PRE-PRINT VERSION

MODELING ENERGY CONSUMPTION OF MOBILE RADIO NETWORKS: AN OPERATOR PERSPECTIVE

ANTONIO CAPONE, STEFANO D'ELIA, ILARIO FILIPPINI, ALESSANDRO E. C. REDONDI, AND MARCO ZANGANI

Antonio Capone, Ilario Filippini, and Alessandro E. C. Redondi are with Politecnico di Milano.

Stefano D'Elia and Marco Zangani are with Vodafone.

ABSTRACT

The exponential growth of mobile traffic is forcing operators to quickly increase the capacity of their networks by means of new technologies and advanced architectures. This capacity expansion not only brings increasing fixed costs for additional network infrastructures, but also inflates operational costs, which are becoming critical, mainly in terms of energy bills. In this perspective, monitoring the energy consumption of network devices and defining their energy profile models are valuable approaches for estimating energy costs and identifying the most efficient configurations. In this article, we propose an energy profiling approach that simplifies the characterization of different base station components and allows the estimation of the network energy efficiency relying only on traffic statistics. We have validated the approach over the extensive dataset of real measurements provided by a probe network for monitoring live energy consumption of Vodafone sites in three different countries.

INTRODUCTION

The explosive growth of mobile data traffic in the last few years is pushing operator networks to their limits. Predictions are that over the next 10 years the traffic volume will skyrocket, increasing dozens of times. In addition, the number of connected devices will reach many billions, including a large number of smart objects [1]. This has strongly stimulated operators' interest in significantly expanding the capacity provided to their users.

Generally speaking, the capacity expansion requires the improvement and the deployment of additional devices. Therefore, facing this traffic growth, mobile operators are challenged by higher costs, both in terms of capital and operational expenditures, and by a constrained carbon footprint caused by the power request of the whole network. This calls for mobile networks that are even more cost-efficient and resource-efficient. Among operational costs, energy costs are becoming a big share, and the reduction of this share is certainly a main concern for mobile operators. In this perspective, cost reduction and environmental sustainability can be seen as two convergent objectives.

Even if information and communication technologies (ICTs), and networks in particular, are generally considered a fundamental instrument to improve the energy efficiency of other economy sectors, their own impact on global warming is no longer negligible. ICTs contribute to 2 percent of global annual greenhouse gas (GHG) emissions [2], which is a value that exceeds the GHG emission of the aviation sector [3]. In particular, telecom infrastructures and devices contribute alone to at least one third of this value. Since ICTs are becoming more and more widely available, their carbon footprint is expected not to fall below 1 $GtCO_{2e}$[1] until at least 2030 (about 2 percent of the global emission), even if current green solutions are taken into account [4]. Furthermore, the worldwide use of electricity for only communication networks has increased from 200 TWh in 2007 to 334 TWh in 2012, representing more than 1.5 percent of the total worldwide electricity use [5].

Among network equipment, radio access nodes are particularly energy-hungry; their energy consumption can reach more than 80 percent of the total energy consumption of the entire access network [6]. This has motivated a large number of studies in the field of green wireless networking, with the aim of developing devices, designing protocols, and proposing network planning and optimization strategies that include energy efficiency aspects in the daily network operation [6]. As a result of this big effort, energy saving mechanisms are now becoming available in several commercial products.

In addition to these studies showing how energy efficiency can be improved, a recent work [1] sets what energy efficiency (measured as the energy spent to serve a traffic unit) can be potentially achieved by mobile networks. It shows that the energy efficiency of mobile networks can be improved with respect to that of the current technology (i.e., first deployment of 4G networks) by more than a factor of 1000, combining together solutions at different levels: hardware, architectures, protocols, and infrastructures. This outcome has been obtained by considering a completely new network infrastructure, mainly based on a large number of small cells. Although deployment costs may limit the feasibility of such high efficiency, it has emerged that large savings require operators to carefully select their equipment and to properly configure the network portions they are expanding or renewing.

---

[1] The unit tCO_2 stands for tonnes of carbon dioxide equivalent, which is a measure that allows to compare different emissions of GHG relative to CO_2.

A rather big fraction of the energy-efficiency improvement is expected from advanced hardware technologies, which can significantly lower the power required by different components, including both power amplifiers for radio transmissions and processing units for baseband signals. Advanced hardware platforms also enable deep-sleep modes with very low energy consumption, which can be exploited in new network architectures to dynamically activate only the network elements necessary to serve the actual traffic demand [7]. Moreover, new transmission technologies and antenna systems characterized by a large number of elements [8] can greatly contribute to the reduction of the energy consumption.

While these advanced solutions start being incorporated in commercial products by manufacturers, it is crucial for operators to identify their potential impact on the network. Indeed, the actual saving achievable by green approaches highly depends on the specific network scenario, in terms of network deployment, configuration, and traffic statistics. Therefore, an accurate study to select the most appropriate solutions is important for successful energy saving. In this perspective, it is fundamental to define simple tools that allow monitoring the energy consumption of the live network and estimating the effects of introducing "green features."

In this article, we present an energy-profiling analysis performed via an energy-consumption monitoring system that has been designed and implemented by Vodafone in 60 sites in three different countries. It uses separate sensors for radio and baseband components, and it is applied to systems of different generations. Sites have been selected to represent a wide range of network scenarios in terms of mix of technologies and devices. Based on a measurement campaign carried out on this monitoring system, we provide the following contributions:

- We have designed a profiling approach that simplifies the energy characterization of radio access equipment and allows us to estimate the network energy-consumption relying on traffic statistics. Different from common approaches based on emitted power, this greatly simplifies operators' activities as it focuses only on network traffic counters, which are already well monitored in network operations.
- We have performed a time and space analysis of the data collected from a real network that shows a fundamental property of the proposed approach, which is that the profiling accuracy is preserved when the model is applied to other sites with similar characteristics or in different days.

The remainder of the article is organized as follows. In the next section, we discuss several approaches for characterizing the energy performance of mobile networks and summarize the radio-access power-profile models proposed in the literature. Following that, we describe the Vodafone energy-consumption monitoring network before introducing the proposed profiling approach together with some representative results. We conclude with some final remarks.

Overview on the Energy Performance Characterization of Mobile Networks

Characterizing the energy performance of mobile networks is not an easy task. Not only does it depend on the performance of individual network devices, but also several other aspects influence energy efficiency [6]: the traffic distribution over time and space determines a different working load-point of each cell; the type of network deployment sets the cell density required in the area; the propagation scenario impacts on achievable rates, and thus on traffic volumes, and so on.

Several performance metrics have been proposed for this purpose. They can be grouped into energy-efficiency metrics at the component level, at the equipment level, and at the network level. Component-level metrics express the individual energy performance of antennas, power amplifiers, power supplies, and so on. Equipment-level metrics assess the energy efficiency of a whole device, like an end-user terminal or a base station. Finally, network-level metrics describe both the energy consumed by the entire network and the corresponding performance measured at the network level (such as coverage, capacity, and delay).

Although operators may be interested in understanding the energy performance of every single component, the final energy figure they can measure, which plays the main role in determining the energy bill, is the overall network-level energy efficiency. However, this high-level performance metric makes the final result depend on several aspects of the network operations, so much so that even deciding what to measure may be an issue. Indeed, although one of the most popular metrics is the energy efficiency (Joule/bit) at nominal network capacity, the metric is appropriate only in the case of networks working close to their capacity (full load), which is rarely the case. Mobile networks undergo strong traffic variations during day hours, and during days of the week. Moreover, they are designed to work far below the full-load level even during peak hours, allowing networks to be robust to both sudden traffic changes and future traffic growth. Finally, the energy efficiency at nominal network capacity does not consider the coverage constraint that forces networks to provide connectivity also in low-load areas, where a metric considering the power consumption over the unit area (like W/m2) would be more appropriate.

Coverage is not the only driver to establish cell density; we must always take into account the type of area where the network is deployed. In rural areas, where large cells typically cover relatively low traffic densities, coverage constraints drive the deployment, while in urban areas the number of required cells per unit area is much higher, as operators must provide sufficient capacity for serving much higher user densities. A metric based only on the spatial density of the power consumption cannot capture the huge difference in urban and rural areas in terms of exchanged traffic volume. In addition, the spatial traffic distribution in urban areas is much less homogeneous than in rural areas. In urban areas we have several traffic hotspots, which force operators to install additional ad-hoc small/femto-cells to boost the network capacity around the hotspot.

This high-level performance metric makes the final result depend on several aspects of the network operations, so much so that even deciding what to measure may be an issue. Indeed, although one of the most popular metrics is the energy efficiency (Joule/bit) at nominal network capacity, the metrics is appropriate only in case of networks working close to their capacity (full load), which is rarely the case.

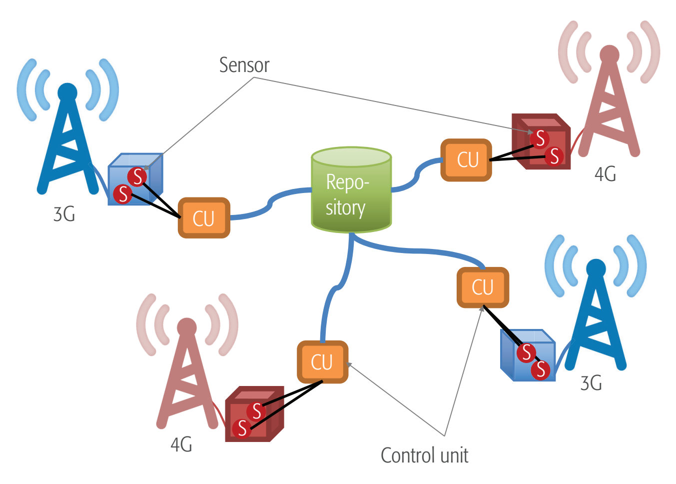

Figure 1. Architecture of the network of energy probes.

We need a sophisticated approach to properly assess the energy efficiency of a mobile network; it must carefully take into account the following aspects:

- Different network layouts across the operator service area.
- Configuration settings in different deployment areas.
- Propagation conditions in considered areas.
- Traffic models and their time and space variations, which may vary in different deployment areas.
- Power consumption profiles of involved devices.
- Operator's targets for user quality of experience.

Ideally, this would require investigating the energy consumption of the whole network at the same time, which cannot be consider a viable solution. A simplified approach consists in splitting the whole network into representative spatial snapshots characterized by a homogeneous traffic profile and network features (e.g., urban, suburban, and rural areas). This makes it possible to perform an energy performance analysis in separate snapshots, which can then be averaged according to the weight of each snapshot in the operator network [9]. Unfortunately, this approach requires additional effort for the characterization of each snapshot and for the estimation of the most efficient configuration of network components and energy saving algorithms in each snapshot.

A fundamental step to evaluate the effect of novel mechanisms and architectures for energy saving is the modeling of the impact of base station (BS) internal components on the aggregate BS energy consumption. The energy efficiency evaluation framework (E^{3}F) proposed in the EARTH project (http://www.ict-earth.eu) investigates the relationship between the BS load and its power consumption: it maps the emitted radio frequency (RF) power (P_{out}) to the power supply of a BS site (P_{in}) [9]. Such a study is based on the analysis of the power consumption of various 4G BS types as of 2010, the effect of the various components of BS transceivers is considered: antenna interface, power amplifier, small-signal RF transceiver, baseband interface, DC-DC power supply, cooling, and AC-DC supply. According to project outcomes, E^{3}F approximates the dependence of the BS power consumption on the cell load with a linear power consumption model:

$$P_{in} = \begin{cases} N_{TRX}P_0 + \Delta_p P_{out} & 0 < P_{out} \le P_{max} \\ N_{TRX}P_{sleep} & P_{out} = 0 \end{cases}$$

where Δ_p is the slope of the load-dependent power consumption, N_{TRX} is the number of transceiver chains, and P_{max} is the RF output power at maximum load. Moreover, P_0 and P_{sleep} indicate the power consumption at minimum non-zero load, and in sleep mode, respectively. More recently, the model has been extended by the GreenTouch consortium (http://www.greentouch.org) on the basis of more detailed hardware profiles [10], so that it can reflect the evolution of hardware characteristics in the next few years (until 2020). As a final remark, the different value between P_0 and P_{sleep} has motivated the research effort toward advanced energy-efficient network architectures [7, 11], which, by separating the control plane and the user plane, allow large parts of the data network to be switched off while maintaining user coverage.

The future trend of energy-aware wireless network development is characterized by small cells and femto cells becoming more widely adopted because of their high energy efficiency. Their energy profiles will be more load-dependent than those of legacy technologies, as shown in [10]: new hardware designs will reduce the power consumption step at minimum non-zero load. Moreover, from a network-level perspective, highly-dynamic networks can be envisioned [7]. In these anticipatory networks, predictive tools will allow the switch-off of redundant resources for the current traffic scenario and reactivate them just in time to cope with the anticipated traffic change, thus providing a further energy saving.

In addition to those in [9] and [10], several other models have been developed for 2G, 3G, and 4G BSs relying on analytical expressions that combine the power consumption of every BS component (e.g., in [12]). However, they result in complex models that are not suitable for the use of mobile operators. In addition, their fine detail level requires a recalibration when the model is applied to a different BS. Alternative approaches based on high-level measurements propose easy-to-use models, which could provide a reasonable accuracy even when applied to a broad set of similar BSs. These works in the literature analyze how the energy consumption of a certain type of BS (indoor femtocell [13], HSPA [14], or GSM/UMTS [15]) depends on some load indicators, like Erlang or bit rate. However, despite their interesting results, all the works consider isolated BSs, trying to provide a customized energy profile for each of them. In this article, instead, we aim at investigating profile robustness to spatial and temporal shifts.

Energy-Consumption Monitoring Network

The system for monitoring and profiling the consumption of the radio access infrastructure is based on a network of remote monitoring probes installed by Vodafone in a subset of radio base station sites.

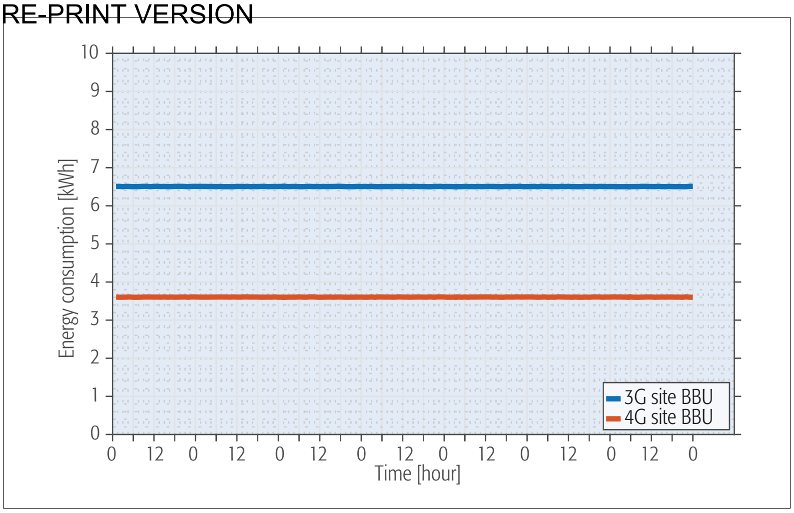

Figure 2. Example of energy consumption of 3G and 4G BBUs in different days and at different time.

Vodafone has a wide network of smart meters installed in several thousand radio base station sites, providing detailed information on energy use, helping to identify and target opportunities to reduce consumption and improving the accuracy of the energy billing. However, these traditional smart meters can measure the overall consumption of a radio base station site, while for a detailed measurement of the energy consumption of a base station, and for its correlation with network counters, a more detailed measurement is needed in order to:

- Monitor the energy consumption of each radio technology installed in the site.
- Separate the contribution of radio and baseband units.
- Track the evolution over time.
- Correlate measures with specific configuration settings.

For such reason, in 2014, Vodafone selected 60 sites in different countries to implement a more sophisticated solution, shown in Fig. 1, aimed at measuring not only the overall energy consumption of the site, but also the energy consumption of each component of the radio base station. In line with ETSI Recommendation 102 706 1, sensors were installed to measure the instantaneous energy absorption of each base station component.

The system implemented in each of the 60 base station sites is based on integrated current sensors (up to 16 per site), having a size of a few millimeters, and its own signal microprocessor, with measured data being transmitted digitally to a single control unit. The control unit of each site is configured to process the sensor data, supply the sensors with power, get access to the measurement data locally, and transmit the data remotely. Data are collected and made available to a single repository, with a time granularity that can be defined in accordance with the granularity available for network counters and key performance indicators (KPI).

A Model for Base Station Energy Consumption

In addition to a detailed energy consumption analysis of monitored sites, the probe network implemented by Vodafone can be used to profile the energy performance of base station components in different configuration settings and traffic load conditions. This not only provides a useful instrument for selecting the most efficient technologies and configurations in monitored sites, but also can support an energy-consumption prediction tool, which, by extending collected results to other sites with similar characteristics, allows the estimation of the impact of new solutions on a much larger network part. However, in order to be effective, the profiling approach must be simple enough to be quickly applied to thousands of sites with limited effort and, as a key feature, it must be robust to time and place shifts. Indeed, once calibrated with monitored sites during a time interval, the model must preserve a good accuracy when applied to other similar sites and/or in different days.

The overall energy consumption of a base station consists of two components: base-band unit (BBU) and the remote radio unit (RRU) energy consumption. BBU performs digital signal processing of received and transmitted data, while RRU implements RF frontend functionalities. We first analyzed the energy profile of BBUs. As shown in Fig. 2, a typical BBU energy profile is almost constant and independent from traffic variations during the day. This suggests that BBU hardware does not implement power-saving strategies and, most important, its energy consumption does not depend on the data traffic load. In addition, the BBU energy consumption of 4G sites is typically much smaller than in 3G systems, because the more recent 4G equipment has been developed with more efficient hardware and software solutions than those adopted for 3G. Moreover, 4G systems have a simplified management that reduces computational complexity. Since BBUs have flat and constant energy profiles, we can focus on modeling RRU power consumption.

Our goal is to derive an alternative model to that presented in the previous section based on RF emitted power. Indeed, counters on RF emitted power are not always available in devices of different vendors, and in some cases they have limitations on measured values. In addition, in order to practically assess the potential impact of new energy-saving techniques for RRUs, it is necessary to define energy consumption models based on variables that can be easily obtained from common network counters. Therefore, we propose a model that replaces the *emitted power* with the *total downlink user traffic volume* measured by network counters over the energy-consumption sampling period. Even if it is rather intuitive that the downlink user is directly correlated with the RF power emitted by a base station, it is also clear that other elements may impact measured energy values, like uplink traffic, cell propagation conditions, traffic distribution, and so on.

In order to assess the impact of these elements, we have validated the linear model comparing it with other higher-order approaches and verifying the achieved accuracy. We carried out an extensive analysis over monitored sites, using traffic volume and emitted power data collected from network databases and correlating them with energy consumption data coming from mon-

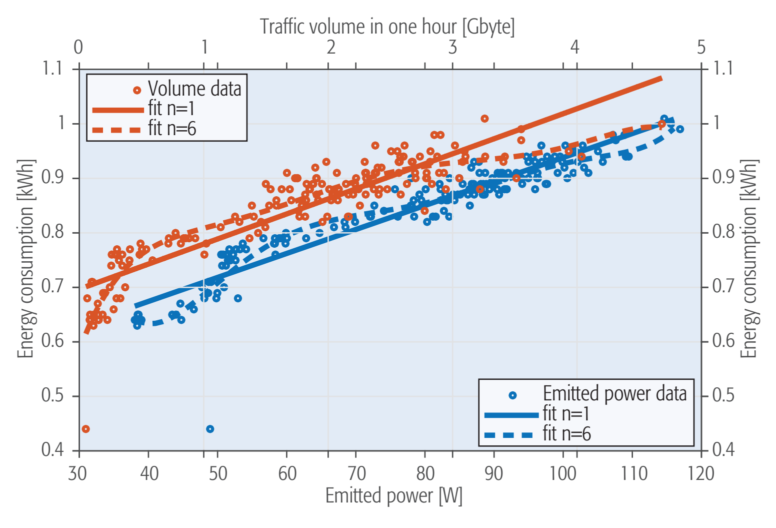

FIGURE 3. Energy consumption plotted against traffic volume and emitted power. Polynomial fitting curves of order n = 1 and n = 6 have been applied as well.

itoring probes. In the following we show a set of representative examples of the energy profiles we derived, focusing on three-sector sites equipped with either 3G or 4G technologies in different European countries. Each dataset consists of a week of measurements with hourly granularity.

As shown in Fig. 3, a strong correlation exists between energy consumption and both traffic volume and emitted power. We have computed fitting polynomial curves with order from 1 to 6 and proceeded in the following way. We first calibrated fitting curves of different orders on a weekly training set at a single site and measured the *training error*, that is, the accuracy of the fitting. Then, we applied calibrated curves both to the same site in different weeks, and to a different site. In these tests we evaluated the *generalization error* between predicted energy values and those measured by monitoring probes, and we compared it against the training error.

Figure 4 shows results for training and generalization errors for models based both on traffic volume and on RF emitted power, considering the case of calibration and application of the model at the same site in different weeks (*different-time*), and the case of model calibration in one site and its application to another site with similar configuration settings (*different-place*).

As shown by the red lines, the training error decreases by increasing the polynomial order. This is somehow expected, as higher-order polynomials better approximate the dataset. Note that the error is below 5 percent and high-order models have only a slightly better accuracy than the linear model. In addition, models based on emitted power have a smaller training error than those based on traffic volume. This is expected as the emitted power is a more direct measure of the energy consumption than the traffic volume.

When models are applied to datasets at different times (dashed lines), we measure the generalization error, which is clearly higher than the training error. Nevertheless, the generalization error is below 10 percent, even when the model based on traffic volume is applied. Interestingly enough, the model based on emitted power, which predicts energy values closer to ground-truth values when the polynomial order is up to four, exhibits a large deviation beyond the fourth order. This behavior is typical of high-order fitting models. Indeed, the increased complexity better follows training dataset results in high stiffness, thus the model badly fits working regions outside those of the training.

Figure 4 shows similar curves (dotted lines) when model calibration and application have been performed at different sites. The two sites have similar hardware features, but they are locat-

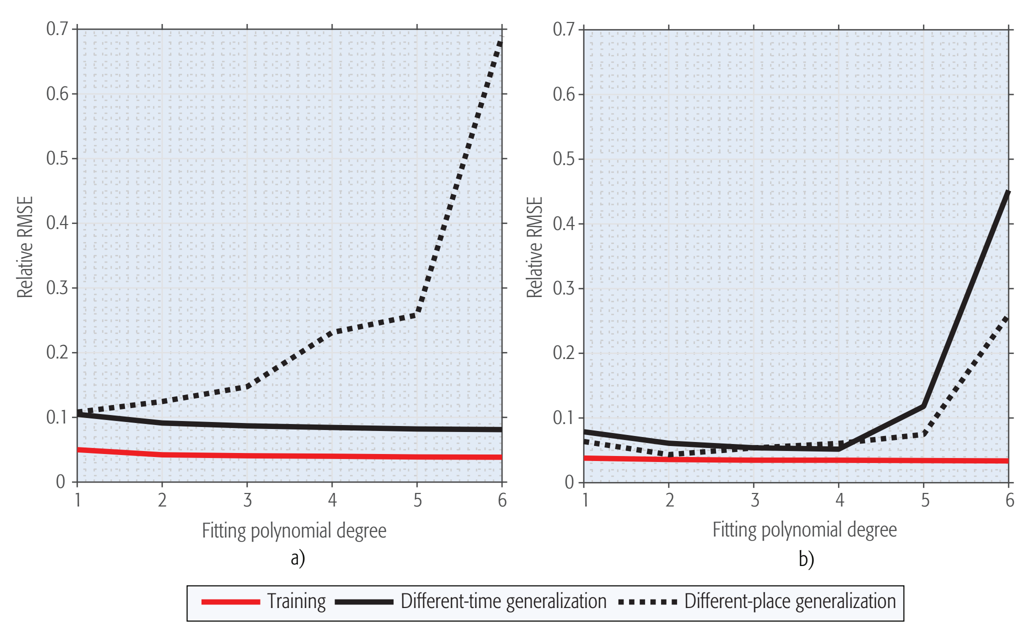

FIGURE 4. Training error and generalization error of the model application at different time and in different place, varying the fitting polynomial order from n = 1 to n = 6: a) Traffic volume ;b) Emitted power.

ed in different environments, that is, urban vs. rural areas. Although the effect of high orders is still evident, we cannot conclude that the generalization error is always larger than in the different-time case. Indeed, there is a strong dependence on the chosen dataset. Vice versa, a true fact is that the linear model produces a generalization error of about 10 percent.

A general conclusion can be drawn: modeling a linear dependence between the emitted power and the energy consumption, as well as between the traffic volume and the energy consumption, is a very good approximation, and it is strongly confirmed by real data. In addition, the linear model based on traffic volume is robust to time and place shifts, therefore it can be applied by mobile operators to the whole network. This allows the placement of the energy measurement probes only in a limited set of sites (to have the model calibrated over representative site types) and straightforwardly have a complete picture of the network energy consumption by using widely available traffic volume counters.

The linear model based on traffic volume, although based on simpler network counters than those related to the emitted power, shows a very accurate prediction quality. Considering its low-complexity and its good accuracy, it is the perfect trade-off for being implemented in deployed networks. Indeed, resorting to more involved modeling approaches not only does not bring significant improvement on the calibration accuracy, but also provides worse results when the model is applied in different places at different times. This has been revealed by the analysis of live data in a real network.

Clearly, the slope and the offset of the linear model depend on the device manufacturer and the type of technology. In Fig. 5 shows an example of 3G sites equipped with devices of different manufacturers. Their curves have different offsets and a different energy consumption slope with respect to the traffic volume. The figure shows the application of the linear model to a 4G site as well. It exhibits a much flatter profile by supporting much higher traffic volumes (one order of magnitude greater) with lower energy consumption than a 3G site.

The linear model can be extended to include base stations operating at different power modes. Several device manufacturers allow operators to run their devices in a low-power mode when the traffic volume is expected to be small, typically during the night. Operators can switch base stations to operational modes where some of the hardware is deactivated and/or the computational power is reduced so as to decrease the energy consumption. An example of this behavior is given in Fig. 6, where the two groups of samples identify a full-power and a low-power mode. The energy profile of the site is equivalent to those of two independent sites, each described by a linear model, operating at different hours of the day according to the activation of the low-power mode. Note that the low-power mode has samples only in the small-volume region, as hardware deactivation usually reduces available capacity. In addition, the low-power mode consumes much less energy than the full-power mode at the same traffic load. Clearly, this behavior can be easily extended to consider multi-mode devices that can finely tune their operational mode in response to the actual traffic load.

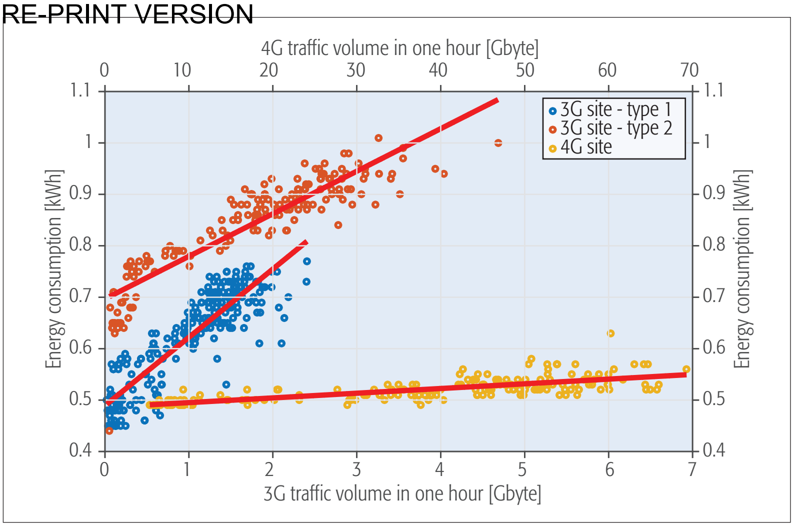

FIGURE 5. Comparison of different-vendor 3G sites and a site with 4G technology.

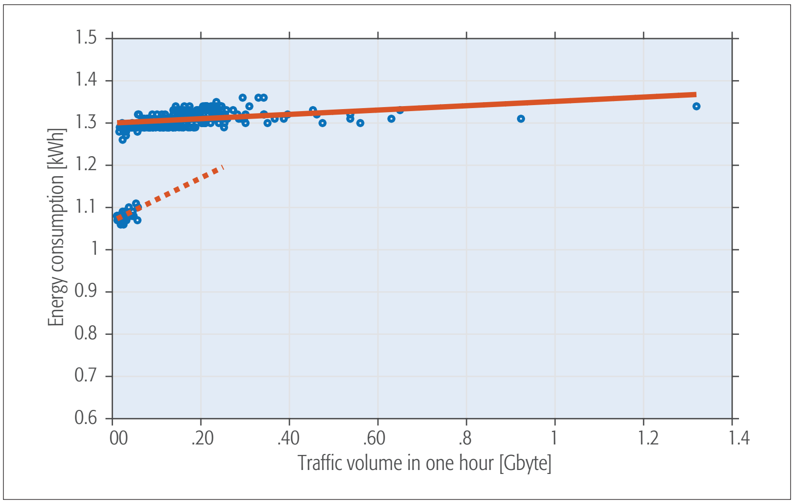

FIGURE 6. Example of dual-mode base station working at full- and low-load operational mode.

CONCLUSION

Operational costs are becoming a critical issue for the energy bill of mobile network operators. For this reason, monitoring the energy consumption of network components and defining energy profile models are valuable instruments to identify the most efficient network configurations and estimate energy costs.

We have presented an energy-consumption monitoring network designed by Vodafone and implemented in three different countries to collect live data from real sites. We have analyzed data coming from this system and compared them with those collected by network counters. This allowed us to derive an energy profiling approach that allows the simplification of the modeling of the different base-station components and the estimation of the energy efficiency on the basis of traffic statistics.

The probes' field data confirm the validity of an energy profile that models a linear dependence between the download traffic volume and the site energy consumption. The slope and the offset of the model depends on the device manufacturer and the technology generation. The model exhibits a good robustness to time and place shifts, proving, once trained on a small set of sites, its applicability to the whole network to predict the energy consumption by relying only on network counters.

Operational costs are becoming a critical issue for the energy bill of mobile network operators. For these reasons, monitoring the energy consumption of network components and defining energy profile models are valuable instruments to identify the most efficient network configurations and estimate energy costs.

REFERENCES

[1] "GreenTouch Green Meter Research Study" in GreenTouch White Paper, http://www.greentouch.org/uploads/documents/GreenTouch_Green_Meter_Research_Study_26_June_2013.pdf, June 2013.

[2] H. Toure, "ICTs and Climate Change — the ITU Perspective," *Climate Action 2008*, www.climateactionprogramme.org, 2008.

[3] A. C. Chan *et al.*, "Methodologies for Assessing the Use-Phase Power Consumption and Greenhouse Gas Emissions of Telecommunications Network Services," *Environmental Science Technology*, vol. 47, no. 1, 2013, pp. 485–92.

[4] "#SMARTer2030, ICT Solutions for 21st Century Challenges," http://smarter2030.gesi.org/, 2015.

[5] W. Van Heddeghem *et al.*, "Trends in Worldwide ICT Electricity Consumption from 2007 to 2012," *Computer Commun.*, vol. 50, 2014, pp. 64–78.

[6] A. De Domenico, E. C. Strinati, and A. Capone "Enabling Green Cellular Networks: A Survey and Outlook," *Computer Commun.*, vol. 37, 2014, pp. 5–24.

[7] A. Capone *et al.*, "Rethinking Cellular System Architecture for Breaking Current Energy Efficiency Limits," *IEEE SustainIT 2012*, Pisa, Italy, Oct. 2012.

[8] T. L. Marzetta, "Massive MIMO: An Introduction," *Bell Labs Technical J.*, vol. 20, Mar. 2015, pp. 11–22.

[9] G. Auer *et al.*, "How Much Energy is Needed to Run a Wireless Network?" *IEEE Wireless Commun.*, vol. 18, no. 4, Oct. 2011.

[10] B. Debaillie, C. Desset, and F. Louagie, "A Flexible and Future-Proof Power Model for Cellular Base Stations," *Proc. IEEE VTC Spring 2015*, Glasgow, Scotland, May 2015.

[11] K. Sakaguchi *et al.*, "Millimeter-Wave Evolution for 5G Cellular Networks," *IEICE Trans. Commun.*, vol. E98-B, no. 3, Mar. 2015.

[12] R. S. Alhumaima, M. Khan, and H. S. Al-Raweshidy, "Modelling the Energy Efficiency of Heterogeneous Cloud Radio Access Networks," *IEEE Int'l. Conf. Emerging Technologies 2015*, Peshawar, Pakistan, Dec. 2015.

[13] R. Riggio and J. L. Douglas, "A Measurement-Based Model of Energy Consumption in Femtocells," *IFIP Wireless Days 2012*, Dublin, Ireland, Nov. 2012.

[14] M. Deruyck, W. Joseph, and L. Martens, "Power Consumption Model for Macrocell and Microcell Base Stations," *Wiley Trans. Emerging Telecommunications Technologies*, vol. 25, no. 3, Aug. 2012, pp. 320–33.

[15] J. Lorincz, T. Garma, and G Petrovic, "Measurements and Modelling of Base Station Power Consumption under Real Traffic Loads," *MDPI Sensors*, vol. 12, no. 4, Mar. 2012, pp. 4281–4310.

BIOGRAPHIES

ANTONIO CAPONE is full professor in the Information, Communication and Bioengineering Technology Department (Dipartimento di Elettronica, Informazione e Bioingegneria), Politecnico di Milano, where he is the director of the Advanced Network Technologies Laboratory (ANTLab). He is a co-founder and CTO of MobiMESH (www.mobimesh.eu), a spin-off company of Politecnico di Milano. His expertise is in networking and his main research activities include protocol design (MAC and routing) and performance evaluation of wireless access and multi-hop networks, traffic management and quality of service issues in IP networks, and network planning and optimization. He received the M.S. and Ph.D. degrees in electrical engineering from the Politecnico di Milano in 1994 and 1998, respectively. He has served as an editor of *IEEE/ACM Transactions on Networking*, and he currently serves as an editor of *IEEE Transactions on Mobile Computing*, *Wireless Communications and Mobile Computing* (Wiley), *Computer Networks* (Elsevier), and *Computer Communications* (Elsevier). He is a Senior Member of the IEEE.

STEFANO D'ELIA is mobile access integrated solutions manager in the Vodafone Group Network Engineering & Delivery team. He graduated in electronic engineering from the University of Rome "La Sapienza." He joined Vodafone Italy (formerly Omnitel Pronto Italia) in 1998. In 2012 he was awarded as a Vodafone Distinguished Engineer, one of the highest steps in the technical career path at Vodafone. He is the author of several papers on electro magnetic fields (EMF), small cells and new technologies, and since 2002 he has been the appointed convenor of the Working Group "Mobile Base Stations" within the Technical Committee 106 of the Italian Electro Technical Committee (CEI), acting as national delegate in several standardization meetings and working groups.

ILARIO FILIPPINI is an assistant professor in the Information, Communication and Bioengineering Technology Department (Dipartimento di Elettronica, Informazione e Bioingegneria), Politecnico di Milano. He received the B.S., M.S., and Ph.D. degrees in electrical engineering from the Politecnico di Milano in 2003, 2005, and 2009, respectively. Since 2008 he has been working as a visiting researcher in the Department of Electrical and Computer Engineering at The Ohio State University in Columbus (OH). His research interests are in networking, and his current activities include planning, optimization, performance evaluation and protocol design of wireless access and multi-hop networks. He serves as area editor of *Elsevier Computer Networks* and *Wiley/Hindawi Wireless Communications and Mobile Computing*. He is a Senior Member of the IEEE

ALESSANDRO E. C. REDONDI is an assistant professor in the Dipartimento di Elettronica, Informazione e Bioingegneria of the Politecnico di Milano, Italy. He received the M.S. in computer engineering in July 2009, and the Ph.D. in information engineering in 2014, both from Politecnico di Milano. From September 2012 to April 2013 he was a visiting student in the Department of Electronic and Electrical Engineerings of the University College of London (UCL). His research activities are focused on network data analysis and on algorithms and protocols for wireless sensor networks, with particular attention on camera networks.

MARCO ZANGANI is head of Mobile Access Engineering in the Vodafone Group Network Engineering & Delivery team. He graduated in electronic engineering from the Politecnico of Milan, where he also completed a post graduate master course at CEFRIEL. He joined Omnitel Pronto Italia, later to become Vodafone Italy, in 1996, where he started his career as an RF engineer. Then, he spent a few years working on the development of the UMTS Network, before moving to Naples as head of Radio Access Network in Southern Italy. In 2008, he returned to Vodafone Italy as head of Radio Access Network, and later he became head of RAN for Southern Europe and later for Europe, being currently responsible for high level radio plans and optimization, as well as the selection, testing and introduction in network of 2G, 3G and 4G radio products.