

An Energy-Flow Model for Self-Powered Routers and its Application for Energy-Aware Routing

Veljko Pejovic
University of California
Santa Barbara
veljko@cs.ucsb.edu

Elizabeth Belding
University of California
Santa Barbara
ebelding@cs.ucsb.edu

Mahesh K. Marina
The University of Edinburgh,
UK
mmarina@inf.ed.ac.uk

ABSTRACT

Self-powered wireless mesh networks have gained popularity as a cheap alternative for providing Internet access in many rural areas of the developed and, especially, the developing world. The quality of service that these networks deliver is often bounded by such rudimentary issues as the unavailability of electrical energy. Dependence on renewable energy sources and variable power consumption make it difficult to predict the available energy and provide guarantees on communication performance. To facilitate energy trend estimation we develop an energy flow model that accounts for communication and energy harvesting equipment hardware specifications; high resolution, time-varying weather information; and the complex interaction among them. To show the model's practical benefits, we introduce an energy-aware routing protocol, Lifetime Pattern-based Routing (LPR), specifically tailored for self-powered wireless networks. LPR's routing decisions are based on energy level estimations provided by the energy flow model. Our protocol balances the available energy budget across all nodes; as a result, power failures are distributed among all participating parties. Using traces captured from a live network, we use simulation to show that LPR outperforms existing work in rural-area wireless network routing.

Categories and Subject Descriptors

C.2 [Computer-communication Networks]; C.4 [Performance of Systems]

General Terms

Algorithms, Design, Measurement, Performance, Reliability

Keywords

Energy efficiency, Rural wireless networks, Energy-flow model, Power-aware routing

1. INTRODUCTION

The map of the information world is highly polarized: only 5% of the African population uses the Internet (compared to 70% in

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NSDR '09 Big Sky, Montana USA

Copyright 2009 ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

North America) [17] and there are more telephones in the city of Montreal, Canada alone than in the whole country of Bangladesh [9]. These examples reveal only the tip of the huge problem of digital divide that can be observed on different levels in our society.

Most rural areas are out of reach of the state-of-the-art networking technologies as low population density and economic factors prevent their deployment. Wireless networks based on the IEEE 802.11a/b/g protocol have emerged as a viable alternative due to their low cost, license-free operation and ease of deployment. Unlike more traditional WiFi network settings, rural wireless deployments often have to face the lack of reliable grid power infrastructure [15], especially for wireless relay nodes that connect distant user clusters. In those situations, network devices are self-powered, i.e., powered by energy harvested from renewable sources such as wind and solar.

Unfortunately, renewable energy sources are inherently intermittent; therefore networks relying on them naturally exhibit intermittent connectivity, making them yet another example of a delay tolerant network (DTN). One way of coping with the problem of time-varying renewable energy supply is to provide a greater safety margin by installing excess energy harvesting hardware. However, such over-engineering significantly increases deployment costs with a powering system costing several orders of magnitude more than the communication equipment it powers. A better way of tackling this problem is to make judicious use of the limited and variable energy supplies via power-aware networking mechanisms, including routing.

Understanding the flow of energy is essential for enabling such intelligent power-aware mechanisms. But to the best of our knowledge, existing work does not model this aspect. Moreover, existing routing protocols suitable for rural wireless networks [2] do not take time-varying node energy budgets into account when making routing decisions.

In this paper, we examine the behavior of energy flow within self-powered routers in rural wireless networks and show that awareness of this behavior can aid in improved routing. Specifically, this work makes the following contributions:

1. We define an energy-flow model that continuously tracks the available energy budget within each self-powered router node (Section 3). It factors in time-varying local weather conditions (solar irradiation, wind speed), node energy consumption and complex interactions among multiple energy sources and the energy storage. This model can be used to predict future up and down times of a node if weather conditions can be forecast at the node.
2. With confidence in our model, we develop a power-aware routing protocol termed LPR (Section 4) that utilizes the model

to select paths that lead to more even consumption of node energy reserves and fewer network partitions, thus improving message delivery. Note that power-aware routing in the presence of self-powered routers is different from the traditional power-aware routing (e.g., [14]) since batteries are replenished in the former case. One could view the self-powered router context as a scenario where nodes and associated links fail temporarily while batteries powering the nodes are recharged by renewable energy sources. To evaluate the effectiveness of LPR (Section 5), we use fine-grained weather monitoring traces collected for the Tegola network [1].

2. NETWORK MODEL

WiFi-based rural area networks have been deployed worldwide from South Africa [6] and India [15] to UK [1]. Although relatively diverse, all the deployments share several common features. Each of these deployments is essentially a network of stationary mesh routers interconnected by long-distance point-to-point links (~ 10 km). A mesh router in these deployments also serves as an access point for a group of users in a location (e.g., village). Node degree is typically small for the mesh network topology. There is usually one mesh router that is connected to the wider Internet, thereby acting as an Internet gateway for the rest of the network. The means of energy supply for mesh routers is diverse with some nodes powered by grid power, others are self-powered using renewable energy sources. A self-powered mesh router consists of communication equipment (a routing board, wireless NICs, antennas, local access AP) and energy harvesting equipment (solar panel and/or wind turbine, rechargeable battery, charging regulator); the former is the energy consumer while the latter is the energy source.

Tegola [1] is a rural wireless mesh network deployed in the north-west of Scotland to serve as a testbed for research into technologies for low-cost, robust broadband wireless access in rural and remote areas. The area where the testbed is deployed is quite remote and known for its harsh weather and mountainous terrain, making it a challenging environment to deploy and operate a network. The Tegola network model closely matches the above generic model. Currently, it consists of five backhaul mesh router nodes (two of which are self-powered) and 17 rooftop nodes, connecting 35-40 households spread across six communities. One of the self-powered routers has been instrumented with a data logger and sensors for power monitoring. Specifically, local wind speed and solar irradiation is measured every minute using an anemometer and a solar radiation sensor attached to the data logger. Additionally, three pairs of voltage and current probes are used to individually measure the power output of the solar panel and wind turbine as well as the overall power consumption of the communication equipment. While the energy flow modeling and power aware routing presented in this work are more generally applicable, we will refer to the Tegola network given that we have immediate access to it.

3. ENERGY FLOW MODEL

Weather conditions are highly variable, so any model that relies on annual (or any other long term average) values is inherently inaccurate. For example, the wind energy harvested is proportional to the cube of the wind speed. Therefore two sites that have the same long-term average wind speed may differ significantly in terms of the actual wind energy harvested. Therefore, we model the system's behavior on a time scale as small as is realistic. The main restriction is the granularity of the weather sensor data.

Our approach to modeling energy flow in a self-powered router node closely follows the physical components of the system and

their types — battery type, solar panel and wind turbine specifications, available power consumption levels. In a self-powered system, the energy harvested by wind turbines and solar panels is converted to electrical energy and stored in a battery. Due to their robustness and low cost, lead-acid carbon-fiber batteries are commonly used for this purpose. Since batteries can store only a limited amount of energy, a charging regulator has to be installed between the battery and energy sources to prevent fatal overcharging of the battery. This necessarily translates to suboptimal harvesting as not all of the wind/solar energy generated can be kept. Moreover, charging regulators transfer the energy with a level of efficiency lower than 100%.

The energy is drained from the battery by the communication equipment. The overall power consumption is influenced by the wireless NIC power consumption, which in turn is dependent on the NIC state (e.g., transmit, receive, idle, sleep) and various other parameters (e.g., TxPower, bit-rate). This is especially true for multi-radio mesh routers, which are common in rural wireless deployments including our Tegola network. Our experience with high-power NICs specially designed for long distance links¹ shows that idle state power consumption is fairly high and that higher bit-rate communication surprisingly can be more energy efficient. Energy flow in a self-powered router over a Δt time period can be described by the following equation:

$$(1) E_{B,t+\Delta t} = E_{B,t} + k * REG_{t+\Delta t} * (E_{WG,t+\Delta t} + E_{PV,t+\Delta t}) - E_{C,t+\Delta t}$$

where $E_{B,t}$ is the battery energy at time t , while $E_{WG,t+\Delta t}$, $E_{PV,t+\Delta t}$ and $E_{C,t+\Delta t}$ are the energy generated by the wind turbine, generated by the solar panel and consumed by the communication equipment at time $t + \Delta t$, respectively. Regulation parameter $REG_{t+\Delta t}$ dictates the amount of energy that is transferred to the battery at that instant of time. Its value depends on the current battery voltage level, while k represents the charging efficiency. In the above equation, Δt is the update interval for battery energy level based on weather sensor measurements and total power consumption.

Table 1: Additional energy flow model equations.

(a)	$E_{WG,t+\Delta t} = P_{WG,t+\Delta t} * \Delta t$
(b)	$P_{WG,t+\Delta t} = f(ws_{t+\Delta t})$
(c)	$E_{PV,t+\Delta t} = P_{PV,t+\Delta t} * \Delta t$
(d)	$P_{PV,t+\Delta t} = g(sit_{t+\Delta t})$
(e)	$REG_{t+\Delta t} = \begin{cases} 0, & V_{B,t+\Delta t} > V_{REG,2} \\ 1, & V_{B,t+\Delta t} < V_{REG,1} \\ \frac{V_{REG,2} - V_{B,t+\Delta t}}{V_{REG,2} - V_{REG,1}}, & \text{otherwise} \end{cases}$
(f)	$E_{C,t+\Delta t} = P_{C,t+\Delta t} * \Delta t$

Δt – update interval
 $P_{WG,t+\Delta t}$ – power generated by the wind turbine at time $t + \Delta t$
 f – wind turbine power curve
 ws_t – wind speed [m/s] at time t
 $P_{PV,t+\Delta t}$ – power generated by the solar panel at time $t + \Delta t$
 g – solar panel specification
 sit_t – solar irradiation [W/m^2] at time t
 $V_{REG,1}, V_{REG,2}$ – charging regulator control voltages

Table 1 lists the additional equations used for the evaluation of expression (1). Harvested energy $E_{WG,t+\Delta t}$ and $E_{PV,t+\Delta t}$ are

¹Ubiquiti Networks XR5 Power Consumption Study, www.ubnt.net

calculated via the wind turbine² and solar panel data sheet³, respectively. Charging regulator characteristics and the battery voltage at time $t + \Delta t$ determine $REG_{t+\Delta t}$, a real number value between zero and one. In our setup, we use a pulse-width modulation (PWM) charging regulator that is suitable for both wind turbines and solar panels⁴. The actual behavior of such a regulator is hard to describe analytically, so we use a linear approximation as given by equation (e) in Table 1.

In the above equations we need to know the battery voltage in order to evaluate the charging regulation function. The battery voltage is directly connected to the state of charge (SOC) we are estimating with the equations in Table 1. Unless no current flows through the battery, the relationship is not a simple one as the internal resistance of the battery influences the voltage read at the battery ends. Under the same SOC, if the battery is being charged, the voltage is higher than if the battery is being recharged; the higher the charging/discharging current, the more observable the impact of the internal resistance. Moreover, the internal resistance is not constant: it changes with the outside temperature, battery age and the SOC itself [3]. Well aware of the problem complexity we decided to simplify the voltage - SOC relationship in our calculations by using the following equation:

$$(2) V_{B,t} = V_{B,80} + \frac{V_{B,0} - V_{B,80}}{\log(E_{B,0} - E_{B,80} + 1)} * \log(E_{B,t} - E_{B,80} + 1)$$

where $V_{B,t}$ is the battery voltage at time t ; $V_{B,0}$ is the battery voltage when fully charged; $V_{B,80}$ is the battery voltage at 80% depth of discharge (DoD); $E_{B,t}$ is the battery energy at time t ; $E_{B,0}$ is the battery energy when fully charged; and $E_{B,80}$ is the battery energy at 80% DoD. We based equation (2) on the manufacturer's data sheet of a new Elecsol⁵ 125Ah deep cycle carbon-fiber battery, the type used in the Tegola network. If a battery bank consists of multiple identical batteries connected in parallel, then equation (2) is still applicable except that $V_{B,t}$ and $E_{B,t}$ then correspond to the overall/aggregate voltage and energy of the battery bank, respectively. Note that currently we do not model the changes in the internal resistance of the battery, but we plan to extend this representation if needed.

The energy flow model captured by equations (1), (2) and Table 1 can be used to estimate the current battery energy given the total power consumption and current readings of solar irradiation and wind speed. If future weather information (solar irradiation and wind speed) can be forecast, then future node up/down times can be predicted using the above model assuming power consumption remains the same. The means of accurate weather forecasting, however, is beyond the scope of this work.

3.1 Model Evaluation

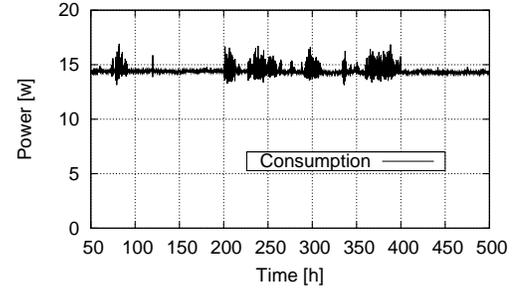
We evaluate the accuracy of the model by comparing the estimated battery voltage with the actual voltage recorded on-site at an instrumented self-powered node in the Tegola network. We provide our model with the measured wind speed and solar irradiation values. We assign an initial battery voltage (charge) and then iterate equation (1) every Δt seconds; the granularity of the available weather measurements restricts Δt to 60s. At time t , $E_{WG,t}$ and $E_{PV,t}$ are calculated from the wind speed and solar irradiation measurements and the hardware specification (wind turbine power

²Rutland 910-3 User Manual, www.marlec.co.uk

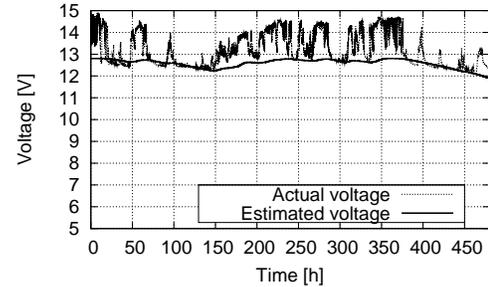
³KC130GHT-2 Data Sheet, www.kyocera.co.uk

⁴Note that there are other types of charging regulators, notably maximum power point tracking (MPPT) regulator specifically designed for use with solar panels.

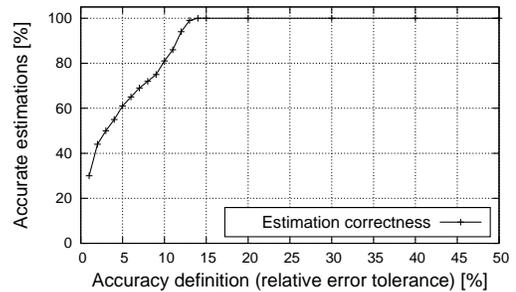
⁵Elecsol, www.elecsolbatteries.com



(a) Tegola network node power consumption



(b) Estimated vs. measured battery voltage



(c) Accuracy with respect to error tolerance

Figure 1: Accuracy of the energy flow model.

curve and solar panel characteristic). For simplicity, we use constant power consumption. This is justified by Figure 1(a) — except for occasional periods of instability, the consumption is largely constant (around 14.5W).

In Figure 1(b), we show the estimated battery voltage at the instrumented self-powered node in our testbed using equation (2) and the measured sensor readings (for solar irradiation and wind speed) over a two week period. Actual, measured voltage over that period is also shown for comparison. We see that the estimated voltage based on the model closely tracks the measured voltage. Discrepancies between the two can be explained by the fact that the measured voltage exhibits sharp rises as the battery is being charged due to the impact of internal resistance even though the actual battery charge does not change so drastically. We do not model the effects of the internal resistance in equation (2).

To quantify the model's performance we need to define what it means for the model to be accurate. As the battery voltage is a real number and small variations from the actual measured values do not play a big role in the end calculations, we define the accuracy as a relative error bracket around the measured value in which the predicted value should fall. For example, if we agree on a 20% error value we are considering the estimated value to be correct if the relative error is less than 20% (compared to the measured

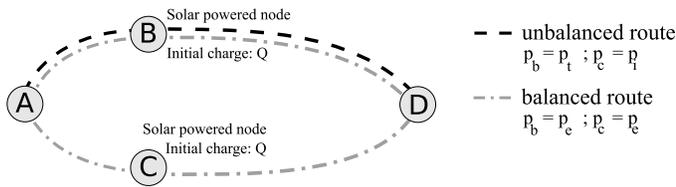


Figure 2: Impact of routing decisions on node availability and network connectivity.

value). Given a sample, we can calculate how many of the values in the sample accurately predict the voltage according to the above definition. Figure 1(c) plots the percentage of values in the sample (of size around 30,000 and corresponding to Figure 1(b)) that are accurate as the definition of accuracy gets looser. For a reasonable tolerance (11.5%), our model is accurate 90% of the time for the sample under consideration. Overall, the model underestimates the actual voltage and the mean modeling error falls in the interval $(-0.5998, -0.5906)V$ with 99% confidence.

Using the above energy flow model for a self-powered router allows higher layer protocols to be energy-aware. In the next section, we demonstrate the benefit of the model through a novel energy-aware routing protocol.

4. ENERGY-AWARE PATH SELECTION

In a rural wireless network with self-powered routers, node energy budgets are time-varying depending on the energy production and consumer patterns. Nevertheless, energy budget distribution within a network of self-powered devices is tightly connected to the routing decisions since the wireless interface activity impacts node energy budgets. As a result, energy-oblivious routing can lead to long periods of path unavailability and/or network partitions, which in turn causes message losses (for lack of a route and buffer space) and large message delays (as a result of having to wait in buffer until the network partition heals).

Consider the example scenario in Figure 2 to illustrate the poor routing decisions resulting from energy-oblivious routing. Suppose nodes B and C are each powered by a solar panel and also suppose that at 7pm each of them holds the same amount of battery charge Q . Node A sends a lengthy packet stream to node D. If the packets are sent via one of the nodes (without loss of generality we choose B), that node's consumption rate is equal to p_t while the other node enjoys idle energy consumption rate p_i . The rate p_t is too high to keep node B up and running until 7am when the sun rises. Instead, B is fully depleted at 4am. If, however, the consumption rate is balanced so that half of the time node B relays A's traffic and half of the time node C relays it ($p_e = \frac{p_i + p_t}{2}$), then both nodes would be up long enough to see the sun and would remain available. Obviously, a wrong routing decision can lower the connectivity of the network. Suppose that node B has data to send around 6am. In that case, the request could not be served until at least an hour later. In this case, a wrong routing decision impacts packet delivery and delay as well. While the routing decision is fairly straightforward in this example, it is not with general topologies and traffic patterns, especially when multiple different types of renewable sources are used (e.g., solar and wind).

4.1 Lifetime Pattern based Routing (LPR)

In our solution, we implicitly control depletion of node energy reserves through routing decisions, and improve the distribution of available energy reserves amongst network nodes. The goal is to avoid (even temporary) network partitioning and node failures so that reliable paths are available in the network most of the time to

successfully route messages without delays. Given the relationship of this scenario to the DTN setting, if a path is unavailable to route a message then we would like to be able to choose a route that becomes available the soonest.

Our protocol, Lifetime Pattern based Routing (LPR), is based on a routing metric that utilizes the energy flow model along with node lifetime predictions described in Section 3. LPR is constructed as a link-state routing protocol. There are multiple reasons why this type of a protocol is the most suitable for long-distance rural area wireless networks; a thorough discussion can be found in [2]. Nodes determine their future up and down times using the energy model, average node power consumption and future weather information. We do not define how far in the future the nodes should be able to predict the weather (we use 24 hour predictions in our evaluation), nor provide the means of getting the accurate prediction. Periodically, each node propagates information about expected future up/down times to its neighbors. This information, combined with the knowledge of its expected future up/down times, allows each node to determine the failure patterns of each of its incident links⁶. If both nodes on either end of a link are up, the weight of the link is proportional to the time left until the first end node failure occurs. If, on the other hand, the link is down, meaning at least one of its end nodes is down, the link weight is inversely proportional to the time left until both end nodes are predicted to be up.

Periodically, each node initiates a network-wide dissemination of its link weights in an epidemic manner as in [7]. Using link weights from all nodes in the network, every node computes routes to every other node by applying Dijkstra's algorithm and using the following path selection criterion: among the possible set of routes to a destination, select the one for which the minimum link weight on the path is the largest.

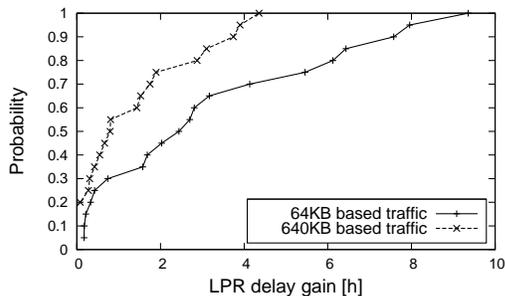
The proposed weighting function has two consequences: 1) the routing algorithm avoids selecting paths that are likely to fail in the near future, balancing the energy reserves equally among the network nodes; 2) the routing algorithm reduces the delay by selecting the routes that are about to be available soon. In a delay tolerant setting a path may consist of some links that are currently up and others that are currently down. In order to compare weights of links with different statuses, we bring links that are up and those that are down to the same scale. Empirically, we arrive at the conclusion that those links that are up and are predicted to go down in one hour should be of the same weight as those links that are down but expected to be up in one hour. The selected point gives a good balance that shows both properties, 1) and 2), of the algorithm. Changing the balance point results in an algorithm that favors one of the above mentioned properties.

5. EVALUATION

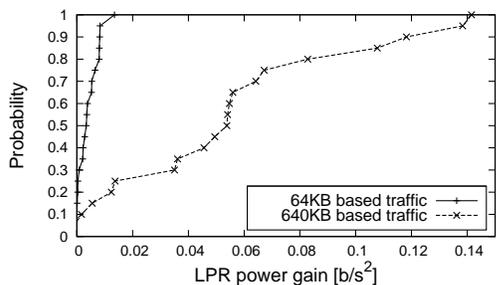
In this section, our goal is to evaluate the routing effectiveness of LPR in rural wireless network scenarios with self-powered routers. We consider three performance metrics for this evaluation: (i) percentage of messages delivered; (ii) end-to-end message delay (average and cdf); and (iii) *power* metric as defined in [11] as a ratio between the throughput and the average delay — higher the better; this metric captures the combined effect of improving delivery ratio and delay.

We compare LPR performance with DTLSR [2], the work that is most similar to ours. The authors specifically target rural-area networks, propose and justify a link-state protocol and select MEED [7]

⁶In this scenario, node locations are usually fixed but nodes (and consequently their incident links) fail temporarily for some periods due to lack of sufficient power.



(a) Delivery delay gain.



(b) Power gain.

Figure 3: LPR gain with varying traffic load.

as the most appropriate metric. Although targeting the deployments where the most probable cause of link/node failure is the lack of sufficient power, the authors do not take into account energy availability behavior. Therefore, we not only compare our work to the state of the art, but also quantify the benefits from power-aware routing in rural wireless networks with self-powered routers, something not previously done.

We briefly elaborate on the MEED metric used in DTLSR [2]. The Minimum Estimated Expected Delay (MEED) metric is a practical version of Minimum Expected Delay (MED), a metric that calculates the expected delay based on the future contact patterns. Instead of the known future contact schedule, MEED uses the observed contact history to estimate the expected delay. This assumes that the future connectivity will be similar to the previously observed connectivity. This makes MEED an attractive option for networks with scheduled power outages and networks that rely on a periodic bus service [4]. We expect MEED to perform reasonably well in a self-powered network scenario as the periodicity is present in at least one environmental factor that determines nodes' up/down patterns — solar irradiation.

5.1 Simulation Setup

We use trace-based simulation as the evaluation methodology. Specifically, we use DTNSim2 [16], which is a Java-based simulator developed for delay-tolerant network (DTN) evaluations. We extended the DTNSim2 with an energy layer which in essence is equation (1). A unique energy layer is instantiated for each node; the energy layer keeps track of a node's available battery capacity and updates this value according to the weather information and communication equipment activity. We set the baseline node power consumption to $16W$. The Tegola network nodes composing of a wireless routing board, local wireless AP, data logger and four NICs consume about the same amount of power. We add a further $2.5W$ ($0.5W$) for each transmitting (receiving) interface.

For our evaluation, we generate twenty different 25-node network topologies that resemble a typical rural area deployment [1]

Table 2: Average packet delay, loss, and power. Note: standard deviation is shown in parentheses.

64KB based traffic load			
	Delay [h]	Delivery ratio [%]	Power [b/s^2]
MEED	18.18 (18.61)	86.01 (8.92)	0.016 (0.005)
LPR	14.91 (15.31)	88.14 (9.19)	0.021 (0.006)
640KB based traffic load			
	Delay [h]	Delivery ratio [%]	Power [b/s^2]
MEED	9.13 (9.39)	51.2 (10.1)	0.198 (0.068)
LPR	7.76 (8.01)	55.9 (9.7)	0.255 (0.077)

[6] [15]: node connection degree is low, one node is elected as a gateway and all communication is between nodes and a gateway. Each of these topologies constitutes an independent simulation run. Before each run, we randomly assign different initial battery charges to the nodes. We use a large trace of weather sensor (wind speed and solar irradiation) data obtained from the instrumented self-powered node in the Tegola network. We slice the trace into four single-day subtraces and assign a subtrace to each of the nodes in the simulated network. The node up and down times are dictated by the specific subtrace assigned to it. Each of these subtraces are long enough for several uptime/downtime transitions to occur during the course of the simulation. At this point we are not considering means of predicting solar irradiation and wind speed; instead, we provide an oracle that gives the exact local weather forecast for the next twenty four hours to each of the nodes.

To exclude the influence of network congestion, all the links have their bandwidth set to $54Mbps$ (the maximum for IEEE 802.11a/b/g networks) and a propagation delay of $1ms$ (an expected value for $20km$ long links). All nodes have identical hardware properties and the same storage buffer sizes of $5MB$.

5.2 Results

In this section, we first present two sets of results corresponding to low and medium traffic intensity, respectively. In the first set, the gateway sends each node $64KB$ of data in $1KB$ chunks (one chunk per second) every hour. This traffic profile is similar to the one used in [2], except that data is split into smaller packets to make it realistic for transmission on 802.11 links and emphasize the influence of changing power consumption. We can expect this low traffic intensity if, for example, the network is used for environmental monitoring. Results are shown in Figures 3(a) and Figure 3(b). It is clear from Table 2 (upper part) that there is only a small improvement with LPR in the percentage of messages delivered compared to MEED but a significant improvement in terms of delay (more than 3 hours). Together this results in an improvement of about 25% in terms of the power metric. Note that with buffering of messages that cannot be routed immediately, the performance differences between protocols would largely manifest in terms of delay differences, which is what we observe from the results. Figure 3(a) shows the cdf of the delay gain with LPR over MEED obtained across all twenty runs. This plot shows that LPR improves the delay by over an hour in more than 70% of the cases compared to MEED. This is mainly because of the fact that MEED naïvely determines the future up/down times based on past patterns without explicitly modeling the energy flow behavior. Moreover diversity of renewable energy sources make it difficult for MEED to make good routing decisions. Figure 3(b) shows the cdf of the gain with LPR over MEED in terms of the power metric. This metric also shows improvement, however, due to small amount of data being sent the difference is not so pronounced.

We now consider the impact of higher traffic intensity on the effectiveness of LPR. $640KB$ of data is now transferred every hour

between the gateway and each of the nodes in the network. This amount of traffic can be observed in a delay tolerant network used for email communication. Figures 3(a) and 3(b) and Table 2 (lower part) show the performance impact. We see a different impact where there is a noticeable improvement in delivery ratio and slightly reduced gain in delay with LPR compared to MEED. This can be explained by the fact that greater traffic intensity leads to greater buffer overflows for a given buffer size, so the packets that would otherwise be delayed for long periods waiting in the buffers are dropped. This explains the different observed effect on delivery ratio and delay relative to the previous case. The net impact on the power metric is greater compared to the previous low traffic intensity case, as there is much more data impacted by the improvements provided by LPR. Like before, we show plots of the delay and power gains with LPR over MEED across all simulation runs as cdfs. Significant benefit from LPR's power-aware routing strategy using the energy-flow model is evident in both these measures more so with the power gain due to the larger differentials in packet delivery ratio with the two protocols.

6. RELATED WORK

To the best of our knowledge, energy flow modeling for self-powered wireless networks has not been considered before. Renewable energy harvesting is covered by a large body of research in wireless sensor networks (WSNs) [12], [5], [13]. The existing WSN energy harvesting techniques and hardware solutions, if used directly, are inappropriate for WiFi networks. Consumption patterns are substantially different: unlike rural area self-powered routers, sensor nodes spend most of their time *sleeping*, when only a small fraction of energy is used, and wake up to receive and send periodically. In addition, sensor networks rarely use more than one environmental energy source - sunlight, as the dimensions of most sensors make wind turbine placing impossible. In our work we concurrently examine both wind and solar generated power as well as the complex interaction between the two. In [8], Kansal et al. propose an energy harvesting oriented method to determine the duty cycling periods of a sensor mote. Their goal is to preserve energy neutrality, i.e. prevent node depletion, and maximize node active time. The authors do not model the energy flow but assume that the energy consumption and generation can be measured. Direct applicability of the above work on our problem is limited since community WiFi networks cannot employ aggressive duty cycling. Wake-on-WLAN [10] uses 802.15.4 sensor motes to register an incoming 802.11 transmission and turn on a self-powered router. In a network with long periods of inactivity this provides substantial energy savings. This approach can be augmented with our energy flow modeling work. With the model in place, before taking any action, network nodes could evaluate the benefits (energy savings) and the drawbacks (unavailability, equipment powering-on delay) of turning off their routers

A number of power-aware routing metrics are examined in [14]. The simulation setup in that work assumes that the nodes are powered by non-rechargeable batteries. Time to failure is then approximated by examining the current battery voltage. Calculating the proposed metrics in a self-powered network, where the batteries are recharged in non-regular intervals, is challenging; we believe that the energy flow model could provide the necessary insight.

7. DISCUSSION AND FUTURE WORK

The goal of our work is to define an energy flow model for self-powered wireless mesh network deployments. Such networks are becoming increasingly important as the communication spreads

to rural areas of the developing world where erratic and sporadic power supply represents a major hurdle for higher ICT penetration. Novel challenges arise in self-powered WiFi networks: energy harvested from renewable sources is available at varying degrees through time, making energy supply difficult to estimate; and the impact of energy-aware network protocols on network performance has not been studied previously.

In this paper we have developed an energy flow model for self-powered wireless network routers. Using the functionality provided by the model and future weather forecast information, we build an energy aware routing protocol called LPR. We evaluated the protocol through simulation based on the measured weather sensor traces obtained from an actual rural wireless network with self powered routers. We have demonstrated that in realistic scenarios, LPR significantly improves delay and provides similar or better packet delivery performance compared to the state of the art approach. In addition, LPR facilitates deployment of heterogeneous networks as it balances routing tasks among nodes while taking into account their individual hardware properties.

LPR protocol is just one example of how reliable energy-flow estimation can be used to improve fundamental network mechanisms. In the future, we see the model's role in novel online energy consumption management and network planning solutions.

8. ACKNOWLEDGEMENTS

This work was funded in part by NSF Career Award CNS-0347886 and The University of Edinburgh Development Trust.

9. REFERENCES

- [1] G. Bernardi, P. Buneman, and M. K. Marina. Tegola Tiered Mesh Network Testbed in Rural Scotland. In *WiNS-DR '08*, San Francisco, CA, Sept. 2008.
- [2] M. Demmer and K. Fall. DTLRSR: Delay Tolerant Routing For Developing Regions. In *NSDR '07*, Kyoto, Japan, August 2007.
- [3] M. Durr, A. Cruden, S. Gair, and J. McDonald. Dynamic model of a lead acid battery for use in a domestic fuel cell system. *Journal of Power Sources*, 161:1400–1411, October 2006.
- [4] S. Guo, M. Falaki, E. Oliver, S. U. Rahman, A. Seth, M. Zaharia, U. Ismail, and S. Keshav. Design and Implementation of the KioskNet System. In *ICTD '07*, Bangalore, India, December 2007.
- [5] J. Jeong, X. Jiang, and D. E. Culler. Design and Analysis of Micro-Solar Power Systems for Wireless Sensor Networks. In *INSS 08*, Kanazawa, Japan, June 2008.
- [6] D. Johnson. Evaluation of a Single Radio Rural Mesh Network in South Africa. In *IEEE/ACM International Conference on Information and Communication Technologies and Development (ICTD2007)*, Bangalore, India, December 2007.
- [7] E. Jones, L. LI, and P. Ward. Practical Routing in Delay-Tolerant Networks. In *WDTN*, Philadelphia, PA, August 2005.
- [8] A. Kansal, J. Hsu, M. Srivastava, and V. Raghunathan. Harvesting Aware Power Management for Sensor Networks. In *DAC '06: Proceedings of the 43rd annual conference on Design automation*, Anaheim, CA, July 2006.
- [9] H. V. Milner. The Digital Divide: The Role of Political Institutions in Technology Diffusion. *Comparative Political Studies*, 39:176 – 199, March 2006.
- [10] N. Mishra, K. Chebrolov, B. Raman, and A. Pathak. Wake-on-WLAN. In *World Wide Web Conference WWW 2006*, Edinburgh, Scotland, May 2006.
- [11] L. L. Peterson and B. S. Davie. *Computer Networks: A Systems Approach*. Morgan Kaufmann Publishers, March 2007.
- [12] V. Raghunathan and P. H. Chou. Design and Power Management of Energy Harvesting Embedded Systems. In *International Symposium on Low Power Electronics and Design, ISLPED'06*, Tegernsee, Germany, October 2006.
- [13] F. Simjee and P. H. Chou. Efficient Charging of Supercapacitors for Extended Lifetime of Wireless Sensor Nodes. *IEEE Transactions on Power Electronics*, 23:1526–1536, May 2008.
- [14] S. Singh, M. Woo, and C. S. Raghavendra. Power-Aware Routing in Mobile Ad Hoc Networks. In *MobiCom '98*, Dallas, TX, October 1998.
- [15] S. Surana, R. Patra, S. Nedeveschi, M. Ramos, L. Subramanian, Y. Ben-David, and E. Brewer. Beyond Pilots: Keeping Rural Wireless Networks Alive. In *ACM USENIX NSDI*, San Francisco, CA, April 2008.
- [16] University of Waterloo. Dtnsim2 - dtn simulator, available at <http://watwire.uwaterloo.ca/dtn/sim/>.
- [17] www.internetworldstats.com. Internet Users and Population Statistics. Technical report, Miniwatts Marketing Group, April 2008.