

Low-Carb: Reducing Energy Consumption in Operational Cellular Networks

M. S. Ilyas*, G. Baig*, M. A. Qureshi*, Q. U. A. Nadeem*, A. Raza*, M. A. Qazi*, B. Rassool[¶]

*SSE, LUMS, Lahore, Pakistan, {saqibm,13100162,13100152,13100094,14100214,14100212}@lums.edu.pk

[¶]Warid Telecom, Lahore, Pakistan bilal.rassool@waridtel.com

Abstract—Electricity costs are a significant fraction of a cellular network’s operations costs. We present Low-Carb, a practical scheme to decrease electrical energy consumption in operational cellular networks by coupling *Base Transceiver Station (BTS) power savings* with call *hand-off*—two features commonly used by cellular operators. Motivated by the practical observation that most callers are in the vicinity of multiple BTSs, Low-Carb presents and solves an optimization problem, allowing calls to *hand-off* from one BTS to another so that *BTS power savings* can be applied to a maximal number of BTSs throughout the cellular network.

We use BTS locations and traffic volume data from a large live GSM network to evaluate the power savings possible using our proposed approach in Low-Carb. Our results indicate that for a GSM 1800 network operator with 7000 sites in an urban setting, a total of up to 35.36 MWh may be saved annually. This is at least 9.8% better than the energy savings obtained by just using BTS power savings alone. Other cellular operators can use the Low-Carb formulation with their own network data to estimate the electricity savings they may achieve on their networks.

I. INTRODUCTION

Cellular networks consume several tens of TWhs of electrical energy every year worldwide [1]. This not only results in significant operational expenditure, which is increasing with rising electricity and fuel prices, but is also a source of concern for ecological reasons. These concerns have motivated a lot of research aimed at reducing energy consumption in cellular networks.

In this paper, our discussion focuses on the legacy 2G GSM cellular networks, which have a significant market share. Our focus on GSM is also due to their dominance in the typically energy-starved developing countries with a large subscriber base. These GSM networks are expected to persist for the foreseeable future due to the upgrade expenses and return on investment (ROI) concerns of the operators. It is, therefore, important to optimize such GSM networks. The energy-saving techniques proposed in this paper may be applicable to 3G and 4G networks, but we make no claims as to the effectiveness of the same.

The major sink of power in a cellular network are Base Transceiver Stations (BTSs), accounting for 60% to 80% of the total power consumption [1–3]. Every BTS is equipped with several transceivers (TRXs), each of which is allocated a single frequency band for transmission and reception of radio signals. Each TRX further uses time multiplexing to handle up to 8 full-rate voice calls over its assigned frequency band in GSM systems. A typical configuration is “6+6+6” depicting a BTS

serving three *sectors* each with six TRXs. Thus, a BTS offers a *fixed* capacity, as determined by the total number of TRXs installed. Sites are deployed such that this fixed BTS capacity can handle the peak traffic load. However, traffic peaks only for a short duration dropping off to a much lower trough each day, which means that the GSM networks are over-provisioned during low-traffic regimes.

Over-provisioned BTSs would be fine if they were also load-proportional, i.e., consumed little power at no traffic load. However, according to [3] the no-load power consumption can be as high as 95% of that at full load. With fixed BTS capacity that is over-provisioned for low traffic loads, today’s cellular networks are highly energy inefficient.

There are generally two approaches to increase cellular network energy efficiency. First, a clean-slate redesign which includes innovations in communication systems, circuits and components. This approach is not attractive for existing GSM operators, which are the most prevalent in the developing world and are expected to stay as such for several years to come, primarily due to the required expensive upgrades. A second approach is to make optimizations to the existing system and equipment to get an improvement in overall energy efficiency. Our present work is aligned with this latter philosophy.

One can improve the energy efficiency of a cellular network by adapting its “online” capacity to changes in traffic load. Recent work has proposed turning off base stations to reduce energy consumption during times of low traffic load [1–4]. In such solutions, to offer the same amount of coverage, the transmit range of some of the remaining BTSs has to be increased. Our conversations with multiple network operators indicate that they are reluctant to employ such techniques citing three reasons:

- Power cycling of entire base stations is expected to reduce equipment life time.
- Turning off some BTSs may require an increased up-link power which may not be handled by many low-cost/power-limited mobile stations (MSs). This raises a risk of customer churn and is not acceptable to the operators in cut-throat competition prevalent in today’s market.
- These techniques of turning off BTSs may underestimate the increase in power needed for indoor MSs.

Our conversations with wireless providers reveal that during low traffic periods, they often use a feature available in most

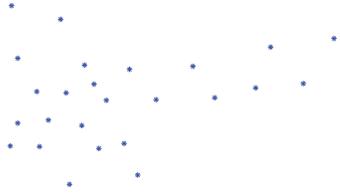


Fig. 1. A depiction of the geographical dispersal of BTSs

vendor’s equipment that power-gates TRX circuits at locations that serve very few customers. Huawei calls this feature *TRX shutdown* while Ericsson calls it *BTS power saving*. We use the latter term generically in this paper.

BTS power consumption’s traffic-independent component depends largely on the number of active TRXs [3]. Therefore, deactivating TRXs reduces the BTS power consumption. For instance, turning off one TRX cuts down BTS power consumption anywhere from 20W to 100W, depending upon the frequency band (900 or 1800) and deployed equipment [5, 6]. Thus, scaling a “6+6+6” to a “2+2+2” configuration, by deactivating 12 TRXs will result in a saving of 240W to 1200W on a single site. The decision to use *BTS power saving* feature is generally local to the BTS without any coordinated effort at the network level.

This paper presents Low-Carb which combines the *BTS power saving* with *hand-off*, another commonly used feature in cellular networks that facilitates user movement from one location to another. Low-Carb proposes to hand-off calls from one BTS to another, without making a negative impact on the network quality of service, such that the *BTS power savings* can be applied to a maximal number of base stations throughout the cellular network. As compared to the use of uncoordinated *BTS power savings*, Low-Carb offers additional power savings as it may allow a larger number of TRXs to be deactivated. In present day deployments, this is possible since most callers receive sufficiently strong signal simultaneously from several nearby BTSs some of which may have relatively lower traffic [3]. Fig. 2 shows coverage diversity evident in the urban data from a large cellular provider that we used in our evaluations; one can see that about half of the callers have 3 or more candidates for serving BTS. Furthermore, Fig. 3 shows normalized traffic at two neighboring sites in our dataset for a 24 hour period, which confirms the presence of geographic diversity in traffic.

We formulate a binary integer program (BIP) to minimize the power consumption in a GSM network by shuffling active calls between nearby BTSs while keeping in check the MS up-link budget. Since BIP is NP-Hard, we also propose a heuristic for Low-Carb and evaluate it’s performance compared to the optimal solution.

Our work is very similar in spirit to the concept of *frequency dimming* in [7] albeit at a different level of abstraction. A similar approach is also proposed in [8] with some rough estimates of expected savings. We, on the other hand, use site

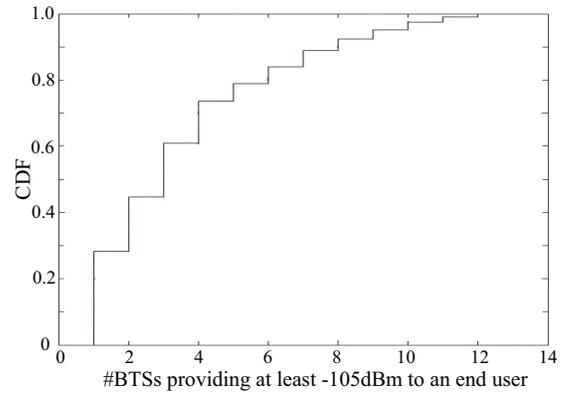


Fig. 2. Empirical CDF of the number of potential serving BTSs for a call in our dataset (large metropolitan area)

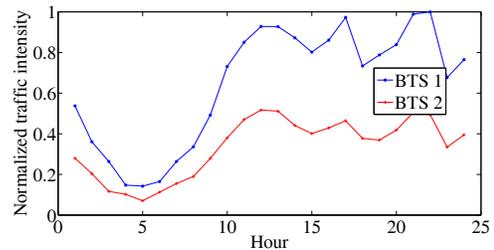


Fig. 3. BTS power consumption model. Low-power (BTS power savings) mode is optional and kicks in at low loads

locations and traffic traces from a large cellular network with more than 13 million subscribers to run a simulation study assessing the benefits of dynamic equipment scaling coupled with call hand-offs. A key benefit of our approach is that it does not require any additional hardware and works within the GSM specifications.

The rest of the paper is structured as follows. The formulation of Low-Carb optimization problem is given in section II. Experimental setup and the results are presented in sections III and IV, respectively. In section V, we draw the conclusions highlighting the power saving strategy for providers.

II. FORMULATION

A. Single Base Transceiver Station (BTS)

Power consumed by a BTS, as a function of traffic load, can be well approximated as a linear curve with a non-zero y-intercept [3] given as $P_1 + l(P_2 - P_1)/t_{max}$. Here P_1 and P_2 are the power consumption at no load and full load, respectively, l is the number of calls presently being handled, and t_{max} is the maximum number of calls that can be handled.

Let δ be the traffic threshold below which the *BTS power savings* may be applied. Since all TRXs are identical, the per call increase in power consumption, and hence the slope of the power consumption profile in Fig. 4, remains the same whether or not some TRXs are deactivated. As also indicated in Fig. 4, the no-load power consumption drops to $P_1 - \gamma$ in the low-power mode, where γ is a constant that depends on the equipment type and the number of TRXs deactivated. If x

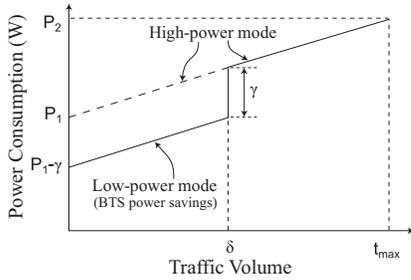


Fig. 4. BTS power consumption model. Low-power (BTS power savings) mode is optional and kicks in at low loads.

is an indicator variable which is 0 when *BTS power savings* is applied, and 1 otherwise, then the BTS power consumption may be given by $P_1 + l(P_2 - P_1)/t_{max} - (1 - x)\gamma$ (indicated in Fig. 4 by the piecewise linear solid line).

B. Multi-BTS Cellular Setting

Consider an area with n active callers being served by m BTSs. We introduce indicator variable $w_{i,j}$, which is 1 if call i is being handled at BTS j and 0 otherwise. We assume availability of an $n \times m$ matrix whose entry $c_{i,j}$ is 1 if caller i can be served through BTS j without exceeding the uplink or downlink budgets. This information can be extracted by the data periodically transmitted by each MS comprising the received signal strength from nearby BTSs during a call. We also introduce indicator variable x_j , which is 1 if BTS j is operating in high-power mode (i.e., without *BTS power savings*) and 0 otherwise. The total power consumption may, therefore, be given as $\sum_{j=1}^m \left[P_1 + \sum_{i=1}^n \frac{w_{i,j}(P_2 - P_1)}{t_{max}} - (1 - x_j)\gamma \right]$. Using this as the objective function, we can formulate the Low-Carb optimization problem, subject to appropriate constraints. However, some of the terms in the above summation are constant and will be omitted from the objective function. The term P_1 in the above summation is a constant, so it can be excluded from the objective function. Furthermore, in order to not affect the grade of service, we will not drop any active calls. Therefore, the summation over $w_{i,j}$ is also constant and can also be excluded from the objective function. After removing the constant additive and multiplier terms from the summation, the Low-Carb optimization may be stated as:

$$\text{minimize } \sum_{j=1}^m x_j \quad (1)$$

subject to the following constraints:

$$\sum_{j=1}^m w_{i,j} = 1 \quad \forall i \quad (2)$$

$$w_{i,j} \leq c_{i,j} \quad \forall i, j \quad (3)$$

$$\sum_{i=1}^n w_{i,j} - \delta \leq Mx_j \quad \forall j \quad (4)$$

$$\sum_{i=1}^n w_{i,j} \leq t_{max} \quad \forall i \quad (5)$$

$$w_{i,j}, x_j \in 0, 1 \quad \forall i, j \quad (6)$$

The first constraint ensures that no active call is dropped just to save on power. The second constraint secures the uplink budget by ensuring that no call is routed to a BTS that is too far away. The third constraint picks the correct value for the decision variable x_j . The fourth constraint is the capacity constraint on all BTSs, while the last constraint is the binary value constraint on the decision variables.

The above optimization problem is a Binary Integer Program (BIP), which is NP-Hard. It is intractable to solve for an operator's entire network. For a large network, it could be applied separately and independently to small disjoint network segments. Alternatively, a heuristic solution to Low-Carb could be deployed over the entire network. We present a heuristic for Low-Carb in Algorithm 1.

Our heuristic first partitions the set of BTSs B into disjoint sets B_1 and B_2 where the latter includes all the BTSs in low-power mode and the former consists of all other BTSs. The heuristic iterates over a random permutation of the BTSs in B_1 . Once a BTS b_j from B_1 is selected (line 3), our heuristic determines the minimum number of calls that must be handed off from b_j before it can be moved to B_2 . Our heuristic iterates over the calls that are presently being handled on b_j and for each such call it attempts to find a candidate serving BTS in B_2 . The choice of the new serving BTS (b_d) is made while ascertaining that the hand off will not cause b_d to move out of B_2 . If such a BTS is found, the call is handed off to it. Calls are handed off from b_j in this manner, until it moves into B_2 , or we exhaust the set of active calls with other candidate BTSs (b_j remains in B_1).

III. DATA AND EXPERIMENTAL SETUP

Our dataset is obtained from a cluster of 26 BTSs operated by a large network operator with more than 7000 sites. These sites are spread over a 31.25 km^2 urban terrain (see Fig. 1). We obtained each site's coverage prediction using a tool popular amongst the operators called Forsk Atoll. With this information, along with a caller's location, we can determine the candidate set of BTSs for the corresponding call (the $c_{i,j}$ parameters). Note that in this work, we do not incorporate user mobility into our model, since we are only interested in instantaneous optimization at small time scales and in determining bounds on the energy savings that Low-Carb can offer.

Also available to us are the hourly cumulative traffic, in Erlang, for each of the sites, spanning two consecutive weekdays. The traffic remained remarkably similar across both days for each site. We have, therefore, only used one day's traffic data in our experiments.

Using the above datasets, we conducted a set of experiments mimicking a 24-hour operation of a subset of a cellular network. Each experiment is a discrete event simulation of the arrival and placement of calls. Since our dataset does not include the arrival times and duration of calls, we synthetically generated this information using the assumption of Poisson

Require: δ : the power-saving traffic threshold,
 B (the set of BTSs) = $\{b_1, b_2, \dots, b_m\}$,
 A (the set of active calls) = $\{a_1, a_2, \dots, a_n\}$,
 W (current call association) = $\{w_{i,j} = 1 \text{ if } a_i \text{ is being served through } b_j, 0 \text{ otherwise}\}$,
 C (Possible call association matrix) = $\{c_{i,j} = 1 \text{ if } a_i \text{ can be served through } b_j, 0 \text{ otherwise}\}$

Ensure: A new and potentially more energy efficient mapping of calls to BTSs

```

1:  $B_1 = \{b_j \mid \sum_{i=1}^n w_{i,j} > \delta\}; \quad B_2 = B - B_1$ 
2:  $B_1 = \text{random\_shuffle}(B_1)$ 
3: for all  $b_j \in B_1$  do
4:    $a = \sum_{i=1}^n w_{i,j} - \delta; \quad d = 1; \quad \text{shuffled} = 0$ 
5:   while  $d < n$  and  $\text{shuffled} \leq a$  do
6:     if  $w_{d,j} = 1$  then
7:        $e = 1; \quad \text{mapped} = 0$ 
8:       while  $e \leq m$  and  $\text{mapped} = 0$  do
9:         if  $e \neq j$  and  $e \in B_2$  and  $c_{d,e} = 1$  then
10:           $w_{d,e} = 1; \quad w_{d,j} = 0$ 
11:           $\text{shuffled}++;$   $\text{mapped}++$ 
12:        end if
13:         $e++$ 
14:      end while
15:       $d++$ 
16:    end if
17:  end while
18:  if  $\sum_{i=1}^n w_{i,j} < \delta$  then
19:     $B_1 = B_1 - \{b_j\}; \quad B_2 = B_2 + \{b_j\}$ 
20:  end if
21: end for

```

Algorithm 1: Heuristic for the Low-Carb problem

call arrivals and exponentially distributed call duration with a mean of 180 seconds [9].

For every hour, the simulator determines the Poisson call arrival rate for each BTS, using Little’s Law and the BTSs traffic intensity for that hour. Using the resulting Poisson process, calls are generated such that it is equally likely for a call to be anywhere in the serving BTSs coverage area.

Our simulator tracks the call volume at every BTS on a minute’s granularity. This enables us to calculate the power consumption level (in Watts) of the BTS during each minute. Accumulating these numbers over the 24 hour period leads to the daily amount of energy consumed (in kWh) if no optimization is used in the network. Our simulator also monitors each BTSs call volume every minute and places the ones with sufficiently low traffic into power-saving mode. This enables us to calculate the possible energy savings using BTS power-saving feature. In addition, our simulator also periodically determines the instantaneous optimal call placement configuration that minimizes the power consumption level by handing-off some calls, thereby placing a maximal number of BTSs in power-saving mode. This allows us to determine the energy savings

possible by combining call hand-offs with BTS power-saving.

The call placement re-optimization may be done at various frequencies. A very aggressive re-optimization regime would keep the network in an optimal state more often than a conservative one, thereby enabling greater energy savings. In order to study the scaling of energy saving with re-optimization frequency, we experimented with a range of intervals between successive optimizations, ranging from a minute to an hour. For a deployment, the re-optimization frequency that can be used would depend on the costs associated with each re-optimization. Let us now consider such costs.

First and foremost, a computational cost is incurred with each optimization. In our case, an optimization run to determine the optimal state over 26 BTSs required an average running time of about 50 seconds on a Core i3 laptop with 4 GB of RAM. An optimization requiring 50 seconds would not be practical to use every minute but may be fine if used less often. For a practical deployment the computational time can be reduced by using a combination of a more powerful machine, distributed optimization and approximation algorithms.

In addition to the computational overhead, for every unit of energy saved some extra energy may be consumed in the network to perform call hand-offs or transitions into and out of BTS power-saving mode. Call hand-offs and TRX (de)activation involve signaling between a Base Station Controller (BSC), BTSs and MSs. The additional energy incurred thus, should be small, because it has been observed that variation in power consumption of network equipment with changes in traffic volume (data or control) is quite small [10]. As far as increased power consumption on MSs due to a greater number of call hand-offs is concerned, we opine that it may be negligible because the MSs energy consumption is far outweighed by that of BTSs.

A. Site Characteristics

All sites in our dataset had three sectors, each equipped with 6 TRXs, for a maximum of 132 simultaneous voice calls¹. The GSM standard includes a provision for half-rate calls, which enables handling greater traffic at the expense of reduced voice quality by allowing a single voice channel to be shared amongst two calls, each using a half-rate codec. In this paper, we only shuffle full-rate calls around, which may be, in reality, two half-rate calls. We do not foresee any significant error arising from using this convention.

We consider a scaling down from a “6 + 6 + 6” site to a “2 + 2 + 2” site, which means that δ should be strictly less than $t_{max}/3$ to avoid quick oscillations into and out of BTS power-saving mode due to short-term traffic variations. We have arbitrarily set δ equal to $\lceil t_{max}/3 \rceil - 5$, because 5 seemed to be a good enough number compared to a sector’s overall capacity and the typical utilization of a site in our datasets.

¹Each TRX’s frequency is shared in time-domain by 8 calls for a total of $3 \times 6 \times 8 = 144$ channels. Four channels in each sector were reserved for control and broadcast purposes.

Parameter	Value		
	Model 1	Model 2	Model 3
P_1	1425	2401.8	2341.5
P_2	1500	3887.5	2973.9
γ	20	50	100

TABLE I
BTS MODEL PARAMETER VALUES

The BTS power consumption model parameters may vary from one BTS model to another. In this paper, we use three different sets of model parameters as listed in table I. We now describe the sources and methods from which we obtained these models.

1) *Model 1*: For the first model, we have used $1.5kW$ as the maximum power consumption [11], a $20W$ per TRX saving when scaling the BTS down [6] and a 5% swing in power consumption between no-load and full-load [3].

2) *Model 2*: Lorincz et. al reported the single sector DC power consumption for a GSM 900 BTS with 7 TRXs [5]. We extrapolate the power consumption for a 6+6+6 site by multiplication of the DC power consumption reported in [5] by $3 \times 6/7$. The DC power consumption does not include the AC power consumed in the power supply units and in air-conditioning. We must, therefore, also compensate for those, to obtain the overall site power consumption. Power supply unit load is negligible compared to air-conditioning, which has a typical power consumption of 1 kW [11]. We used this scaling and AC load correction to obtain the values for P_1 and P_2 using the minimum and maximum reported power consumption in [5]. Furthermore, the authors measured a drop of $50W$ in power consumption when a TRX is disable, which gives us the value of γ as listed in table I.

3) *Model 3*: Using the same method as for model 2 in III-A2, we derived the values for P_1 and P_2 based on the measurements for a GSM 1800 BTS reported in [5]. As for the value of γ , the paper reported a $100W$ cut in power consumption when deactivating a single TRX. The parameter values for this model are given in Table I

IV. RESULTS

First, we consider the benefit of BTS power-saving alone, resulting from traffic diversity at each BTS compared to running the network in the default configuration. The percentage reduction in energy consumption is listed in TABLE II. The results indicate that a saving of between 4% and 12% can be achieved in a network just by enabling BTS power savings. We note here that these results are in agreement with Ericsson's claim of saving 10-20% energy by using BTS power-saving on Germany's Vodafone network [12].

In absolute terms, this represents a cumulative saving of between 43 kWh and 217 kWh per day over a set of 26 BTSs. Now, consider that there are five cellular operators in Pakistan: Mobilink with more than 8500 sites [13], Ufone with more than 8000 sites [14], Zong with more than 5500 sites [14], Telenor with more than 7000 sites [15] and Warid with more than 4500 sites [14]. Overall, there were more than 31000 sites

Energy saving	Model 1	Model 2	Model 3
Percentage	4.73%	5.43%	12.89%
Daily absolute saving over 26 BTSs (in kWh)	43.28	109.68	217.12
Country-wide daily saving over 31000 sites (in MWh)	51.6	130.77	258.87

TABLE II
ENERGY SAVINGS BY USING BTS POWER SAVINGS ONLY

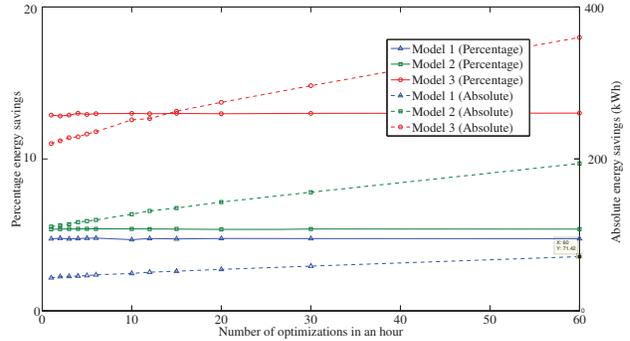


Fig. 5. Percent and absolute reduction in energy consumption vs re-optimization interval

in Pakistan at the end of 2011. Extrapolating our experimental results suggests that a significant reduction in energy consumption is possible country-wide in Pakistan (and any other country with similar deployment and traffic characteristics) by just deploying BTS power-saving. This is especially significant for an energy-starved developing country. As we shall see next, greater energy savings are possible if we couple periodical call shuffling with BTS power savings in the network.

Fig. 5 shows the energy savings achievable by coordinated call handoff and BTS power saving. Along the left-side y-axis, it shows percentage energy saving compared to the default deployment using solid lines for all three BTS models considered in this paper. Since the three BTS models have quite different power consumption characteristics, percentage energy savings do not provide a useful metric for comparison of the relative utility of Low-Carb for the three BTS models. We, therefore, also plot the absolute energy savings (in kWh) for each of the BTS models along the right-side y-axis using dashed lines. For all three BTS models, we see an almost linear increase in energy saving as the duration of the re-optimization interval is decreased. In terms of absolute energy savings (kWh) as well, we see the same linear trend and relative order of the three BTS models as is the case for percentage energy savings.

Re-optimizing at an interval less than the mean call duration should offer greater savings than a less frequent re-optimization. This is because the former regime has the ability to optimally hand off most of the calls at least once. Our results confirm this intuition. For model 1 BTS, for instance, the gain in energy savings while going from a 60 minutes re-optimization interval to 30 minutes is merely 0.0506 kWh per minute, whereas it is 15.5421kWh when decreasing the

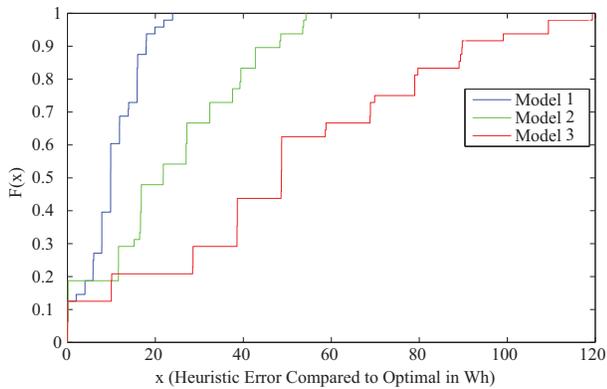


Fig. 6. Empirical CDF of the difference between the cost offered by our heuristic compared to the optimal

re-optimization interval from 2 minutes to 1 minute.

Let us now interpret what these results mean physically in terms of ecological impact. An extrapolation of our results indicates that for Pakistan, Low-Carb offers a total energy saving of 60.72 MWh, 156.84 MWh or 301.61 MWh daily, respectively for each of the BTS models. For a small and energy-starved developing country, these energy savings are significant. Since network deployments and traffic patterns are similar in different countries, the same extrapolation can be applied to other countries as well.

In the above extrapolation, we have assumed that the same amount of energy saving would be applicable in rural as well as urban settings. One may argue that in rural settings, due to sparse deployments the energy savings potential would be lower because few calls would have multiple candidate BTSs. A counter-argument, however, is that in rural settings, call traffic is already low, which implies that BTS power-saving is applicable to most BTSs most of the time.

We also ran experiments for each BTS model in which the electricity cost for the optimal as well as the heuristic solution (Algorithm 1) was computed. We assessed the performance of our heuristic by computing the difference (error) in the electricity cost of the two solutions. For statistical significance, we computed the error in our heuristic relative to the optimal solution over 48 different experiment runs for each BTS model. The resulting CDF of the heuristic error (in Wh) is plotted in Fig. 6. We can see in Fig. 6 that our heuristic is quite close to the optimal solution most of the time, especially for the Model 1 and Model 2 BTS. For Model 3 BTS, while the error is comparatively larger, since the amount of savings with the optimal solution is quite high (Fig. 5), the heuristic will still result in significant energy savings.

V. CONCLUSION

BTSs account for most of a cellular network's energy consumption. Motivated by the non load-proportionality of BTS energy consumption, prior work proposed shutting down some BTSs when traffic is low. However, network operators are reluctant to do so for a variety of reasons.

To reduce energy consumption, operators often use a feature called BTS power savings that deactivates some TRXs at BTSs that have low traffic. This is typically done without any *intelligence*. Using real network topology and traffic traces in a simulation study, we found that merely using BTS power saving in an urban setting can result in considerable energy savings.

We formulate an optimization problem that intelligently re-routes active calls to place a maximal number of BTSs in power-saving mode, thereby minimizing the electricity consumption. When re-routing active calls, our optimization keeps the MS uplink budget in check and does not drop any active calls either. Since our formulation is NP-Hard, we also propose a heuristic. Experiments using topologies and datasets from a live GSM operator indicate that Low-Carb can significantly reduce BTS power consumption and that our heuristic's performance is quite comparable to the optimal solution.

ACKNOWLEDGEMENTS

The authors wish to acknowledge Zartash Afzal Uzmi, Ihsan Ayub Qazi, Fahad Rafique Dogar, Omer Butt, Ammar Qadir, Anum Ajmal and Adnan Ahmed for their invaluable feedback during this research project.

REFERENCES

- [1] E. Oh, B. Krishnamachari, X. Liu, and Z. Niu, "Toward dynamic energy-efficient operation of cellular network infrastructure," in *IEEE Communications Magazine*, June 2011.
- [2] J. T. Louhi, "Energy efficiency of modern cellular base stations," in *IEEE INTELEC '07*. New York, NY, USA: IEEE, 2007, pp. 475–476.
- [3] C. Peng, S.-B. Lee, S. Lu, H. Luo, and H. Li, "Traffic-driven power saving in operational 3g cellular networks," in *ACM MobiCom '11*, 2011, pp. 121–132.
- [4] H. He, J. Sun, Y. Wang, Q. Liu, and J. Yuan, "Cluster analysis based switching-off scheme of base stations for energy saving," *Journal of Networks*, vol. 7, no. 3, pp. 494–501, 2012.
- [5] J. Lorincz, T. Garma, and G. Petrovic, "Measurements and modelling of base station power consumption under real traffic loads," *Sensors*, vol. 12, no. 4, pp. 4281–4310, 2012.
- [6] N. S. Networks, "Flexi BSC (BSC14)," <http://bit.ly/QMD8O5>, June 2011.
- [7] D. T. A. Rezgui, Krishnamurthy, and P. Pacharintanaku, "Dimming cellular networks," in *IEEE GLOBECOM'10*, December 2010.
- [8] O. Blume, "Energy savings in mobile networks based on adaptation to traffic statistics," *Bell Labs Technical Journal*, vol. 15, no. 2, pp. 77–94, September 2010.
- [9] M. Gerla and J. T.-C. Tsai, "Multicenter, mobile, multimedia radio network," *Wireless Networks*, vol. 1, no. 3, pp. 255–265, Aug. 1995.
- [10] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, "Power awareness in network design and routing," in *In Proc. IEEE INFOCOM*, 2008.
- [11] S. Mbakwe, M. T. Iqbal, and A. Hsaio, "Design of a 1.5kw hybrid wind/photovoltaic power system for a telecoms base station in remote location of benin city nigeria," in *IEEE NECEC*, November 2011.
- [12] Ericsson, "Vodafone Germany first to launch Ericsson's power-saving feature to reduce energy consumption and cut CO₂ emissions," <http://bit.ly/SIs1v1>, December 2007.
- [13] Mobilink, "Mobilink Network," <http://bit.ly/UWXYzB>, September 2012.
- [14] Pakistan Telecommunication Authority, "Pakistan Telecommunication Authority, 2011 Annual Report," <http://bit.ly/S8YtWS>, December 2011.
- [15] Wikipedia, "Telenor Pakistan - Wikipedia," <http://bit.ly/R1Uyqm>, September 2012.