

Expanding Rural Cellular Networks with Virtual Coverage

Kurtis Heimerl
kheimerl@cs.berkeley.edu
UC Berkeley

Kashif Ali
kashif@cs.berkeley.edu
UC Berkeley

Joshua Blumenstock
joshblum@uw.edu
University of Washington

Brian Gawalt
bgawalt@eecs.berkeley.edu
UC Berkeley

Eric Brewer
brewer@cs.berkeley.edu
UC Berkeley

Abstract

The cellular system is the world's largest network, providing service to over five billion people. Operators of these networks face fundamental trade-offs in coverage, capacity and operating power. These trade-offs, when coupled with the reality of infrastructure in poorer areas, mean that upwards of a billion people lack access to this fundamental service. Limited power infrastructure, in particular, hampers the economic viability of wide-area rural coverage.

In this work, we present an alternative system for implementing large-scale rural cellular networks. Rather than providing constant coverage, we instead provide *virtual coverage*: coverage that is only present when requested. Virtual coverage powers the network on-demand, which reduces overall power draw, lowers the cost of rural connectivity, and enables new markets.

We built a prototype cellular system utilizing virtual coverage by modifying a GSM base station and a set of Motorola phones to support making and receiving calls under virtual coverage. To support the billions of already-deployed devices, we also implemented a small radio capable of adding backwards-compatible support for virtual coverage to existing GSM handsets. We demonstrate a maximum of 84% power and cost savings from using virtual coverage. We also evaluated virtual coverage by simulating the potential power savings on real-world cellular networks in two representative developing counties: one in sub-Saharan Africa and one in South Asia. Simulating power use based on real-world call records obtained from local mobile operators, we find our system saves 21-34% of power draw at night, and 7-21% during the day. We expect even more savings in areas currently off the grid. These results demonstrate the feasibility of implementing such a system, particularly in areas with solar or otherwise-intermittent power sources.

1 Introduction

No recent technology has had a greater impact on economic development than mobile phones, which comprise the largest networks on Earth and cover over five billion subscribers [18]. Unfortunately, many people still lack this fundamental service. Although it is difficult to know the total number of potential users currently without coverage, it is likely more than a billion.

Nearly 95% of this uncovered population live in rural areas without grid power [12]. The primary reason for their lack of service is economic; operators are unwilling to make the large infrastructure investments (or pay the large operating costs) required to operate in areas without enough users to cover expenses. Contacts in rural areas have reported prices between five hundred thousand to one million USD for the installation of a cell tower in an area without existing power or network [25]. Our theory of change is simple: by reducing the cost of infrastructure we make mobile phones viable for new areas and users. With lower costs, we believe operators will naturally extend their reach into more rural areas, and new rural entrants become economically viable.

Recent advances have dramatically lowered the price of cellular equipment [23] and backhaul networks [26]. Unfortunately, power remains a fundamental cost in any deployment, dominating both the capital and operating costs of rural cellular networks. The International Telecommunications Union has indicated that 50% of the OPEX cost for a rural network is power [17]. Commercial equipment providers have attempted to address this fact [2, 28, 33], reducing the power consumption of small-scale cellular equipment to less than 90 watts. Unfortunately, further reducing power draw is impossible without causing service interruptions; the power amplifier quickly dominates total power used (65%-84% of total draw) and is directly tied to coverage radius and capacity. Each watt of power drawn is another watt needing generation (usually diesel [12]) and

storage; a severe limitation with expensive rural power.

One obvious way to save more power is to turn off portions of the network for a period of time (typically at night). This is common in rural and developing regions [17]. For instance, in the Punjab province of Pakistan, citizens were without grid power for over eighteen hours a day [1]. Unfortunately, this means that users cannot make important calls (such as emergency calls, a critical use case [15]) while the network is off.

Our solution, *virtual coverage*, resolves this concern. Virtual coverage also powers down individual cellular towers, but only when not in use, which we demonstrate to be a substantial fraction of the time. When network is needed, as signaled by a user initiating or receiving a call, the power is restored and the tower is available for communications. This design allows us to save a large amount of power in the largest networks on Earth while still providing consistent coverage at all times.

Specifically, individual cellular towers are powered-down (i.e., in “idle” mode) during periods of prolonged idleness. Although powering down a tower is conceptually simple, waking one is much harder; when a user wishes to make an outbound call, they are able to wake the network by sending a burst from either a modified GSM phone (our Wake-up Phone) or via a small, low-cost, push-button transmitter (our Wake-up Radio, functionally similar to a garage-door opener). Mobile-terminated calls require no changes to user behavior; the tower simply turns on and holds the call until the requested user connects. This solution also benefits from economies of scale by leveraging existing handset and radio equipment.

We implemented this design by modifying a Range Networks 5150 2G GSM base transceiver station (BTS). We first show that “idle” mode saves between 65% to 84% of the power on a BTS, depending on the coverage and capacity required. We also demonstrate that the user’s experience is not dramatically affected, with the setup of all calls increasing by only two seconds for the Wake-up Phone and an average of at most 25 seconds in standard, unmodified phones using the Wake-up Radio. We show that an installation using virtual coverage in a low-density area could operate with less than one-sixth of the solar panels, batteries, and price of a traditional setup. Lastly, we demonstrate that a virtual coverage network’s power requirements scale sub-linearly with the total number of calls. This allows smaller operators to invest in their network as it grows, rather than with one large capital expenditure, reducing their risk.

In addition, we simulated the use of our technology in two existing developing world cellular networks. We gathered a week of tower-level call activity from one country in sub-Saharan Africa (roughly 15 million calls) and one country in South Asia (roughly 35 million

calls). Using these records, we calculated the exact amount of “idle time” (time where no calls were active) in each network and combined them with estimates of the power draw for each tower. We show that, by utilizing virtual coverage, we could reduce the network power draw by 34% at night (21% during the day) for our South Asian operator. In sub-Saharan Africa, where towers are more heavily utilized, we are able to save 21% of the power at night (and just 7% during the day). Although this simulation uses only existing networks, we expect calling patterns in currently unserved areas to have more available idle time and thus, power savings. This reduction in operating power would dramatically reduce the operating cost of an off-grid, renewable powered, rural wide-area cellular network, enabling cheaper, greener telecommunications for people currently without network connectivity.

Summary of contributions:

- The concept of *Virtual Coverage*: on demand wide-area cellular networking;
- A working cellular base station implementing virtual coverage using OpenBTS;
- The implementation of a handset capable of waking a virtual coverage enabled BTS;
- The design and implementation of a low-cost wake-up radio capable of waking the network, for supporting existing handsets;
- A technical evaluation of virtual coverage, showing “idle” power savings, user experience impact, and sub-linear growth in power needs as a function of total calls; and
- An *in situ* evaluation of the potential power savings based on trace data from all the cellular towers operated by both an sub-Saharan African and South Asian telecommunication firm.

2 Design

Our system is targeted at a specific set of users: those in rural areas too sparse or poor to support traditional cellular infrastructure. This encompasses more than a billion people, in both developed and developing areas. We start with use cases that inform our design.

Alaska Large swaths of rural America live in areas without network coverage. Often these are agricultural or mountainous areas where the population density is too low or coverage too difficult to warrant the (expensive) deployment. One example is Hatcher Pass in South Central Alaska. Though surrounded by large towns (Palmer, Wasilla, and Houston), the mountains are too far from these population centers for network coverage,

especially in the valleys that define the range.

Our system would enable lightweight, autonomous, solar-powered cellular equipment that can be easily deployed on a central mountain top, providing “on demand” network coverage for travelers in need of communication. As communications in these areas will be rare, we want to enable wide-area coverage, amortizing the cost of deploying and maintaining the cellular equipment.

Uganda Another primary use case is in the developing world, where potentially billions of wireless consumers live in areas without coverage. This is almost always in rural areas, as the high population density of cities (despite the low per capita earnings) ensures that wireless providers will see a return on their investment. Similarly, even though the population in rural areas may be higher in the developing world than in the rural U.S., their lower income discourages investment.

One potential target for virtual coverage is rural Uganda. One example is city of Mpigi, an hour from the capital of Kampala. As a major trading hub, multiple carriers provide coverage in the center of town. However, in the hills less than 20 kilometers away, there’s no coverage at all. Heimerl et al. [15] showed how populations in these areas manage to communicate with their limited cellular coverage. They found that users heavily valued emergency services, and were unwilling to accept a fully asynchronous telecommunication system for this reason. Thus we aim for low-cost but continuous operation, explicitly supporting emergency services.

2.1 Design Goals

With the above use cases in mind, we have generated a set of design goals:

- Enable solar-powered cellular infrastructure capable of wide-area coverage;
- Provide continuous network availability (with a small start-up delay) to support emergency communications;
- Reduce infrastructure cost, enabling coverage for areas currently too poor or sparse for traditional cellular; and
- Utilize existing economies of scale by building off of existing GSM handsets and base stations and minimizing hardware changes.

We also note that our design explicitly does not support mobile handoff, nor does it ensure that wake-up requests are made in good faith. We address these concerns in Section 7.

3 Background

Cellular telephony is an enormous field, with multiple standards deployed across nearly every nation. In this section, we detail the specific wireless standards and hardware most suited to virtual coverage.

3.1 Cellular Telephony

The 2G (GSM) standard was officially launched in Finland in 1991. Subsequent 3G (UMTS) and 4G (LTE) standards first appeared in 2001 and 2009, respectively. As the standards progressed, the effective bit-rate for channels increased, primarily to support high-bandwidth data services such as streaming video.

Unfortunately, these superior encodings are more sensitive to errors and loss [29], which limits their propagation and usefulness in rural areas. With this in mind, most 3G/4G deployments are smaller in scale, primarily targeting dense urban areas. The 2G GSM standard, especially in the lower 900-MHz band, delivers the most consistent network propagation.

For these reasons, the 2G standard is still present in almost all commercial cellular hardware. Similarly, it is also present in almost all cellular handsets, with many brand new handsets in the developing world supporting only the 2G standard.

There have been recent advances in open-source GSM telephony, specifically the OpenBTS [23] project. This is an open implementation of the 2G GSM standard. As such, this is the wireless technology we use to prototype virtual coverage. However, there is no fundamental difficulty applying Virtual Coverage to 3G/4G networks.

3.2 Cellular Base-Station Hardware

In a modern software-defined GSM base station there are three core pieces of hardware that draw power, while the rest is passive. These pieces are the computer, the radio, and the power amplifier. Figure 1 shows our base station.

Unlike many other wireless systems (like 802.11), GSM coverage range is inherently limited by the *uplink* power (phone to BTS) and not the tower’s transmit power. The handset must be able to reach the tower, and increasing our broadcast power does not make that any easier. Though there is equipment on the BTS that improves this slightly, any transmit amplifications over 10W will not improve the range of the system; the GSM standard sets the maximum handset power to be 2W.

However, past the 10W limit, increasing the transmit power does allow for *more* communications at the same range. The extra power can amplify other channels, increasing the totally capacity of the tower. Table 1 provides an example.

	Range (km)	Capacity (Calls)
2W Tower	7	7
10W Tower	35	7
50W Tower	35	35

Table 1: Range Networks [28] cellular tower propagation and capacity specifications.

In any such wide-area setup (greater than 1 kilometer), the amplifier will dominate the power consumed by the unit. A 10W amplifier draws 45W of power (65% of the total) in a low-capacity BTS. A 50W amplifier draws 130W (84% of the total) for a high-capacity, 35 concurrent call, BTS. This amplifier operates continuously, amplifying the beacon channel.

Because of these properties, any attempt to save power in a wide-area cellular network must change the behavior of the power amplifier [3]. Unfortunately, this will also affect the user experience; amplification is the mechanism by which the tower broadcasts long distances and increases capacity. This is the core problem we address with virtual coverage: meaningfully covering a sparse population is currently energy-intensive.

4 System Implementation

Enabling virtual coverage requires a holistic rethinking of the base station itself. First, the BTS must be modified to enable programmatic control of the power amplifier. This will allow us to enter an “idle” mode in which the power amplifier is turned off. As a byproduct of this, the network is unavailable during this period.

Second, we must implement a mechanism for allowing users to wake the cellular tower remotely and promptly, thus enabling coverage on demand. We implemented two models of virtual coverage wake up: 1) implementing software-only handset modifications to send special wake-up bursts, and 2) developing a custom autonomous low-cost radio that sends the same message, allowing the system to work with existing, unmodified handsets. After detection of this burst, the network exits the idle state and resumes normal operation.

4.1 Enabling Low-Power Modes in Cellular Infrastructure

Virtual coverage requires the base station to have a low-power mode when the network is not in use. There are two core changes needed create a low-power mode for a GSM base station. First, the hardware must be modified to provide a mechanism for programmatic control of the power amplifier, the primary power draw. Second, the software must actually cease broadcasting during idle times while still listening to detect wake-up bursts.




Figure 1: Our Range Networks GSM BTS.

Hardware Figure 1 shows the internals of our revised Range Networks 5150 cellular base station. The key pieces of equipment are the radio, computer, duplexer, and power amplifier (PA). We added a USB-controlled high-current switch and connected it directly to the power amplifier, allowing us to control the PA’s status via serial commands from the computer. When the BTS enters idle mode, the PA is turned off.

Software There are two key software modifications. First, we implement the idle mode and drop all transmissions (including the beacon) while the power amplifier is off. Second, we implement a mechanism for the BTS to receive wake-up signals from user radios.

We implemented idle mode with a service that sends messages to the switch controlling the power amplifier. This daemon, which has access to the GSM and switch state, controls entering and exiting idle mode. Instead of naively returning to idle when all calls have terminated, we use a number of heuristics to improve the user experience. First, we require that the network be active for a minimum of 90 seconds, approximately double what we found to be the worst-case time necessary for a handset to connect to and communicate with the tower (i.e., *camping*) (Table 3). This ensures that all handsets waiting to camp will have ample time to do so should a tower return from idle mode. Second, the BTS only transitions to idle if there has been no cellular traffic for 30 seconds. This enables serialized actions, e.g., redialing a dropped call.

Originally, we had hoped to provide a “low coverage mode” (i.e., signal transmission without amplification),

where the BTS could still provide coverage for people physically near the radio. In testing, however, we discovered that the radio *must* be disabled when in idle, even to the exclusion of transmitting with the amplifier in pass-through mode. If the BTS broadcasts (even at low power), handsets nearby will attempt to camp (a process near-all handsets perform automatically and periodically). As our burst-detection is a simple power measurement, this legitimate network traffic would be indistinguishable from wake-up bursts.

The BTS provides two primary functions: incoming (mobile-terminated) and outgoing (mobile-originated) calls. Mobile-originated calls are simple. The tower must be in active mode, as only a camped handset can initiate a call. Mobile-terminated calls are more complicated. If the tower is “active” when it receives a call, the call is simply routed to the appropriate handset. If the tower is “idle”, the caller either leaves a voicemail (and the tower remains “idle”) or they are put on hold and the tower immediately wakes and waits for the handset to camp. According to our measurements (Table 3), this can be up to 40 seconds (with most being under 25 seconds). When the handset eventually camps, the callee is immediately connected to the caller by bridging to the held call or initiating a new call if they hung up.

The basic mechanism for detecting bursts is implemented in the transceiver of the radio. If the radio is in idle mode, any high enough power burst on the tower’s Absolute Radio Frequency Channel Number (ARFCN, basically just the frequency the tower listens and transmits on) will cause a message to be sent to the daemon, waking the system. The power required depends on the current noise level as determined by the transceiver. This technique is similar to ones used in sensor networks [13].

As OpenBTS utilizes voice-over-IP (VoIP) as it’s interconnect, there are no changes required to any other network services. Were we to interconnect using more traditional protocols (i.e., SS7/MAP) the name database (HLR) would have to allow longer registrations from users on virtual coverage enabled BTSs. This is the only change required for inter-operation.

With this system, we are able to provide on-demand voice services for rural networks at the cost of increased call-connection latency. SMS and data traffic are assumed to sync during periods of active voice traffic.

4.2 Waking Up in Virtual Coverage

Virtual coverage is not just a change in the cellular tower; it also requires a device capable of sending a “wake up” message. As mentioned in the previous section, we implemented two mechanisms for sending this message: from a handset or via an autonomous radio.

4.2.1 Cellular Handset

We have implemented our base station wake-up mechanism using an osmocomBB compatible mobile phone. We call this the *Wake-up Phone (WUP)*. OsmocomBB [24] is an open-source GSM baseband software implementation, which simplifies changes to the GSM protocol. However, every GSM handset should be able to send a wake-up burst with a software change from the manufacturer. The mechanism for waking up the BTS is sending a burst packet on the BTS’s ARFCN. The BTS, though not transmitting, receives this message and exits the idle state, allowing the handset to camp.

Each BTS broadcasts its ARFCN number (as well as the ARFCNs of similar nearby towers) on the beacon channel, which details the exact frequency used to communicate with the BTS. A handset periodically scans the network for towers to camp on, and gathers these numbers. In our system, the handset stores these numbers when the network idles, and then uses them to send “wake up” messages (as above) during periods without network availability.

Mobile-Originated (MO) Call In order to initiate a call, the Wake-up Phone will transmit “wake up” bursts on a selected set of ARFCNs. These ARFCNs are either a list of previously detected base stations or a static configured list. The “wake up” packets are random packets that are transmitted on the selected ARFCN. After transmission, handset scans for a tower broadcasting on the ARFCN just awoken, instead of scanning the whole cellular band (as in normal cell selection). If discovered, the handset camps to this tower and the user is able to communicate.

If a WUP is unsuccessful in camping to the recently awakened base station, the handset will proceed to the next ARFCN in its list, if any, and perform similar operations. This mechanism repeats until the handset is successfully camped or it runs out of available ARFCNs. At this point it will default back to the standard GSM protocol, which scans the entire band looking for available towers.

Mobile-Terminated (MT) Call As stated above, when the BTS receives a mobile-terminated call it immediately exits idle mode and waits for the handset to camp. The WUP scans the stored ARFCNs much more frequently (10:1), reducing the average time to camp. However, this does not affect the worst case analysis, which is 7s. When found, the phone camps and the call is connected.

4.2.2 Wake-up Radio

We have also designed and implemented a system to wake-up our BTS, the *Wake-up Radio (WUR)*,




Figure 2: Prototype implementation of WUR.

Component	Draw (W)	% (10W)	% (50W)
Computer	12W	17.4%	7.8%
Radio	12W	17.4%	7.8%
10W Amp	45W	65.2%	
50W Amp	130W		84.4%

Table 2: Power draw of the components of our Range Networks 5150 BTS.

nicknamed the *garage door opener*. The WUR transmits wake-up bursts, similar to our modified GSM handset, on a specific ARFCN.

The radio is designed to be both cheap and low power. The primary user interface is just a single button. When pressed, this button triggers a burst on the configured ARFCN. The radio produces a signal at approximately 500mW, the minimum power required for a handset in the GSM standard. The WUR uses an on-board battery pack, but provides interfaces for other power sources (e.g., solar) as well. The WUR, with two AAA batteries is capable of 5000+ bursts. The WUR can be configured to produce different ARFCNs with dip switches.

The WUR is only needed for mobile-originated communications. For mobile-terminated communications, the tower simply wakes and waits for the recipient's handset to camp.

5 Technical Evaluation

5.1 BTS Power Savings

We begin by evaluating the performance of a single modified Range Networks 5150 BTS. This unit has two power amplifiers available: 10 Watt and 50 Watt. The 10W unit supports just one channel and seven concurrent calls. It is commonly used for low-density areas. The 50W unit is designed for denser areas, produces up to five channels, and is capable of providing 35 concurrent calls. Both towers cover up to 35 kilometers, depending on configuration and geography. Areas with buildings or dense foliage will have worse signal propagation.

Model	Time (Avg)	Time (Max)
WUP (2G) (MT)	2s	7s
WUP (2G) (MO)	2s	2s
HTC Dream (3G)	12.1s	41.8s
Samsung Nexus S (3G)	23.6s	37.6s
Nokia 1202 (2G)	10.8s	14.6s

Table 3: Measurements on how long a handset has to wait, on average, to camp to a specific tower. This is the additional connect time if the network is idle.

Table 2 shows the relative power draw for each component of the BTS. As expected, the power amplifier dominates the overall power draw, consuming 65% of the power for a 10W unit and 84% for a 50W.

In our system, we added a USB-controlled switch to programmatically control the power amplifier. This switch draws negligible power (less than 1W). We also saw no change in the power draw of the computer, as expected with the BTS handling no calls in “idle” mode. As such, we are able to reduce the overall power draw of our BTS by over 65%.

5.2 Handsets

Wake-up Radio With the wake-up radio (WUR), the user has access to a device tuned to the particular frequency of their local cellular tower. Depending on local cost and logistic constraints, this device can be either widely deployed as an attachment to each local user's individual cellular phone, or singly deployed at some central location as a “phone booth”. The user presses the button, sending the wake-up message to the tower, taking the station out of idle mode. When the tower wakes, it broadcasts a beacon signaling its location and ownership.

Traditional GSM handsets periodically scan the airwaves looking for beacons. As such, a user's handset will eventually camp to the newly awoken BTS. We measured the time to camp, after waiting to ensure the initial network search failed, for three different handsets: the Samsung Nexus S (Android), the HTC Dream (Android), and the Nokia 1202 (Symbian S30) over thirteen trials. The results are shown in Table 3. The first two phones are quad-band 3G phones, meaning that they scan a wider band than a dual-band 2G-only feature phone (e.g., our Nokia 1202) commonly used in developing regions.

Users must wait for their handsets to camp in order to communicate using the network. Our results mean that using the WUR increases the setup time for all calls by a maximum of approximately 40 seconds. The average wait measured is less than 25 seconds. Both the maximum and average time to camp are highly dependent on the specific phone used. Though this




Figure 3: Solar Power (a), Battery Power (b), Individual Batteries (c), and total spending (d) required to operate a virtual coverage tower at a certain amount of idleness for a week in an area with 5 hours of sun. Note that 0% idle is equivalent to a traditional tower.

potential wait is a nontrivial amount of time, we believe this is acceptable to rural users who have limited alternatives for communication.

Wake-up Phone In the embedded solution, the WUP is able to camp on the BTS almost immediately after sending the wake-up burst. Since the same hardware device both delivers the burst and camps to the tower, the timing of the two tasks can be optimized for minimal delay. We measured the amount of time needed to register with the BTS¹ and found that, in all cases, it took exactly two seconds from wake-up burst to “camped normally”. This is shown in Table 3. The practical impact of this result is that, with a virtual coverage network, every mobile-originated communication takes two seconds longer to set up when the BTS starts idle.

For mobile-terminated calls, the handset does not generate the wake-up burst and must instead listen for a cellular beacon. Fortunately, the handset still knows what cellular towers are in the area. Instead of periodically scanning the entire band (as in standard GSM), we can scan just the beacons that are present in the area. This again takes no more than two seconds. Other operating system functions occasionally delay this scan, causing a maximum wait time of seven seconds.

5.3 Deployment Example

To begin to understand the impact of virtual coverage in a real-world situation, we calculate the approximate amount of infrastructure (solar panels and batteries) required to support a 50W (drawing 155W) cellular station year-round using only solar power.

We first frame our “real-world situation”: Providing winter-time network coverage in the South Asian country profiled in our later evaluation. During this time period, the country receives 5 hours of usable sun [10], as solar panels deliver minimal power when the sun is on the horizon. Using this, we are able to calculate the amount of power optimally tiled solar panels generate.

¹GSM state A1_TRYING_RPLMN to C3_CAMPED_NORMALLY

We assume an operating temperature of 40 degrees Fahrenheit and 24V batteries. Batteries are priced at 442 USD for 200 Amp-Hours and solar panels are priced at 1.07 USD per Watt. Lastly, the BTS draws 155W at full power (3720 Watt-hours/day) and 25W at idle (600 Watt-hours/day). As there is often inclement weather, we calculate the requirements for powering the station over a week without any power generation. The results are shown in figure 3.

The actual impact of virtual coverage is large; a completely idle tower requires one sixth of the batteries, solar panels, and total infrastructure cost of a traditional tower. As expected, these variables scale linearly with increasing idleness. We later (Section 6.3) show that idleness scales sub-linearly with respect to total calls (and thus users), meaning that the price of infrastructure required to support a virtual coverage tower scales sub-linearly with the total number of calls handled. Contrast this with a traditional cellular tower that must install the same amount of solar panels and batteries regardless of the number of calls and users serviced.

We wish to note that these costs are not only monetary. A single traditional BTS wanting week-long backup requires a kilowatt of solar panels and seventeen deep-cycle batteries (each weighing 68 pounds!). This equipment will be hiked into rural areas, an enormous load. Compare that with a virtual coverage station in an 80% idle area. There, just 300W of solar panels and 6 batteries must go up the hill. Lastly, virtual coverage also allows for growth; as an area moves from 80% to 70% idleness, new batteries and panels can be installed.

6 Real-World Evaluation

In addition to the micro-benchmarks just presented, it is important to understand how our system would perform *in situ*. For although we’ve demonstrated that our modified base station saves power and provides a consistent user experience, it requires periods of idle time in individual base stations, and this important variable is not known.

To resolve this, we use traces from existing cellular networks to provide a “worst case” analysis of virtual coverage’s benefits to areas without coverage. Specifically, we run two simulations using log data obtained from two mobile telecommunications firms in developing countries. Our simulations calculate the amount of power that the operator would save if they implemented our system of virtual coverage, based on actual patterns of idleness. Although these networks service millions of users (and not the rural areas we target), we show that our technology has the potential to save substantial amounts of power by utilizing the available idle time.²

6.1 Data

To assist in evaluating our research question, we acquired data from two large mobile telecommunications operators, one in sub-Saharan Africa and one in South Asia. The data we utilize contains a detailed log of tower activity over the span of one week, including information on: the start time of each mobile-originated and mobile-terminated call, the duration of the call (in seconds), and the approximate location of each tower. Note that all cellular activity is *per-tower* and not *per-cell*. We do not know how many cells were located on each tower; a negative bias as our system is capable of idling per-cell.

Although we are contractually bound not to disclose the identity of the mobile operators whose data we analyze, some basic facts are relevant.

Sub-Saharan Africa (SSA) The first operator’s data comes from a sub-Saharan African nation. Towers in this country are heavily utilized, although overall mobile phone penetration during the week we analyze was roughly 20 percent. We observe 15 million calls from a representative random sample of approximately 150 unique towers. There are an average of 90000 calls per tower, with a median of near 70000.

South-Asia (SA) The second operator’s data comes from a large country in South Asia. Mobile penetration here was roughly 55 percent during the week we analyze, but average tower utilization was lower than SSA. We observe 35 million calls from a representative random sample of around 5000 unique towers. There are an average of 6000 calls per tower, with a median of 4000.

²It is worth noting that we do not advocate replacing existing telecommunications networks with our equipment. As stated above, we are designing for autonomous rural networks. However, we believe it is an instructive demonstration of the value of virtual coverage.

6.2 Analysis

Combining the call log data with our user experience extensions (Section 4.1), we determine the amount of time where the network can be put into an “idle” mode while still providing complete cellular service to all users. We separately compute results during daytime (6am-6pm) and nighttime (6pm-6am), as user behavior and power generation are different during these periods.

We begin by determining the relative amount of idleness in the cellular network by using the detailed logs to determine when users initiate actions on the network. Of course, the network cannot power down for every idle period; users must have time to camp on the network and they may wish to communicate multiple times without having to wake the tower up repeatedly. With this in mind, we instead model a realistic user experience of the network. In this model, we assume the tower will remain awake for some time after a logged call in case a new call is placed or received. In our network, the towers remain active for 30 seconds following any user-initiated action. Each tower must also stay available for a minimum of 90 seconds at a time; this ensures that any phones within virtual coverage range waiting to initiate an action have enough time to detect and connect to the network (noting that the prior section found a maximum of 42 seconds to camp when scanning in disconnected mode). Any idle periods seen in the logs greater than 24 hours are removed from consideration (rather than being “idle”), as they are almost certainly a power outage. We found 18 such periods in SSA (in 25000 hours of coverage) and 59 such periods in SA (in 940000 hours).

6.3 Results: Idleness

Figure 4 shows the relative amount of idleness vs total number of calls per base station in both sub-Saharan Africa and South Asia. We note that the amount of idle time in the network (rural or urban) scales *sub-linearly* with the total number of calls. Fitting a logarithmic trend line (occupied time vs number of calls) results in $y = -0.077\ln(x) + 1$ for SSA, and $y = -2.04\ln(x) + 2.8$ for SA. This result means that each new call causes less occupied network time than the last, on average. This makes sense; as the number of calls on a tower increases, the amount of overlapping network activity also increases. Overlapping calls have zero marginal cost (in terms of power), and so we see this sub-linear benefit. This result, combined with our previous result showing that the cost of power infrastructure scales linearly with the idle percentage (Section 5.3) means that virtual coverage allows infrastructure cost to scale sub-linearly with the total number of calls serviced, and presumably the number of users serviced. As a byproduct of this




Figure 4: The idle time for both South Asia (a) and sub-Saharan Africa (b).




Figure 5: The difference in day and night idle time for each of the cellular towers in SA (a) and SSA (b), sorted by difference.

result; rural entrepreneurs can invest in more power infrastructure as demand grows, rather than requiring a large capital expenditure during installation.

Figure 5 next compares the day and night idle time in the network. This comparison is done for each BTS in the study, and over five thousand towers are difficult to represent graphically. To resolve this, we instead graph the *difference* between the night and day idle time. For instance, a tower that is idle 25% at night and 5% during the day would be represented by a point at $(.25-.05)=.2$.

For sub-Saharan Africa, we see more idle time at night. Over 98% of the towers are more idle at night than during the day. This result is diminished in South Asia, with just under 89% of the towers being more idle at night than during the day. We believe this is primarily due to a differential pricing scheme used in SA to encourage nighttime communications.

These results suggest that we should target night idle periods for saving power. This also works well with solar, since night power requires storage and thus costs considerably more due to both lower efficiency (10–20% loss), and the ongoing cost of battery replacement.

Lastly, our data also demonstrates the enormous amount of idle time available in these networks at night. 86% of the towers in SA are over 20% idle at night, while 53% of those in SSA pass the same metric. There is a significant opportunity for virtual coverage to reduce power consumption expenditures in both networks.

6.4 Results: Power Savings

We begin by noting the power measurements in Section 5.1. These measurements indicate that high-




Figure 6: Comparison of the power saved by every BTS in SA (a) and SSA (b) by using virtual coverage.

	Power Draw	Savings %
SA Original	1483 kW	0%
SA Day	584 kW	21.3%
SA Night	488 kW	34.3%
SA Total	1071 kW	27.7%

Table 4: The final results of our network simulation in South Asia. There are always significant power savings.

capacity towers (greater than 7 concurrent calls) require larger amplifiers, drawing 155W at full power and 25W at idle (a savings of 84%). Smaller capacity towers (supporting less than 7 concurrent calls) draw 70W at full power and 25 at idle (a savings of 65%). We use these measurements, combined with the measures of the maximum number of concurrent calls observed on each tower, to estimate the power draw of each individual BTS on both networks.

Figure 6 shows the results of this calculation. Each tower’s original power draw (yellow) is compared directly against the same tower using virtual coverage (blue). There is significant power savings in each network. We now move to calculate the exact amount of power saved.

Tables 4 and 5 show the final results if we sum the power drawn (and saved) by all tower equipment observed. Using virtual coverage, we are able to reduce the total night power budget by 34% for our South Asian network, and 21% in sub-Saharan Africa. During the day, when solar power is more available, we found that the network power draw could be reduced by 21% in South Asia, and just 7% in sub-Saharan Africa.

We wish to remind the readers that this analysis is on two *existing* cellular networks with broad coverage

	Power Draw	Savings %
SSA Original	45.6 kW	0%
SSA Day	21 kW	7.2%
SSA Night	18 kW	20.7%
SSA Total	39 kW	13.9%

Table 5: The final results of our network simulation of the sub-Saharan Africa network. Significantly more power is saved at night.

on nation-wide scales. Our system is designed to provide a mechanism for covering parts of the world currently without coverage, and gathering call records from such a place is impossible. We hope this proxy measurement, demonstrating massive possible power savings in networks that are economically feasible is a suitable demonstration that there would be similar (likely better) savings in areas currently without network.

These power savings (34% SA, 21% SSA), when combined with our earlier technical evaluation showing limited impact on the user experience (adding an average of 25 seconds to each call) and sublinear scaling of power infrastructure cost with regards to total calls, demonstrate a compelling system. It is low-cost, low-power, and well-suited for wide-area cellular communications in off-grid areas dependent on renewable energy sources.

7 Discussion

Data/SMS Services An obvious critique of our work is that we do not actively support SMS or data services on the base stations. In particular, as presented we do not wake up base stations for those messages alone.

SMS as a protocol already incorporates delay-handling, and functions asynchronously with respect to sender and recipient. There exist opportunities to improve SMS user experience in a virtual coverage cell. For instance, the carrier could set a maximum message delay: if a tower receives a message for one of its handsets when idle, the tower must activate and deliver the message within one hour. This has the potential to limited wasted active network time and the carrier is free to establish this interval based on their own preferences.

If a user requests data service while the tower is idle, the BTS can handle this in a similar fashion to placing a phone call: the tower is sent a wake-up burst (via either phone handset or wake-up radio). If idle time is significantly reduced by data requests (e.g., by data-channel apps seeking updates from the web), users can be incentivized to turn off these features when negotiating service rates with the provider.

Mobility As stated in Section 2, we explicitly avoid the issue of mobility in this paper. Our equipment is designed to create “islands” of coverage, simplifying the architecture dramatically. It is assumed that a rural virtual coverage cell will not intersect any other covered region. However, as this work moves forward, we recognize that the issue of mobility should be addressed.

The GSM tower broadcasts not only its own ID, but also those of nearby towers. This helps in two ways. First, a handset can try waking up towers in succession to increase its chances of successfully waking a BTS, at the cost of extra delay. Second, during a call the

handset could try to wake-up nearby towers proactively, either due to low signal from the current tower or just in case. Once awake hand-off to neighboring towers works as usual. Finally, on higher-end phones, a GPS-indexed database could inform the handset of exactly what tower(s) to wake in a specific location.

Inter-operation with existing infrastructure The GSM specification assumes constant coverage; a by-product of a system designed for developed, urban areas with strong power infrastructure. Virtual coverage changes this, turning cellular towers into dynamic agents. The interaction of these networks is complex.

Our system handles this already: Modified handsets *always* connect and call through existing static systems if possible: we are only capable of waking a tower if we are not attached to any existing tower. This is done primarily to save power; the static tower is on (whether or not it fields an additional call), but we’d prefer to keep our local tower idle. Similarly, if the other tower is dynamic but active, we’d prefer to send two calls through the powered tower, rather than waking an idle tower.

Security The system, as designed, has no mechanism for authenticating users before they wake the BTS. This means that one dedicated attacker could launch a power-based denial of service (DOS) attack by constantly sending the wake-up bursts to the tower. For example, a dedicated user may DOS the tower to leverage an information asymmetry they have developed.

The two mechanisms for waking up the BTS have different trade-offs for preventing such DOS attacks. For the WUR it is impossible to identify the user sending the wake-up burst, as the device is totally separate from the phone. Instead, we could add a cost to the use of the physical radio; perhaps by placing it in a phone-booth-like structure and charging a fee. This way, DOSing the tower could be made prohibitively expensive.

For the modified handset, we can identify the users by changing the protocol slightly. Currently, users send the burst, camp, and then wait. They must choose separately to make a call. We can instead enforce that the user must make a call (or any communication) immediately after camping to the station. This would allow us to know who did the waking. Unfortunately, this is still susceptible to another, similar attack; a user could send the wake-up burst and immediately pull the battery. In this case, we could modify the BTS to note there has been no traffic and immediately re-enter the idle state, reducing the impact of the attack.

Lastly, it is possible to send identifying information, such as the IMSI (unique SIM ID), in the burst message. This would allow us to charge users for waking the tower and prevent DOS attacks. We leave this to future work.

8 Related Work

Virtual coverage utilizes ideas first developed for shorter-range wireless sensor networks to save power in wide-area cellular networks. It also makes use of the recent advances in open-source telephony to implement these power-saving techniques. Lastly, our focus is on the problem of rural connectivity.

8.1 Saving Power

As the world goes “green”, reducing power consumption has become a critical goal of system designers [4]. Researchers (and system implementers) have focused on many different mechanisms for saving power, including disabling specific pieces of hardware. Gabriel et al. [11] investigated how to save power by utilizing the “idle time” in data center networks. Zheng et al. [34] used asynchronous wake-up mechanisms to save power in ad-hoc wireless networks. Wake-on-Wireless-LAN [21, 30] enables behavior similar to wake-on-LAN [22] across wireless networks. We similarly save power by utilizing idle time and a wake-up mechanism in cellular networks. However, the potential power savings in rural GSM cellular networks far outweigh smaller wireless deployments, as the range and power consumption of cellular networks are orders of magnitude larger.

8.2 Sensor Networks

Efficient use of power is of critical importance in wireless sensor networks due to battery-based operation. The core technique we use to save power in cellular networks is very similar to one developed by Gu et al. [13] for wireless sensor networks. Other researchers have explored similar designs [7]. As in our system, the nodes sit idle until communication is needed, and wake each other up using large radio bursts that are distinguishable from noise through some mechanism. Similarly, other researchers also created standalone devices for creating wake-up signals [8]. This is unsurprising, as these technologies inspired our design. The key differences are in scale, intent, and mechanisms.

8.3 GSM/Cellular

The OpenBTS project [23] is a full-featured GSM base station using software-defined radios. It bridges GSM handsets to voice-over-IP systems, enabling cheaper, lower-power cellular equipment. The OsmocomBB project [24] is an open-source GSM handset, capable of interfacing with any 2G GSM station. Our research is built on these two pieces of technology that enable us to

modify the GSM standard and implement a low-power GSM telephony solution.

Lin et al. investigated multihop cellular networks [20] to broaden the range of GSM cellular towers. In their solution, each handset could potentially act as a router for other handset’s messages, directing them to a centralized base station. Others expanded this idea, investigating how one might properly incentivize users to share their cellular connections [19]. Our work utilizes a different mechanism, virtual coverage, to achieve the goal of increased effective network range. These two designs are not mutually exclusive. A solution using both could have a dramatic impact on rural telephony.

A few groups have explicitly investigated reducing power consumption in cellular networks. Bhaumik et al. [5] proposed varying the size of cells; saving power by turning a subset of the cellular towers off during times of low load. Peng et al. [27] took these ideas even farther, demonstrating significant over-provisioning in cellular networks. Unfortunately, their proposed method for saving power is not possible in most rural areas as there is often just one tower providing service.

8.4 Developing Regions/Rural Networks

Network connectivity in developing and rural areas is an area of active research [6]. Wireless telecommunications are a common idiom [31], as the cost of deploying these networks is significantly less than traditional wired networks in areas with limited infrastructure. Researchers have investigated long-distance wireless [26] and sustainability in these areas [32].

Seeing the need for rural cellular communications, commercial providers have begun to develop products optimized for these areas. Both Vanu [33] and Range Networks [28] have developed “low power” 2G cellular equipment capable of running off of entirely renewable energy and using only around 90 watts of power. Altobridge’s [2] lite-site product varies its capacity in order to save power. As all of these products provide constant network coverage (rather than virtual coverage), they cannot reduce their power draw below 90 watts.

The rural/urban divide is also an area of active work. Eagle et al. [9] investigated how calling patterns changed as users migrated between rural and urban areas. Heimerl et al. [14, 15, 16] researched how users in developing, developing, urban, and rural areas viewed and made use of their cellular infrastructure. He found that rural users had a better understanding of network properties, including coverage patterns. This work informs our decision to include users in network power provisioning, as we expect them to quickly understand and make use of the basic primitives provided by our system.



Figure 7: Installing the Village Base Station in rural Papua, Indonesia.

9 Future Work

We're currently deploying the Village Base Station [14], complete with virtual coverage, in rural Papua, Indonesia (Figure 7). This is a longitudinal study investigating not only the impact of virtual coverage, but also the general feasibility of off-grid, off-network, rural, community-owned and supported cellular installations. We hope to have this work completed by mid 2013.

10 Conclusion

The positive impact of mobile phones on the poor is one of the great ongoing success stories of development, with more users and more impact than any other advanced technology. The arrival of low-cost devices, and even low-cost smart phones, make mobile phones the best platform for current and future interventions, including mobile banking, education, health care, and governance. Yet much of the rural world lacks coverage due to the high cost of infrastructure in rural areas; reaching the next billion+ users will be much harder now that most urban areas have coverage.

In this work we presented virtual coverage, a mechanism for dramatically reducing the power draw of cellular infrastructure in rural areas. This is done by introducing an “idle” mode to the network, similar to work done in wireless sensor networks. Instead of providing constant coverage (wasted in times of low communication, such as at night), we provide coverage

only when needed. A user demonstrates their need to communicate with one of two mechanisms. First, we modified the baseband of a cellular phone to send a “wake up” message when a user wants coverage. Second, recognizing that modifying the baseband of the billions of cellular phones already deployed is likely infeasible, we developed a custom low-cost radio capable of producing the same signal. These changes utilize existing manufacturing economies, requiring just a small hardware modification to the BTS, a software change to the handset, and the manufacturing of a \$14 device.

We validated this design by implementing the system and demonstrating the power savings. We showed that, with proper use, our equipment saves between 65%-84% of the power at idle. We measured the impact on users, who would see an average of less than 25 seconds added to any call. Users of our custom firmware would see just two seconds delay. We showed that a virtual coverage installation could be built with one-sixth of the power infrastructure of a traditional tower. We demonstrated that the power requirements for a virtual coverage tower scale sub-linearly with the total number of calls (and presumably callers) serviced. This allows smaller operators to invest in their network as it grows, rather than having the entire expenditure be up front.

We also simulated both an sub-Saharan African and South Asian cellular carrier using our system. We found that we are able to save 34% of the night power (21% during the day) in South Asia. For the denser sub-Saharan African country, we can save 21% of the power at night and 7% during the day. This reduction in power consumption enables more use of solar power and makes cellular system more economically viable in rural areas far from grid power or network.

11 Code

All of our software and hardware designs are open source. Our modified versions of OpenBTS, OsmocomBB, and the schematics for the Wake-up Radio are in the following repositories:

- <https://github.com/kheimerl/openbts-vbts>
- <https://github.com/kheimerl/osmocom-bb-vbts>
- <https://github.com/kheimerl/VBTS>

All of our systems can be run on open hardware, primarily the Ettus USRP line of products.

12 Acknowledgments

Thanks to Tapan Parikh, Kelly Buchanan, Stephen Dawson-Haggerty, Andrew Krioukov, Nathan Eagle and Amy Wesolowski.

References

- [1] Unsheduled loadshedding irks people in punjab. *The Nation (Pakistan)*, October 2011.
- [2] Altobridge. <http://www.altobridge.com/>. Retrieved 8/2012.
- [3] G. M. Amdahl. Validity of the single processor approach to achieving large scale computing capabilities. In *Proceedings of the April 18-20, 1967, spring joint computer conference, AFIPS '67 (Spring)*, pages 483–485, New York, NY, USA, 1967. ACM.
- [4] L. A. Barroso and U. Hölzle. The case for energy-proportional computing. *Computer*, 40(12):33–37, Dec. 2007.
- [5] S. Bhaumik, G. Narlikar, S. Chattopadhyay, and S. Kanugovi. Breathe to stay cool: adjusting cell sizes to reduce energy consumption. In *Proceedings of the first ACM SIGCOMM workshop on Green networking, Green Networking '10*, pages 41–46, New York, NY, USA, 2010. ACM.
- [6] E. Brewer, M. Demmer, B. Du, M. Ho, M. Kam, S. Nedeveschi, J. Pal, R. Patra, S. Surana, and K. Fall. The case for technology in developing regions. *Computer*, 38(6):25–38, 2005.
- [7] I. Demirkol, C. Ersoy, and E. Onur. Wake-up receivers for wireless sensor networks: benefits and challenges. *Wireless Commun.*, 16(4):88–96, Aug. 2009.
- [8] B. V. d. Doorn, W. Kavelaars, and K. Langendoen. A prototype low-cost wakeup radio for the 868 mhz band. *Int. J. Sen. Netw.*, 5(1):22–32, Feb. 2009.
- [9] N. Eagle, Y.-A. de Montjoye, and L. M. A. Bettencourt. Community computing: Comparisons between rural and urban societies using mobile phone data. In *Proceedings of the 2009 International Conference on Computational Science and Engineering - Volume 04, CSE '09*, pages 144–150, Washington, DC, USA, 2009. IEEE Computer Society.
- [10] Global Energy Network Institute. <https://www.geni.org/>. Retrieved 8/2012.
- [11] S. Gabriel, C. Maciocco, and T.-Y. Tai. Long idle: Making idle networks quiet for platform energy-efficiency. In *Systems and Networks Communications (ICSNC), 2010 Fifth International Conference on*, pages 340–347, aug. 2010.
- [12] GSM Association. Powering Telecoms: East Africa Market Analysis Sizing the Potential for Green Telecoms in Kenya, Tanzania and Uganda. <http://www.gsma.com/mobilefordevelopment/powering-telecoms-east-africa-market-analysis>. Retrieved 2/2013.
- [13] L. Gu and J. A. Stankovic. Radio-triggered wake-up for wireless sensor networks. *Real-Time Syst.*, 29(2-3):157–182, Mar. 2005.
- [14] K. Heimerl and E. Brewer. The village base station. In *Proceedings of the 4th ACM Workshop on Networked Systems for Developing Regions, NSDR '10*, pages 14:1–14:2, New York, NY, USA, 2010. ACM.
- [15] K. Heimerl, R. Honicky, E. Brewer, and T. Parikh. Message phone: A user study and analysis of asynchronous messaging in rural uganda. In *SOSP Workshop on Networked Systems for Developing Regions (NSDR)*, 2009.
- [16] K. Heimerl and T. S. Parikh. How users understand cellular infrastructure. Technical Report UCB/EECS-2012-42, EECS Department, University of California, Berkeley, Apr 2012.
- [17] International Telecommunications Union. Green Solutions to Power Problems (Solar & Solar-Wind Hybrid Systems) for Telecom Infrastructure. www.itu-apt.org/gtas11/green-solutions.pdf. Retrieved 1/2013.
- [18] International Telecommunications Union. The World in 2010: ICT Facts and Figures. <http://www.itu.int/ITU-D/ict/material/FactsFigures2010.pdf>. Retrieved 2/2011.
- [19] M. Jakobsson, J.-P. Hubaux, and L. Buttny. A micro-payment scheme encouraging collaboration in multi-hop cellular networks. In W. Hunt and F. Somenzi, editors, *Computer Aided Verification*, volume 2742 of *Lecture Notes in Computer Science*, pages 15–33. Springer Berlin / Heidelberg, 2003.
- [20] Y.-D. Lin and Y.-C. Hsu. Multihop cellular: a new architecture for wireless communications. In *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies.*, volume 3, 2000.
- [21] N. Mishra, K. Chebrolu, B. Raman, and A. Pathak. Wake-on-wlan. In *Proceedings of the 15th international conference on World Wide Web*,

- WWW '06, pages 761–769, New York, NY, USA, 2006. ACM.
- [22] S. Nedeveschi, J. Chandrashekar, J. Liu, B. Nordman, S. Ratnasamy, and N. Taft. Skilled in the art of being idle: reducing energy waste in networked systems. In *Proceedings of the 6th USENIX symposium on Networked systems design and implementation*, NSDI'09, pages 381–394, Berkeley, CA, USA, 2009. USENIX Association.
- [23] OpenBTS. <http://openbts.sourceforge.net/>. Retrieved 4/2012.
- [24] OsmocomBB. <http://bb.osmocom.org/trac/>. Retrieved 4/2012.
- [25] PapuaCom, July 2012. Personal Communication.
- [26] R. Patra, S. Nedeveschi, S. Surana, A. Sheth, L. Subramanian, and E. Brewer. Wildnet: Design and implementation of high performance wifi based long distance networks. In *4th USENIX Symposium on Networked Systems Design and Implementation*, 2007.
- [27] C. Peng, S.-B. Lee, S. Lu, H. Luo, and H. Li. Traffic-driven power saving in operational 3g cellular networks. In *Proceedings of the 17th annual international conference on Mobile computing and networking*, MobiCom '11, pages 121–132, New York, NY, USA, 2011. ACM.
- [28] Range Networks. <http://www.rangenetworks.com/>. Retrieved 8/2012.
- [29] Range Networks, September 2011. Personal Communication.
- [30] E. Shih, P. Bahl, and M. J. Sinclair. Wake on wireless: an event driven energy saving strategy for battery operated devices. In *Proceedings of the 8th annual international conference on Mobile computing and networking*, MobiCom '02, pages 160–171, New York, NY, USA, 2002. ACM.
- [31] L. Subramanian, S. Surana, R. Patra, S. Nedeveschi, M. Ho, E. Brewer, and A. Sheth. Rethinking wireless in the developing world. In *In Proc. of the 5th Workshop on Hot Topics in Networks (HotNets)*, 2006.
- [32] S. Surana, R. Patra, S. Nedeveschi, and E. Brewer. Deploying a rural wireless telemedicine system: Experiences in sustainability. *Computer*, 41(6):48–56, 2008.
- [33] Vanu. <http://www.vanu.com/index.html>. Retrieved 8/2012.
- [34] R. Zheng, J. C. Hou, and L. Sha. Asynchronous wakeup for ad hoc networks. In *Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing*, MobiHoc '03, pages 35–45, New York, NY, USA, 2003. ACM.