In the context of TREND, we have developed solutions to save energy, starting from the observation that BSs capacity and locations are normally planned to satisfy the peak traffic demand. However, traffic heavily varies during 24 hours, with a peak during the day and a minimum during the night. This results in a waste of BS capacity when traffic is low. Intuition says that it is possible to put into sleep mode BSs during periods of low traffic, and consequently save energy. When some base stations are put into sleep mode, radio coverage and service provisioning are taken over by the BSs that remain active, so as to guarantee that service is available over the whole area at all times. This is a realistic assumption in the case of the dense BSs layouts of urban areas, which consume most of the network energy. However, putting a BS into sleep mode, while guaranteeing user performance, is a complex task. Therefore, TREND has considered different aspects of the problem, including: the optimal management of cell switch-off, the transient analysis of a single cell switch-off, the possible cooperation among different operators to save energy, and finally the impact of joint planning and management.

A. Optimal Management of Cell Switch-off

We have first developed an analytical model to identify the optimal set of active BSs on over an area, as a function of the daily traffic pattern [3]. In particular, we have initially considered the case in which only a single set of BSs in sleep mode per day is admitted, by bringing the network from a high-power, fully-operational configuration, to a low-power reduced configuration. We have then extended our analysis to the case in which several sets are permitted, by progressively reducing the number of active base stations and consequently the network power. First of all, we computed potential energy savings up to almost 50%. This is an important signal that sleep modes can indeed be a useful

tool for energy-efficient networking. Second, our results show that few sets (typically two) applied during the late night and for periods of several hours are at most 10% far from the upper bound, in which a new set of BSs in sleep mode is computed at each traffic variation. This is also an important indication, because it shows that most of the gains can be obtained with limited effort in network management.

B. Transient Analysis of Sleep Modes

We have considered the times required to put a BS into sleep mode, and then to return back to the full power state. In particular, we have considered different case studies of BS deployments driven by operators [4] and manufacturers [5] feedback. Before putting a BS into sleep mode, all the users connected to the BS have to handover to other BSs which remain powered on, to prevent data losses and call dropping. In particular, the BS power has to be gradually decreased in order to shift users from the current BS to the neighbor BSs. This process takes an amount of time which we estimated to be around one minute. Similarly, the reactivation phase from the sleep mode to the full power state cannot be instantaneous. In particular, the BSs have to gradually increase power to prevent disconnections of users that are already connected to the neighbor BSs. We have estimated that also this process takes around one minute to be completed. Thus, we can conclude that the time required to put into sleep mode a BS and recover back has a marginal impact on the energy saving achievable with sleep mode schemes (typically lasting in the order of hours).

C. Cooperation Among Operators to Save Energy

We also considered the case in which different operators provide cellular service over an area. In such context, there is a redundancy in terms of equipment, since each operator manages a cellular network which is completely independent from the other operators. Therefore, a natural question is if it is possible to reduce energy consumption by considering the whole set of operators. In this context, we have considered the case in which operators cooperate to reduce the overall energy consumption [6]. Our idea is to progressively put into sleep mode entire cellular networks during periods when traffic decreases, and eventually using the cellular network of a single operator during night-time. When a network is switched off, its customers are allowed to roam over the networks that remain powered on. Our results show that up to 20% of energy can be saved by adopting a cooperative approach, suggesting that energy-aware cooperative attitudes should be encouraged, for example with incentives regulated by national governments.

D. Joint Planning and Management

Finally, we have estimated the impact of jointly taking into account planning and management of cellular networks, to reduce energy consumption [7].

In particular, we have proposed and evaluated algorithms to put into sleep mode BSs, taking into account different planning strategies. We first solve the planning problem, by considering different cases: minimization of the number of transmitters, minimization of power, and a hybrid case

which is a mixture of those two cases. We then propose two sleep mode strategies that are based either on the cell load or the BS coverage overlap. Results, obtained over a realistic case study, show that sleep modes are able to save energy even if the deployment has been already planned to minimize power consumption, proving that sleep modes are essential to save energy in cellular networks.

IV. SAVING ENERGY IN THE CORE

Network devices in the core consume non-negligible amount of power due to the huge amount of data that they transport. Similarly as in the cellular networks, the daily variation of traffic can be used to switch off (or put into sleep mode) idle devices in the core. While switching on and off whole routers is time consuming, activation and deactivation of some of their parts (in particular line cards) has been shown to have high potential for saving energy [8]. The computation of the optimal configuration of the network for each period of traffic is unfeasible, because of the time needed to solve complex Mixed Integer Linear Problems. Therefore, several heuristic approaches to reduce power consumption of backbone networks were proposed within TREND. We briefly mention three of them applied to IP-over-WDM networks, namely the Least Flow Algorithm (LFA), the Genetic Algorithm (GA) and the Energy Watermark Algorithm (EWA) presented in [9] and [10].

All the algorithms adapt the IP (logical) topology and routing of IP traffic demands to the current traffic. LFA starts from the static base network with all devices switched on. The algorithm first sorts the logical links according to increasing amount of flow, and then iteratively tries to switch off logical links and corresponding line cards, respecting the constraints on network connectivity and link utilization.

GA takes into account both the power consumption of the network, and the reconfiguration cost in terms of rerouted IP traffic. In the context of IP-over-WDM networks and GA, an individual represents the logical topology supporting the given traffic matrix, and fulfilling the constraints on the number of installed devices. GA searches (through the generation of offsprings) for an individual with the best fitness value (weighted sum of power consumption and reconfiguration cost).

EWA utilizes the network configuration from the previous observation period, and uses thresholds on the utilization of the last lightpath of a logical link to trigger the establishing of additional lightpaths, or releasing the unnecessary ones, while respecting constraints on connectivity, utilization of logical links and number of installed devices.

Simulation results presented in [9] and [10] show that all heuristics reduce power consumption in IP-over-WDM networks. Looking at the Abilene network loaded with traffic originating from measurements, and assuming 15-minute observation periods and realistic power values of single components, the results show that GA manages to reduce the power consumption of the network by more than 55% during the night, while the corresponding value for

LFA is 37% [9]. Power savings achieved by EWA are comparable with the ones obtained with GA. This is due to the fact that LFA attempts to switch off whole logical links, while GA and EWA can establish and release single lightpaths (with corresponding line cards), which constitute logical links. This corresponds to capacities of finer granularity, and therefore better adaptivity of the network to the changing traffic. Moreover, GA and EWA start from the network configuration found in the previous time period. This allows EWA to compute the network configuration within a few (sub)seconds (comparable with LFA), while GA needs more time (range of seconds to minutes) due to numerous iterations searching for the best individual.

Significant power savings are achieved in the low demand-hours with respect to the high-demand hours (e.g. up to 33% for EWA with aggressive setting [10]), which is a metric independent of the static base network. Moreover, GA and EWA outperform LFA in terms of the reconfiguration cost.

V. ENERGY EFFICIENCY IN PROTECTED ELASTIC OPTICAL TRANSPORT NETWORKS

Guaranteeing high resilience is a must for operators due to the importance of telecommunication services in our society. Many protection schemes have been proposed so far, but in most cases energy efficiency was not taken into account in performance evaluation.

Recently, the concept of elastic optical network, based on, for example, Orthogonal Frequency Division Multiplexing (OFDM) and coherent detection, was introduced as a promising candidate for the operation of future optical transport networks. This elastic OFDM-based network allows for a more flexible allocation of network resources compared to classical wavelength division multiplexing (WDM), in which the coarse granularity of a wavelength and the rigid channel spacing specified by the International Telecommunication Union (ITU-T) grid may lead to an inefficient use of the spectral resources, and to low energy efficiency. The OFDM modulation technique allows for two levels of flexibility to better adjust the transmission rate to the actual demand: 1) an elastic transmission bandwidth by selecting a variable number of subcarriers, and 2) the use of different modulation formats for subcarriers (distance-adaptive modulation possibility). The main distinction in the deployment of this elastic approach with respect to conventional WDM networks lies in the presence of bandwidth-variable transponders (OFDM transponders), and bandwidth-variable OXCs (Optical Cross Connects), instead of WDM transponders and fixed-grid OXCs, respectively. The performance of this innovative network approach with different protection schemes was compared in terms of energy efficiency to that of conventional fixed-grid WDM networks, operating with Single Line Rate (SLR 10, 40 or 100 Gbps) and Mixed Line Rate (MLR – a mix of 10, 40 and 100 Gbps).

In [11], an energy efficiency comparison of conventional path protection schemes (*dedicated protection* – DP – 1+1,

DP 1:1, and shared protection – SP) for WDM, and elastic OFDM-based networks is reported. More specifically, the energy efficiency per GHz (bits/Joule/GHz) of the elastic OFDM-based network is compared with that of WDM in a Spanish network topology under different traffic load conditions. The simulation results in Figure 1. show that the Elastic network can generally offer a superior performance with respect to conventional WDM networks, thanks to the fine granularity offered by the different modulation formats and the flexible-grid operation. More specifically, a shared protection scheme can offer significantly better results in terms of energy efficiency than the dedicated ones (1+1, 1:1). Furthermore, the elastic approach offers a much better spectral efficiency, which may actually allow for the accommodation of more traffic in a single fiber, reducing thus the number of required network elements (i.e. decrease in cost and also in power consumption).

Even though dedicated protection 1:1 and shared protection schemes commonly show better spectral and energy efficiency than dedicated protection 1+1, the latter is still the most widely used, due to its higher resilience and shorter recovery time. In this scheme, resources are allocated to both working and protection paths according to the peak traffic demand, and the transmission on both paths remains active, independent of the current traffic requirements, which may lead to energy wastage. Thus, in order to reduce the energy consumed by backup resources, a protection scheme that exploits the hourly traffic fluctuations by adapting the rate of the backup transponders to the current required bandwidth requirements is proposed in [12]. This protection scheme can bring considerable reduction in total energy consumption, compared to the conventional dedicated protection 1+1, as can be seen in the average energy savings presented in Figure 2. These savings are especially significant for the elastic network scenario, where up to 11.4% and 18.5% of energy can be saved on a working day and a weekend day, respectively. This is explained by its better adaptability to different traffic conditions, thanks to its multiple modulation format possibilities and its elastic bandwidth transmission.

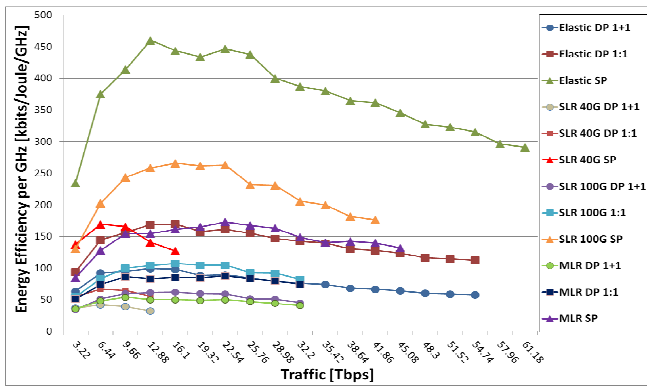


Figure 1. Energy Efficiency per GHz (kbits/Joule/GHz) for the different technologies and traffic load conditions without blocking.

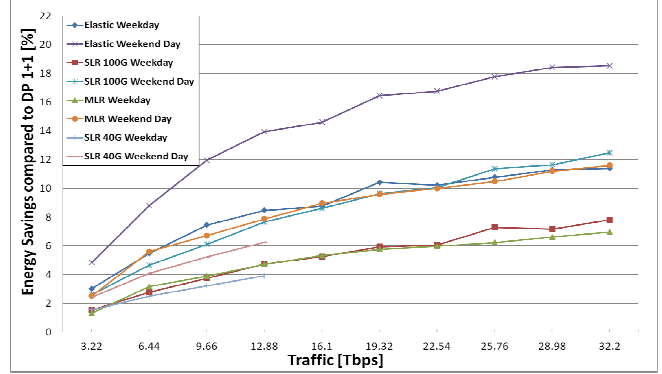


Figure 2. Average energy savings with a Multi-hour protection scheme on a weekday and weekend day with respect to the conventional 1+1 dedicated protection scheme for different traffic load conditions.

VI. ENERGY EFFICIENCY IN CONTENT DISTRIBUTION

A key role of today's Internet is to provide efficient content distribution among users, from distributing multimedia content (e.g. IPTV or Netflix) to sharing user generated data (e.g. Facebook, Twitter, and Youtube). Major content service providers like Google or Microsoft are expanding their data centers to adapt to the increasing demand. The energy consumption in data centers is a key factor to consider in the content distribution infrastructures. In [13] the authors make a study on the environmental foot print of data centers infrastructures.

Besides the content service providers infrastructures, backbone network providers and ISPs in general have increasing network capacity to meet the traffic demands. The increase in the networking infrastructures at ISPs level brings also an increase in the energy cost of the corresponding infrastructures. A distributed approach between ISPs and content service providers to cooperate to minimize the total power consumption for content and service management is proposed in [14].

Another possible approach to optimize energy consumption is to push content to the network edge, like in content distribution networks (CDNs) such as Akamai. Originally, CDNs were designed to improve end user performance in terms of delay and throughput. However, there is also interest in the research community to propose energy-efficient CDN infrastructures considering the relevance of this kind of infrastructures in the current and future Internet. In this field, the TREND consortium is working on the optimization in terms of energy of the distributed networking infrastructures of services like Twitter and BitTorrent. The TREND research is based on network measurements to infer the current ecosystems of these popular services [15][16][17]. The next steps are directed to combine different techniques to optimize the energy consumption of this kind of content distribution infrastructures, including data centers, ISPs, content providers and CDNs.

The third possible solution studied within TREND is based on the possibility of turning off a portion of the equipment in data centers when load is low. For this purpose, we have studied the message arrival pattern in Twitter, identifying a clear day/night pattern, as can be seen in Figure 3. This figure shows that the load at 9:30 is about half of the load at 15:00.

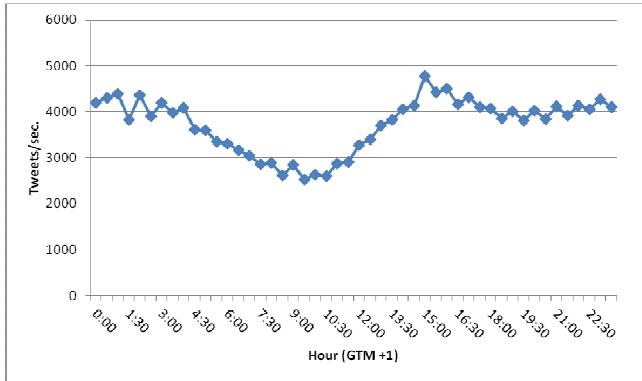


Figure 3. Daily Twitter message arrival pattern

VII. CONCLUSIONS

In this paper we have provided a short overview of the activities of the TREND FP7 NoE in energy-efficient networking, and we have summarized some of the main research results obtained so far. More details on the TREND results can be found in the technical papers listed at <http://www.fp7-trend.eu/content/publications>

The success of TREND is reflected in the growing number of international research groups which are joining the project as Collaborating Institutions, and by the high quality of its research output. Researchers interested in the activities of TREND are invited to visit the TREND web site at <http://www.fp7-trend.eu/> and to contact the project office or the project coordinator.

ACKNOWLEDGMENT

The research reported in this paper has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n. 257740 (Network of Excellence TREND).

REFERENCES

[1] S. Buzzi, G. Colavolpe, D. Saturnino, A. Zappone, "Potential games for energy-efficient power control and subcarrier allocation in uplink multicell OFDMA systems," *IEEE Journal on Selected Topics in Signal Processing (special issue on Game Theory in Signal Processing)*, Vol. 6, No. 2, pp. 89 - 103, April 2012.

[2] A. Zappone, G. Alfano, S. Buzzi, M. Meo, "Energy-efficient non-cooperative resource allocation in multicell OFDMA systems with multiple base station antennas," *2011 IEEE Online Conference on Green Communications (GREENCOM 2011)*, Dresden, Germany, September 2011.

[3] M. Ajmone Marsan, L. Chiaraviglio, D. Ciullo, M. Meo, "Multiple Daily Base Station Switch-Offs in Cellular Networks," *4th International Conference on Communications and Electronics (ICCE 2012)*, Hue, Vietnam, August 2012.

[4] M. Ajmone Marsan, L. Chiaraviglio, D. Ciullo, M. Meo, "Switch-Off Transients in Cellular Access Networks with Sleep Modes," *4th IEEE ICC 2011 Workshop on Green Communications (GreenComm4)*, Kyoto, Japan, June 2011.

[5] A. Conte, A. Feki, L. Chiaraviglio, D. Ciullo, M. Meo, M. Ajmone Marsan, "Cell Wilting and Blossoming for Energy Efficiency," *IEEE Wireless Communications Magazine*, Vol. 18, No. 5, pp. 50-57, October 2011.

[6] M. Ajmone Marsan, M. Meo, "Energy efficient wireless Internet access with cooperative cellular networks," *Computer Networks Journal, Special Issue: Wireless for the Future Internet*, Vol. 55, No. 2, February 2011.

[7] L. Chiaraviglio, D. Ciullo, G. Koutitas, M. Meo, L. Tassiulas, "Energy-Efficient Planning and Management of Cellular Networks," *9th International Conference on Wireless On-demand Network Systems and Services (WONS)*, Courmayeur, Italy, January 2012.

[8] F. Idzikowski, S. Orłowski, C. Raack, H. Woesner, A. Wolisz, "Dynamic routing at different layers in IP-over-WDM networks -- Maximizing energy savings," *Optical Switching and Networking*, Vol. 8, No. 3, pp. 181 - 200, July 2011.

[9] E. Bonetto, L. Chiaraviglio, D. Cuda, F. Idzikowski, and F. Neri, "Exploiting Traffic Dynamics in Power-Aware Logical Topology Design," *ECOC 2011*, Geneva, Switzerland, September 2011.

[10] F. Idzikowski, L. Chiaraviglio, and E. Bonetto, "EWA: an Adaptive Algorithm for Energy Saving in IP-over-WDM Networks," *NOC 2012*, Vilanova i la Geltru, Spain, June 2012.

[11] J. López, et al., "On the Energy Efficiency of Survivable Optical Transport Networks with Flexible-grid," *ECOC 2012, September 2012, to appear*.

[12] J. López, et al., "Traffic and Power-Aware Protection Scheme in Elastic Optical Networks," *Networks 2012, September 2012, to appear*.

[13] W. Vereecken, W. Van Heddeghem, D. Colle, M. Pickavet, B. Dhoedt, P. Demeester, "The Environmental Footprint of Data Centers: The Influence of Server Renewal Rates on the Overall Footprint," *GCN 2011*, Chongqing, China, July 2011.

[14] L. Chiaraviglio, I. Matta, "An Energy-Aware Distributed Approach for Content and Network Management," *IEEE INFOCOM 2011 Workshop on Green Communications and Networking*, Shanghai, China, April 2011.

[15] R. Gonzalez, R. Cuevas, A. Cuevas and C. Guerrero, "Where are my followers? Understanding the Localitty Effect in Twitter," *arXiv:1105.3682v1*, May 2011.

[16] M. Kryczka, R. Cuevas, A. Cuevas, C. Guerrero, A. Azcorra, "Measuring the Bittorrent Ecosystem: Techniques, Tips, and Tricks," *IEEE Communications Magazine*, Vol. 49, No. 9, pp. 144-152, USA, September 2011.

[17] R. Cuevas, N. Laoutaris, Y. Xiao, G. Siganos, P. Rodriguez, "Deep Diving into BitTorrent Locality," *The 30th IEEE International Conference on Computer Communications, IEEE INFOCOM 2011*, Shanghai, China, April 2011.