Toward Dynamic Energy-Efficient Operation of Cellular Network Infrastructure

Eunsung Oh and Bhaskar Krishnamachari, University of Southern California
Xin Liu, University of California at Davis
Zhisheng Niu, Tsinghua University

ABSTRACT

The operation of cellular network infrastructure incurs significant electrical energy consumption. From the perspective of cellular network operators, reducing this consumption is not only a matter of showing environmental responsibility, but also of substantially reducing their operational expenditure. We discuss how dynamic operation of cellular base stations, in which redundant base stations are switched off during periods of low traffic such as at night, can provide significant energy savings. We quantitatively estimate these potential savings through a first-order analysis based on real cellular traffic traces and information regarding base station locations in a part of Manchester, United Kingdom. We also discuss a number of open issues pertinent to implementing such energy-efficient dynamic base station operation schemes, such as various approaches to ensure coverage, and interoperator coordination.

INTRODUCTION

With increasing awareness of the potential harmful effects to the environment caused by CO_2 emissions and the depletion of non-renewable energy sources, there is a growing consensus on the need to develop more energy-efficient telecommunication systems. It has been estimated that 3 percent of the world's annual electrical energy consumption and 2 percent of CO_2 emissions are caused by the information and communication technology (ICT) infrastructure. Moreover, it is estimated that ICT energy consumption is rising at 15–20 percent per year, doubling every five years [1]. According to one estimate, about a tenth of this can be attributed to cellular mobile communication systems. As of 2008, this corresponded to 60 billion kWh of electricity usage annually, about 40 million metric tons of CO_2 emissions each year. To put it in perspective this is equivalent to annual greenhouse gas emissions from about 8 million cars. Another consistent estimate states that there were over 600 thousand base stations in China deployed by three major operators which consumed about 20 billion kWh in 2007.

From the perspective of cellular network operators, reducing electrical energy consumption is not only a matter of being green and responsible, it is also very much an economically important issue. A significant portion of the operational expenditure (OPEX) of a cellular network goes to pay the electricity bill. From the above, it can be estimated that the mobile network OPEX for electricity globally is more than $10 billion dollars today.

For all these reasons, cellular network operators have been exploring ways to increase energy efficiency in all components of cellular networks, including mobile devices, base stations, and core (backhaul) networks. There has been a tremendous amount of work on mobile device energy efficiency with the objective of prolonging battery life time. Similarly, green operation of Internet has been considered and some of the techniques can be extended to the cellular backhaul networks. However, the key source of energy usage in cellular networks is the operation of base station equipment. It has been estimated that base stations contribute to 60–80 percent of the total energy consumption [2]. In this article, we therefore focus on energy-efficient operation of cellular base stations, which is referred to as green cellular operation.

Energy efficiency with respect to base stations has been considered in all stages of cellular networks, including hardware design and manufacture, deployment, and operation. A number of these efforts have focused on hardware improvements. For instance, next-generation base stations are designed to be substantially more energy efficient, for example, using more energy-efficient power amplifiers and natural resources for cooling. Others have considered collocating cellular base stations with renewable energy sources such as solar power and wind energy. In addition, cellular operators have evaluated deployment strategies that minimize the energy expenditure on base stations [3]. We primarily discuss operation, and touch on deployment.

There is room for significant improvement in cellular operation. Even when a site is experienc-

0163-6804/11/$25.00 © 2011 IEEE

ing little or no activity, the base station consumes more than 90 percent of its peak energy; for example, a typical universal mobile telecommunications system (UMTS) BS consumes 800–1500 W and has a radio frequency (RF) output power of 40–80 W. While turning off some of the radio transceivers on a BS during low traffic periods can provide some relief, this is still not sufficient. For significant energy savings, what is called for is a more carefully coordinated dynamic approach that allows the system to shut entire base stations and transfer the corresponding load to neighboring cells during periods of low utilization.

In this article, we explore such dynamic base station management mechanisms for cellular networks to understand quantitatively the scope for potential energy savings, and also the technical challenges that arise in implementing these mechanisms, particularly to ensure that energy efficiency does not come at the expense of reduced quality of service for mobile customers. Note that although similar approaches can also be considered for reducing the energy consumption in wireless local area networks such as 802.11-based WLANs, we do not explore these in this work as the electrical energy usage and expense in such networks is an order of magnitude lower than in cellular networks (e.g., Cisco estimates that its wireless control system adaptive power management software can save 7 W/h per Aironet access point).

The article is organized as follows. We present a first-order analysis on the scope of potential energy saving through green cellular operation. We use a simple greedy algorithm to estimate the energy saving based on real cellular traffic traces and actual base station location information in the analysis. We discuss related challenges and potential solutions in green cellular operation, including maintaining coverage, enabling cooperation between operators, and providing E911 service. Last, we conclude the article.

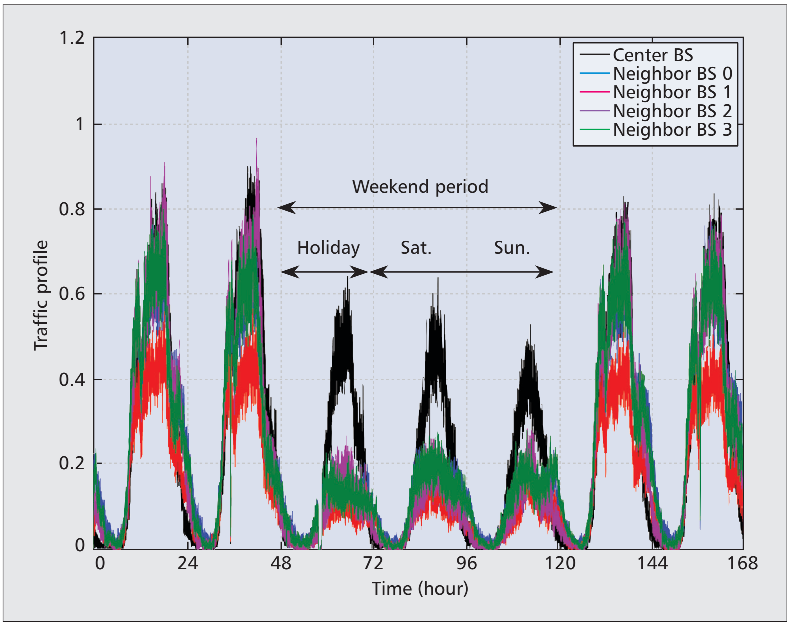

Figure 1. *Normalized traffic profile during one week from real voice call information.*

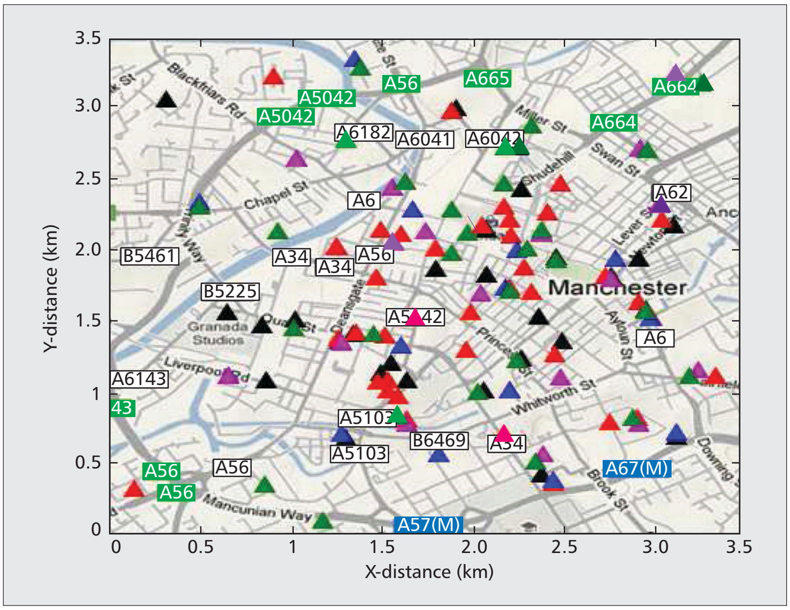

Figure 2. *Base station location data from a part of Manchester, United Kingdom. (Some base stations are collocated, in which case just one triangle is shown on the map.)*

A CASE STUDY

The basic idea behind green cellular operation is to alleviate the inefficiencies resulting from the fact that today's base stations are typically deployed and operated continuously based on peak traffic estimates. Intuitively, it saves energy to carefully turn off unutilized and underutilized base stations during off-peak times, while maintaining coverage. Next, we use real cellular network data to estimate the degree of energy savings possible in urban areas using such a dynamic base station operation.

We look at two sets of real data. The first is a temporal traffic trace shown in Fig. 1, which was obtained from a cellular operator (the source has to be anonymized to preserve confidentiality) in a metropolitan area; and the second is base station location data from a central portion of the city of Manchester, United Kingdom, as illustrated in Fig. 2, obtained from a UK government-sponsored website [4]. We discuss these traces in detail below.

TEMPORAL TRAFFIC TRACE

The data trace we have contains voice call information from five base stations — one central base station and its neighboring four base stations — over the course of one week. The data is plotted with a resolution of 1 s in Fig. 1. We normalize the values of load by a constant value to show the relative change. We observe a periodic sinusoidal profile of the traffic in each cell and notice that the traffic during the daytime (11 a.m.–9 p.m.) is much higher than that at nighttime (10 p.m.–9 a.m.). In addition, the traffic profile during a weekend/holiday period, even during peak hours, is much lower than that of a normal weekday.

SPATIAL DEPLOYMENT

Figure 2 is a snapshot of the base station locations (shown as colored triangles) in part of the city of Manchester, United Kingdom. This infor-

Threshold	Weekday	Weekend	Average per week
5% of peak	23.2	29.8	25.1
10% of peak	30.2	43.3	34.0
20% of peak	38.6	75.6	49.2

Table 1. *Analysis of sample cellular traffic load profiles: percentage of time the traffic is below* x *percent weekday peak during weekdays and weekends, for* x *= 5, 10, 20.*

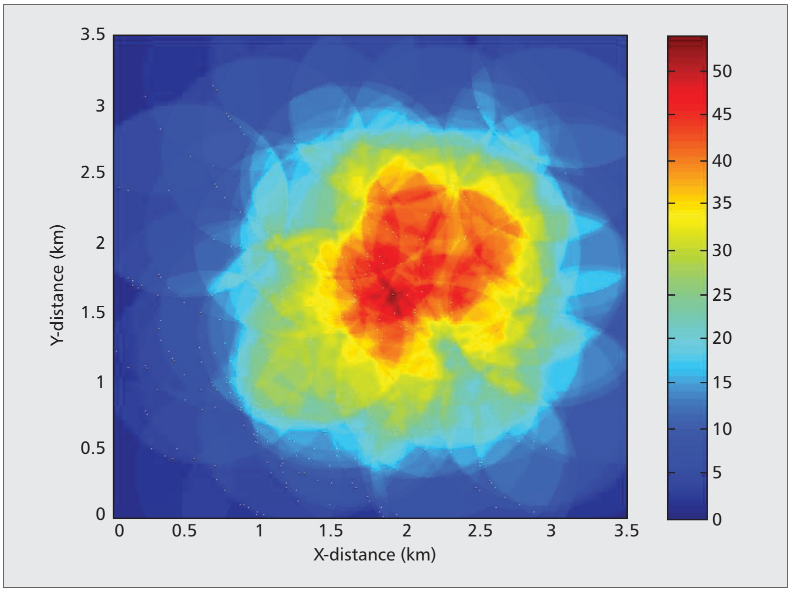

Figure 3. *Redundancy in the cellular coverage of the Manchester area at base station coverage.*

mation is obtained from a U.K. government sponsored interactive website [4] based on information provided voluntarily by mobile network operators on the location and operating characteristics of individual base stations. The information on this site includes location information, operator information, maximum transmission power, and cellular tower height and operating frequency. The information is updated by operators every three months or so [4]. We manually extract base station data from a part of Manchester in the United Kingdom from this website. This area is of size 3.5 × 3.5 km, and has a total of 139 base stations in 128 locations.

ENERGY SAVING ESTIMATE

We now use these real datasets to estimate the energy savings that are possible using dynamic base station operation in this part of Manchester. Note that the anonymized temporal trace data we have does not actually correspond to the base station locations in Manchester. However, our methodology for estimating the energy savings from dynamic base station management decouples the temporal and spatial components, as follows. We use the temporal traffic trace to generate rough estimates of the fraction of time when cells show low activity during weekday and weekend/holiday periods. Independently we analyze the spatial distribution of base stations (on a per-operator basis, as well as jointly) assuming uniform circular coverage zones to estimate the level of redundancy that can be eliminated while maintaining more than 95 percent of the coverage, for different communication ranges. We assume for now that the remaining 5 percent of coverage can be made up through the kinds of techniques discussed later. We combine these two intermediate estimates to generate the final estimate of the expected energy savings.

Based on the traffic load profiles shown in Fig. 1, we obtain the percentage of time that the traffic is below *x* percent of weekday peak, during weekdays and weekends, where $x = 5, 10, 20$. The numbers are presented in Table 1. In the following discussion, for ease of exposition, we shall focus on time when traffic is less than 10 percent of the peak (further work is needed to better understand exactly what fraction of peak traffic should be chosen as the threshold for dynamic base station operation). We find that during weekdays, about 30 percent of the time the traffic is less than 10 percent of the peak. During weekends and holidays, this low activity period increases to 43 percent of the time. Assuming two weekend days in a typical week, we estimate that the average low activity period is about 34 percent of the time.

We next consider the level of redundancy present in the spatial base station deployment. As shown in Fig. 2, the deployment shows non-homogeneous density for each operator, with greater density near the city center, which is a primarily commercial district. In the interest of a simplified first-order analysis, we assume that all cells have a common maximum coverage range. (In order to minimize intercell interference, some base stations may be operated below this range when all base stations are activated. We assume that the range of these base stations can be extended to the maximum level when necessary.) If all base stations were to operate at maximum range, there would be significant redundant overlap in coverage. For instance, Fig. 3 shows the coverage of all base stations when a common range of 700 m is chosen for each base station, with colors indicating the extent of redundancy. Near the city center, we find that more than 50 base stations from various operators overlapping if operated at this maximum range.

We first treat each operator separately. For a given maximum range, we use a greedy algorithm to identify a minimal number of base stations such that the net area coverage of these base stations is more than 95 percent of the area covered when all base stations are active. In the greedy algorithm, we sequentially switch off the base station with the minimum distance to its nearest active base station so long as the coverage condition is met. (We discuss later how any remaining coverage holes may be filled by techniques such as range extension, multihop relay, and multipoint coordination). This is plotted in Fig. 4. As expected, the level of redundancy is higher when the maximum range is higher. For example, when the maximum range is 700 m, depending on the operator, from 25 to more than 65 percent of the base stations can be deactivated while maintaining more than 95 percent of the original coverage.

We also plot in this figure the percentage of base stations that could be switched off while preserving 95 percent of the original coverage if all operators were able to share the base station resources. Clearly the overall energy savings are highest in this case (note that from the dataset we have, it appears that some sites are already jointly shared by multiple operators; in this analysis we have assumed that there is a single base station shared by all active operators at such sites). For the 700 m range, we observe that more than 85 percent of the base station can be shut down while keeping greater than 95 percent coverage with minimal overlap in this case. Comparing the total turned-off base stations in the case when operators act individually, we see that base station sharing increases the energy saving by about 35 percent during off-peak hours.

Putting together the temporal and spatial analysis above, we estimate that individual operators can save between about 8 percent to 22 percent of energy in such an urban deployment. Sharing base station resources together, we get a total reduction of about 29 percent for the energy expended on base station operation. Thus an important collateral finding of our analysis is that greater cooperation among operators is essential for substantial savings. We shall discuss this issue further next.

This percentage of energy saving corresponds to between 32 and 60 kWh of absolute energy savings for the roughly 12 km^2 area of Manchester we have considered (assuming the single base station power is between 800 and 1500 W). This in turn translates to about \$42,000 to \$78,000 annually for the electricity bill for this set of base stations, or about 200 to 375 metric tons of annual CO$_2$ emissions. This is a substantial reduction in greenhouse gas emission as well as cost of operation.

Some caveats are in order. Our study has not considered heterogeneous networks. Also, an analysis of savings requires better radio propagation modeling. These can be considered in future work.

CHALLENGES AND FUTURE DIRECTIONS

In the above analysis we focused on obtaining a coarse-grained estimate of energy savings if "redundant" base stations could be switched off during periods of low activity. However, this analysis has glossed over many important details pertaining to how such dynamic base station operation can be implemented in practice. A number of questions arise: What are the mechanisms by which coverage can be maintained when subsets of the base stations are switched off? What changes may be needed to make the mobile units more "cognitive" in order to enable more agile base station management across multiple spectrum bands? What is the temporal granularity at which base station operation decisions should be made? At what locality level of the cellular network hierarchical architecture should these decisions be implemented? Is there really an incentive for different operators to cooperate in base station operation? If so, how can such cooperation be realized? We now explore these issues in greater detail, in the process identifying avenues for future research and development in this area. In particular, we focus on coverage extension, then address other technical issues, and consider the location estimation problem that arises in some cellular deployments.

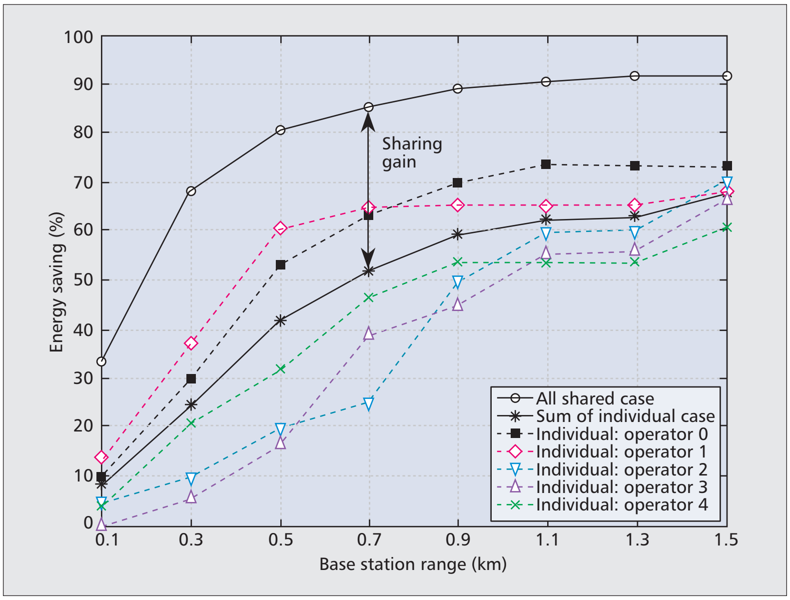

Figure 4. *Energy saving during off-peak times vs. the base station range in the Manchester area.*

MAINTAINING COVERAGE USING POWER CONTROL

The primary constraint of green cellular operation is to preserve coverage and service quality when certain cellular base stations are turned off in the case of low load. Techniques for preserving coverage and maintaining quality of service have much similarity. To be concise, we focus the following discussion on maintaining coverage.

Power control will play an important role in balancing coverage and interference. Power control has been extensively used in cellular networks for interference management and quality of service provisioning. In the case of low traffic load, the main challenge is to increase coverage, while interference management is less critical. Therefore, one can potentially increase transmission power when some base stations are turned off to increase the coverage area of the remaining base stations. Although a simple idea, it needs to be carefully evaluated for both uplink and downlink transmissions.

COGNITIVE MULTIFREQUENCY OPERATION

Multi-frequency-band operation can potentially be explored in green operation. Lower frequency bands have better penetration capability and can provide better coverage under the same transmission power constraint. An example of such a band is the vacant 700 MHz TV band in the United States, auctioned in February 2009 and acquired mainly by large cellular operators. Such a lower frequency band could potentially be used by larger cells, overlapping with smaller cells on a higher frequency band. It is worth noting that frequency selection requires more advanced physical layer technology, which is available in current commercial mobile devices to some degree, with more advanced features being developed as part of the cognitive radio paradigm.

MULTIHOP RELAY FOR COVERAGE EXTENSION

Various techniques have been considered in cellular networks to improve network coverage

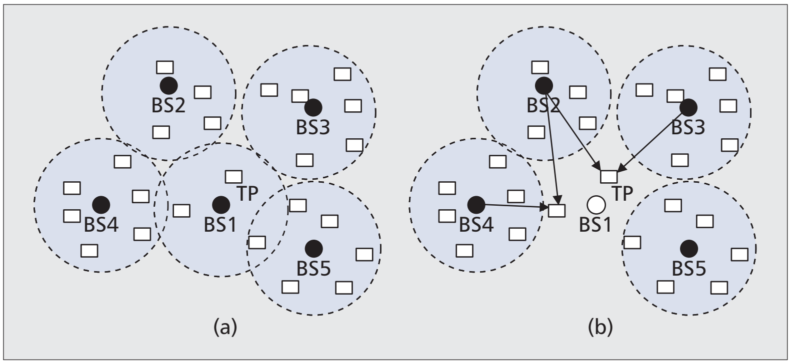

Figure 5. *Coverage extension with coordinated transmissions.*

and/or increase network throughput. Some of these techniques can be applied in the case of green cellular operation with a focus on covering the cells with shut-down base stations. The basic idea is simply to allow other devices to relay the traffic of a cellular user that receives weak or no signal directly from its nearby cellular base stations. This will clearly be useful to ensure that the dynamic shutting down of base stations, while saving energy, does not leave coverage holes. However, finding such relays can be a challenge.

COORDINATED MULTIPOINT TRANSMISSION AND RECEPTION

Another novel approach that can play a role in extending coverage is the use of coordinated multipoint transmission and reception (CoMP) being developed in the context of Long Term Evolution (LTE)-Advanced cellular networks [5]. The basic idea in CoMP, a macro-diversity scheme, is to coordinate transmission and reception at neighboring cells. As shown in Fig. 5, there are many users and five candidate base stations. The dashed lines demonstrate the transmission ranges of the base stations. If intercell cooperation is not allowed, all five base stations should operate to guarantee the coverage requirement (Fig. 5a). However, if the neighboring base stations can transmit cooperatively (e.g., BS2 and BS4, BS2 and BS3) to enlarge their coverage, all three users served by BS3 can be covered, and BS3 can be switched off to save energy (Fig. 5b). Similarly on the uplink, coordinated receptions at multiple base stations can effectively reduce the received power requirement at each individual base station. The increased power on the downlink and lowered power requirement at the uplink together provide coverage with acceptable service quality for mobile users in nearby cells whose base stations have been shut down during low activity periods [6].

HANDLING HETEROGENEOUS CELL SIZES

In our simplified estimate of energy savings, we considered uniform cell sizes. In third-/fourth-generation (3G/4G) systems, because of the required high data rate, cellular towers will be more dense and have varying coverage, with more heterogeneous cells, such as microcells, picocells, and femtocells. While most existing deployment strategies focus on quality of service during peak periods, energy savings during off-peak periods should also be taken into account. In particular, heterogeneous cellular architecture should be considered in the deployment stage so that umbrella macrocells can provide overall coverage with lower data rate, and smaller microcells/picocells can provide better data rates for congested areas [7]. This can potentially be exploited for energy efficiency, since in the case of light load we can turn off (most of) the smaller cells. On the other hand, considering energy efficiency introduces additional complexity in network planning and deployment, and thus needs to be further studied.

GRANULARITY OF CONTROL

The issue in dynamic base station switching is the temporal granularity of the control. While most of the early work in this area has primarily focused on switching off base stations once a day, it remains to be seen whether finer-grained operation (e.g., on an hourly or even more frequent basis) can improve energy efficiency, particularly in settings where traffic patterns are not predictable from day to day. However, this may require more active online monitoring of traffic and increase the complexity of coordination with other cells to ensure coverage.

Another related question is the spatial granularity of coordination. Base stations can be switched off and on in a distributed manner at either a base station controller or mobile switching center level. This granularity can affect different trade-offs between timeliness, complexity of coordination, and efficiency of the base station switching policy [8].

COOPERATION BETWEEN OPERATORS

As shown in our case study earlier, there is room for greater energy savings if different operators can pool their resources. This is particularly helpful in metropolitan areas where each operator sets up a dense deployment. Such a scenario has also previously been examined by Marsan and Meo [9], who consider the setting of a single cell with two operators that cooperate with each other by accepting each others' traffic (as roaming traffic) when they are shut down and show that considerable energy savings can be obtained in such a setting. In implementing such cooperation, physical sharing may not be a substantial concern. In many cases, in urban areas operator base stations tend to be closely situated, or even on the same tower. The main challenge is the complexity in network operation, regarding issues such as cross-operator user authentication and billing, introduced by fine-granularity roaming.

This is also an interesting problem from a game theoretic perspective. Under what conditions would self-interested operators agree to cooperate with others? What kind of profit-sharing agreements will provide an adequate incentive for all participants? Cooperative game theory formulations that consider coalition formation may provide insights on these questions. From a policy perspective, however, it also needs to be examined whether such operator agreements can potentially be abused to create oligopolies that hurt customers, and how these can be ameliorated in the interest of green operation that offers other benefits to society.

MAINTAINING EMERGENCY LOCATION SERVICE

There are other practical considerations related to the green operation of cellular infrastructure. One such issue in the United States is E911 (Enhanced 911). E911 is a North American telecommunications-based system that automatically associates the physical address with a caller's phone number, which enables the call to be routed to the most appropriate public safety answering point for that address. In cellular networks, this requires location determination of a cellular caller. In E911 Phase II, it is required that 95 percent of a network operator's in-service phones be E911 compliant, with a granularity of 300 m. Therefore, if cellular towers are used to triangulate caller locations, the 1-coverage requirement needs to be extended to 3-coverage (i.e., each caller can reach three cellular towers). However, as shown in a large-scale cellular network study by a major cellular operator in the United States, over 50 percent of the time, a caller receives information from only one base station instead of the three required for triangulation or trilateration [10]. Therefore, to achieve 3-coverage in the original cellular network, much higher cellular tower density is required. To maintain 3-coverage at low traffic volume, we might expect a similar percentage of energy saving. More detailed analysis is needed. Note that there are also alternative and complimentary techniques that address the issue of location determination, including using GPS data, combining with WiFi location information, and using Bayesian inference. Such location techniques apply in the case of green cellular operation as well.

CONCLUDING COMMENTS

In this article we have discussed the current trends in green cellular technology. We have focused on green cellular operation. Using real data traces, we derived a first-order approximation of the percentage of power saving one can expect by turning off base stations during low traffic periods while maintaining coverage. Our coarse-grained analysis shows promising potential. We have also presented and discussed a number of relevant challenges and solutions, such as maintaining coverage, enabling cooperation between operators, and providing E911 service.

In summary, we argue that energy efficiency is an important metric that should be used in the design and analysis of cellular networks, along with the metrics that have been considered traditionally, such as blocking/dropping probability, throughput, and delay. While a few recent papers, such as [2, 6–9], have started to address this topic, there are still many rich open problems in this domain, and we encourage the community to give it greater attention.

We argue that energy efficiency is an important metric that should be used in the design and analysis of cellular networks, along with the metrics that have been considered traditionally, such as blocking/dropping probability, throughput, and delay.

ACKNOWLEDGMENTS

We thank Thomas Lee, Yucho Ho, and Sophia Tsang, undergraduate students at the University of California, Davis, for their help in extracting the base station information used in this article. We thank Mr. Raj Savoor from AT&T Labs for a helpful discussion regarding E911 and Mr. Donald Draper Campbell from FCC for helpful input regarding cellular coverage. The work of Dr. Oh was supported by a National Research Foundation of Korea Grant funded by the Korean Government (NRF-2009-352-D00205). The work of Dr. Niu is partially supported by the National Science Fund for Distinguished Young Scholars (60925002).

REFERENCES

[1] G. Fettweis and E. Zimmermann, "ICT Energy Consumption — Trends and Challenges," *Proc. IEEE WPMC*, Lapland, Finland, Sept. 8–11, 2008.

[2] M. A. Marsan *et al.*, "Optimal Energy Savings in Cellular Access Networks," *Proc. IEEE ICC*, Dresden, Germany, June 18, 2009.

[3] "Sustainable Energy Use in Mobile Communications," White Paper, Ericsson, Aug. 2007.

[4] Ofcom, "Sitefinder: Mobile Phone Base Station Database," http://www.sitefinder.ofcom.org.uk/.

[5] S. Parkvall and D. Astely, "The Evolution of LTE Towards IMT-Advanced," *J. Commun.*, vol. 4, no. 3, Apr. 2009, pp. 146–54.

[6] D. Cao *et al.*, "Energy Saving Performance Comparison of Coordinated Multi-Point Transmission and Wireless Relaying," *Proc. IEEE GLOBECOM*, Miami, FL, Dec. 6–10, 2010.

[7] A. J. Fehske, F. Richter, and G. P. Fettweis, "Energy Efficiency Improvements through Micro Sites in Cellular Mobile Radio Networks," *Proc. IEEE GLOBECOM*, Honolulu, HI, Dec. 4, 2009.

[8] S. Zhou *et al.*, "Green Mobile Access Network with Dynamic Base Station Energy Saving," *Proc. ACM MobiCom* (Poster Paper), Beijing, China, Sept. 20–25, 2009.

[9] M. A. Marsan and M. Meo, "Energy Efficient Management of two Cellular Access Networks," *Proc. ACM SIGMETRICS*, Seattle, WA, June 15, 2009.

[10] Z. Hui, F. Baccelli, and J. Bolot, "Bayesian Inference for Localization in Cellular Networks," *Proc. IEEE INFOCOM*, San Diego, CA, Mar. 15–19, 2010.

BIOGRAPHIES

EUNSUNG OH (eunsung.oh@usc.edu) received his B.S., M.S., and Ph.D. degrees in electrical engineering from Yonsei University, Seoul, Korea, in 2003, 2006, and 2009, respectively. He is currently a post-doctoral researcher in the Department of Electrical Engineering at the University of Southern California's Viterbi School of Engineering. His main research interests include the design and analysis of algorithms for green communication networks and next-generation wireless networks.

BHASKAR KRISHNAMACHARI (bkrishna@usc.edu) received his B.E. in electrical engineering at The Cooper Union, New York, in 1998, and his M.S. and Ph.D. degrees from Cornell University in 1999 and 2002, respectively. He is currently an associate professor and a Ming Hsieh Faculty Fellow in the Department of Electrical Engineering at the University of Southern California's Viterbi School of Engineering. His primary research interest is in the design and analysis of algorithms and protocols for next-generation wireless networks.

XIN LIU (liu@cs.ucdavis.edu) received her B.S. and M.S. degrees in electrical engineering from Xi'an Jiaotong University, People's Republic of China, in 1994 and 1997, respectively; and her Ph.D. degree from Purdue University in 2002. She is currently an associate professor in the Computer Science Department at the University of California, Davis. Before joining UC Davis, she was a postdoctoral research associate in the Coordinated Science Laboratory at UIUC. Her research interests include resource allocation and scheduling mechanisms in wireless networks, integration of wireless and wireline networks, sensor networks, ultra-wideband communication, and timing channels.

ZHISHENG NIU [SM] (niuzhs@tsinghua.edu.cn) graduated from Northern Jiaotong University, Beijing, China, in 1985, and got his M.E. and D.E. degrees from Toyohashi University of Technology, Japan, in 1989 and 1992, respectively. After spending two years at Fujitsu Laboratories Ltd., Kawasaki, Japan, he joined Tsinghua University, Beijing, China, in 1994, where he is now a full professor in the Department of Electronic Engineering. His current research interests include teletraffic theory, mobile Internet, radio resource management of wireless networks, and cognitive and multiple radio networks. Currently, he is a Fellow and Councilor of the IEICE, Director for Conference Publications of IEEE Communications Society, and council member of the Chinese Institute of Electronics (CIE).